Title

Characterizing regularity in semantic shift of individuals

Permalink

https://escholarship.org/uc/item/00q8v7qt

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 48(0)

Authors

Henderson, Kieran
Xu, Yang

Publication Date

2026

Characterizing regularity in semantic shift of individuals

Kieran Henderson (hende325@cs.toronto.edu)
Department of Computer Science
University of Toronto

Yang Xu (yangxu@cs.toronto.edu)
Department of Computer Science, Cognitive Science Program
University of Toronto

Abstract

Semantic shift, or the diachronic variation in word meaning, is a topic commonly discussed in historical and cognitive linguistics. Previous work has typically focused on characterizing semantic change in a population, but how word meanings shift over time in individuals is underexplored. We propose a computational framework for analyzing diachronic semantic shift at the individual level by adapting techniques for population-level semantic change. Using speaker-labelled text corpora, our framework utilizes contextualized word embeddings to identify predictable patterns in semantic shift among individual speakers. We discover that frequency predicts the rate of semantic shift over time across speakers, while polysemy does not. Additionally, nouns exhibit higher semantic stability over adjectives and verbs. These patterns partially mirror existing findings on regularity in historical semantic change, suggesting shared tendencies across individual and population levels. Our work bridges computational approaches to historical semantics with the diachronic modeling of personal lexicons.

Keywords: personal lexicon; semantic shift; semantic change; contextualized word embedding

Introduction

The lexicon is a symbolic system that maps meanings to words and it plays a fundamental role in cognition and communication. Although the lexicon is typically viewed as a shared system within a population, every individual possesses their own version of the lexicon, or *personal lexicon* (Clark, 1998). A personal lexicon is pertinent to an individual's lifelong development: it emerges in childhood and evolves constantly through adolescence and adulthood. How do personal lexicons evolve in different individuals: Are there regularities or predictable patterns underlying these individual trajectories? In the current study, we explore this question computationally with an initial investigation on the diachronic shift in word meaning among individual speakers.

The diachronics of word meaning has been a topic of significant interest in historical and cognitive linguistics. A long-standing focus of this research is characterizing the regularity in the semantic shift of words (also known as semantic change) over historical periods (e.g., Bréal, 1904; Geeraerts, 1997; Traugott & Dasher, 2001). This traditional focus has been propelled by recent work that uses historical text corpora and computational methods to uncover law-like principles in semantic change. For instance, it has been suggested that word frequency is a factor that predicts the historical stability of word meaning: more commonly used words tend to resist meaning change over time (e.g., Pagel, Atkinson, & Meade, 2007; Hamilton, Leskovec, & Jurafsky, 2016). Although such principles have been established in a historical setting, the level of analysis often applies to a population where word usages and meanings are considered to be shared or in aggregate (i.e., from all speakers). Here we focus on a more granular level of analysis by asking whether there is regularity in semantic shift over time across different speakers, therefore uncovering general principles in the temporal dynamics of personal lexicons across individuals.

Our study is also related to research on the lexicon in psychology and cognitive science. For example, extensive work in social psychology has studied individual variation in language use and linguistic style (e.g., Pennebaker & King, 1999; Slatcher, Chung, Pennebaker, & Stone, 2007), while more recent work in cognitive science has demonstrated how concepts may be perceived differentially across individuals (Marti, Wu, Piantadosi, & Kidd, 2023). In the area of semantic development, existing research has analyzed structural changes in the mental lexicon through lifespan with an emphasis on understanding how semantic networks constructed from word association differ in young and elderly populations (Cain & Ryskin, 2025; Dubossarsky, De Deyne, & Hills, 2017; Wulff, Hills, & Mata, 2022). Our current work differs from this line of research in two respects. First, we work with longitudinal and speaker-labelled text corpora based on natural language use of individual speakers rather than data drawn from lab-based experiments. Second, we present a computational framework that models word semantics of personal lexicons over time and supports the analysis of shared patterns in the semantic trajectories across individual speakers. To our knowledge, this high-resolution approach to individual diachronic semantics remains underexplored in the study of semantic shift, which in its closet form has explored variation in sub-populations (e.g., Kamath, Yang, Reddy, Sonderegger, & Card, 2025) or across communities (e.g., Noble, Sayeed, Fernández, & Larsson, 2021).

