

JESEME: A Website for Exploring Diachronic Changes in Word Meaning and Emotion

Johannes Hellrich^{1,2} Sven Buechel² Udo Hahn²

¹ Graduate School ‘The Romantic Model’
modellromantik.uni-jena.de

² Jena University Language & Information Engineering (JULIE) Lab
julielab.de

Friedrich-Schiller-Universität Jena, Jena, Germany

Abstract

We here introduce a substantially extended version of JESEME, an interactive website for visually exploring computationally derived time-variant information on word meanings and lexical emotions assembled from five large diachronic text corpora. JESEME is designed for scholars in the (digital) humanities as an alternative to consulting manually compiled, printed dictionaries for such information (if available at all). This tool uniquely combines state-of-the-art distributional semantics with a nuanced model of human emotions, two information streams we deem beneficial for a data-driven interpretation of texts in the humanities.

1 Introduction

Historical, manually compiled dictionaries are central to many kinds of studies in the humanities, since they provide scholars with information about the lexical meaning of terms in former time periods. Yet, this traditional approach is limited in many ways, coverage being perhaps the most pressing issue: Is a dictionary for the specific time period a scholar is investigating really available and, if so, does it cover all of the lexical items of interest?

Word embeddings have been proposed as a technical vehicle to increase lexical coverage (Kim et al., 2014). However, they require locally installed software and time-consuming calculations, thus being ill-suited for mostly non-technical users in the humanities. As an alternative, we here present an extended version of JESEME, a user-friendly open source website¹ for accessing embedding-derived diachronic information on lexical meaning and emotion. The first release of JESEME (Hellrich and Hahn, 2017b) mainly provided time-variant diachronic lexical semantic information. Its second version, the focus of this paper, excels with the unique capability to additionally track the diachronic *emotional* connotation of words in parallel with their lexical semantics. Such a functionality is widely considered beneficial for the data-driven interpretation of literary text genres (Kim et al., 2017).

Measuring affective information on the lexical level is an active field of research in computational linguistics (Liu, 2015). Yet, most contributions focus on contemporary language and are limited to shallow representations of human emotions, mainly distinguishing between *positive* and *negative* feelings. Current research in sentiment analysis either starts to include historical trends in word polarity (Hamilton et al., 2016a) or incorporates more nuanced models of emotions, such as Valence-Arousal-Dominance (Buechel and Hahn, 2018). This contribution integrates both lines of work in a unique way based on our prior research activities (Buechel et al., 2016; Buechel et al., 2017). To the best of our knowledge, only few systems share similarities with JESEME. Alternative websites for tracking diachronic word meaning yet offer far less diverse collections of corpora compared to JESEME and neither of them incorporates emotion values attached to lexical entries. For example, Arendt and Volkova (2017) provide only short term trends in word similarity in two social media corpora in their ESTEEM system.² The system³ by

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>

¹Website available at: jeseme.org; sources available at: github.com/JULIELab/JeSemE

²esteem.labworks.org/

³embvis.flovis.net/s/neighborhoods.html

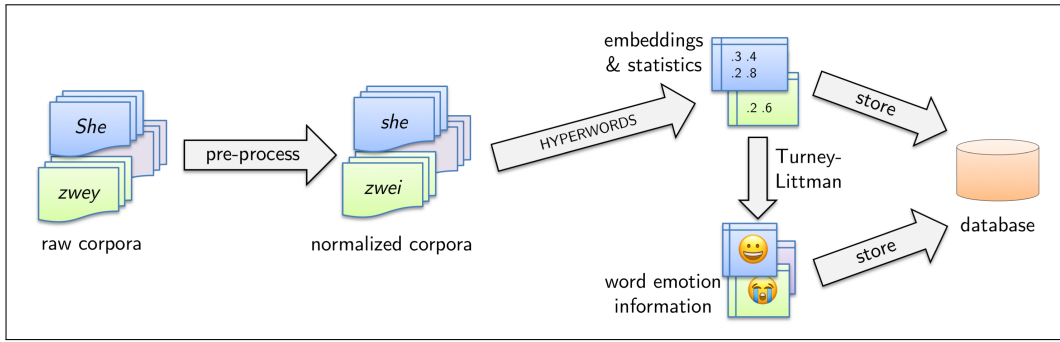


Figure 1: JESEME’s text processing pipeline.

Heimerl and Gleicher (2018) is intended as a mere showcase for a novel visualization technique and re-uses SGNS embeddings trained on the English Google Books corpus by Hamilton et al. (2016b). The DIACHRONIC EXPLORER⁴ which uses sparse vector representations instead of word embeddings to calculate lexical similarity is limited to the Spanish Google Books corpus (Gamallo et al., 2018).

