To appear in Cognition.

Evolution of emotion semantics

Aotao Xu[a,*], Jennifer E. Stellar[b], Yang Xu[a,c]

[a]*Department of Computer Science, University of Toronto*
[b]*Department of Psychology, University of Toronto*
[c]*Cognitive Science Program, University of Toronto*

Abstract

Humans possess the unique ability to communicate emotions through language. Although concepts like anger or awe are abstract, there is a shared consensus about what these English emotion words mean. This consensus may give the impression that their meaning is static, but we propose this is not the case. We cannot travel back to earlier periods to study emotion concepts directly, but we can examine text corpora, which have partially preserved the meaning of emotion words. Using natural language processing of historical text, we found evidence for semantic change in emotion words over the past century and that varying rates of change were predicted in part by an emotion concept's prototypicality—how representative it is of the broader category of "emotion". Prototypicality negatively correlated with historical rates of emotion semantic change obtained from text-based word embeddings, beyond more established variables including usage frequency in English and a second comparison language, French. This effect for prototypicality did not consistently extend to the semantic category of birds, suggesting its relevance for predicting semantic change may be category-dependent. Our results suggest emotion semantics are evolving over time, with prototypical emotion words remaining semantically stable, while other emotion words evolve more freely.

Keywords: emotion; semantic change; semantic stability; prototype theory; word embedding

*Corresponding author
Email address: `a26xu@cs.toronto.edu` (Aotao Xu)

1. Introduction

Much like emotion concepts vary in their meaning across cultures [1, 2, 3], it is possible emotion words can take on different meanings over time.[1] For instance, the English word *awe* in the 18th century may not represent the same feeling it does today, after a century of evolving perspectives on power and beauty [4]. Although we cannot travel to earlier historical periods to study emotion concepts directly, we do have access to text corpora which have partially preserved the meaning of emotion words. These words do not reflect the entirety of an emotion concept, which includes expressive, experiential, and physiological components, but they do offer insight into its shared meaning within a society. Here we use computational linguistic analyses to investigate the evolution of emotion semantics.

If the meaning of different emotion words like *awe* or *joy* are changing over time, are they changing at the same rate or are there features of an emotion word that might predict its rate of semantic change? We propose that an emotion's conceptual prototypicality is one such feature. Prototypicality is a graded measure of the goodness of a concept's membership in a semantic category [5, 6]. In the case of emotions, *joy* is considered a more prototypical concept than *optimism*. Prototypical emotion concepts may have clearer biological and cultural functions and more distinct features than less prototypical ones. For instance, prototypical concepts like *fear* and *disgust* are particularly suited to solving evolutionary challenges or taking advantage of opportunities that faced early humans [7], and they may have particularly strong social or cultural scripts [8, 5, 9]. These emotion concepts are often more clearly marked by distinctive expressions, experience, and patterns of activation in the body [10], and prototypical members may even help define the meaning of their less proto-

[1] In our study, we use the terms "emotion concept" and "emotion word" interchangeably to refer to emotions that are lexicalized in natural language.

typical counterparts [11] (see *Supplementary Information* for further evidence). We hypothesize that these well-defined functions and features of prototypical emotion concepts could promote semantic stability. As a result, the meaning of words for more prototypical concepts like *joy* may tend to resist change, more so than words for less prototypical ones like *optimism*; see Figure 1 for an illustration.

Although prototypicality has been discussed in other semantic categories, we do not expect prototypicality to predict semantic stability in every category. The basis of prototypicality and thus its ability to predict semantic change may differ in the classic example of birds [12]. The prototypicality of a bird name is primarily based on differences in biological taxonomies [13] and features grounded in sensory or visual perception [14]. As such, the features that define more (e.g., *sparrow*) or less prototypical birds (e.g., *penguin*) are equally well-defined, so the meanings of prototypical bird names do not help define the meanings of less prototypical bird names (see *Supplementary Information*), in contrast to the category of emotion words. We expect that while prototypicality should correlate with semantic stability of emotion words, it should not correlate with semantic stability of bird names.

Our hypothesis augments general principles of lexical evolution and semantic change. It has been shown that usage frequency is a general determiner of stability in English verb regularization [15], lexical replacement [16, 17, 18], loan word borrowing [19], and semantic change [20, 21]. If explained through the lens of communication, frequency should predict semantic stability: when speakers change the meaning of a highly frequent lexical item, they would face a higher number of misunderstandings than if they change a low-frequency item [22, 23]. As a result, we expect frequent emotion words to change less in meaning than other emotion words. We examine the prototypicality of emotion concepts as an additional predictor of semantic stability beyond usage frequency.

Our hypothesis differs from diachronic prototype semantics [24], which states that more prototypical senses of a word tend to stay prototypical over time and exhibit more stability than peripheral senses. Although this theory is consistent

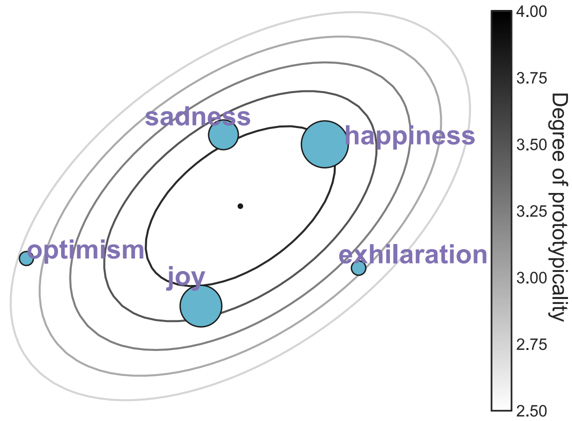

Figure 1: An illustration of the relation between prototypicality and semantic stability of emotion words. Each blue dot represents an emotion word, and the size of the dot is proportional to its predicted semantic stability; the smaller the dot, the higher its rate of semantic change over time. The contours indicate degrees of prototypicality. Visually, an emotion word close to the center has high prototypicality, and vice versa.

with our hypothesis regarding the pattern that prototypicality offers stability, we focus on explaining rates of semantic change among concepts in a lexical field, as opposed to characterizing principles of change among senses of an individual word [24, 25]. Previous studies have examined the theory of diachronic prototype semantics over the whole lexicon and found the prototypicality of words in statistical clusters (formed in meaning space) to negatively correlate with rates of semantic change [26, 21]. However, these studies do not explain how semantic change relates to prototypicality in the scope of a specific category such as emotions or birds.