Figure 1 illustrates our framework by visualizing the semantic trajectory of a single English word *deal* based on its usage by an individual speaker in recent decades. The semantic clusters in the left panel signify the different meanings or senses (detected automatically) the word has taken on, and the plot in the right panel shows the historical shift in the probability distribution of these word senses over the years. In this case, the word has become more predominantly used

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

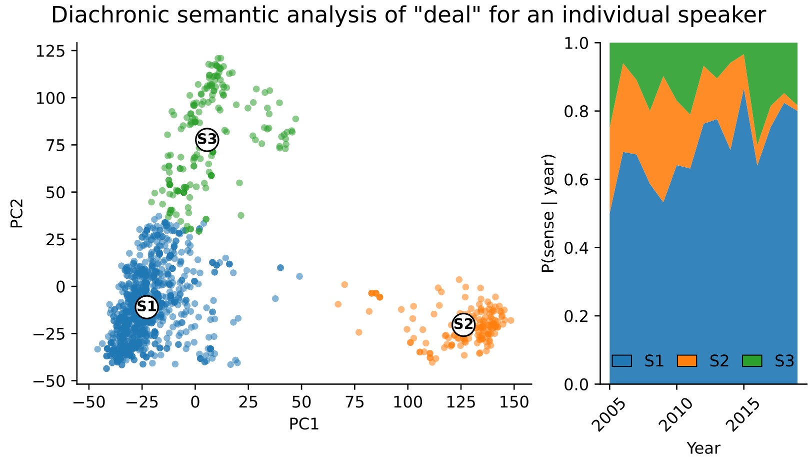

Figure 1: An illustration of diachronic semantics of an individual tracked over the recent decades. **Left:** Principal-component (PC) projection of contextual embeddings for the word *deal* based on its usages from a speaker available in the dataset, colored by inferred semantic or sense clusters (S1-S3); labelled points correspond to centroids. **Right:** Historical distribution of word senses over time for that individual, colors matched to semantic clusters. Example usages of the semantic clusters: S1 - agreement/negotiation “a peace [deal] could be reached”; S2 - address a situation “try to [deal] with this”; S3 - transaction “a good [deal] on your new house”.

for the sense “agreement/negotiation”, while the sense “address a situation” has diminished over time. Moving beyond this single example, individual speakers can vary in both the vocabulary they use (see review by Enfield, 2023) and the semantic trajectories of the (same) words over time, but our goal is to characterize the general principles that might give rise to predictable patterns in the diachronic shift of word meanings in the personal lexicons despite this individual variation.

Hypothesis formulation

Although individuals may vary substantially in their semantic trajectories, we hypothesize that such processes of word meaning shift are unlikely to be random. In particular, we evaluate the possibility that the known factors shown to predict semantic change in a historical setting might also predict semantic shift in the lexicons of individuals. Our proposal is motivated by the notion that individual-level language change might mirror population-level language evolution (e.g., Chater & Christiansen, 2010). This idea received empirical support, including how lexical evolution is predicted by age of word acquisition in child development (Monaghan, 2014), and how word meanings evolve drawing on similar principles in language change and language development (Brochhagen, Boleda, Gualdoni, & Xu, 2023).

We build on existing findings in historical semantic change to test the following hypotheses about regularity in the semantic shift of individual speakers. First, we postulate that word frequency should predict rate of semantic shift across individuals. Based on prior work (Hamilton et al., 2016; Pagel et al., 2007), we expect frequently used words should resist semantic shift more so than rarer words. Existing research has also suggested that more polysemous words (i.e., words containing more senses) should experience a higher rate of semantic shift compared to words with fewer senses (Hamilton et al., 2016). However, follow-up work has demonstrated that the effect of polysemy is minimal and may be attributable to statistical artifacts in the methodology that measures semantic change (Dubossarsky, Weinshall, & Grossman, 2017). Here we evaluate the roles of frequency and polysemy in predicting rate of semantic shift in individuals. Additionally, we expect different word classes to exhibit differential rates of semantic shift: verbs should show a higher rate of semantic shift than nouns, motivated by previous findings in historical semantics (Dubossarsky, Weinshall, & Grossman, 2016). We acknowledge that our hypotheses are meant to reflect statistical tendencies, but not absolute rules. Importantly, to guard against spurious findings, we use a rigorous evaluation scheme with a controlled analysis adopted by existing work on historical semantic change (Dubossarsky, Weinshall, & Grossman, 2017).

