

Chapter 10

Computational approaches to lexical semantic change: Visualization systems and novel applications

Adam Jatowt^a, Nina Tahmasebi^b & Lars Borin^b

^aUniversity of Innsbruck ^bUniversity of Gothenburg

The purpose of this chapter is to survey visualization and user interface solutions for understanding lexical semantic change as well as to survey a number of applications of techniques developed in computational analysis of lexical semantic change. We first overview approaches aiming to develop systems that support understanding semantic change in an interactive and visual way. It is generally accepted that computational techniques developed for analyzing and uncovering semantic change are beneficial to linguists, historians, sociologists, and practitioners in numerous related fields, especially within the humanities. However, quite a few non-professional users are equally interested in the histories of words. Developing interactive, visual, engaging, and easy-to-understand systems can help them to acquire relevant knowledge.

Second, we believe that other fields could benefit from the research outcomes of computational approaches to lexical semantic change. In general, properly representing the meaning of terms used in the past should be important for a range of natural language processing, information retrieval and other tasks that operate on old texts. In the latter part of the chapter, we then focus on current and potential applications related to computer and information science with the underlying question: “How can modeling semantic change benefit wider downstream applications in these disciplines?”



1 Visualization systems supporting manual analysis

It is already evident that computational approaches greatly speed up the research and the resulting discoveries related to semantic change and lexical replacement. The applications of this research progress still remain to be seen, both inside and outside the academic realm. Providing effective methods for detecting changes, their characteristics, timing, and causal factors are all important for our understanding of languages and their evolution. Computational methods support the acquisition of such knowledge and the formation and validation of various kinds of hypotheses. However, other uses of the technology outside the academic field are less discussed. Word meaning change is not only interesting to professionals (e.g., linguists, historians, or librarians) but also to the wider public. For example, many books discussing the origins and evolution of word meaning have been published aiming at a wider readership, suggesting significant interest by average users in the histories of words. We think that computational approaches and especially online interactive systems are important to help to further disseminate knowledge about etymology.

Visual analytics have been increasingly applied in historical linguistics (Schätzle & Butt 2020) by combining automated algorithms with interactive visual components to let us perform effective investigations through data manipulation and presentation. In this context, several online visualization systems and demonstrations supporting manual analysis have been proposed to complement the research methods developed for detecting diachronic conceptual change. They allow for verification of the results obtained from automatic methods or provide novel means for supporting manual determination of diachronic conceptual change and its characteristics. In these systems the level of interactivity and user freedom in querying the data, as well as the provision of features enabling multidimensional analysis play crucial roles. Evaluation of these systems focuses on usability criteria and their user interfaces. Visualization systems tend to be attractive as they provide either a complement to automatic analysis or serve as the main tool for analysis, and not only for professionals and scientists. Many of them are also particularly suited to lay users, especially if the systems are intuitive and highly usable. Below, we discuss representative systems designed for learning about semantic change and we highlight several new directions based on their analysis. Due to copyright restrictions, we can only show the screenshots of few, selected examples.

1.1 Description- and frequency-based approaches

Starting with a simple example, the enhancement of definitions generated by the Google word search engine for “definition queries” is an effective way to disseminate basic information on words’ origins and patterns of popularity over time. When given an input word, a standard definition can be complemented by an additional click with a brief textual description of the word’s origin as well as its frequency plot over time (see Figure 10.1 for an example). This service provides historical context for word definitions that can support better understanding of queried words and also trigger user interest in word histories by appending complementary basic knowledge about the temporal evolution of word meaning. Although users can see the change in the frequency of a word over time and read brief information on the word’s origin, they still essentially have to guess about the word’s meaning change over the entire span of time. Nevertheless, this application is worth mentioning because, thanks to the popularity of Google search services, information on word change (albeit rather superficial) can be widely disseminated to the public through this service.

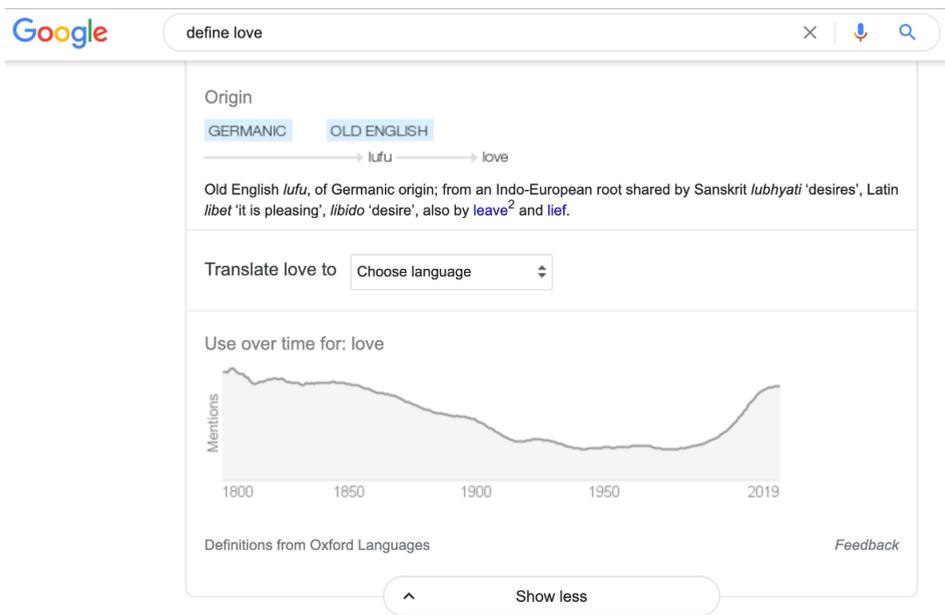


Figure 10.1: Snapshot of an example output from Google search engine for the query `define love` where simple information on the origin and popularity of word *love* over time is given (image captured on 2021-01-24).

The *Online Etymology Dictionary*¹ is an easy-to-use system that draws on the etymology of numerous words compiled from several dictionaries. The service returns a short etymological description of an input word with dates indicating the earliest year for which there is a surviving written record of its use.

The *Google Books Ngram Viewer*² (Michel et al. 2011) – another notable contribution from the search engine company – is a powerful online application for observing and analyzing the frequency of words or Ngrams over time. It is based on the Google Books Ngrams datasets. The service has frequently been used for digital humanities research and the like (e.g. Michel et al. 2011, Acerbi et al. 2013, Bentley et al. 2014, Pechenick et al. 2015, Iliev et al. 2016). The temporal frequency plots of several words or Ngrams can be contrasted with each other. Users can choose a wildcard search (by putting * in place of a word in a given phrase to obtain the top ten substitutions) or do a case-insensitive search. Analysis based on parts of speech (POS) tags is possible (e.g., plotting frequencies of *tackle* as either a verb or as a noun). It is also possible to plot the frequency based on five composition operators (e.g., summing or subtracting the frequencies of several expressions). Inflection-oriented search can be done (e.g., searching with `book_INF a hotel` returns results for *book*, *booked*, *books*, and *booking a hotel*). The Ngram viewer allows the identification of words at the start or end of sentences to be plotted. It provides dependency relations using the => operator. For instance, to understand how often *tasty* was used to modify the word *dessert* one would input (`tasty => dessert`). This search combines frequencies of all instances in which the word *tasty* modifies *dessert*, including *tasty frozen dessert*, and *tasty yet expensive dessert*. Dependencies can be further combined with wildcards (e.g., `drink => *_NOUN` to track frequencies of expressions containing different kinds of beverages as nouns). Nevertheless, because the viewer is mainly based on the frequency signals of words (i.e., probabilities of seeing a given Ngram or a set or composite of Ngrams in a given year), it does not provide a direct means for portraying exactly how a term was used in the past or when its meaning transitions occurred. The viewer is thus best suited to cultural-omics or cultural text mining studies and is similar to other tools available for general purpose interactive exploration of diachronic corpora (e.g. Michel et al. 2011, Odijk et al. 2014, van Eijnatten et al. 2014, Jatowt & Bron 2016). Despite this limitation, this application is worth mentioning because it uses extremely large datasets coupled with basic manipulation capabilities, even though it does not explicitly support analysis of semantic change.

¹<https://www.etymonline.com/>

²<https://books.google.com/ngrams>

The online interfaces of several diachronic corpora created by Mark Davies at Brigham Young University³ provide effective options for users. Users can perform simple analyses without the need to write any code. For example, they can generate frequency plots over time, examples of keywords in context (KWIC) at different time points, and listings of collocates. However, the amount and level of detail of the displayed data make it rather difficult to draw broader conclusions about the semantics of words from a longitudinal perspective.