We present a methodology for modeling emotion semantics and its evolution by building on work from machine learning and natural language processing in word embedding [27, 28, 29] and its historical extensions [30, 20, 31, 32]. We model emotion semantics using a vector-space representation trained on historical text corpora of natural language use, and we use this representation to model human judgments of prototypicality and semantic change of emotion words. Vector-space models of word meaning have been used within affective science for reconstructing human emotion ratings on dimensions such as valence and arousal [33], sentiment analysis [34], and analyzing emotion categories in documents [35], but not for investigating the open question of the evolution of emotion semantics.

2. Methodologies for quantifying rates of semantic change

Quantifying the rate of semantic change for a word requires records of its meaning from two distinct time periods and a quantitative metric that compares these records. One type of methods that constructs word meanings and enables comparisons over time is based on word embeddings [27, 28]. The embedding of a word is a real-valued vector that represents its meaning through a high-dimensional space; vectors for words with similar meanings tend to be close in this space, such as *compassion* and *sympathy*. Word embeddings are constructed from co-occurrence statistics in large text corpora. We thus obtain meaning representations from two distinct time periods by constructing word embeddings based on historical text corpora from the corresponding periods [20].

Existing methods for computing rates of semantic change often rely on the cosine distance between two embeddings [20, 21]. According to this metric, a large cosine distance between historical and contemporary embeddings implies a high rate of semantic change, and vice versa. However, this metric tends to bias the correlation between rate of semantic change and frequency [21]. For this reason, we use an alternate neighbourhood-based metric to compare word embeddings across time [36]. This metric quantifies the rate of semantic change for a word w between periods t_1 and t_2 via the Jaccard distance between sets of k-nearest neighbours in meaning space:

$$rate(w, t_1, t_2) = 1 - \frac{|kNN(w, t_1) \cap kNN(w, t_2)|}{|kNN(w, t_1) \cup kNN(w, t_2)|} \tag{1}$$

where $kNN(w, t)$ contains the k words whose embeddings are the closest to the embedding of w in terms of cosine similarity. Intuitively, we say a word underwent semantic change if the composition of its semantic neighbourhood has changed. Following [36], the part of speech (POS) of the members of $kNN(w,t)$ is always the same as the POS of w, and we also set k to 100. In *Supplementary Information*, we show that this measure is robust to variations in k. Compared to the cosine metric, this metric enables more transparent interpretation on rates of change because we can inspect and evaluate the sets of semantic neighbours (see *Supplementary Information* for examples of emotion semantic change).

To implement this metric at scale, we used pretrained historical word embeddings and POS tags from HistWords [20]. Specifically, we used 300-dimensional Word2Vec (SGNS) embeddings obtained from the Skip-Gram model [28] and trained on the corpora Google Books Ngrams English and French. We used historically most frequent POS tags from the same sources. This provided us with historical word embeddings and most frequent POS tags for 100,000 English words and 100,000 French words, for every decade between 1800 and 2000.

3. Analyses of emotion concepts

In the first set of analyses, we provide evidence for our hypothesis that the well-defined features and functions of prototypical emotion words promote semantic stability. Specifically, we test against the null hypothesis that prototypicality does not predict semantic stability in English and French emotion words over the past century.[2] We describe resources that provide us with lists of English and French emotion words, emotion prototypicality ratings, and historical frequency estimates. We then describe our methods for estimating prototypicality ratings historically and for hypothesis testing, which is followed by a presentation of our results.

[2]We focused on these two languages because 1) we want to test if our analysis generalizes beyond a single language, and 2) there is a limited cross-linguistic variety of empirical studies on emotion prototypicality and of the historical data provided by HistWords.

3.1. Materials

We obtained a list of English emotion words and their corresponding prototypicality ratings from [6]. The authors produced the list by obtaining 213 emotion nouns from a collection of emotion concepts. They produced emotion prototypicality ratings by asking 112 American university students to rate each of these nouns on a 4-point scale, where 4 means the noun is definitely an emotion, and 1 means the noun is definitely not an emotion. Following this work, our analyses focused on nouns that have prototypicality ratings at least 2.75 with the addition of *surprise* and exclusion of *abhorrence*, *ire*, *malevolence*, and *titillation*; we additionally included the word *awe*. We also obtained the valence of these emotion words from the study, which was originally derived from applying multidimensional scaling to similarity judgments [6].

We also obtained a list of French emotion words with their corresponding prototypicality ratings [37]. The authors produced the list by translating 237 Italian emotion words from an earlier study into French. They produced emotion prototypicality ratings by asking 319 French university students to rate each of these words on a 10-point scale, where 10 means the word is certainly an emotion, and 1 means it is not an emotion. To be consistent with the English list, we kept emotion words whose most frequent POS tag is noun in the final decade of our historical POS data. We also obtained the valence of these emotion words from the study, which was originally obtained by asking 300 French university students to rate the words on a scale of -5 (very unpleasant) to 5 (very pleasant) [37].