Materials and methods

Dataset and preparation of individual text corpora

For our analysis, we work with the INTERVIEW: NPR Media Dialog Transcripts dataset (Majumder, Li, Ni, & McAuley, 2020a, 2020b). This dataset includes interview transcripts from seven programs on National Public Radio (NPR) spanning two decades from 1999 to 2019. It contains 184K speakers, of which 287 are interviewers. Interviewers are well-suited for the purpose of this study, because they exhibit consistent speaking roles and produce relatively large volumes of naturalistic speech over a stretched period of time where semantic shift can take place.

To ensure sufficiently dense data for the individual-based diachronic analysis, we work with a gender-balanced subset of 12 interviewers (6 men and 6 women). Each of the selected speakers contributed at least 17,000 utterances, for a total of 430K utterances across all speakers. This averages to 36K utterances per year for the period of interest. The list of speakers and information about their gender and individualized text corpora are provided in Table 1. We also perform an analysis of a single person with the most utterances, Neal C., among the other 12 speakers, because they have considerably more utterances than all other speakers.

Table 1: Information of speakers and their individual corpora.

Speaker	Gender	Period	Corpus size	Words
Neal C	Male	2004-13	256524	1932
Ira F	Male	2005-13	73393	586
Steven I	Male	2004-19	59487	631
Robert S	Male	2005-18	49663	385
David G	Male	2009-19	28876	90
Guy R	Male	2007-13	27524	196
Alex C	Male	2005-08	21397	334
Renee M	Female	2004-19	37108	89
Farai C	Female	2005-09	36900	141
Rachel M	Female	2010-19	30582	189
Madeleine B	Female	2005-10	27003	143
Michele N	Female	2005-13	22366	334
Liane H	Female	2005-11	17323	182

Obtaining dense longitudinal data at the individual level remains a significant challenge for research on diachronic personal semantics. We also considered alternative datasets, including US congressional speeches (Gentzkow & Taddy, 2018) and the SAYCam corpus (Sullivan, Mei, Perfors, Wojcik, & Frank, 2021). Congressional speech data was not feasible for our analysis due to uncertainty associated with speech identification, since speeches may be written or edited by staff rather than the speaker. We also considered SAYCam, a longitudinal audiovisual dataset centered on early language acquisition. Unfortunately, this dataset has very limited number of speakers (3) and shorter temporal spans making it unsuitable for studying semantic shift in the long term.

For each individual corpus, we perform consistent processing by removing rows with empty fields, lower-casing all text, and performing Part Of Speech (POS) tagging on each utterance. Afterwards, we obtain occurrences of word tokens for each year to determine which words have sufficient data for meaningful analysis.

We filter word types according to the following criteria: (1) At least 5 appearances per year; (2) Total occurrence count of at least 100; (3) POS falls under the 3 major word classes: noun, verb, or adjective; (4) Absent for no more than 2 years in the period of interest.

These processing steps ensure each target word has enough observations per year while also maximizing the individuals' vocabulary sizes for our analysis. Finally, we prepare an individual corpora for each speaker, containing all utterances with target words, associated POS tags, and timestamps. These corpora are then used as input for the computational analysis.

Computational framework

Contextualized word embedding. Representing word meanings accurately and longitudinally is a prerequisite for our computational analysis of diachronic semantics of individuals. There are several established approaches to modeling word meanings over time. The earlier approach using word vectors or word embeddings per time point (e.g., Kulkarni, Al-Rfou, Perozzi, & Skiena, 2015; Hamilton et al., 2016) is sensible for very large text corpora, but it is not suitable for individual corpora due to diminished corpus size compared to aggregate or population-level corpora. Here we use contextualized word embeddings based on the BERT language model (Devlin, Chang, Lee, & Toutanova, 2019), shown to be accurate in representing word senses (Reif et al., 2019) and in capturing phenomena relevant to historical semantic change (Giulianelli, Del Tredici, & Fernández, 2020).