Hilpert & Gries (2008) apply a variant of hierarchical clustering called variability-based neighbor clustering. The idea is to cluster adjacent time units (hence the name “neighbor clustering”) if the frequency of a target term does not change much. The resulting dendrogram allows for visual identification of time points of large frequency change, which may indicate increased possibility of diachronic sense shifts (e.g., due to sudden triggers like large events). No context is used for a target word because the method relies only on the the frequency information of a query word, which limits the applicability of this approach in representing diachronic conceptual change of words.

Odijk et al. (2014) demonstrate an interactive environment that visualizes information on the volumes and correlations of words and documents across time. Similar to Michel et al. (2011), their focus is more on understanding historical and social aspects than on shifts in word meaning.

1.2 Context-based semantic approaches

Rohrdantz et al. (2011) use latent Dirichlet application (LDA) to represent the different senses of words and track their intensity of change over time. Twenty-five words before and twenty-five words after the target word are used as the context of the term, following the suggestion given by Schütze (1998) for automatic sense discrimination. This approach allows one to notice various kinds of semantic change in words, such as the broadening or narrowing of senses and the first occurrences of senses, especially as all the topics are shown over time in a single view. According to Rohrdantz et al. (2011), their interactive visualization approach provides the possibility of detecting key patterns at a glance, while at the same time observing the details of the data by zooming in on the occurrences of particular words in their contexts. Additionally, the results of the pairwise comparisons of word senses with respect to their shared contexts are also displayed. The authors, however, restrict their system to only a short time period, demonstrating results on the New York Times Annotated corpus, which spans roughly two decades.

³<https://corpus.byu.edu/>

Heylen et al. (2012) propose using a multidimensional scaling (Cox & Cox 2008) technique with a window length of 4 words before and after the target word and pointwise mutual information for weighting context terms. Heylen et al. (2012) took this approach because they had observed that earlier automatic approaches which use distributional models use them in an indirect, black-box fashion, failing to indicate particular semantic properties and relations that play key roles. Motion charts from the Google Chart tools are then used to visualize occurrences of nouns in a 2D representation of their semantic distances. Hovering a mouse pointer over the bubbles denoting nouns shows the text in which each noun occurs so that users can interpret the precise meaning of the occurrence of the noun. In their case study, the authors focus on Dutch words extracted from Dutch newspaper articles published between 1999 and 2005, which were organized in 218 synsets containing 476 nouns in total. Although they do not use the motion feature of the charts, the authors note that it should be possible to track the centroid of the tokens of a target word over time in the semantic space, and also to show the dispersion of the tokens around the centroid.

Hilpert & Perek (2015) use animations in the form of animated scatterplots to portray change in patterns over time using the metaphor of a petri dish. The authors focused on a single pattern, “many a [noun]” as a case study. Spots on the graphs represent nouns involved in the same pattern, and are plotted next to each other if they have high similarity. The size of the spot is linked to the frequency of a noun or noun type in a particular time unit. During the animation, the changes in the size and distances of spots provide knowledge of different uses of the pattern over time.

Dimensionality reduction techniques such as principal components analysis (PCA), latent semantic analysis (LSA) or the popular t-distributed stochastic neighbor embedding (t-SNE, van der Maaten & Hinton 2008)⁴ have been frequently used to plot “trajectories” of word meaning over time in vector spaces using 2D plots. By showing points that represent the meaning of the same words at different years or decades on the same 2D plot (see, e.g., Hamilton et al. 2016a, Kulkarni et al. 2015), and optionally connecting them with arrows, a single static view can show how the words changed their meaning over time, by simply following their “trajectories”. Typically some background reference terms are added along these “trajectories” to ground and explain the meaning.

Martinez-Ortiz et al. (2016) introduce a system called *ShiCo* for visualizing shifting concepts of Dutch words over time. It measures change in words used to refer to concepts based on a model previously introduced by Kenter et al. (2015).

⁴<https://lvdmaaten.github.io/tsne/>

This model requires a series of semantic spaces that are constructed by training word embeddings (e.g., word2vec) for different units of time (typically, each unit spans 10 years). It is based on two steps, generation and aggregation. The generation step works in an iterative fashion. An initial seed set is selected that typically consists of a small number of user-provided terms. Then, words most semantically similar to the seed set are found based on similarity values between word embeddings. A semantic graph is constructed from these terms, and the central terms are extracted using graph centrality measures. Next, the central terms are used as the seed set for the next iteration of the generation step. In the aggregation step, the lists of words produced in the generation step are aggregated to generate the final word lists to be presented to the user. The visualization by Martinez-Ortiz et al. (2016) is composed of two kinds of complementary graphs: a stream graph and a series of network graphs. The former shows color-coded streams for each term; the stream sizes represent the relative importance of the term in a period. This importance is measured either as a term count in each time unit or as a sum of the similarities of the term to the seed terms. The network graphs for each time unit display the relations between terms in the time unit.

Xu & Crestani (2017) use term clouds and a heatmap to visualize semantic shifts of target words by utilizing sequentially trained embedding vectors with initialization based on previous time periods, as proposed by Kim et al. (2014). They use *The New York Times* and *National Geographic Magazine* articles for the underlying datasets, which span about 110 years. Following Martinez-Ortiz et al. (2016), a temporal semantic similarity word cloud is used to show terms most similar to a target query for a given time unit. As in standard term clouds, the font size of the terms is linked to their similarity to the query word. Heatmap views let users see the similarity values of the terms most similar to the target term in each year, using colors. The y-axis of the heatmap is a list of words and the x-axis is a list of temporal periods such that for each given word (each row) one can understand the pattern of the change in the similarity of this term to the target term (also called anchor term). The results from the *The New York Times* and *National Geographic Magazine* are then contrasted with each other.

Jatowt & Duh (2014) describe an analytical framework that incorporates different types of similarity plotting. The plots include across-time self-similarity, a decade-to-decade similarity heatmap, across-time sentiment analysis, diachronic comparative word analysis, and key context term listing. They use both the Google Books Ngrams and COHA datasets. The signals from the different views (e.g. frequency analysis, semantic analysis, and sentiment analysis) can be combined to allow for multi-evidence-based reasoning on diachronic conceptual change. Two different word representations are used: a simple bag-of-words, and

a distance-aware bag-of-words, where the distance used is the relative position of a context term to the target term. For simplicity, the sentiment values of context terms are assumed to remain stable over time when computing the change of the sentiment of the target term's context.

Their work resulted in the development of an online interactive system for diachronic conceptual change analysis (Jatowt et al. 2018) (see Figures 10.2, 10.3 and 10.4).⁵ The system enables detailed analysis of diachronic conceptual change from different viewpoints, such as word's context comparison across-time, temporal term cloud, and temporal term tree generation. It can also perform contrastive analysis of word pairs (see Figure 10.2) of larger groups of words such as synonyms. The results can be generated using both the Google Books Ngrams and COHA datasets. Pearson correlation, cosine similarity, and Jaccard similarity can be used as similarity measures of word representations from different time points. Jatowt et al. (2018) recommend that diachronic conceptual change over time should be contrasted with term frequency plots (as also suggested by Kim et al. 2014), since together both provide a more informed view on how often and in what sense a term was used in the past. Any conclusions drawn from semantic change plots should be treated with caution when the frequency plot of a target term shows a low utilization rate. To better visualize the word change in relation to the average change of other words, the degree of the target word's change over time is also displayed with reference to the average change of words in the same frequency bin as the target word.

The system has another novel feature that allows the change in individual context terms over time to be investigated in the form of a *time-enhanced term cloud* (see Figure 10.3) and *time-enhanced term tree* (see Figure 10.4). Finally, the framework provides a unique functionality to track the semantic shifts of entire concepts represented as word sets, for example, the concept of a vehicle represented by words like *auto*, *automobile*, *car*, *truck*, and so on.

Hellrich et al. (2018) proposed JeSemE, the Jena Semantic Explorer,⁶ which is an interactive website for visually exploring temporal information on word meanings and lexical emotions on the basis of five large diachronic text corpora in English and German, including COHA and Google Books English fiction. A unique feature of this system is the provision of predicted emotion values of words over time based on the valence-arousal-dominance (VAD) scheme of Bradley & Lang 1994. It also shows similar words and specific contexts of a query word. Figure 10.5 and 10.6 show example outputs.