We obtained historical frequency data from HistWords [20], which is based on the corpora Google Books Ngrams English and French. This yielded historical frequency data for 682,459 English words and 213,686 French words, for every decade between 1800 and 2000. We intersected the word lists with historical word embeddings, POS tags, and frequency from HistWords. We noticed that more emotion words were unavailable when we increased the span between flanking decades than otherwise: if $t_1 = 1890$ and $t_2 = 1990$, only 9 words from the English list are unavailable in HistWords and the HTE, but if we used $t_1 = 1800$, the number increased to 28; similarly in French, the shorter time span resulted in 32 unavailable words, but the longer one resulted in 58 unavailable words. Consequently, we decided to use the decades of 1890 and 1990 as the flanking decades for our analysis (i.e. $t_1 = 1890$, $t_2 = 1990$), and we used historical frequency data from the 1890s. After the intersection, we had a total of 123 English emotion words and 112 French emotion words.

3.2. Methods

Since we cannot go back in time to measure the prototypicality of emotion concepts in the past, we needed a method for estimating historical prototypicality. Let x represent the word embedding of a concept in category c. Following previous work in prototype theory [38, 39], we estimated the prototypicality of x as the unnormalized conditional probability $p(c|x)$, which can be computed using an isotropic Gaussian via Bayes rule:

$$p(c|x) \propto p(x|c) \sim N(\mu, I) \tag{2}$$

where $\mu = \frac{1}{|E_c|} \sum_{v \in E_c} v$ and E_c is the set of embeddings for members of c; I is an identity matrix. Intuitively, we estimated the prototypicality of x by computing its distance from the category centroid μ; the closer they are, the higher its estimated prototypicality. To estimate the prototypicality of an emotion concept in history, we used its historical embedding x and the embeddings of other emotion concepts to compute $p(x|c = \textit{emotion})$. We evaluated this method by computing the correlation between our empirical prototypicality ratings obtained from [6, 37] and our estimated prototypicality based on embeddings from the 1980s and 1990s, the decades closest to the publication of those studies.

To test against the null hypothesis, we computed the rate of change for every emotion concept x, $rate(x, 1890, 1990)$ using Equation (1) and historical embeddings and POS tags from HistWords. Separately for English and French, we then computed the Pearson correlations between the emotion concepts' rates of change and prototypicality estimated for the 1890s. To evaluate whether the prototypicality of emotion concepts predicts rates of change beyond frequency, we performed multiple linear regressions for English and French using the following regression formula:

$$rate(x, 1890, 1990) \sim p(x|c = emotion) + freq(x) + val(x) \tag{3}$$

where for every concept x, we denote its usage frequency as $freq(x)$ and its valence as $val(x)$, which we added to control for unequal numbers of negative and positive emotion concepts in our datasets. We fitted the model using ordinary least squares implemented by `statsmodel` [40]; we also used this package to compute relevant test statistics. Following previous work [20], we performed a log transformation on frequency.

3.3. Results

Figure 2 shows the Pearson correlation between estimated prototypicality from English word embeddings and ratings from English speakers [6]: $\rho = 0.428$, $p < 0.001$, $n = 123$. We obtained similar results with French word embeddings for a set of French emotion concepts [37]: $\rho = 0.438$, $p < 0.001$, $n = 112$. These initial results show our estimated degrees of prototypicality for emotion concepts capture human judgments to some extent. For this reason, we used the same method to estimate historical prototypicality ratings and evaluated them as predictors of semantic stability.

Figure 3 shows a significant negative correlation between emotion prototypicality and degree of semantic change: $\rho = -0.580$, $p < 0.001$, $n = 123$. On average, emotion concepts rated prototypical such as *anger*, *joy*, *fear* underwent less change in meaning compared to words denoting less prototypical concepts such as *zest, exhilaration* and *hysteria* (see annotated word samples in Figure 3). Similar results hold for French: $\rho = -0.576$, $p < 0.001$, $n = 112$. *Supplementary Information* provides additional examples of English and French emotion concepts from the most changing to the most semantically stable, along with their semantic neighbours retrieved from our methods.

Figure 4 shows our results for multiple regression. The adjusted R^2 of the model for English is 0.680, with $p < 0.001$, $n = 123$; mean regression coeffi-

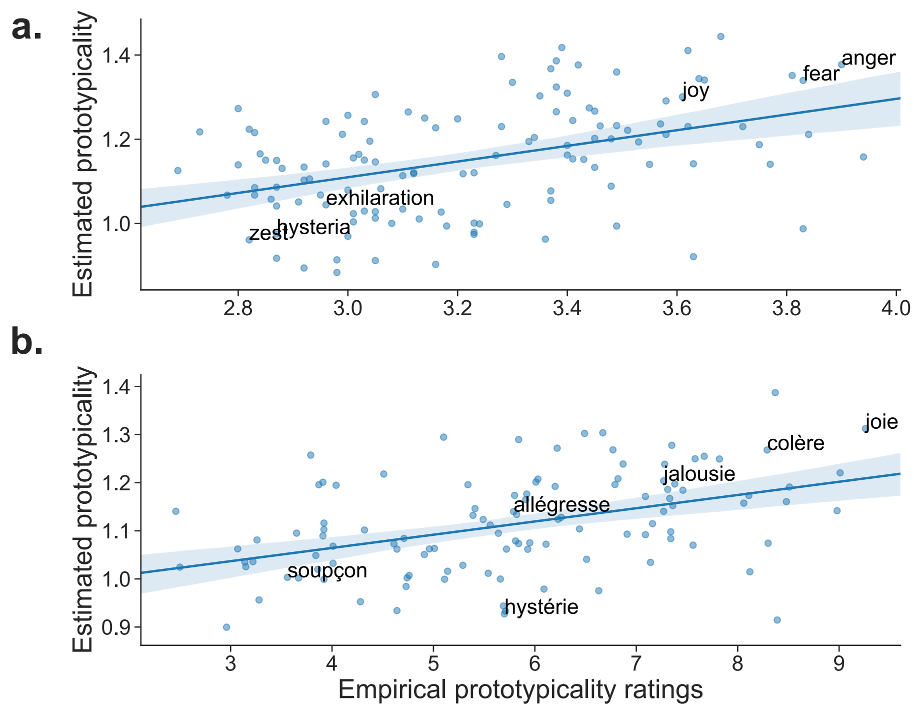

Figure 2: Word embedding reconstruction of emotion prototypicality in a) English and b) French. Scatter plots compare estimated prototypicality computed from Equation 2 against empirical ratings. Each dot corresponds to an emotion concept (a sample of concepts annotated), and each band shows a 95% confidence interval for the line of best fit.