To generate each embedding, we first identify all sentences containing utterances of the target word in a speaker's individual corpus. Each utterance is then tokenized into subword units using the pretrained BERT-base-uncased tokenizer. As a result of tokenization, some target words may be split into multiple subword tokens, so we identified all tokens corresponding to the target word. For each target word, we followed the procedures in existing work (Giulianelli et al., 2020) by extracting the model hidden states from all layers, summed across layers, and averaged over the subword tokens associated with each word. These procedures yield a single fixed length vector representing the contextualized embedding of the word within its utterance. For each target word, we collect all embeddings for each utterance into a matrix, which can be considered the usage matrix for that word.

Quantification of semantic shift. To quantify semantic shift of a word for an individual speaker, we calculate average pairwise distance (APD) between contextualized usage embeddings across two time points, a measure shown to account for human judgment of semantic change with better precision compared to alternative measures (Giulianelli et al., 2020). APD measures the within-speaker temporal variation of a word independent of other speakers as follows:

$$\Delta^{(t)}(w_i) = \mathrm{APD}\left(U_w^t, U_w^{t^+}\right) = \frac{1}{N^t \cdot N^{t^+}} \sum_{x_i \in U_w^t,\, x_j \in U_w^{t^+}} d(x_i\,,\, x_j), \tag{1}$$

Here t^+ represents the immediate next future time, $d(\cdot,\cdot)$ is cosine distance, U_w^t is the usage matrix for word w at time t, and N^t is the number of occurrences of w in year t.

Predictors of semantic shift. For the individual speakers, we examine how semantic shift into the next future time point may be predicted by frequency and polysemy of words in the current time point. This analysis is similar to existing work that studies laws of historical semantic change (Hamilton et al., 2016). Since our analysis is longitudinal and includes individuals of different gender, we also incorporate year and gender as fixed effects in a mixed-effects analysis.

Frequency We define frequency at a time point as the relative proportion of usage for a word with respect to the total usage of all words from a speaker:

$$f^{(t)}(w_i) = \frac{C_{w_i}^{(t)}}{T^{(t)}} \tag{2}$$

Here $C_{w_i}^{(t)}$ is the number of occurrences of word w_i in year t, and $T^{(t)}$ is the total number of words which occur in year t. We use the natural logarithm of frequency as a predictor to account for the skewed distribution of counts.

Polysemy via semantic clustering We compute polysemy using normalized entropy of a words usage distribution across semantic clusters:

$$d^{(t)}(w_i) = \log K_w \cdot \exp\left(-\sum_{k=1}^{K_w} u_w^t[k] \log u_w^t[k]\right) \tag{3}$$

Here $d^{(t)}(w_i)$ is the proportion of occurrences of word w in cluster k at time t, K_w is the fixed number of clusters, and k indexes clusters. Values of $d^{(t)}(w_i)$ are non-negative and increase with semantic diversity, meaning that words whose usage is concentrated in a single cluster have lower values, whereas words with usage spread more evenly across clusters have higher values. We use silhouette scores to determine optimal cluster number and apply K-means clustering on the contextualized embeddings to obtain semantic clusters or senses of a word identical to the procedures used in existing work (Giulianelli et al., 2020) and apply a log-transform on the polysemy measure to stabilize the distribution, consistent with our treatment of frequency.

Mixed-effects analysis. We fit a mixed-effects regression model to predict the rate of semantic shift given the predictors described above, with random intercepts for both word and speaker and fixed effects for year (β_t), log-frequency (β_f), log-polysemy (β_d), and gender (β_g) such that the coefficients of the fixed effects indicate the population level effect of a predictor variable on the result:

$$\tilde{\Delta}^{(t)}(w_i) = \tag{4}$$
$$\beta_f \log f^{(t)}(w_i) + \beta_d \log d^{(t)}(w_i) + \beta_t + \beta_g + Z_{w_i} + S_{w_i} + \varepsilon_{w_i}^{(t)}$$

where

$$\begin{aligned} Z_{w_i} &\sim \mathcal{N}(0, \sigma_w) && \text{random intercept for word } w_i, \\ S_{w_i} &\sim \mathcal{N}(0, \sigma_s) && \text{random intercept for speaker,} \\ \varepsilon_{w_i}^{(t)} &\sim \mathcal{N}(0, \sigma) && \text{residual error term.} \end{aligned}$$

This formulation allows us to test regularity in individual-level diachronic semantics while controlling for frequency, polysemy, time, and speaker-level variability.