⁵<http://www.okayama.silk.jp/WordEvolution/>

⁶<http://jeseme.org/>

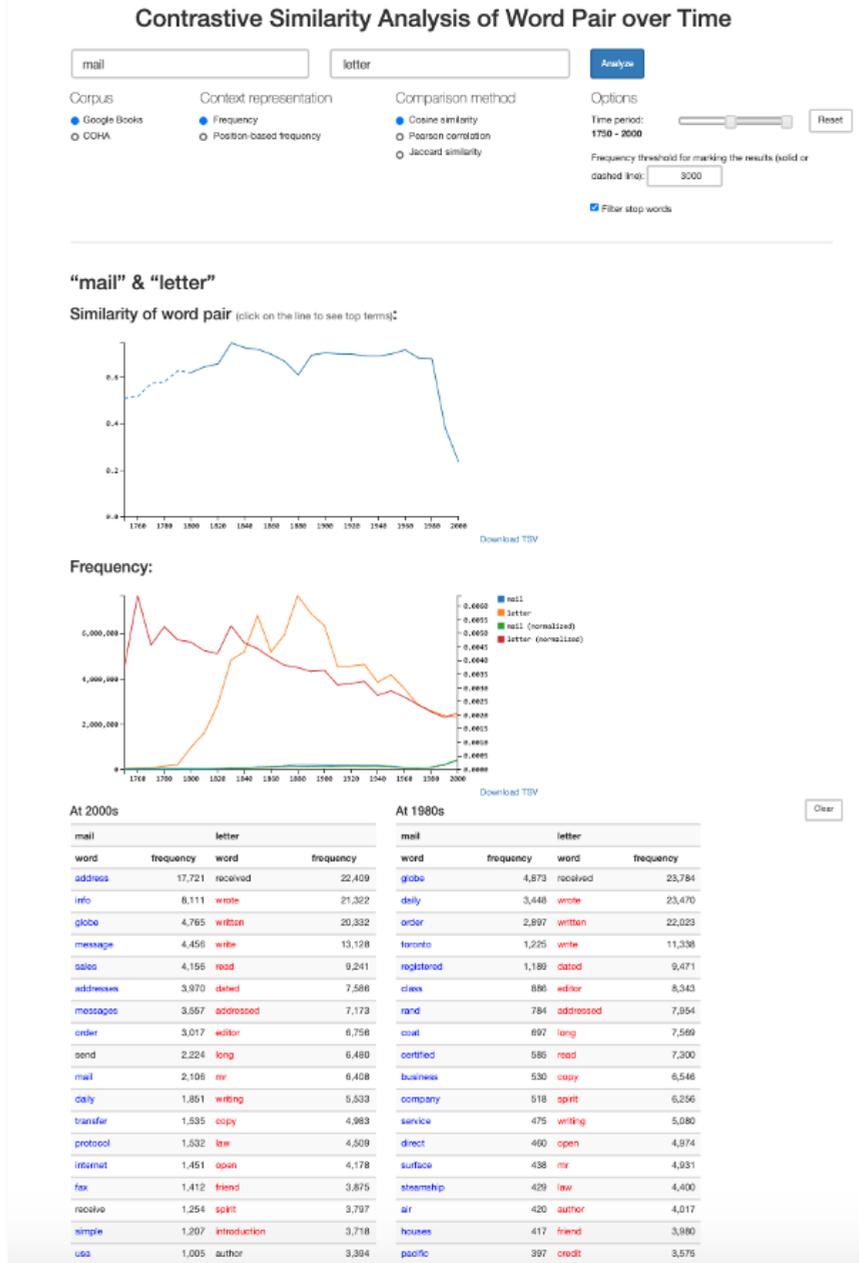


Figure 10.2: Snapshot of an example output from the diachronic conceptual change analysis system for comparing the words *mail* and *letter* that displays their similarity plot over time, frequency plots and the contrasted lists of top-frequent context term at 1980s and 2000s (image captured on 2021-01-24).

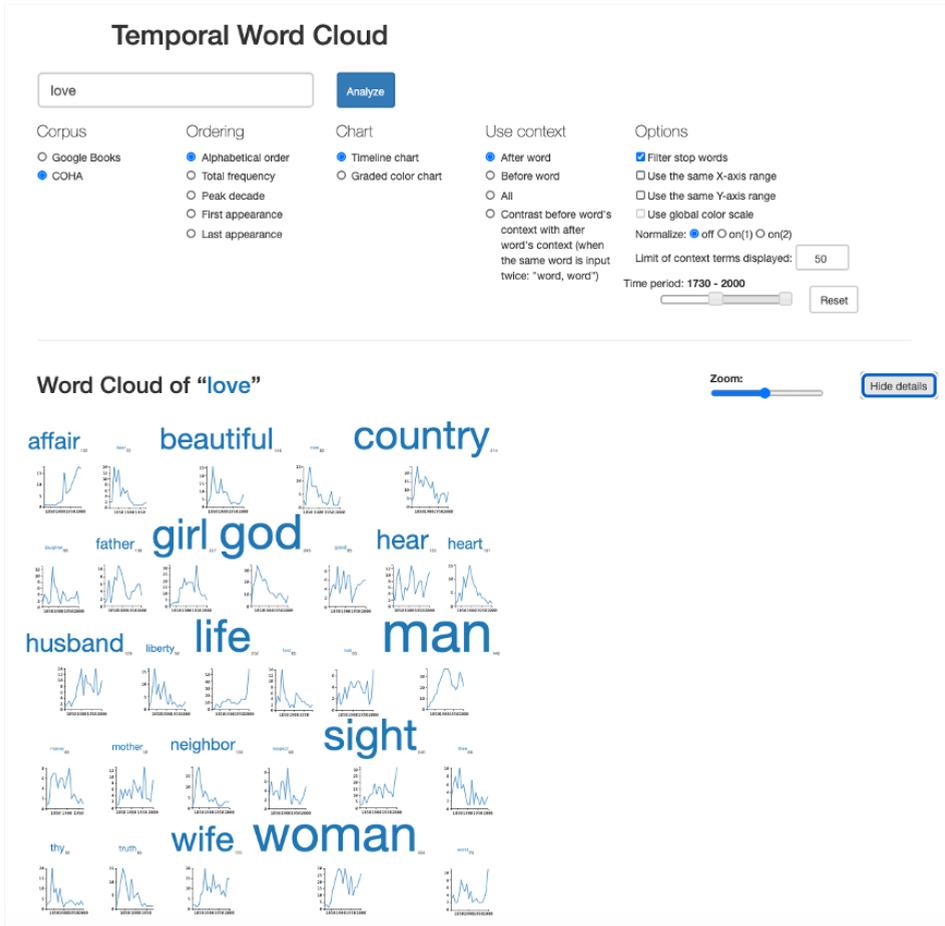


Figure 10.3: Snapshot of an example output from the diachronic conceptual change analysis system which shows a part of temporal term cloud of context words for the input word *love* (image captured on 2021-01-24).

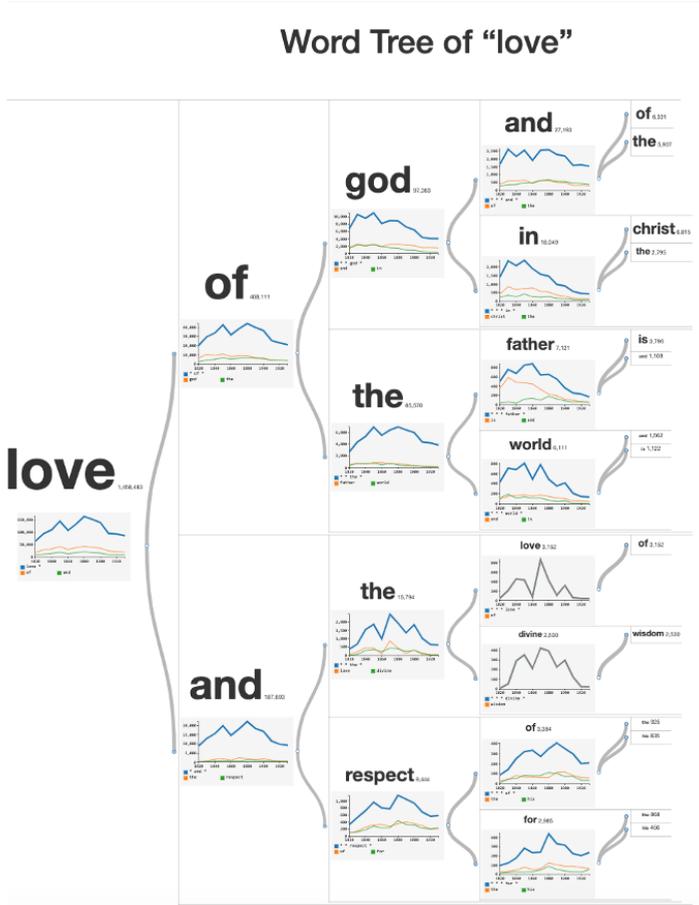


Figure 10.4: Snapshot of an example output from the diachronic conceptual change analysis system which shows a part of temporal term tree for the input word *love* (image captured on 2021-01-24).