cients for prototypicality ($\beta = -0.417$, $p < 0.001$) and frequency ($\beta = -0.0451$, $p < 0.001$) are significant, but for valence ($\beta = 0.0053$, $p = 0.208$) it is insignificant. For French, the adjusted R^2 of the model is 0.538, with $p < 0.001$, $n = 112$; mean regression coefficients for prototypicality ($\beta = -0.6363$, $p < 0.001$) and frequency ($\beta = -0.0331$, $p < 0.001$) are significant, but for valence ($\beta = 0.0019$, $p = 0.454$) it is insignificant. These results show that frequency predicts semantic stability, which confirms the previous findings [20, 21]. Beyond frequency, we find that prototypicality plays an important role in predicting semantic stability of emotion words, manifested in its significant and negative effect. This provides evidence for our hypothesis that prototypical emotion words tend to

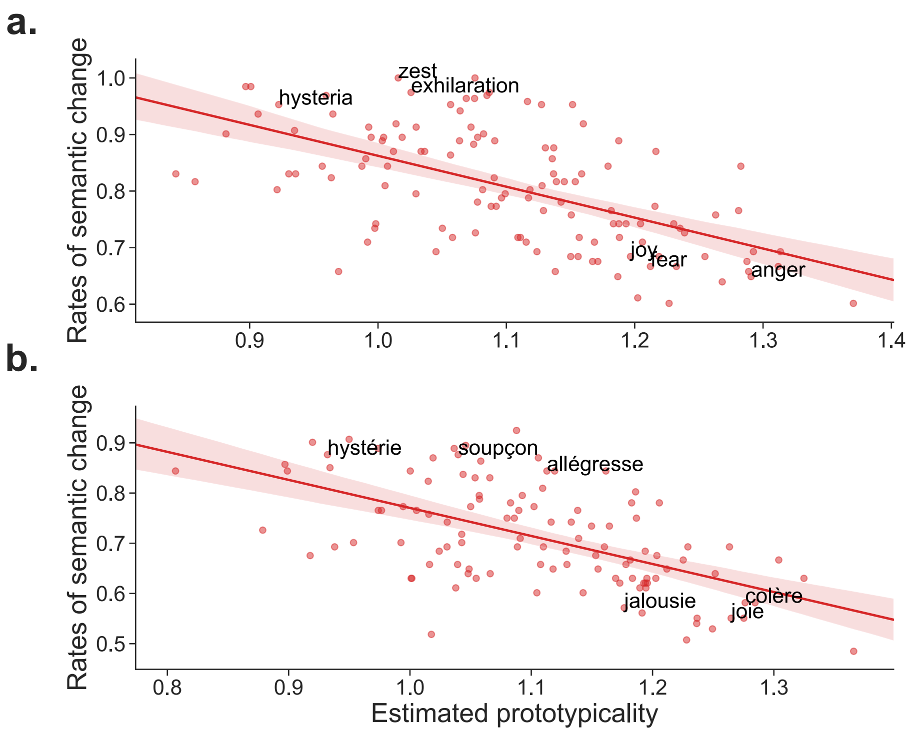

Figure 3: Scatter plots showing the negative correlations between emotion prototypicality and rates of emotion semantic change between the 1890s and 1990s, in a) English and b) French. Each dot corresponds to an emotion word (with a sample set of words annotated), and each band shows a 95% confidence interval for regressions between emotion prototypicality and rates of semantic change.

be semantically stable over time.

Supplementary Information includes three more analyses that further corroborate our findings. The first analysis repeats the multiple regression but restricts the neighbourhoods to emotion concepts only when computing $rate(w, 1890, 1990)$; the results rule out the possibility that our findings are an artifact of the non-emotion senses of polysemous emotion concepts (e.g., *zest*). The second analysis extends the multiple regression for English by including additional predictors based on hypernymy-hyponymy, age of acquisition, and degrees of polysemy, which could potentially subsume the effects of prototypicality; our results

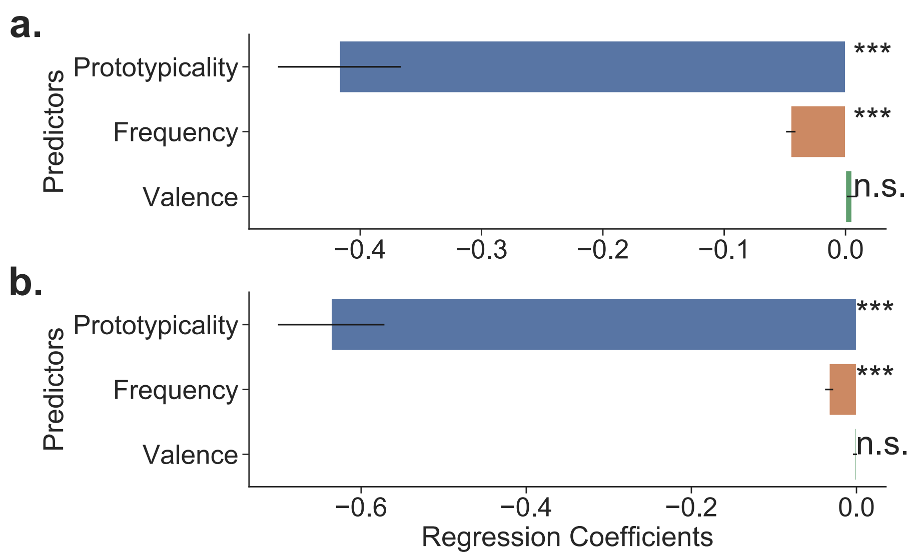

Figure 4: Predictor coefficients from multiple regressions on rates of emotion semantic change. Error bars show standard error, and "n.s.", "*", "**", "***" denote no significance at $p < 0.05$, and $p < 0.05$, $p < 0.01$, $p < 0.001$ respectively. a) shows results for English, and b) shows results for French.