Results

We present our results in three steps. First, we evaluate the roles of frequency and polysemy in predicting rate of semantic shift among the individuals through the mixed-effects analysis. Second, we offer a controlled version of this analysis by repeating it in time-shuffled corpora of the same individuals. Third, we examine how the rate of semantic shift varies in different word classes.

Predictive factors of semantic shift in individuals

To evaluate our first hypothesis, we examine the role of word frequency along with polysemy in the mixed-effects model described. Figure 2 shows that there is a clear variation in word frequencies among the individual speakers, although our frequency estimates tend to correlate significantly with word frequencies in Brown corpus ($p < 0.05$) which are independently estimated from population-level statistics.

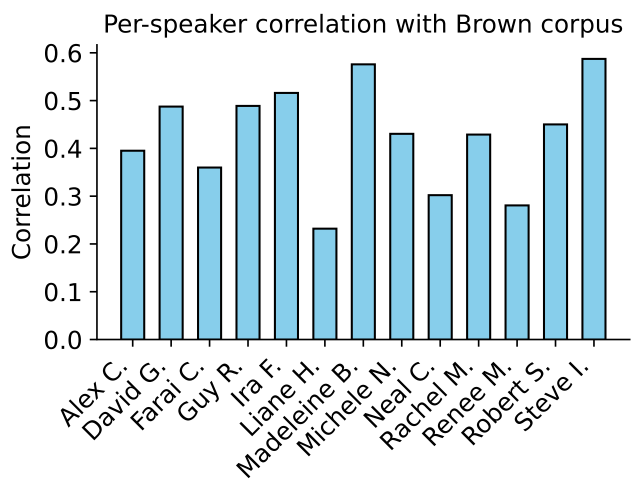

Figure 2: Word frequency variation among individual speakers and correlation with population-based frequency.

Table 2 summarizes the results of the mixed-effects regression analysis that takes into account all the 12 individual speakers. We found that the effect of word frequency is statistically significant ($p < 0.05$) with a negative coefficient, which indicates that more frequent words are semantically more stable than rarer words across the individual speakers. This finding is consistent with the existing view in historical semantics that frequent words tend to resist meaning change over time (Pagel et al., 2007; Hamilton et al., 2016). We also found that polysemy shows a positive effect on semantic shift ($p < 0.05$), with more polysemous words exhibiting greater semantic shift across the individual speakers. Although this result also mirrors an earlier finding that polysemy tends to positively correlate with rate of historical meaning change (Hamilton et al., 2016), such an effect has been later disputed and found to be minimal (Dubossarsky, Weinshall, & Grossman, 2017), an issue that we will revisit in the next section. Moreover, we found significant coefficients for the majority of the years in the period, indicating consistent meaning shift patterns over the years. Additionally, we found gender to have no significant effect ($p = 0.125$) suggesting that these findings are generally applicable to the speakers regardless of gender.

Table 2: Mixed-effects regression results predicting individual-level semantic shift for the genuine condition and the temporally shuffled control condition. Statistically significant results ($p < 0.05$) are indicated with *.

Predictor	Genuine		Shuffled	
	Coeff.	Std.Err.	Coeff.	Std.Err.
Intercept	0.241*	0.004	0.235*	0.002
Year (2006)	-0.003*	0.001	-0.001	0.000
Year (2007)	-0.004*	0.001	-0.001	0.000
Year (2008)	-0.003*	0.001	0.000	0.000
Year (2009)	-0.004*	0.001	0.000	0.000
Year (2010)	-0.005*	0.001	0.000	0.000
Year (2011)	-0.003*	0.001	0.000	0.001
Year (2012)	-0.003*	0.001	-0.001	0.001
Year (2013)	-0.002*	0.001	0.000	0.001
Year (2014)	-0.001	0.001	0.000	0.001
Year (2015)	-0.004*	0.001	0.000	0.001
Year (2016)	-0.005*	0.001	0.001	0.001
Year (2017)	-0.005*	0.001	0.000	0.001
Year (2018)	-0.005*	0.001	-0.001	0.001
Gender	0.005	0.003	0.006	0.003
Log frequency	-0.005*	0.000	0.000	0.000
Log polysemy	0.013*	0.001	0.012*	0.001
Speaker variance	0.003	0.010	0.003	0.016

Table 3: Mixed-effects regression results predicting individual-level semantic shift for the genuine condition and the temporally shuffled control condition for Neal C. Statistically significant results ($p < 0.05$) are indicated with *.