2 Architecture and Website

JESEME uses five diachronic corpora: the Google Books N-Gram Corpus for German and its English fiction register (Michel et al., 2011), the Corpus of Historical American English (COHA; Davies (2012)), the Deutsches Textarchiv [‘German Text Archive’] (Geyken, 2013) and the Royal Society Corpus (Kermes et al., 2016). To ensure high embedding quality, these corpora are divided into temporal slices of similar size covering between 10 to 50 years each.

JESEME’s processing pipeline is depicted in Figure 1. It starts with orthographically normalizing the corpus slices, i.e., lower casing only for English and a historical spelling-aware lemmatization for German (Jurish, 2013). We then use a modified version of HYPERWORDS⁵ to calculate slice-specific

embedding models with SVD_{PPMI} (Levy et al., 2015). This algorithm was chosen for its superior reliability which is essential for interpreting local neighborhoods in embedding spaces as is done in the remainder of this paper (Hellrich and Hahn, 2016; Hellrich and Hahn, 2017a). Apart from word vectors, we also calculate word-based co-occurrence statistics, frequency information and emotion values for each slice (see Section 3). All this information is stored in a relational database. Compared to Hellrich and Hahn (2017b), our current version also reduces the database size from approximately 120GB to 40GB. This is achieved by storing word vectors instead of pre-computed similarity scores. Unlike the previous version, semantic similarity between most words will be computed on the fly. Only the most




Figure 2: JESEME in operation I: Meaning change of “heart” relative to reference words since the 1830s in the COHA.

⁴tec.citius.usc.es/explorador-diacronico

⁵github.com/hellrich/hyperwords

similar ones for each word (automatically picked as references) are cached for fast retrieval.

JESEME’s website prompts a search form for selecting the word under scrutiny as well as one of the five corpora we supply. Its result page then provides graphs depicting the development of semantic similarity to automatically chosen reference words over time as an indicator for semantic change, as well as information on diachronic affective meaning (see Figure 2). These two main sources of information are complemented with information on word co-occurrence and relative frequency, thus providing scholars with additional information to increase interpretability and rule out measurement artifacts. Users may also add further reference words to the analysis on demand. Besides this graphical interface JESEME also offers a REST API.⁶

3 Representing and Computing Emotions

We represent emotions following the Valence-Arousal-Dominance (VAD) scheme (Bradley and Lang, 1994), one of the major models of emotion in psychology (for an illustration, see Figure 3). The VAD model describes affective states relative to three dimensions, namely, Valence (degree of displeasure vs. pleasure), Arousal (degree of calmness vs. excitement) and Dominance (degree of perceived control in a social situation).

We used a modified version of the emotion induction algorithm by Turney and Littman (2003) based on evidence that it outperforms alternative methods for historical emotion lexicon creation (Buechel et al., 2017; Hellrich et al., 2018). In this work, each word’s predicted emotion value $\hat{e}(w)$ is calculated by averaging the emotion values $e(s)$ for each member s of a seed set S , with $sim(w, s)$, the similarity between w and s , serving as a weight:

$$\hat{e}(w) := \frac{\sum_{s \in S} sim(w, s) \times e(s)}{\sum_{s \in S} sim(w, s)}$$

For the emotion scores stored in JESEME, we used the emotion lexicons by Warriner et al. (2013) and Schmidtke et al. (2014) as seed sets for English and German corpora, respectively. Word emotions were induced independently for each temporal corpus slice, using the respective embedding model to retrieve similarity scores. Hence, the similarity between the seed words and the target word reflects word usage at a given language stage, thereby infusing historical emotion information into the resulting emotion ratings (Buechel et al., 2017).

4 Examples

The new insights provided by diachronic emotion models can be demonstrated by re-visiting the example of “heart” we used in Hellrich and Hahn (2017b) as shown in Figure 2. This lexeme is often used metaphorically or metonymically despite the fact that the heart’s anatomical function was already known for a long time. Results for our novel emotion tracking functionality match a move from metaphorical to anatomical usage we previously observed in the genre-balanced COHA. Around 1900, the similarity of “heart” to words such as “stroke” increases, while Dominance and Valence ratings drop sharply in




Figure 3: Affective space spanned by the Valence-Arousal-Dominance (VAD) model, together with the position of six basic emotion categories. Adapted from Buechel and Hahn (2016).