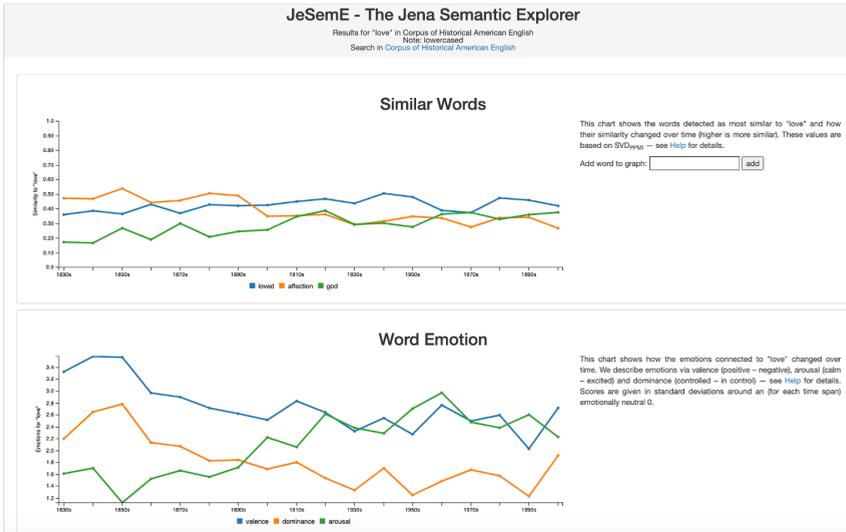


Figure 10.5: Snapshot of an example output from JeSemE for the word *love* which shows the similarity of similar words to the input word and associated emotions over time (image captured on 2021-01-24).

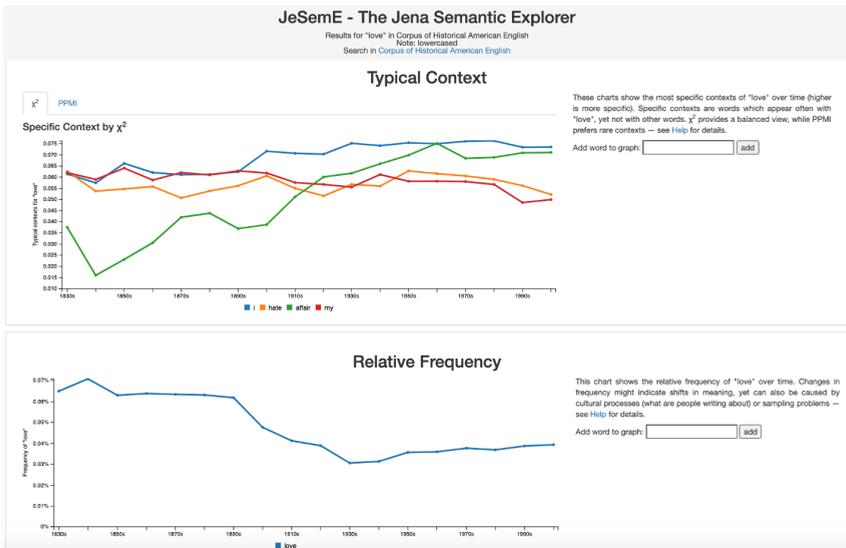


Figure 10.6: Snapshot of an example output from JeSemE for the word *love* which shows specific context words and relative word frequency over time (image captured on 2021-01-24).

1.3 Dictionary-based approaches

Theron & Fontanillo (2015) demonstrate an interactive visual tool for advanced analysis of the data in Spanish historical dictionaries. Their approach is unique as they utilize different editions of Spanish language dictionaries over time: the 1780, 1817, 1884, 1925, 1992 and 2001 editions provided by the Royal Spanish Academy. In this method the dictionary editions are arranged in a matrix in columns (right to left in chronological order), while the meanings of a word are placed on the rows (top to bottom in ascending order). Lines are drawn to connect related meanings across time, with the connection computed using NIST or BLEU metrics (Zhang et al. 2004), which are frequently utilized in evaluating machine translation or summarization accuracy. Starting from the most recent dictionary, a particular meaning is connected to its closest meaning in the previous dictionary; if there is nothing that satisfies the predefined similarity threshold, then the procedure is repeated with the older dictionary. Connecting lines can have branches in cases of bifurcation or merging of meanings. Theron & Fontanillo (2015) call the resulting diagrams *diachronlex diagrams*. The diagrams can be further improved by collapsing nearby lines with similar temporal patterns or by simplifying branches. Furthermore, users with editing rights can annotate meanings or change their associations.

1.4 Systems for analyzing lexical replacement

Mazeika et al. (2011) focus on semantically similar entities from different time periods to provide visual support in analyzing lexical replacement, particularly for entities. They extracted named entities from the Yet Another Great Anthology (YAGO) database to provide a visual analytics tool to analyze the evolution of named entities in the New York Times Annotated Corpus. Name changes are not tracked, but the tool offers a visualization of the evolution of an entity in relation to other entities.

Investigation of lexical replacement and temporal analogy are also possible in the aspect-based temporal analog retrieval (ATAR) system (Zhang et al. 2019) which uses perceptrons to compute transformations between present and past vector spaces trained on the present and past participles, respectively. For a user query (e.g., *euro*) and its defining aspect or sense (e.g., *currency*) a list of analogical terms is produced based on the analysis of its past document collection, together with an extracted representative sentence for each output term (see Figure 10.7). The sentence provides a typical context in which an analogous word was used in the past, using the output style dubbed KWECT (keyword in exemplar context at time), similar to the traditional KWIC (keyword in context) style.

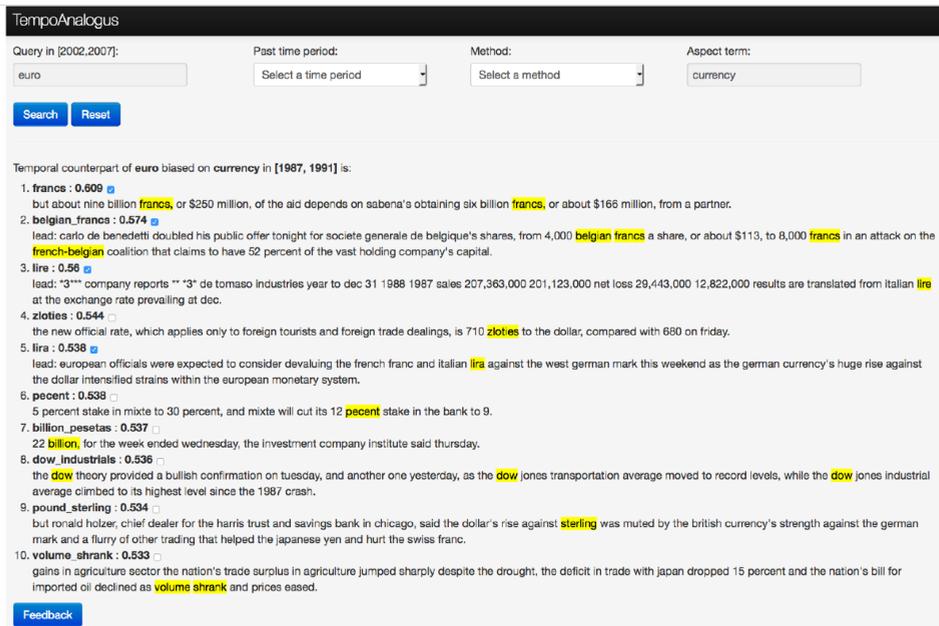


Figure 10.7: Snapshot of an example output from aspect-based temporal analog retrieval system (ATAR) for the query word *euro* under the aspect of currency (image captured on 2021-01-24).