show that this is not the case. The third analysis repeats the multiple regression for English emotion concepts, except the rates of change are computed as $rate(w, 1980, 1990)$ and empirical prototypicality from [6] were used; these results provide evidence that the effect of prototypicality is not caused by potential artifacts in our estimation of prototypicality based on Equation 2.

Figure 5 illustrates our main finding with two example words: *disgust* and *awe*. These words had similar usage frequencies over time, but *disgust* is rated as a more prototypical emotion word than *awe* [6]. Over time, *awe* has shifted meaning more substantially than *disgust*. In particular, both words were in the neighbourhood of negative emotion words (e.g., *sadness*, *anger*, and *fear*) in the 1890s. However, while *disgust* still remained close to these words in the 1990s, *awe* moved closer to positive emotion words (e.g., *love* and *happiness*).

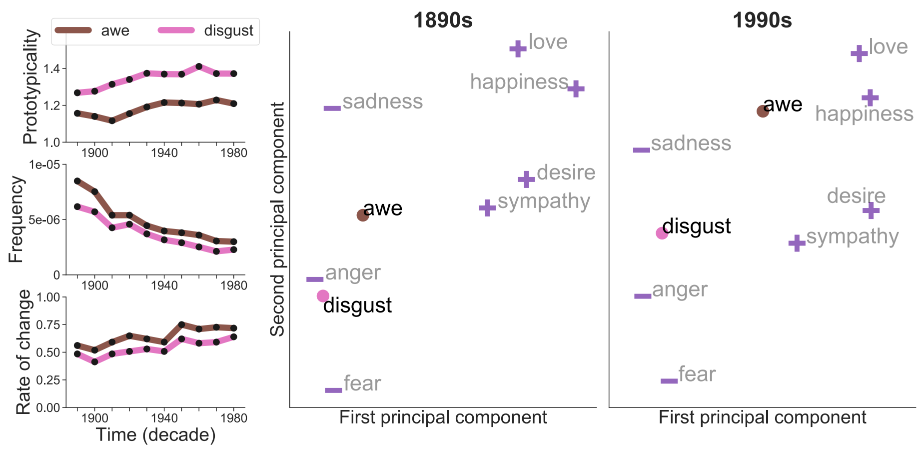

Figure 5: An illustrative comparison of prototypicality, frequency, and semantic stability in emotion words *awe* and *disgust*. Left panels show the embedding-based prototypicality, frequency and degree of semantic change of *awe* and *disgust* over time. Right panels visualize the rates of change in these words by placing them in the two principal components [41] of meaning space, along side prototypical emotion concepts which are annotated based on their valence ("+" for positive, "-" for negative).

4. Analyses of bird names

In this set of analyses, we demonstrate that the sources of prototypicality do not always provide semantic stability as we have shown for emotion concepts. Here we repeat our previous analyses on a case study of birds, a frequently investigated category in prototype theory [12, 42]. As we will see, our embedding-based estimation of prototypicality does not work well with bird names, and we will focus our analysis on using empirical ratings from the 1970s.

4.1. Materials

We obtained a list of English bird names with prototypicality ratings from [12]. The author produced the list by consulting previous work so that the selected names were relatively frequent. They produced bird prototypicality ratings by asking 209 American university students to rate each of these names on a 7-point scale, where 1 means the name refers to a very good example of a bird, and 7 means the name refers to a very poor example. Note that the scale operates in the opposite direction of our prototypicality ratings for emotion concepts. For clarity, we multiplied these ratings by -1 so the direction is the same as our emotion data. Focusing on the 1970s and 1990s, we used historical data from HistWords [20], which was intersected with the word list and provided us with 41 bird names.

4.2. Methods

Similar to the previous section, we attempted at estimating bird prototypicality using Equation 2. We then computed the rates of change for every bird name w, $rate(w, 1970, 1990)$ using Equation 1. We computed the Pearson correlation between rates of change and prototypicality ratings obtained from the 1970s, and we performed a multiple regression using the following formula:

$$rate(w, 1970, 1990) \sim proto(w) + freq(w) \tag{4}$$

where we denote the empirical prototypicality rating of every bird name w as $proto(w)$.

4.3. Results

Figure 6a shows the Pearson correlation between estimated prototypicality and empirical ratings from [12]: $\rho = 0.153$, $p = 0.340$, $n = 41$. While the same method reconstructs prototypicality for emotion concepts to some extent, our text-based method does not explain a significant amount of variance in the prototypicality of birds which depends more on sensory features [14]. It has been shown that prototypical birds in our dataset tend to be passerines, small perching birds that sing (e.g., *robin*), and less prototypical ones tend to be non-passerines (e.g., *penguin*) [13], which our text-based methodology did not capture. For this reason, we chose to focus on empirical prototypicality ratings for birds in our analyses.

