Cultural Shift or Linguistic Drift? Comparing Two Computational Measures of Semantic Change

William L. Hamilton, Jure Leskovec, Dan Jurafsky

Department of Computer Science, Stanford University, Stanford CA, 94305

`wleif,jure,jurafsky@stanford.edu`

Abstract

Words shift in meaning for many reasons, including cultural factors like new technologies and regular linguistic processes like subjectification. Understanding the evolution of language and culture requires disentangling these underlying causes. Here we show how two different distributional measures can be used to detect two different types of semantic change. The first measure, which has been used in many previous works, analyzes global shifts in a word's distributional semantics; it is sensitive to changes due to regular processes of linguistic drift, such as the semantic generalization of *promise* ("I promise."→"It promised to be exciting."). The second measure, which we develop here, focuses on local changes to a word's nearest semantic neighbors; it is more sensitive to cultural shifts, such as the change in the meaning of *cell* ("prison cell" → "cell phone"). Comparing measurements made by these two methods allows researchers to determine whether changes are more cultural or linguistic in nature, a distinction that is essential for work in the digital humanities and historical linguistics.

1 Introduction

Distributional methods of embedding words in vector spaces according to their co-occurrence statistics are a promising new tool for diachronic semantics (Gulordava and Baroni, 2011; Jatowt and Duh, 2014; Kulkarni et al., 2014; Xu and Kemp, 2015; Hamilton et al., 2016). Previous work, however, does not consider the underlying causes of semantic change or how to distentangle different types of change.

We show how two computational measures can be used to distinguish between semantic changes caused by cultural shifts (e.g., technological advancements) and those caused by more regular processes of semantic change (e.g., grammaticalization or subjectification). This distinction is essential for research on linguistic and cultural evolution. Detecting cultural shifts in language use is crucial to computational studies of history and other digital humanities projects. By contrast, for advancing historical linguistics, cultural shifts amount to noise and only the more regular shifts matter.

Our work builds on two intuitions: that distributional models can highlight syntagmatic versus paradigmatic relations with neighboring words (Schutze and Pedersen, 1993) and that nouns are more likely to undergo changes due to irregular cultural shifts while verbs more readily participate in regular processes of semantic change (Gentner and France, 1988; Traugott and Dasher, 2001). We use this noun vs. verb mapping as a proxy to compare our two measures' sensitivities to cultural vs. linguistic shifts. Sensitivity to nominal shifts indicates a propensity to capture irregular cultural shifts in language, such as those due to technological advancements (Traugott and Dasher, 2001). Sensitivity to shifts in verbs (and other predicates) indicates a propensity to capture regular processes of linguistic drift (Gentner and France, 1988; Kintsch, 2000; Traugott and Dasher, 2001).

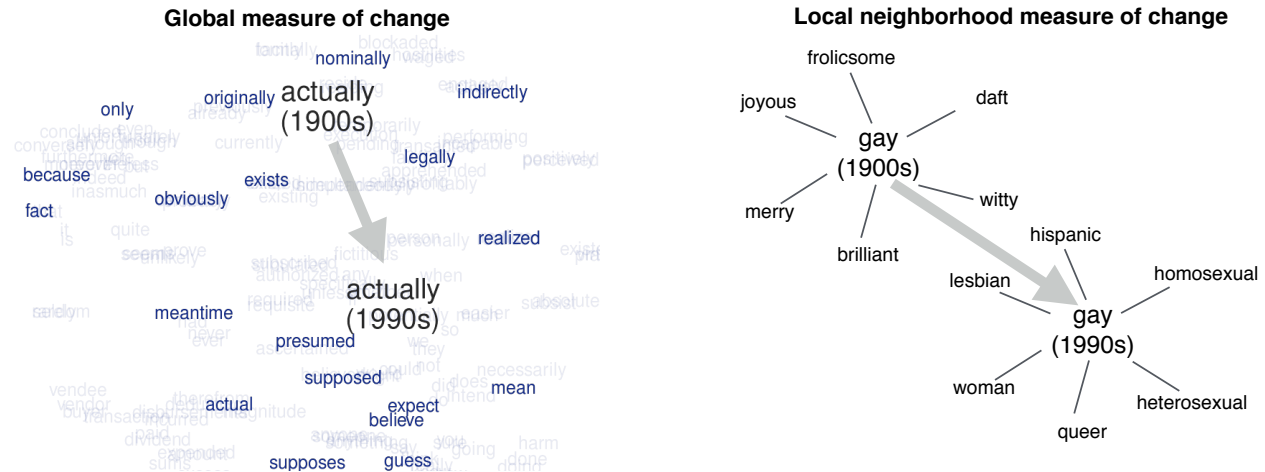The first measure we analyze is based upon changes to a word's local semantic neighborhood;