Predictor	Genuine		Shuffled	
	Coeff.	Std.Err.	Coeff.	Std.Err.
Intercept	0.252*	0.007	0.246*	0.006
Year (2005)	-0.015*	0.004	-0.011*	0.002
Year (2006)	-0.016*	0.004	-0.011*	0.002
Year (2007)	-0.016*	0.004	-0.011*	0.002
Year (2008)	-0.015*	0.004	-0.011*	0.002
Year (2009)	-0.018*	0.004	-0.011*	0.002
Year (2010)	-0.017*	0.004	-0.011*	0.002
Year (2011)	-0.018*	0.004	-0.011*	0.002
Year (2012)	-0.016*	0.004	-0.011*	0.002
Log frequency	-0.005*	0.001	0.000	0.000
Log polysemy	0.027*	0.004	0.008*	0.002
Speaker variance	0.004	0.047	0.004	0.126

A control analysis using time-shuffled corpora

To assess whether the observations we made in the mixed-effect analysis reflect genuine temporal trends rather than spurious variation, we also performed a rigorous control analysis where the time stamps of the individual corpora were randomly shuffled, while all other factors were kept equal including the corpus size per year. This time-shuffled analysis is motivated by existing research on historical semantic analysis suggesting that measures of semantic change (such as cosine distance) may introduce statistical artifacts when examined under predictive factors relevant to our analysis (Dubossarsky, Weinshall, & Grossman, 2017).

Table 2 summarizes the results based on the time-shuffled control analysis. In comparison to the main analysis in the genuine condition, none of the year coefficients are statistically significant in the shuffled condition, indicating that the diachronic trends in word meaning depend on correct temporal ordering of the data. Interestingly, we observed that the effect of frequency is greatly diminished in the shuffled condition (about 7-fold decrease) although it remains significant ($p = 0.042$). This result shows that the frequency effect observed in the genuine condition is not attributable to statistical artifacts. Conversely, we observed that polysemy remains a strong predictor of rate of semantic shift with minimal change in the coefficient ($p < 0.05$). This result indicates that the effect of polysemy is comparable in the genuine and shuffled conditions, and therefore may actually be minimal. Table 4 summarizes the partial R^2 that further describe the variances explained by these variables. Complementing these analyses, we also performed similar analyses with Neal C based on their individual corpora. Table 3 shows that these results resemble those with the 12 individuals, suggesting that the effects hold robustly at the level of an individual.

Table 4: Partial R^2 of fixed effects predicting rate of semantic shift in genuine and shuffled data. Partial R^2 values represent the approximate proportion of total variance (fixed + random + residual) attributable to each predictor.

Condition	**log_freq**	**log_polysemy**
Genuine	0.0089	0.0460
Shuffled	0.0000	0.0387

Overall, the negative effect of frequency is consistent with findings in historical semantics showing that frequent exposure reinforces stable semantic representations (Pagel et al., 2007; Hamilton et al., 2016). Our finding parallels principles of semantic change identified at the population level, which suggests that semantic shift in individuals may operate under similar principles to semantic change at a population level. However, the (stronger) effect of polysemy may be inflated (more prominently than the effect of frequency). One potential cause of this inflation is that polysemy is directly associated with high APD by definition, since words which are more polysemous should have more separable clusters identified in semantic space. This means that such words should yield high values in cosine-based semantic distance, and as a result APD should remain high regardless of whether the data is shuffled across time or not.

Rates of semantic shift across word classes

Our results thus far show that the diachronic semantic shift at the individual level exhibits regular patterns beyond random. In particular, word frequency appears to be a shared factor that predicts semantic stability in both individual and population levels.

In this final analysis, we examine how rate of semantic shift might vary across the main word classes. In particular, we test the hypothesis that certain word classes might show greater semantic instability compared to other classes.