Figure 6b shows a significant positive correlation between bird prototypicality and degree of semantic change: $\rho = 0.428$, $p = 0.005$, $n = 41$. This finding suggests that the relation between semantic change and prototypicality in bird names is opposite to our previous findings for emotion words. Figure 7a shows the results for multiple regression. The adjusted R^2 is 0.508, with $p < 0.001$, $n = 41$; mean regression coefficients for empirical prototypicality ($\beta = 0.0283$, $p = 0.011$) and frequency ($\beta = -0.0454$, $p < 0.001$) are significant. We observe frequency still predicts semantic stability, suggesting it is indeed a general predictor of semantic change. Interestingly, prototypicality of birds not only failed to predict stability as in the case of emotion concepts, but also pointed to the opposite trend: in the category of birds, names of prototypical birds tend to undergo more change than other names.

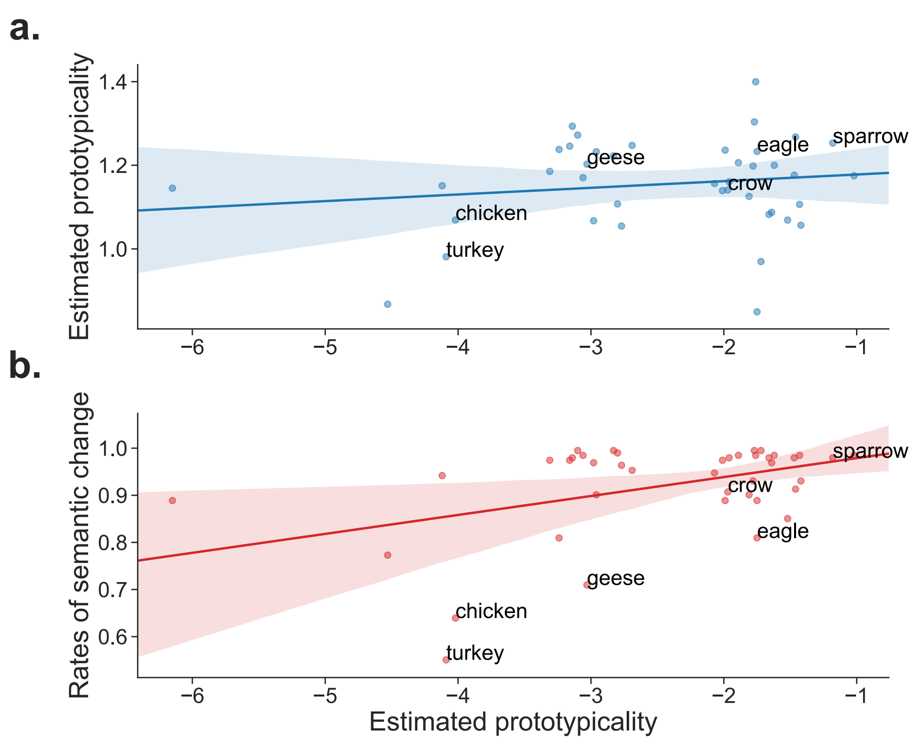

Figure 6: Analyses of bird names: a) word embedding reconstruction of bird prototypicality and b) correlations between bird prototypicality and rates of semantic change between the 1970s and 1990s. Each dot corresponds to a bird name, and each band shows a 95% confidence interval for the line of best fit.

To better understand the implications of this variability to our finding about

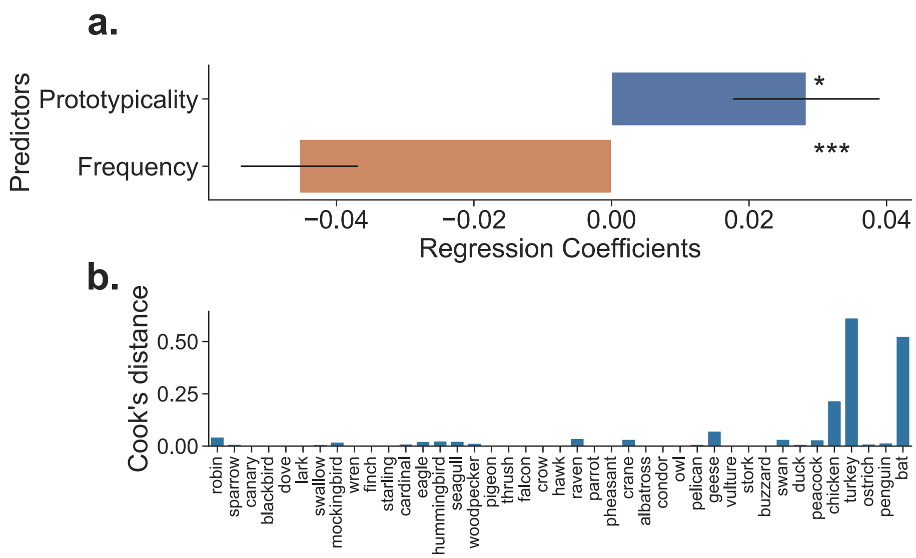

Figure 7: Multiple regression analysis of bird names: a) predictor coefficients from multiple regressions on rates of semantic change, following the same layout as Figure 4; b) Cook's distance for every bird name, showing the influence of individual data points on the regression result.

bird names, we performed a more in-depth analysis of the data. Unlike the case of emotion, we observe bird names exhibit high variability in Figure 6b, which is reflected in the wide confidence region.[3] This suggests the opposite trend in bird names is influenced by only a handful of less prototypical birds. We estimated the influence of each bird name using Cook's distance, which takes into account the data point's residual and leverage. Figure 7b shows the influence of each bird in the regression analysis: the most influential points are *turkey*, *bat* and *chicken*. We can observe *bat* is likely to be influential as it has a much higher rating (not prototypical) than other bird names; this might be because subjects in the original study, being university students, were familiar with the scientific classification of bats. More importantly, despite

[3]Note that the number of available bird terms for our analysis is substantially lower than that of emotion terms.

not being prototypical birds, *turkey* and *chicken* could have important cultural roles (festive or culinary) in North America so that they provided anchors for their meaning, thereby contributing to the significant correlation between bird prototypicality and semantic change.