⁶See online documentation: jeseme.org/help.html#api

tandem (see Figure 2; y-axis values are centered and scaled). This simultaneous drop seems plausible, since we can “change our heart” in a metaphorical sense, yet have little control over our anatomical heart. Also, with its increasing anatomical usage, “heart” becomes less positive, since we are under mortal threat by cardiovascular diseases such as a “stroke”.

Changes in emotion can also be traced for items with a more constant meaning, e.g., for “woman” as shown in Figure 4. Here similarity scores for the most similar words—“man” and “girl”—remain rather static. Yet, emotion values are highly dynamic and seem to match turning points in women’s rights movement, e.g., women’s suffrage in the US is connected with an increase in all VAD dimensions for the 1920s.

5 Conclusion

We introduced a substantially extended version of JESEME, an interactive website for tracking diachronic changes in word meaning and, as a novel and unique feature, word emotion. To the best of our knowledge, no other system combines these two traits. JESEME allows users with a limited technical background to interactively explore semantic evolution based on

five large diachronic corpora for two languages, German and English. We believe that JESEME will be most useful for diachronic linguists and scholars within the digital humanities. We see two major applications: First, it can be used to generate hypotheses by querying words of interest to get a first impression of their semantic evolution. Second, scholars can first shape a hypothesis using traditional means and then query JESEME for testing its plausibility based on diachronic statistical evidence.

Acknowledgements

This research was partially funded by grant GRK 2041/1 from *Deutsche Forschungsgemeinschaft* within the Graduate School “*The Romantic Model. Variation–Scope–Relevance*”.

References

- Dustin Arendt and Svitlana Volkova. 2017. ESTEEM: A novel framework for qualitatively evaluating and visualizing spatiotemporal embeddings in social media. In *ACL 2017: System Demonstrations*, pages 25–30.
- Margaret M. Bradley and Peter J. Lang. 1994. Measuring emotion: The self-assessment manikin and the semantic differential. *Journal of Behavior Therapy and Experimental Psychiatry*, 25(1):49–59.
- Sven Buechel and Udo Hahn. 2016. Emotion analysis as a regression problem—Dimensional models and their implications on emotion representation and metrical evaluation. In *ECAI 2016*, pages 1114–1122.
- Sven Buechel and Udo Hahn. 2018. Word emotion induction for multiple languages as a deep multi-task learning problem. In *NAACL 2018: Long Papers*, pages 1907–1918.

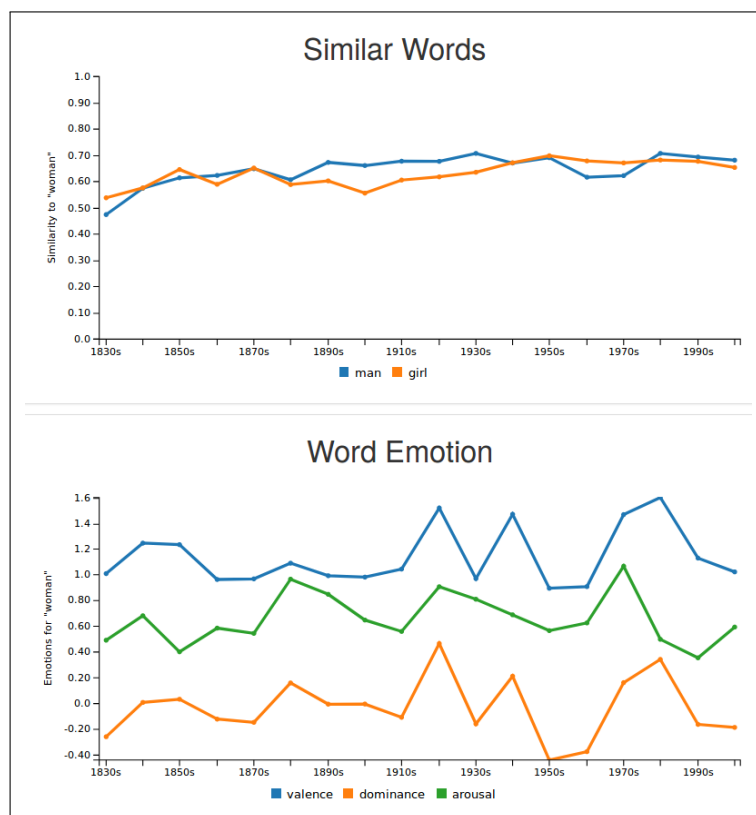


Figure 4: JESEME in operation II: Meaning of “woman” since the 1830s in the COHA.