2116

Global measure of change

Local neighborhood measure of change

Figure 1: Two different measures of semantic change. With the global measure of change, we measure how far a word has moved in semantic space between two time-periods. This measure is sensitive to subtle shifts in usage and also global effects due to the entire semantic space shifting. For example, this captures how *actually* underwent subjectification during the 20th century, shifting from uses in objective statements about the world ("actually did try") to subjective statements of attitude ("I actually agree"; see Traugott and Dasher, 2001 for details). In contrast, with the local neighborhood measure of change, we measure changes in a word's nearest neighbors, which captures drastic shifts in core meaning, such as *gay*'s shift in meaning over the 20th century.

we show that it is more sensitive to changes in the nominal domain and captures changes due to unpredictable cultural shifts. Our second measure relies on a more traditional global notion of change; we show that it better captures changes, like those in verbs, that are the result of regular linguistic drift.

Our analysis relies on a large-scale statistical study of six historical corpora in multiple languages, along with case-studies that illustrate the fine-grained differences between the two measures.

2 Methods

We use the diachronic word2vec embeddings constructed in our previous work (Hamilton et al., 2016) to measure how word meanings change between consecutive decades.[1] In these representations each word w_i has a vector representation $\mathbf{w}^{(t)}$ (Turney and Pantel, 2010) at each time point, which captures its co-occurrence statistics for that time period. The vectors are constructed using the skip-gram with negative sampling (SGNS) algorithm (Mikolov et al., 2013) and post-processed to align the semantic spaces between years. Measuring the distance between word vectors for consecutive decades allows us to compute the rate at which the different words

---

[1] http://nlp.stanford.edu/projects/histwords/. This URL also links to detailed dataset descriptions and the code needed to replicate the experiments in this paper.

change in meaning (Gulordava and Baroni, 2011).

We analyzed the decades from 1800 to 1990 using vectors derived from the Google N-gram datasets (Lin et al., 2012) that have large amounts of historical text (English, French, German, and English Fiction). We also used vectors derived from the Corpus of Historical American English (COHA), which is smaller than Google N-grams but was carefully constructed to be genre balanced and contains word lemmas as well as surface forms (Davies, 2010). We examined all decades from 1850 through 2000 using the COHA dataset and used the part-of-speech tags provided with the corpora.

2.1 Measuring semantic change

We examine two different ways to measure semantic change (Figure 1).

Global measure

The first measure analyzes global shifts in a word's vector semantics and is identical to the measure used in most previous works (Gulordava and Baroni, 2011; Jatowt and Duh, 2014; Kim et al., 2014; Hamilton et al., 2016). We simply take a word's vectors for two consecutive decades and measure the cosine distance between them, i.e.

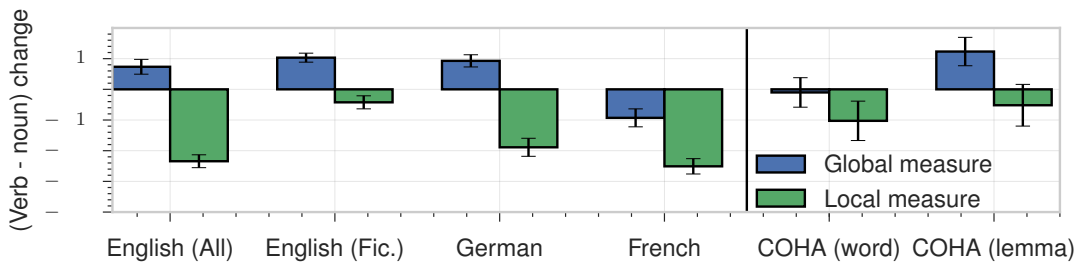$$d^G(w_i^{(t)}, w_i^{(t+1)}) = \text{cos-dist}(\mathbf{w}_i^{(t)}, \mathbf{w}_i^{(t+1)}). \quad (1)$$