1.5 Summary and observations

Visualization and analysis of diachronic conceptual change belong to an emerging and powerful research field of interactive visualization for computational linguistics (Collins et al. 2008). Its purpose is to let users understand models of language and their abstract representations, and to visually uncover patterns in language. In view of the inherent complexity of tracking word senses and understanding their shifts over time, we expect an increase in the availability and popularity of visual, interactive approaches to diachronic corpora. A similar conclusion was reached by Tang (2018), who considers further use of data visualization techniques to prove hypotheses as one of the core issues to be solved. Below we list several approaches and future directions.

- Easy to use and attractive services should be built to allow non-professional users to freely investigate histories of any words they are interested in, and to appreciate their language. Entertaining visualizations and explanations would attract many interested visitors. Automatically

generating accessible explanations and suggesting interesting words to explore would be beneficial (see Section 2.4 for the discussion of example query recommendation techniques).

- It is often difficult to precisely determine the exact time of sense change, let alone accurately determine the nature of the change. Sometimes several conflicting conclusions or hypotheses can appear simultaneously valid, prompting scientists and professionals to look at the results from different angles and use different datasets as well as visualization techniques. Hence, frameworks providing multiple views or analysis angles, and using parallel datasets, should be especially useful (e.g. Jatowt et al. 2018, Kalouli et al. 2019, Hellrich et al. 2018). Related to this is the role of a word's frequency over time as calculated for the particular corpus used in the analysis. This can work as a confidence measure for observed semantic change. It is also helpful to contrast semantic analysis results of similar or related words, or words associated with the same concept (e.g. Jatowt et al. 2018).
- Previous analyses (Hamilton et al. 2016b, Pagel et al. 2007, Lieberman et al. 2007) have revealed differing average degrees of change of very frequent and less frequent words, although the influence of frequency on the semantic change degree was later found to be smaller than expected (Dubossarsky et al. 2017). Investigating historical data such as the degree of the target word's change over time could be referenced to the average degree of change of words in the same frequency bin as the bin of the target word.⁷
- When evaluating these systems, it is common practice to investigate particular cases to determine whether the results support expectations or existing knowledge about diachronic conceptual change. The same type of evaluation is done with general purpose information visualization systems (Carpendale 2008). The investigation can extend to checking whether novel kinds of information can be obtained. We think that more systematic and extensive evaluation frameworks should be applied to determine whether new systems really help to find changes. New systems should also be compared with other systems to determine their strengths as well as their weaknesses. The unsupervised lexical semantic change detection task (Schlechtweg et al. 2020) at SemEval-2020 is an example of a standardization initiative to evaluate algorithms based on a shared dataset and evaluation metrics.

⁷As, for example, in Jatowt et al. (2018).

- Finally, systems designed primarily for visualization and interactive analysis of syntactic change in historical linguistics such as HistoBankVis (Schätzle & Butt 2020) and ParHistVis (Kalouli et al. 2019) can provide novel insights for building interactive tools for semantic analysis. Comprehensive visual analytics frameworks for historical linguistics could embrace both semantic and syntactic change and their interrelation to enable the comprehensive study of change in linguistic phenomena over time.

2 Applications of computational analysis of semantic change and lexical replacement

The remainder of this chapter deals with several applications of approaches designed for computational modeling, analysis of semantic change, and lexical replacement. In particular, we focus on the use of the technologies outside the core objective of analyzing the change in word meaning per se, that is, aiming to reveal knowledge of a word's history. The techniques developed for diachronic conceptual analysis and the findings from their use can be beneficial for various applications and services that deal with old texts or documents in long-term document archives. We discuss some current as well as promising future applications, especially within computer and information sciences. We note that our overview is in no way definitive and exhaustive, as many computational processes applied to old texts could benefit from the techniques and discoveries in the field of computational approaches to semantic change.

2.1 Semantics-aware culturomics

Michel et al. (2011) coined the term *culturomics* – the study of cultural and historical phenomena based on large textual data. In their seminal paper the authors demonstrate changes in the frequencies of selected words that reveal high-level cultural or abstract change occurring in a society over time. As one example, they contrast the popularity plots of the words *men* and *women* to provide evidence for the increasing social role and emancipation of women in recent decades. The work by Michel et al. (2011) inspired many similar studies using the Google Books Ngrams datasets (e.g. Acerbi et al. 2013, Bentley et al. 2014, Pechenick et al. 2015, Iliev et al. 2016) or other diachronic corpora (e.g. Hills & Adelman 2015, Sneffjella et al. 2018, Kutuzov et al. 2017), as well as other languages (e.g. Viklund & Borin 2016, Hengchen et al. 2019, Marjanen et al. 2020).

While the approach of culturomics relies on investigating change in the usage intensity of words, and especially the data around their first appearances, it

should be extended to also consider fluctuations in the meaning that words represent. Tahmasebi & Risse (2017) discuss the utility of automatic sense detection for archive users and digital humanities research. They propose a sense-based approach to capture changes related to the usage and culture of a word. We also believe that correctly recognized shifts in term meanings should be accounted for in order to produce reliable data in any cultural study based on the analysis of the aggregate statistics obtained from term occurrences and term relations over time. Fridlund et al. (2019) attempted to estimate the number of certain types of events (in particular, terror attacks) in the past, portraying the societal responses based on a diachronic document collection (e.g., news archives spanning a longer time period). Inspired by their work, we take as an example a political demonstration. Simply issuing direct queries such as *political demonstration* to a search engine indexing a document archive would be insufficient, and would likely produce inaccurate results. This is because the term *demonstration* and its close derivatives were probably not used in the past to indicate a public show of feelings in support of or against something, or at least one cannot assume this was always the case. Past meanings of *demonstration*, *political demonstration* or *public demonstration* probably did not exactly correspond to their contemporary meanings. Many events that would currently be regarded and labeled as such would be missed during the data collection. In addition, some false negatives can be identified if one does not properly take into consideration diachronic semantic change.⁸ On the other hand, equipped with knowledge of actual terms used in the past and accounting for semantic variations in the known term and related ones, the researchers could more accurately collect data and more credibly represent the true frequency of target types of incidents over time. In general, semantic change awareness should improve trustworthy, precise collection building and, by extension, culturomics studies in general.

While the approaches were usually developed for long-term sense tracking and analysis (over decades or centuries), recently researchers have also focused on analyzing diachronic change over shorter time spans such as a few years (Dodds et al. 2011, Danescu-Niculescu-Mizil et al. 2013, Eisenstein et al. 2014, Goel et al. 2016, Del Tredici & Fernández 2018). Short-term change is intensified nowadays due to the popularity of the Web, the high dynamics of social media, and the dramatic increase in the speed of information exchange brought about by novel communication and Web technologies. These all mean that lexical change can

⁸In our example, the sense of *demonstration* as a ‘public show of feeling by a number of persons in support of some political or social cause’ dates back to 1839 (<https://www.etymonline.com/search?q=demonstration>).

now materialize in much shorter time frames than in the past. The technologies developed for semantic change analysis over long-span diachronic corpora could be adapted for cultural studies drawing on temporal and social aspects of social media.

2.2 Natural language processing

We expect that many text processing tasks ranging from POS tagging, grammatical dependency detection, semantic role labeling, named entity extraction and linking, and sentiment analysis to language inference could benefit from correct estimation of word senses present in past documents. Currently, post-OCR error detection and correction are among the most common text processing procedures applied to old texts. Automatically detecting and correcting errors in OCR-processed historical texts (Chiron et al. 2017) could also benefit from the research on diachronic conceptual change. This is because knowledge of a word sense that is expected at a given position in a text should help to determine whether the word at that position is erroneous or not (especially in the case of so-called “real-word errors” which are misspellings that result in valid words). This process should also help to generate the most plausible substitutes if the word is deemed an error.

2.3 Document analysis and understanding

We expect that knowledge of change in the diachronic semantics of words constituting a document created at a certain time in the past should help in the analysis of the document (Tahmasebi 2013). Below we discuss three examples.

2.3.1 Providing temporal context to support analysis of past documents

Jatowt et al. (2019) proposed viewing a past document through the lens of its time by utilizing knowledge of the change in the frequency and semantics of words contained in the document (e.g., based on a large diachronic corpus such as Google Books Ngrams). The document in context of its time (DICT) visualization style lets users (e.g., professionals such as historians or other humanities researchers studying old literature) observe whether the words in the document were frequently used or were rather rare at the time of the document’s creation. This helps to locate neologisms and archaisms used by the document’s author. Furthermore, words that have changed their meaning when compared to a given specified date (e.g., a present time) are identified in text. Together, these functions

let users better understand the writing style of the author, which can be important in literature studies, and let them “connect” the document to the words and their senses commonly used at the time the document was created.