- Sven Buechel, Johannes Hellrich, and Udo Hahn. 2016. Feelings from the past: Adapting affective lexicons for historical emotion analysis. In *LT4DH @ COLING 2016*, pages 54–61.
- Sven Buechel, Johannes Hellrich, and Udo Hahn. 2017. The course of emotion in three centuries of German text: A methodological framework. In *Digital Humanities 2017*, pages 176–179.
- Mark Davies. 2012. Expanding horizons in historical linguistics with the 400-million word Corpus of Historical American English. *Corpora*, 7:121–157.
- Pablo Gamallo, Iván Rodríguez-Torres, and Marcos Garcia. 2018. Distributional semantics for diachronic search. *Computers & Electrical Engineering*, 65:438–448.
- Alexander Geyken. 2013. Wege zu einem historischen Referenzkorpus des Deutschen: Das Projekt Deutsches Textarchiv. In Ingelore Hafemann, editor, *Perspektiven einer corpusbasierten historischen Linguistik und Philologie*, pages 221–234.
- William L. Hamilton, Kevin Clark, Jure Leskovec, and Dan Jurafsky. 2016a. Inducing domain-specific sentiment lexicons from unlabeled corpora. In *EMNLP 2016*, pages 595–605.
- William L. Hamilton, Jure Leskovec, and Daniel Jurafsky. 2016b. Diachronic word embeddings reveal statistical laws of semantic change. In *ACL 2016: Long Papers*, pages 1489–1501.
- F. Heimerl and M. Gleicher. 2018. Interactive analysis of word vector embeddings. In *EuroVis 2018 — Eurographics Conference on Visualization*. [to appear]. Preprint: graphics.cs.wisc.edu/Papers/2018/HG18.
- Johannes Hellrich and Udo Hahn. 2016. Bad company—Neighborhoods in neural embedding spaces considered harmful. In *COLING 2016: Technical Papers*, pages 2785–2796.
- Johannes Hellrich and Udo Hahn. 2017a. Don’t get fooled by word embeddings: Better watch their neighborhood. In *Digital Humanities 2017*, pages 250–252.
- Johannes Hellrich and Udo Hahn. 2017b. Exploring diachronic lexical semantics with JESEME. In *ACL 2017: System Demonstrations*, pages 31–36.
- Johannes Hellrich, Sven Buechel, and Udo Hahn. 2018. Inducing affective lexical semantics in historical language. arxiv.org/abs/1806.08115.
- Bryan Jurish. 2013. Canonicalizing the Deutsches Textarchiv. In Ingelore Hafemann, editor, *Perspektiven einer corpusbasierten historischen Linguistik und Philologie*, pages 235–244.
- Hannah Kermes, Stefania Degaetano-Ortlieb, Ashraf Khamis, Jörg Knappen, and Elke Teich. 2016. The Royal Society Corpus: From uncharted data to corpus. In *LREC 2016*, pages 1928–1931.
- Yoon Kim, Yi-I Chiu, Kentaro Hanaki, Darshan Hegde, and Slav Petrov. 2014. Temporal analysis of language through neural language models. In *LT-CSS 2014 @ ACL 2014*, pages 61–65.
- Evgeny Kim, Sebastian Padó, and Roman Klinger. 2017. Investigating the relationship between literary genres and emotional plot development. In *LaTeCH-CLfL 2017 @ ACL 2017*, pages 17–26.
- Omer Levy, Yoav Goldberg, and Ido Dagan. 2015. Improving distributional similarity with lessons learned from word embeddings. *Transactions of the Association of Computational Linguistics*, 3:211–225.
- Bing Liu. 2015. *Sentiment Analysis: Mining Opinions, Sentiments, and Emotions*. Cambridge University Press.
- Jean-Baptiste Michel, Yuan Kui Shen, Aviva Presser Aiden, Adrian Veres, Matthew K. Gray, The Google Books Team, Joseph P. Pickett, Dale Hoiberg, Dan Clancy, Peter Norvig, Jon Orwant, Steven Pinker, Martin A. Nowak, and Erez Lieberman Aiden. 2011. Quantitative analysis of culture using millions of digitized books. *Science*, 331(6014):176–182.
- David S. Schmidtke, Tobias Schröder, Arthur M. Jacobs, and Markus Conrad. 2014. ANGST: Affective norms for German sentiment terms, derived from the affective norms for English words. *Behavior Research Methods*, 46(4):1108–1118.
- Peter D. Turney and Michael L. Littman. 2003. Measuring praise and criticism: Inference of semantic orientation from association. *ACM Transactions on Information Systems*, 21(4):315–346.
- Amy Beth Warriner, Victor Kuperman, and Marc Brysbaert. 2013. Norms of valence, arousal, and dominance for 13,915 English lemmas. *Behavior Research Methods*, 45(4):1191–1207.