Figure 3 summarizes the results based on rates of semantic shift estimated for the three main word classes from the 12 individuals in our dataset. We observed that verbs exhibit significantly greater rate of semantic shift than nouns ($p < 0.05$ from paired t-test) across these individual speakers; the verb class also shows higher rate compared to the adjective class, but this difference is not significant. This finding is consistent with previous work and suggests that verbs show more fluidity in historical semantic change in comparison to nouns (Dubossarsky et al., 2016). However, contrary to the existing finding that nouns change meaning more than adjectives, we found the opposite pattern to hold true in our dataset where adjectives tend to show less stability than nouns in individuals ($p < 0.05$). The same pattern holds for the individual case of Neal C. These findings suggest there exist some differences in semantic shift viewed across individuals and in a population.

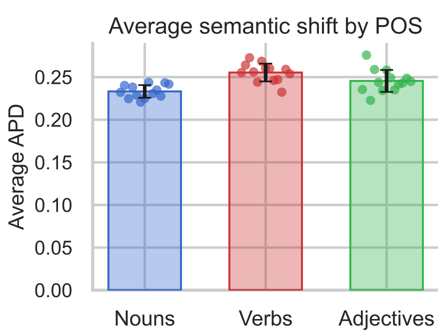

Figure 3: Rates of semantic change across the three main word classes. Each dot represents rate estimated for an individual.

Discussion

We presented a computational framework for quantifying regularity in the diachronic semantic development of individuals. Unlike existing approaches that examine semantic development based on psychological experimentation, our approach draws on individual text corpora based on natural language use of individuals.

Applying our methodology to the INTERVIEW: NPR Media Dialog Transcripts dataset, we found systematic speaker-level variance in a mixed-effects model, suggesting that semantic shift occurs differently across individuals while also showing interesting shared patterns. More specifically, we identified a parallel in the effect of frequency with previous findings in historical semantic change, whereby frequently used words tend to be more robust against semantic shift over time. We also found that the effect of polysemy is equally strong in the genuine and time-shuffled versions of the analysis, suggesting that polysemy may play a minimal role.

Together, these findings support the view that individual semantic development follows regular trajectories shaped by the principle of utility in both population-level and individual-level word meaning dynamics. Additionally, we found only partial evidence that semantic shift across word classes in individuals resembles that in a historical setting, which suggests that meaningful differences may lie in these diachronic processes.

While our framework offers a novel way for quantifying individual-level semantic shift, several limitations should be acknowledged in this initial exploration. One important issue is to what extent the current computational methodology differentiates (or conflates) semantic shift, or temporal change in meaning, with referential switching or sampling (e.g., see Del Tredici, Fernández, & Boleda, 2019).

In addition, the dataset used for our analysis is also limited in several aspects. First, the use of the *NPR Media Dialog Transcripts* dataset introduces restrictions due to its limited temporal span and the number of observations at the individual level, both of which reduce observable semantic shift over time. Second, the sample of speakers is restricted to radio hosts, which increases the likelihood of shared occupation-specific vocabulary while limiting the diversity of language use. Expanding the analysis to a broader range of individuals may yield greater speaker-level variation, although this remains an open question requiring larger and more diverse datasets. Third, the effects of gender and age are not comprehensively examined in the current study and warrant further research once the framework is applied to more densely collected longitudinal data. Fourth, the dataset exclusively contains english text, which means that all findings are specific to the English language. Finally, evaluating semantic change in individuals is inherently challenging and computationally intensive. However, our work adopts reliable methods from computational historical semantics and applies them to an individual setting as a grounded starting point.

Conclusion

Our study highlights a parallel between semantic development at the individual level and semantic change observed in a population in the process of language evolution. By adapting methods from computational historical semantics to model individual semantic trajectories, our work offers the first step toward quantifying regularity in the diachronic development of personal lexicons. We hope that the framework presented here will stimulate future research on how word meanings develop within individuals across time, communities, and cultures.

References

Bréal, M. (1904). *Essai de sémantique (science des significations)*. Hachette.

Brochhagen, T., Boleda, G., Gualdoni, E., & Xu, Y. (2023). From language development to language evolution: A unified view of human lexical creativity. *Science*, *381*(6656), 431–436.

Cain, E., & Ryskin, R. (2025). Semantic representations are updated across the lifespan reflecting diachronic language change. *Open Mind*, *9*, 2114–2148.

Chater, N., & Christiansen, M. H. (2010). Language acquisition meets language evolution. *Cognitive Science*, *34*(7), 1131–1157.

Clark, H. H. (1998). Communal lexicons. In K. Malmkjær & J. Williams (Eds.), *Context in language learning and language understanding* (pp. 63–87).