2.3.2 Comprehensibility of past texts

Another possible application is to support comprehensibility of past documents. Methods designed to estimate the reading difficulty of past documents could then be incorporated into archival retrieval engines and recommendation systems, so that relevant past texts are provided that current users can understand. Many of the methods described in this book could then be useful, as awareness of changes in word meaning over time could lead to increase the ease of reading and comprehension, as suggested by Tahmasebi & Risse (2013). One application would be to design extensions to traditional readability indexes to cover the additional difficulty caused by the semantic change. Other examples of initiatives in this direction are highlighting words in old texts that have undergone considerable change, as suggested by Jatowt et al. (2019), and/or clarifying their actual senses to improve comprehension by the average reader. The work of Tran et al. (2015) is related to this idea of COMPREHENSION-FOCUSED DOCUMENT ENRICHMENT. They propose recontextualizing past texts by enriching them with explanatory content extracted from Wikipedia, although Wikipedia does not explicitly focus on diachronic conceptual change.

2.3.3 Diachronic text evaluation

Diachronic conceptual change detection and exploration can also be applied to support the date of origin detection. This task is also called DIACHRONIC TEXT EVALUATION (DTE) (Popescu & Strapparava 2015). Many of the DTE solutions rely on information about word occurrence in the past, with the underlying hypothesis that if a document contains many words that were common at a given time point in the past, it is likely that it was created/published at that time point. This is especially the case if the words were rarely used at other time points (see, e.g. Kanhabua & Nørvgå 2009, Chambers 2012, Szymanski & Lynch 2015, Jatowt & Campos 2017). Including information on diachronic conceptual change could further improve the performance of DTE, as demonstrated by Frermann & Lapata (2016) through task-based evaluation of a Bayesian model of diachronic meaning change. This is because additional information on a word’s sense (or the probability distribution over its known senses) can be utilized alongside the frequency-based signal to more precisely determine age.

2.4 Information access and recommendation

2.4.1 Query suggestion and content ranking

Semantic change detection techniques can enhance information access and retrieval for users of digital libraries and digital archives. For example, at the query-level, effective query suggestion and correction techniques can be provided based on the computation of across-time analogies and present-to-past semantic relations. These could help users who lack the specific vocabulary of things common in that era to select appropriate query terms. Berberich et al. (2009) and Holzmann et al. (2012) discussed direct application of a method for finding analogical entities across time to information retrieval, mainly for query suggestions.

Recently, some work has aimed at across-time content retrieval and matching to enable novel information access approaches in news archives. For example, semantic term matching was used for extracting and summarizing comparative sentences (Duan & Jatowt 2019), computing temporal analogies (Zhang et al. 2015, Szymanski 2017) or estimating the contemporary relevance of past news articles, defined as the degree of the utility and attractiveness of old news articles to present-day users (Sato et al. 2019). Morsy & Karypis (2016) also propose using information about language change in document similarity computation.

2.4.2 Recommending words with interesting change history

As mentioned earlier, etymological knowledge is not only interesting to professionals such as linguists, historians, or librarians, but also to the wider public. Educators could use it to make students aware of etymological developments and arouse their interest in learning about language and history. Computational approaches and particularly online interactive systems could help to further disseminate knowledge of word etymologies. For such systems to be effective and attractive, it would be beneficial to recommend interesting words to be explored by non-professional users. Explanation of past meanings of words like *gay* or *nice*, for example, tends to surprise lay users who are not aware of them, and triggers questions on the reasons for the change. Existing systems for exploring diachronic conceptual change require users to provide words as the input. As average users may not know what words to search for, recommending sample queries to explore and learn about semantic change could be a useful option to attract or entertain users. Unique or specific input words could be recommended based on the shapes of their self-similarity plots over time (e.g., words that retained stable senses over a long time, or that underwent significant semantic shifts within short time frames) (Jatowt et al. 2018). Finally, suggesting words

that show interesting or unexpected semantic changes could be used to provide attractive content on specialized websites, or even might inspire new books directed at the average reader.

2.5 Temporal summarization and trend detection in domain-specific diachronic collections

Detecting shifts in word senses can also be limited to specific domains such as scholarly or legal documents. For example, Degaetano-Ortlieb & Strötgen (2017) analyze differences in frequency, meanings and the underlying temporal scopes of temporal expressions used in scientific writing from 1665 to 2007.

From an application viewpoint, semantic analysis of specialized terminology could help in detecting emerging trends (Dridi et al. 2019) and in summarizing entire domain-specific collections (Mohd Pozi et al. 2020). For example, Dridi et al. (2019) propose detecting emerging trends in scholarly publication collections in computer science and bioinformatics. Rather than employing citation analysis or straightforward frequency-based trend assessment, as has been usually done, the authors use temporal word embeddings to observe shifts in scientific language over several decades. A simple improvement of this approach would be to use contextualized embedding models pre-trained on domain-specific collections such as SciBert, which was trained on scholarly corpora (Beltagy et al. 2019). Another option is to consider specialized term extraction techniques such as ones based on recognizing meaning shifts between general and domain-specific language (Hätty et al. 2019).

A further extension is debiasing semantic change of analyzed words by considering the overall change direction of the collection. Such *temporal normalization in domain-constrained collections* would remove the overall, average drift that the collection underwent over time based on the evolution trajectory of the words studied. This will help to better represent the specific semantic change of these words. Techniques for gender-specific and other kinds of debiasing of word embeddings could be adapted (Bolukbasi et al. 2016, Kaneko & Bollegala 2019).

In general, change detection, temporal summarization, emerging trend detection, and other similar tasks in domain-constrained document collections are promising applications for the computational tools of diachronic semantic change detection and analysis.

3 Conclusions

Recently, we have witnessed many developments and advances in methods for recognizing, analyzing and understanding diachronic semantic change and lexical replacement. In this chapter, we discussed examples and applications of these methods besides the usual purpose of supporting research in historical linguistics by revealing unknown change and improving understanding of known change. We began by surveying representative visual systems that can help the wider public and non-professional users investigate evidences of semantic change and so learn about word etymology and evolution. Finally, we discussed the possibilities of enhancing and improving downstream applications in NLP, information retrieval, and related fields.

Acknowledgements

Thanks go to the Center for Digital Humanities (CDH), University of Gothenburg and to Kyoto University. A first version of this work was written while Nina Tahmasebi was employed at CDH. Adam Jatowt was employed by Kyoto University when the first version was completed. This work has in part been funded by the project *Towards Computational Lexical Semantic Change Detection* supported by the Swedish Research Council (2019–2022; contract 2018-01184), the project *South Asia as a linguistic area?* funded by the Swedish Research Council (2015–2020; contract 2014-00969), and *Nationella språkbanken* (the Swedish National Language Bank), a research infrastructure jointly financed by the Swedish Research Council (2018–2024; contract 2017-00626) and the 10 partner institutions. In addition we are grateful for the strong and enduring commitment that the University of Gothenburg, its Faculty of Humanities, and its Department of Swedish have shown to our R&D unit Språkbanken Text for almost 50 years. The authors would also like to acknowledge Dominik Schlechtweg, Haim Dubossarsky, and Simon Hengchen for fruitful discussions and insights. The authors wish to thank the Department of Literature, History of Ideas, and Religion at the University of Gothenburg for providing financial support for language editing services.

Abbreviations

ATAR	aspect-based temporal analog retrieval
IR	information retrieval
KWIC	keyword in context
LSA	latent semantic analysis
NLP	natural language processing
OCR	optical character recognition
PCA	principal component analysis
POS	part of speech
YAGO	Yet Another Great Anthology

References

- Acerbi, Alberto, Vasileios Lampos, Philip Garnett & R. Alexander Bentley. 2013. The expression of emotion in 20th century books. *PLOS ONE* 8(3). e59030.
- Beltagy, Iz, Kyle Lo & Arman Cohan. 2019. SciBERT: A pretrained language model for scientific text. *arXiv preprint arXiv:1903.10676*.
- Bentley, R. Alexander, Alberto Acerbi, Paul Ormerod & Vasileios Lampos. 2014. Books average previous decade of economic misery. *PLOS ONE* 9(1). e83147.
- Berberich, Klaus, Srikanta J. Bedathur, Mauro Sozio & Gerhard Weikum. 2009. *Bridging the terminology gap in Web archive search*. Paper presented at the Twelfth International Workshop on the Web and Databases (WebDB 2009). <http://webdb09.cse.buffalo.edu/papers/Paper20/webdb2009-final.pdf>.
- Bolukbasi, Tolga, Kai-Wei Chang, James Y. Zou, Venkatesh Saligrama & Adam T. Kalai. 2016. Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. In *Advances in neural information processing systems*, 4349–4357.
- Bradley, Margaret M. & 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.
- Carpendale, Sheelagh. 2008. Evaluating information visualizations. In Andreas Kerren, John T. Stasko, Jean-Daniel Fekete & Chris North (eds.), *Information Visualization: Human-centered issues and perspectives*, 19–45. Berlin: Springer.
- Chambers, Nathanael. 2012. Labeling documents with timestamps: Learning from their time expressions. In *ACL 2012*, 98–106.
- Chiron, Guillaume, Antoine Doucet, Mickaël Coustaty & Jean-Philippe Moreux. 2017. ICDAR2017 competition on post-OCR text correction. In *Document analysis and recognition (ICDAR), 2017*, vol. 1, 1423–1428.