Figure 2: The global measure is more sensitive to semantic changes in verbs while the local neighborhood measure is more sensitive to noun changes. Examining how much nouns change relative to verbs (using coefficients from mixed-model regressions) reveals that the two measures are sensitive to different types of semantic change. Across all languages, the local neighborhood measure always assigns relatively higher rates of change to nouns (i.e., the right/green bars are lower than the left/blue bars for all pairs), though the results vary by language (e.g., French has high noun change-rates overall). 95% confidence intervals are shown.

Local neighborhood measure

The second measure is based on the intuition that only a word's nearest semantic neighbors are relevant. For this measure, we first find word w_i's set of k nearest-neighbors (according to cosine-similarity) within each decade, which we denote by the ordered set $\mathcal{N}_k(w_i^{(t)})$. Next, to measure the change between decades t and $t+1$, we compute a "second-order" similarity vector for $w_i^{(t)}$ from these neighbor sets with entries defined as

$$\mathbf{s}^{(t)}(j) = \text{cos-sim}(\mathbf{w}_i^{(t)}, \mathbf{w}_j^{(t)})$$
$$\forall w_j \in \mathcal{N}_k(w_i^{(t)}) \cup \mathcal{N}_k(w_i^{(t+1)}), \quad (2)$$

and we compute an analogous vector for $w_i^{(t+1)}$. The second-order vector, $\mathbf{s}_i^{(t)}$, contains the cosine similarity of $\mathbf{w}_i$ and the vectors of all w_i's nearest semantic neighbors in the the time-periods t and $t+1$. Working with variants of these second-order vectors has been a popular approach in many recent works, though most of these works define these vectors against the full vocabulary and not just a word's nearest neighbors (del Prado Martin and Brendel, 2016; Eger and Mehler, 2016; Rodda et al., 2016).

Finally, we compute the local neighborhood change as

$$d^L(w_i^{(t)}, w_i^{(t+1)}) = \text{cos-dist}(\mathbf{s}_i^{(t)}, \mathbf{s}_i^{(t+1)}). \quad (3)$$

This measures the extent to which w_i's similarity with its nearest neighbors has changed.

The local neighborhood measure defined in (3) captures strong shifts in a word's paradigmatic relations but is less sensitive to global shifts in syntagmatic contexts (Schutze and Pedersen, 1993). We

Dataset	# Nouns	# Verbs
Google English All	5299	2722
Google English Fic.	4941	3128
German	5443	1844
French	2310	4992
COHA (Word)	4077	1267
COHA (Lemma)	3389	783

Table 1: Number of nouns and verbs tested in each dataset.

used $k = 25$ in all experiments (though we found the results to be consistent for $k \in [10, 50]$).

2.2 Statistical methodology

To test whether nouns or verbs change more according to our two measures of change, we build on our previous work and used a linear mixed model approach (Hamilton et al., 2016). This approach amounts to a linear regression where the model also includes "random" effects to account for the fact that the measurements for individual words will be correlated across time (McCulloch and Neuhaus, 2001).

We ran two regressions per datatset: one with the global d^G values as the dependent variables (DVs) and one with the local neighborhood d^L values. In both cases we examined the change between all consecutive decades and normalized the DVs to zero-mean and unit variance. We examined nouns/verbs within the top-10000 words by frequency rank and removed all words that occurred <500 times in the smaller COHA dataset. The independent variables are word frequency, the decade of the change (represented categorically), and variable indicating

Word	1850s context	1990s context
actually	"...dinners which you have actually eaten."	"With that, I actually agree."
must	"O, George, we must have faith."	"Which you must have heard ten years ago..."
promise	"I promise to pay you...'	"...the day promised to be lovely."
gay	"Gay bridals and other merry-makings of men."	"...the result of gay rights demonstrations."
virus	"This young man is...infected with the virus."	"...a rapidly spreading computer virus."
cell	"The door of a gloomy cell..."	"They really need their cell phones."