Del Tredici, M., Fernández, R., & Boleda, G. (2019). Short-term meaning shift: A distributional exploration. In *Proceedings of the 2019 Conference of the North American chapter of the Association for Computational Linguistics: Human Language Technologies.*

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Cinguistics: Human Language Technologies* (pp. 4171–4186).

Dubossarsky, H., De Deyne, S., & Hills, T. T. (2017). Quantifying the structure of free association networks across the life span. *Developmental Psychology*, *53*(8), 1560–1570.

Dubossarsky, H., Weinshall, D., & Grossman, E. (2016). Verbs change more than nouns: A bottom-up computational approach to semantic change. *Lingue e linguaggio*, *15*(1), 7–28.

Dubossarsky, H., Weinshall, D., & Grossman, E. (2017). Outta control: Laws of semantic change and inherent biases in word representation models. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* (pp. 1136–1145).

Enfield, N. J. (2023). Scale in language. *Cognitive Science*, *47*(10), e13341.

Geeraerts, D. (1997). *Diachronic prototype semantics: A contribution to historical lexicology*. Oxford University Press.

Gentzkow, J. M. S., Matthew, & Taddy, M. (2018). *Congressional record for the 43rd-114th congresses: Parsed speeches and phrase counts.* Stanford Libraries. Retrieved from https://data.stanford.edu/congress$_t$*ext*

Giulianelli, M., Del Tredici, M., & Fernández, R. (2020). Analysing lexical semantic change with contextualised word representations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics.*

Hamilton, W. L., 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.*

Kamath, G., Yang, M., Reddy, S., Sonderegger, M., & Card, D. (2025). Semantic change in adults is not primarily a generational phenomenon. *Proceedings of the National Academy of Sciences*, *122*(31), e2426815122.

Kulkarni, V., Al-Rfou, R., Perozzi, B., & Skiena, S. (2015). Statistically significant detection of linguistic change. In *Proceedings of the 24th International Conference on World Wide Web.*

Majumder, B. P., Li, S., Ni, J., & McAuley, J. (2020a). *Interview: A large-scale open-source corpus of media dialog.* Retrieved from https://arxiv.org/abs/2004.03090

Majumder, B. P., Li, S., Ni, J., & McAuley, J. (2020b). Interview: Large-scale modeling of media dialog with discourse patterns and knowledge grounding. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing.*

Marti, L., Wu, S., Piantadosi, S. T., & Kidd, C. (2023). Latent diversity in human concepts. *Open Mind*, *7*, 79–92.

Monaghan, P. (2014). Age of acquisition predicts rate of lexical evolution. *Cognition*, *133*(3), 530–534.

Noble, B., Sayeed, A., Fernández, R., & Larsson, S. (2021). Semantic shift in social networks. In *Proceedings of SEM 2021: The Tenth Joint Conference on Lexical and Computational Semantics* (pp. 26–37).

Pagel, M., Atkinson, Q. D., & Meade, A. (2007). Frequency of word-use predicts rates of lexical evolution throughout Indo-European history. *Nature*, *449*(7163), 717–720.

Pennebaker, J. W., & King, L. A. (1999). Linguistic styles: Language use as an individual difference. *Journal of Personality and Social Psychology*, *77*(6), 1296–1312.

Reif, E., Yuan, A., Wattenberg, M., Viegas, F. B., Coenen, A., Pearce, A., & Kim, B. (2019). Visualizing and measuring the geometry of BERT. *Advances in Neural Information Processing Systems*, *32*.

Slatcher, R. B., Chung, C. K., Pennebaker, J. W., & Stone, L. D. (2007). Winning words: Individual differences in linguistic style among US presidential and vice presidential candidates. *Journal of Research in Personality*, *41*(1), 63–75.

Sullivan, J., Mei, M., Perfors, A., Wojcik, E., & Frank, M. C. (2021). SAYCam: A large, longitudinal audiovisual dataset recorded from the infant's perspective. *Open Mind*, *5*, 20-29.

Traugott, E. C., & Dasher, R. B. (2001). *Regularity in semantic change* (Vol. 97). Cambridge University Press.

Wulff, D. U., Hills, T. T., & Mata, R. (2022). Structural differences in the semantic networks of younger and older adults. *Scientific Reports*, *12*(1), 21459.