- Collins, Christopher, Gerald Penn & Sheelagh Carpendale. 2008. Interactive visualization for computational linguistics. In *Proceedings of HLT 2008: Tutorial abstracts*, 6–6. ACL.
- Cox, Michael A.A. & Trevor F. Cox. 2008. Multidimensional scaling. In Chunhouh Chen, Wolfgang Karl Härdle & Antony Unwin (eds.), *Handbook of data visualization*, 315–347. Heidelberg: Springer.
- Danescu-Niculescu-Mizil, Cristian, Robert West, Dan Jurafsky, Jure Leskovec & Christopher Potts. 2013. No country for old members: User lifecycle and linguistic change in online communities. In *Proceedings of WWW 2013*, 307–318. ACM.
- Degaetano-Ortlieb, Stefania & Jannik Strötgen. 2017. Diachronic variation of temporal expressions in scientific writing through the lens of relative entropy. In *International Conference of the German Society for Computational Linguistics and Language Technology*, 259–275. Springer.
- Del Tredici, Marco & Raquel Fernández. 2018. The road to success: Assessing the fate of linguistic innovations in online communities. In *Proceedings of COLING 2018*, 1591–1603. ACL.
- Dodds, Peter Sheridan, Kameron Decker Harris, Catherine A. Bliss & Christopher M. Danforth. 2011. Temporal patterns of happiness and information in a global social network: Hedonometrics and Twitter. *PLOS ONE* 6(12). e26752.
- Dridi, Amna, Mohamed Medhat Gaber, R. Muhammad Atif Azad & Jagdev Bhogal. 2019. Leap2trend: A temporal word embedding approach for instant detection of emerging scientific trends. *IEEE Access* 7. 176414–176428.
- Duan, Yijun & Adam Jatowt. 2019. Across-time comparative summarization of news articles. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, 735–743.
- Dubossarsky, Haim, Daphna Weinshall & Eitan Grossman. 2017. Outta control: Laws of semantic change and inherent biases in word representation models. In *Proceedings of EMNLP 2017*, 1136–1145. Copenhagen: ACL. DOI: 10.18653/v1/D17-1118.
- Eisenstein, Jacob, Brendan O’Connor, Noah A. Smith & Eric P. Xing. 2014. Diffusion of lexical change in social media. *PLOS ONE* 9(11). e113114.
- Frermann, Lea & Mirella Lapata. 2016. A Bayesian model of diachronic meaning change. *Transactions of the ACL* 4. 31–45. DOI: 10.1162/tacl_a_00081.
- Fridlund, Mats, Leif-Jöran Olsson, Daniel Brodén & Lars Borin. 2019. Trawling for terrorists: A big data analysis of conceptual meanings and contexts in Swedish newspapers. In *5th International Workshop on Computational History, HistoInformatics@TPDL 2019* (CEUR Workshop Proceedings 2461), 30–39. http://ceur-ws.org/Vol-2461/paper_5.pdf.

- Goel, Rahul, Sandeep Soni, Naman Goyal, John Paparrizos, Hanna Wallach, Fernando Diaz & Jacob Eisenstein. 2016. The social dynamics of language change in online networks. In *International Conference on Social Informatics*, 41–57. Springer.
- Hamilton, William L., Jure Leskovec & Dan Jurafsky. 2016a. Cultural shift or linguistic drift? Comparing two computational measures of semantic change. In *Proceedings of EMNLP 2016*, 2116–2121. Austin: ACL. DOI: 10.18653/v1/D16-1229.
- Hamilton, William L., Jure Leskovec & Dan Jurafsky. 2016b. Diachronic word embeddings reveal statistical laws of semantic change. In *Proceedings of ACL 2016 (Volume 1: Long papers)*, 1489–1501. Berlin: ACL. DOI: 10.18653/v1/P16-1141.
- Hätty, Anna, Dominik Schlechtweg & Sabine Schulte im Walde. 2019. SURel: A gold standard for incorporating meaning shifts into term extraction. In *Proceedings of the Eighth Joint Conference on Lexical and Computational Semantics (*SEM 2019)*, 1–8.
- Hellrich, Johannes, Sven Buechel & Udo Hahn. 2018. JeSeme: A website for exploring diachronic changes in word meaning and emotion. *CoRR* abs/1807.04148. <http://arxiv.org/abs/1807.04148>.
- Hengchen, Simon, Ruben Ros & Jani Marjanen. 2019. A data-driven approach to the changing vocabulary of the *nation* in English, Dutch, Swedish and Finnish newspapers, 1750–1950. In *Proceedings of the Digital Humanities (DH) conference 2019*.
- Heylen, Kris, Dirk Speelman & Dirk Geeraerts. 2012. Looking at word meaning: An interactive visualization of semantic vector spaces for Dutch synsets. In *Proceedings of the EACL 2012 Joint Workshop of LINGVIS & UNCLH*, 16–24. ACL.
- Hills, Thomas T. & James S. Adelman. 2015. Recent evolution of learnability in American English from 1800 to 2000. *Cognition* 143. 87–92.
- Hilpert, Martin & Stefan Th. Gries. 2008. Assessing frequency changes in multi-stage diachronic corpora: Applications for historical corpus linguistics and the study of language acquisition. *Literary and Linguistic Computing* 24(4). 385–401.
- Hilpert, Martin & Florent Perek. 2015. Meaning change in a petri dish: Constructions, semantic vector spaces, and motion charts. *Linguistics Vanguard* 1(1). 339–350.
- Holzmann, Helge, Gerhard Gossen & Nina Tahmasebi. 2012. Fokas: Formerly known as, a search engine incorporating named entity evolution. In *Proceedings of COLING 2012: Demonstration papers*, 215–222.

- Iliev, Rumen, Joe Hoover, Morteza Dehghani & Robert Axelrod. 2016. Linguistic positivity in historical texts reflects dynamic environmental and psychological factors. *PNAS* 114(2). E7871–E7879.
- Jatowt, Adam & Marc Bron. 2016. HistoryComparator: Interactive across-time comparison in document archives. In *Proceedings of COLING 2016*, 84–88.
- Jatowt, Adam & Ricardo Campos. 2017. Interactive system for reasoning about document age. In *Proceedings of CIKM 2017*, 2471–2474. ACM.
- Jatowt, Adam, Ricardo Campos, Sourav S. Bhowmick & Antoine Doucet. 2019. Document in context of its time (DICT): Providing temporal context to support analysis of past documents. In *Proceedings of CIKM 2019*, 2869–2872. Beijing: ACM.
- Jatowt, Adam, Ricardo Campos, Sourav S. Bhowmick, Nina Tahmasebi & Antoine Doucet. 2018. Every word has its history: Interactive exploration and Visualization of word sense evolution. In *Proceedings of CIKM 2018*, 1899–1902. ACM.
- Jatowt, Adam & Kevin Duh. 2014. A framework for analyzing semantic change of words across time. In *Proceedings of Joint Conference on Digital Libraries (JCDL '14)*, 229–238. <http://dl.acm.org/citation.cfm?id=2740769.2740809>.
- Kalouli, Aikaterini-Lida, Rebecca Kehlbeck, Rita Sevastjanova, Katharina Kaiser, Georg A. Kaiser & Miriam Butt. 2019. ParHistvis: Visualization of parallel multilingual historical data. In *Proceedings of the 1st International Workshop on Computational Approaches to Historical Language Change*, 109–114.
- Kaneko, Masahiro & Danushka Bollegala. 2019. Gender-preserving debiasing for pre-trained word embeddings. *arXiv preprint arXiv:1906.00742*.
- Kanhabua, Nattiya & Kjetil Nørvåg. 2009. Using temporal language models for document dating. In *Machine learning and knowledge discovery in databases*, 738–741. Springer.
- Kenter, Tom, Melvin Wevers, Pim Huijnen & Maarten De Rijke. 2015. Ad hoc monitoring of vocabulary shifts over time. In *Proceedings of the 24th ACM international conference on information and knowledge management*, 1191–1200.
- Kim, Yoon, Yi-I Chiu, Kentaro Hanaki, Darshan Hegde & Slav Petrov. 2014. Temporal analysis of language through neural language models. In *Proceedings of the ACL 2014 workshop on language technologies and computational social science*, 61–65. Baltimore: ACL. DOI: 10.3115/v1/W14-2517.
- Kulkarni, Vivek, Rami Al-Rfou, Bryan Perozzi & Steven Skiena. 2015. Statistically significant detection of linguistic change. In *Proceedings of the 24th international conference on the World Wide Web*, 625–635. Florence: ACM. DOI: 10.1145/2736277.2741627.