Table 2: Example case-studies of semantic change. The first three words are examples of regular linguistic shifts, while the latter three are examples of words that shifted due to exogenous cultural factors. Contexts are from the COHA data (Davies, 2010).
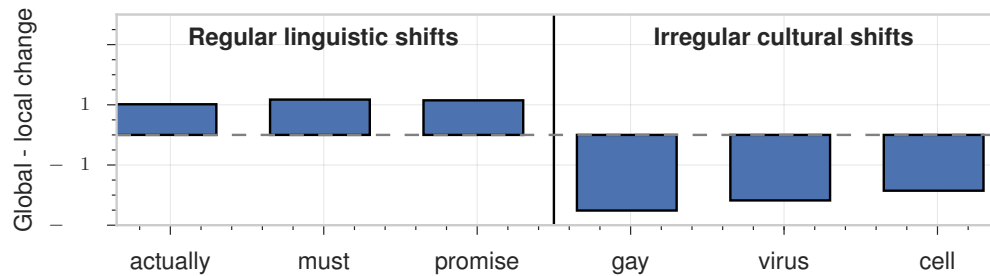

Figure 3: The global measure captures classic examples of linguistic drift while the local measure captures example cultural shifts. Examining the semantic distance between the 1850s and 1990s shows that the global measure is more sensitive to regular shifts (and vice-versa for the local measure). The plot shows the difference between the measurements made by the two methods.

whether a word is a noun or a verb (proper nouns are excluded, as in Hamilton et al., 2016).[2]

3 Results

Our results show that the two seemingly related measures actually result in drastically different notions of semantic change.

3.1 Nouns vs. verbs

The local neighborhood measure assigns far higher rates of semantic change to nouns across all languages and datasets while the opposite is true for the global distance measure, which tends to assign higher rates of change to verbs (Figure 2).

We focused on verbs vs. nouns since they are the two major parts-of-speech and previous research has shown that verbs are more semantically mutable than nouns and thus more likely to undergo linguistic drift (Gentner and France, 1988), while nouns are far more likely to change due to cultural shifts like new technologies (Traugott and Dasher, 2001). However, some well-known regular linguistic shifts include rarer parts of speech like adverbs (included in our case studies below). Thus we also confirmed

---

[2]Frequency was included since it is known to strongly influence the distributional measures (Hamilton et al., 2016).

that the differences shown in Figure 2 also hold when adverbs and adjectives are included along with the verbs. This modified analysis showed analogous significant trends, which fits with previous research arguing that adverbial and adjectival modifiers are also often the target of regular linguistic changes (Traugott and Dasher, 2001).

The results of this large-scale regression analysis show that the local measure is more sensitive to changes in the nominal domain, a domain in which change is known to be driven by cultural factors. In contrast, the global measure is more sensitive to changes in verbs, along with adjectives and adverbs, which are known to be the targets of many regular processes of linguistic change (Traugott and Dasher, 2001; Hopper and Traugott, 2003)

3.2 Case studies

We examined six case-study words grouped into two sets. These case studies show that three examples of well-attested regular linguistic shifts (set A) changed more according to the global measure, while three well-known examples of cultural changes (set B) change more according to the local neighborhood measure. Table 2 lists these words with some representative historical contexts (Davies, 2010).

Set A contains three words that underwent attested regular linguistic shifts detailed in Traugott and Dasher (2001): *actually*, *must*, and *promise*. These three words represent three different types of regular linguistic shifts: *actually* is a case of subjectification (detailed in Figure 1); *must* shifted from a deontic/obligation usage ("you must do X") to a epistemic one ("X must be the case"), exemplifying a regular pattern of change common to many modal verbs; and *promise* represents the class of shifting "performative speech acts" that undergo rich changes due to their pragmatic uses and subjectification (Traugott and Dasher, 2001). The contexts listed in Table 2 exemplify these shifts.

Set B contains three words that were selected because they underwent well-known cultural shifts over the last 150 years: *gay*, *virus*, and *cell*. These words gained new meanings due to uses in community-specific vernacular (*gay*) or technological advances (*virus*, *cell*). The cultural shifts underlying these changes in usage — e.g., the development of the mobile "cell phone" — were unpredictable in the sense that they were not the result of regularities in human linguistic systems.

Figure 3 shows how much the meaning of these word changed from the 1850s to the 1990s according to the two different measures on the English Google data. We see that the words in set A changed more when measurements were made using the global measure, while the opposite holds for set B.