Figure 8 compares the degrees of semantic change that took place in emotion concepts and bird names between the 1970s and 1990s. Many prototypical emotion concepts tend to lie at the lower tail of the density distribution and show high stability, mirroring the results we have seen previously, but the same pattern does not hold for birds. We observe that overall bird names tend to undergo greater change than emotion concepts do. It is possible that prototypical birds possess the most representative features of the bird category, which could provide points of attachment for meaning change via processes such as chaining, in which a word for one object is extended to be used for another, or metaphor [43, 25]. This general pattern of more rapid change among bird names together with the additional semantic stability of a handful of influential bird exemplars may be responsible for the positive correlation between degrees of bird prototypicality and rates of semantic change.

5. Conclusion

Language offers a lens into the history of emotion semantics. Our computational linguistic analyses of semantic change suggest that a new view of emotion concepts in language may be warranted. Rather than perceiving of emotion concepts as static, their meaning is evolving over time. The exact cultural or societal factors responsible for semantic change in emotion words are difficult to pinpoint, and they may be different for each emotion term. For example, semantic change in *awe* may reflect a movement away from its use in religious contexts, in which it reflects more of a fearful respect, towards greater use in beautiful artistic and natural contexts that followed the emergence of romanticism and transcendentalism in the early to middle nineteenth century. We assessed semantic change over a relatively short timescale, suggesting that in

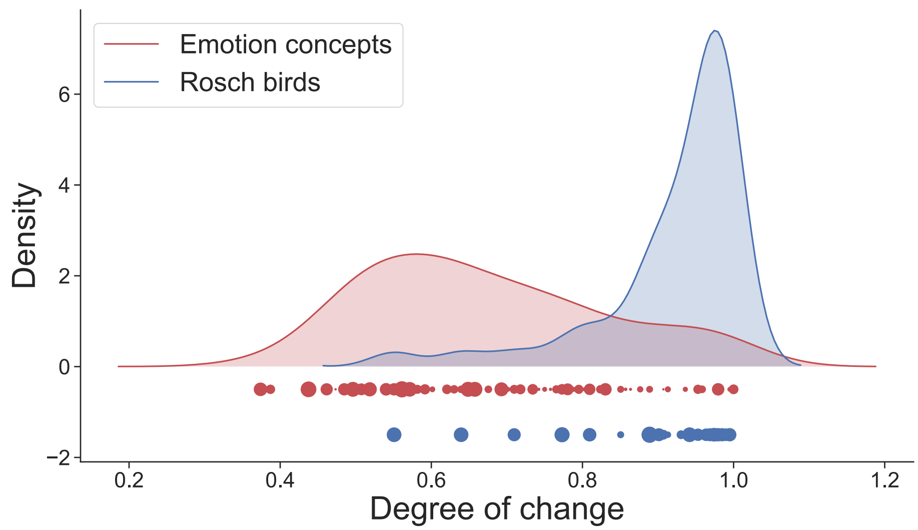

Figure 8: Distributions of semantic change in emotion and bird categories. Each dot corresponds to a word, and the size of the dot is proportional to its degree of prototypicality. The density plots were obtained using kernel density estimation; although degrees of change given by Equation 1 are technically bounded between 0 and 1, we did not bound the support of this figure for illustrative reasons.

the centuries to come it is possible that words like *awe* may continue to evolve and mean something very different than they do today.

Further, we found in two languages that more prototypical emotion words [6, 42] showed greater semantic stability than other emotion words over time. The relation between prototypicality and semantic change depends on its exact sources, as we observed opposite trends for emotions and birds. The importance of prototypicality as a predictor in semantic change for other semantic categories remains an open question and future work should investigate what features affect the importance of prototypicality. Our study extends research on emotions to its historical development and offers a computational cognitive characterization of evolving emotion semantics from natural language use.

Acknowledgements

We thank Brett Ford and two anonymous reviewers for providing constructive comments on our manuscript. This work was supported by a University of Toronto Graduate Entrance Scholarship to AX, a Canada Foundation for Innovation grant #502426 to JES, a NSERC Discovery Grant RGPIN-2018-05872 and a SSHRC Insight Grant #435190272 to YX.

Appendix A. Supplementary material

Code and data used for our analyses are available on GitHub at `https://github.com/johnaot/Emotion_Semantic_Change`.

References

[1] S. E. Kitayama, H. R. E. Markus, Emotion and culture: Empirical studies of mutual influence, American Psychological Association, 1994.

[2] A. Wierzbicka, Emotions across Languages and Cultures: Diversity and Universals, Cambridge University Press, 1999.

[3] J. C. Jackson, J. Watts, T. R. Henry, J.-M. List, R. Forkel, P. J. Mucha, S. J. Greenhill, R. D. Gray, K. A. Lindquist, Emotion semantics show both cultural variation and universal structure, Science 366 (6472) (2019) 1517–1522.

[4] D. Keltner, J. Haidt, Approaching awe, a moral, spiritual, and aesthetic emotion, Emotion and Cognition 17 (2) (2003) 297–314.

[5] B. Fehr, J. A. Russell, Concept of emotion viewed from a prototype perspective, Journal of Experimental Psychology: General 113 (3) (1984) 464.

[6] P. Shaver, J. Schwartz, D. Kirson, C. O'connor, Emotion knowledge: Further exploration of a prototype approach., Journal of Personality and Social Psychology 52 (6) (1987) 1061.

[7] L. Cosmides, J. Tooby, Evolutionary psychology and the emotions, in: M. Lewis, H.-J. J. M (Eds.), Handbook of Emotions 2nd Edition, Guilford, NY, 2000, pp. 91–115.

[8] L. F. Barrett, Are emotions natural kinds?, Perspectives on Psychological Science 1 (1) (2006) 28–58.