- Kutuzov, Andrey, Erik Velldal & Lilja Øvrelid. 2017. Tracing armed conflicts with diachronic word embedding models. In *Proceedings of the events and stories in the news workshop*, 31–36. ACL.
- Lieberman, Erez, Jean-Baptiste Michel, Joe Jackson, Tina Tang & Martin A. Nowak. 2007. Quantifying the evolutionary dynamics of language. *Nature* 449(7163). 713.
- Marjanen, Jani, Jussi Kurunmäki, Lidia Pivovarova & Elaine Zosa. 2020. The expansion of isms, 1820–1917: Data-driven analysis of political language in digitized newspaper collections. *Journal of Data Mining & Digital Humanities*. <https://jdmhd.episciences.org/6728>.
- Martinez-Ortiz, Carlos, Tom Kenter, Melvin Wevers, Pim Huijnen, Jaap Verheul & Joris Van Eijnatten. 2016. Design and implementation of ShiCo: Visualising shifting concepts over time. In *HistoInformatics 2016*, vol. 1632, 11–19.
- Mazeika, Arturas, Tomasz Tylanda & Gerhard Weikum. 2011. Entity timelines: Visual analytics and named entity evolution. In *Proceedings of the 20th ACM international conference on information and knowledge management*, 2585–2588. ACM.
- Michel, Jean-Baptiste, 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 & Erez Lieberman Aiden. 2011. Quantitative analysis of culture using millions of digitized books. *Science* 331. 176–182.
- Mohd Pozi, Muhammad Syafiq, Adam Jatowt & Yukiko Kawai. 2020. Temporal summarization of scholarly paper collections by semantic change estimation: Case study of CORD-19 dataset. In *JCDL 2020*, 459–460. Virtual Event, China: ACM.
- Morsy, Sara & George Karypis. 2016. Accounting for language changes over time in document similarity search. *ACM Transactions on Information Systems* 35(1).
- Odijk, Daan, Giuseppe Santucci, Maarten de Rijke, Marco Angelini, Guido Lorenzo Granato, et al. 2014. Time-aware exploratory search: Exploring word meaning through time. In *Sigir 2012 workshop on time-aware information access*.
- Pagel, Mark, Quentin D. Atkinson & Andrew Meade. 2007. Frequency of word-use predicts rates of lexical evolution throughout Indo-European history. *Nature* 449(7163). 717.
- Pechenick, Eitan Adam, Christopher M. Danforth & Peter Sheridan Dodds. 2015. Characterizing the Google Books corpus: Strong limits to inferences of socio-cultural and linguistic evolution. *PLOS ONE* 10(10). e0137041.

- Popescu, Octavian & Carlo Strapparava. 2015. Semeval 2015, task 7: Diachronic text evaluation. In *SemEval 2015*, 870–878.
- Rohrdantz, Christian, Annette Hautli, Thomas Mayer, Miriam Butt, Daniel A. Keim & Frans Plank. 2011. Towards tracking semantic change by visual analytics. In *Proceedings of HLT 2011: Short papers*, 305–310. ACL.
- Sato, Mari, Adam Jatowt & Masatoshi Yoshikawa. 2019. Present relatedness estimation of archival documents using external knowledge base. In *DEIM forum 2019*, 1–5. <https://db-event.jpn.org/deim2019/post/papers/275.pdf>.
- Schätzle, Christin & Miriam Butt. 2020. Visual analytics for historical linguistics: Opportunities and challenges. working paper or preprint. <https://hal.inria.fr/hal-02914284>.
- Schlechtweg, Dominik, Barbara McGillivray, Simon Hengchen, Haim Dubossarsky & Nina Tahmasebi. 2020. SemEval-2020 task 1: Unsupervised lexical semantic change detection. In *Proceedings of SemEval 2020*, 1–23. Barcelona: ACL. <https://www.aclweb.org/anthology/2020.semeval-1.1>.
- Schütze, Hinrich. 1998. Automatic word sense discrimination. *Computational Linguistics* 24(1). 97–123.
- Sneffjella, Bryor, Michel Génereaux & Victor Kuperman. 2018. Historical evolution of concrete and abstract language revisited. *Behavior Research Methods* online first. 1–13. DOI: 10.3758/s13428-018-1071-2.
- Szymanski, Terrence. 2017. Temporal word analogies: Identifying lexical replacement with diachronic word embeddings. In *Proceedings of ACL 2017*. ACL.
- Szymanski, Terrence & Gerard Lynch. 2015. UCD: Diachronic text classification with character, word, and syntactic n-grams. In *SemEval 2015*, 879–883.
- Tahmasebi, Nina. 2013. *Models and algorithms for automatic detection of language evolution*. Gottfried Wilhelm Leibniz Universität Hannover. (Doctoral dissertation). <http://edok01.tib.uni-hannover.de/edoks/e01dh13/771705034.pdf>.
- Tahmasebi, Nina & Thomas Risse. 2013. The role of language evolution in digital archives. In *SDA 2013*, vol. 1091, 16–27. <http://ceur-ws.org/Vol-1091/paper2.pdf>.
- Tahmasebi, Nina & Thomas Risse. 2017. On the uses of word sense change for research in the digital humanities. In *TPDL 2017*, 246–257. Berlin: Springer.
- Tang, Xuri. 2018. Survey paper: A state-of-the-art of semantic change computation. *Natural Language Engineering* 24(5). 649–676.
- Theron, Roberto & Laura Fontanillo. 2015. Diachronic-information visualization in historical dictionaries. *Information Visualization* 14(2). 111–136.
- Tran, Nam Khanh, Andrea Ceroni, Nattiya Kanhabua & Claudia Niederée. 2015. Back to the past: Supporting interpretations of forgotten stories by time-aware re-contextualization. In *Proceedings of WSDM 2015*, 339–348. Shanghai: ACM. DOI: 10.1145/2684822.2685315.

- van der Maaten, Laurens & Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of Machine Learning Research* 9. 2579–2605.
- van Eijnatten, Joris, Toine Pieters & Jaap Verheul. 2014. Using texcavator to map public discourse. *Tijdschrift voor Tijdschriftstudies* 35. 59–65.
- Viklund, Jon & Lars Borin. 2016. How can big data help us study rhetorical history? In *Selected Papers from the CLARIN Annual Conference 2015*, 79–93. Linköping: LiU EP. <http://www.ep.liu.se/ecp/123/007/ecp15123007.pdf>.
- Xu, Zaikun & Fabio Crestani. 2017. Temporal semantic analysis and visualization of words. In *IIR*, 52–62.
- Zhang, Yating, Adam Jatowt, S. Sourav Bhowmick & Yuji Matsumoto. 2019. ATAR: Aspect-based temporal analog retrieval system for document archives. In *WSDM2019*. Melbourne: ACM.
- Zhang, Yating, Adam Jatowt, Sourav Bhowmick & Katsumi Tanaka. 2015. Omnia mutantur, nihil interit: Connecting past with present by finding corresponding terms across time. In *Proceedings of ACL/IJCNLP 2015 (Volume 1: Long papers)*, 645–655. Beijing: ACL. DOI: 10.3115/v1/P15-1063.
- Zhang, Ying, Stephan Vogel & Alex Waibel. 2004. Interpreting bleu/nist scores: How much improvement do we need to have a better system? In *Proceedings of LREC 2004*. ELRA.