4 Discussion

Our results show that our novel local neighborhood measure of semantic change is more sensitive to changes in nouns, while the global measure is more sensitive to changes in verbs. This mapping aligns with the traditional distinction between irregular cultural shifts in nominals and more regular cases of linguistic drift (Traugott and Dasher, 2001) and is further reinforced by our six case studies.

This finding emphasizes that researchers must develop and use measures of semantic change that are tuned to specific tasks. For example, a cultural change-point detection framework would be more successful using our local neighborhood measure, while an empirical study of grammaticalization would be better off using the traditional global distance approach. Comparing measurements made by these two approaches also allows researchers to assess the extent to which semantic changes are linguistic or cultural in nature.

References

Mark Davies. 2010. The Corpus of Historical American English: 400 million words, 1810-2009. http://corpus.byu.edu/coha/.

Fermin Moscoso del Prado Martin and Christian Brendel. 2016. Case and Cause in Icelandic: Reconstructing Causal Networks of Cascaded Language Changes. In *Proc. ACL*.

Steffen Eger and Alexander Mehler. 2016. On the Linearity of Semantic Change: Investigating Meaning Variation via Dynamic Graph Models. In *Proc. ACL*.

Dedre Gentner and Ilene M. France. 1988. The verb mutability effect: Studies of the combinatorial semantics of nouns and verbs. *Lexical ambiguity resolution: Perspectives from psycholinguistics, neuropsychology, and artificial intelligence*, pages 343–382.

Kristina Gulordava and Marco Baroni. 2011. A distributional similarity approach to the detection of semantic change in the Google Books Ngram corpus. In *Proc. GEMS 2011 Workshop on Geometrical Models of Natural Language Semantics*, pages 67–71. Association for Computational Linguistics.

William L. Hamilton, Jure Leskovec, and Dan Jurafsky. 2016. Diachronic Word Embeddings Reveal Statistical Laws of Semantic Change. In *Proc. ACL*.

Paul J Hopper and Elizabeth Closs Traugott. 2003. *Grammaticalization*. Cambridge University Press, Cambridge, UK.

Adam Jatowt and Kevin Duh. 2014. A framework for analyzing semantic change of words across time. In *Proc. 14th ACM/IEEE-CS Conf. on Digital Libraries*, pages 229–238. IEEE Press.

Yoon Kim, Yi-I. Chiu, Kentaro Hanaki, Darshan Hegde, and Slav Petrov. 2014. Temporal analysis of language through neural language models. *arXiv preprint arXiv:1405.3515*.

Walter Kintsch. 2000. Metaphor comprehension: A computational theory. *Psychon. Bull. Rev.*, 7(2):257–266.

Vivek Kulkarni, Rami Al-Rfou, Bryan Perozzi, and Steven Skiena. 2014. Statistically significant detection of linguistic change. In *Proc. 24th WWW Conf.*, pages 625–635.

Yuri Lin, Jean-Baptiste Michel, Erez Lieberman Aiden, Jon Orwant, Will Brockman, and Slav Petrov. 2012. Syntactic annotations for the google books ngram corpus. In *Proc. ACL System Demonstrations*.

Charles E McCulloch and John M Neuhaus. 2001. *Generalized linear mixed models*. Wiley-Interscience, Hoboken, NJ.

Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S. Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *NIPS*.

Martina Rodda, Marco Senaldi, and Alessandro Lenci. 2016. Panta rei: Tracking Semantic Change with Distributional Semantics in Ancient Greek. In *Italian Conference of Computational Linguistics*.

Hinrich Schutze and Jan Pedersen. 1993. A vector model for syntagmatic and paradigmatic relatedness. In *Proc. 9th Annu. Conf. of the UW Centre for the New OED and Text Research*, pages 104–113. Citeseer.

Elizabeth Closs Traugott and Richard B Dasher. 2001. *Regularity in Semantic Change*. Cambridge University Press, Cambridge, UK.

Peter D. Turney and Patrick Pantel. 2010. From frequency to meaning: Vector space models of semantics. *J. Artif. Intell. Res.*, 37(1):141–188.

Yang Xu and Charles Kemp. 2015. A computational evaluation of two laws of semantic change. In *Proc. 37th Annu. Conf. Cogn. Sci. Soc.*