[9] J. A. Russell, L. F. Barrett, Core affect, prototypical emotional episodes, and other things called emotion: Dissecting the elephant., Journal of Personality and Social Psychology 76 (5) (1999) 805.

[10] P. Ekman, D. Cordaro, What is meant by calling emotions basic, Emotion Review 3 (4) (2011) 364–370.

[11] P. N. Johnson-Laird, K. Oatley, The language of emotions: An analysis of a semantic field, Cognition and Emotion 3 (2) (1989) 81–123.

[12] E. Rosch, Cognitive representations of semantic categories., Journal of Experimental Psychology: General 104 (3) (1975) 192.

[13] J. S. Boster, Natural sources of internal category structure: Typicality, familiarity, and similarity of birds, Memory & Cognition 16 (3) (1988) 258–270.

[14] P. Garrard, M. A. Lambon Ralph, J. R. Hodges, K. Patterson, Prototypicality, distinctiveness, and intercorrelation: Analyses of the semantic attributes of living and nonliving concepts, Cognitive Neuropsychology 18 (2) (2001) 125–174.

[15] E. Lieberman, J.-B. Michel, J. Jackson, T. Tang, M. A. Nowak, Quantifying the evolutionary dynamics of language, Nature 449 (7163) (2007) 713–716.

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

[17] M. Pagel, Q. D. Atkinson, A. S. Calude, A. Meade, Ultraconserved words point to deep language ancestry across Eurasia, Proceedings of the National Academy of Sciences 110 (21) (2013) 8471–8476.

[18] S. Vejdemo, T. Hörberg, Semantic factors predict the rate of lexical replacement of content words, PLOS ONE 11 (1) (2016) e0147924.

[19] P. Monaghan, S. G. Roberts, Cognitive influences in language evolution: Psycholinguistic predictors of loan word borrowing, Cognition 186 (2019) 147–158.

[20] W. L. Hamilton, J. Leskovec, D. Jurafsky, Diachronic word embeddings reveal statistical laws of semantic change, in: Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, 2016.

[21] H. Dubossarsky, D. Weinshall, E. Grossman, 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, 2017, pp. 1136–1145.

[22] R. Boyd, P. J. Richerson, Culture and the evolutionary process, University of Chicago press, 1988.

[23] J. Bybee, et al., Frequency of use and the organization of language, Oxford University Press on Demand, 2007.

[24] D. Geeraerts, Diachronic Prototype Semantics: A Contribution to Historical Lexicology, Oxford University Press, 1997.

[25] C. Ramiro, M. Srinivasan, B. C. Malt, Y. Xu, Algorithms in the historical emergence of word senses, Proceedings of the National Academy of Sciences 115 (10) (2018) 2323–2328.

[26] H. Dubossarsky, Y. Tsvetkov, C. Dyer, E. Grossman, A bottom up approach to category mapping and meaning change., in: NetWordS, 2015, pp. 66–70.

[27] T. Mikolov, K. Chen, G. S. Corrado, J. Dean, Efficient estimation of word representations in vector space, in: Workshop Papers in International Conference on Machine Learning, 2013.

[28] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, J. Dean, Distributed representations of words and phrases and their compositionality, in: Advances in Neural Information Processing Systems, 2013, pp. 3111–3119.

[29] A. Caliskan, J. J. Bryson, A. Narayanan, Semantics derived automatically from language corpora contain human-like biases, Science 356 (6334) (2017) 183–186.

[30] V. Kulkarni, R. Al-Rfou, B. Perozzi, S. Skiena, Statistically significant detection of linguistic change, in: Proceedings of the 24th International Conference on World Wide Web, 2015.

[31] J. Xie, R. Ferreira Pinto Jr., G. Hirst, Y. Xu, Text-based inference of moral sentiment change, in: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing, 2019.

[32] Y. Li, T. Engelthaler, C. S. Siew, T. T. Hills, The macroscope: A tool for examining the historical structure of language, Behavior Research Methods 51 (4) (2019) 1864–1877.

[33] S. Buechel, U. Hahn, Word emotion induction for multiple languages as a deep multi-task learning problem, in: Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, 2018, pp. 1907–1918.

[34] S. M. Mohammad, Sentiment analysis: Detecting valence, emotions, and other affectual states from text, in: Emotion Measurement, 2016, pp. 201–237.

[35] R. A. Calvo, S. Mac Kim, Emotions in text: dimensional and categorical models, Computational Intelligence 29 (3) (2013) 527–543.

[36] Y. Xu, C. Kemp, A computational evaluation of two laws of semantic change., in: Proceedings of the 37th Annual Meeting of the Cognitive Science Society, 2015.

[37] P. Niedenthal, C. Auxiette, A. Nugier, N. Dalle, P. Bonin, M. Fayol, A prototype analysis of the french category "émotion", Cognition and Emotion 18 (3) (2004) 289–312.

[38] S. K. Reed, Pattern recognition and categorization, Cognitive Psychology 3 (3) (1972) 382–407.

[39] F. G. Ashby, L. A. Alfonso-Reese, Categorization as probability density estimation, Journal of Mathematical Psychology 39 (2) (1995) 216–233.

[40] S. Seabold, J. Perktold, statsmodels: Econometric and statistical modeling with python, in: 9th Python in Science Conference, 2010.

[41] S. Wold, K. Esbensen, P. Geladi, Principal component analysis, Chemometrics and Intelligent Laboratory Systems 2 (1-3) (1987) 37–52.

[42] E. Rosch, Principles of categorization, in: E. Rosch, B. B. Lloyd (Eds.), Cognition and Categorization, Lawrence Erlbaum Associates, Hillsdale, NJ, 1978, pp. 27–48.

[43] Y. Luo, Y. Xu, Stability in the temporal dynamics of word meanings., in: Proceedings of the 40th Annual Meeting of the Cognitive Science Society, 2018.