Introducing Traveling Word Pairs in Historical Semantic Change: A Case Study of Privacy Words in 18th and 19th Century English

Thora Hagen[1,*], Erik Ketzan[2]

[1]Chair of Computational Philology, University of Würzburg, Germany
[2]Department of Digital Humanities, King's College London, UK

Abstract

In recent years, Lexical semantic change detection (LSCD) has become a central task of NLP. Because most studies in LSCD only consider the semantic change of words in isolation, in this paper, we propose a new direction for the analysis of semantic shifts: *traveling word pairs*. First, we introduce shift correlation to find pairs of words that semantically shift together in a similar fashion. Second, we propose word relation shift to analyze how the relationship between two words has changed over time. As a test case, we investigate the word *privacy* (and related words identified by a pre-existing dictionary), as an example of a word that has shifted semantics historically and remains vibrantly explored as a concept in contemporary humanistic discourse. We report that the term *privacy* in comparison shows relatively little change initially – with correlation analysis revealing more about how key terms surrounding *privacy* have shifted in tandem, and explore nuanced changes through word pair analysis, suggesting a shift toward concreteness in particular.

Keywords

semantic change, language models, computational semantics

1. Introduction

There is a growing body of recent work in computational approaches to lexical semantic change detection in historical linguistics research [27, 25, 22, 21, 33, 8, 4]. There is recognition that such research will eventually evolve from tracing semantic shifts in individual words or lexemes to larger groups of words [16, 13]. Eventually, the investigation of semantic historical shifts may uncover groups of words which shift in correlation, or larger, tectonic shifts of meaning which we have not yet perceived. As McGillivray writes, "Truly cutting-edge computational research in historical semantics should involve the development of innovative and impactful methods, which are build to answer questions relevant to humanists" [21].

In this paper, we explore two methods by investigating semantic shift around the term *privacy* as a test case, using a pre-existing dictionary of words relating to *privacy*. Privacy is selected, firstly, because it has had shifting semantics over time. The *Oxford English Dictionary*

CHR 2023: Computational Humanities Research Conference, December 6 – 8, 2023, Paris, France
*Corresponding author.
thora.hagen@uni-wuerzburg.de (T. Hagen); erik.ketzan@kcl.ac.uk (E. Ketzan)
0000-0002-3731-6397 (T. Hagen); 0000-0003-0264-2157 (E. Ketzan)

entry[1] begins with word sense 1 — ”The state or condition of being alone, undisturbed, or free from public attention, as a matter of choice or right; seclusion; freedom from interference or intrusion” — and of the seven remaining word senses and sub-word senses, six are ”obsolete” and/or ”rare”. In the long nineteenth century, ”privacy as a concept underwent important evolutions in society, both in its literature, which increasingly explored privacy as a theme, and in law, with the foundation of the recognition of privacy as a legal right separate from copyright or defamation” [15]. Privacy as a theory is so manifold that Tavani [34] proposes a taxonomy of theories of privacy with four categories: nonintrusion, seclusion, control, and limitation. A consistent, uniform theory of privacy has proven elusive [37], and this historical uncertainty over the definition of the lexical item and theories of privacy, as well as the voluminous discourse around privacy today (for instance, in the widely-used metric of differential privacy, which captures the increased risk to one’s privacy incurred by participating in a database [7]), support *privacy* as a word worth exploring for semantic change, and perhaps related words shifting semantics in some correlation.

To investigate the semantic shift of *privacy*, we employ contextualized word embeddings derived from a historical English BERT model, MacBERTh [20]. We first generate token embeddings for all of our dictionary words from sentences from CLMET [5], a corpus of historical English. We then follow the standard procedure of calculating semantic change for all dictionary words via centroid distances between token embedding sets of two time slices. Finally, we propose two novel methods to further break down semantic shifts in the context of *privacy*: correlated semantic word shifts and shifting word relations, with the latter being a reframing of the nearest neighbor approach.

We report that the term *privacy* shows relatively little semantic change by our models, a surprising result given the presumed manifold semantics of this term, and expand the investigation of *privacy* by detecting *traveling word pairs*; we report, for example, that *revealing* and *protecting* shift in tandem semantically, which could be due to both of these terms shifting from more figurative to more concrete over time.

2. Related Work on Semantic Change

Lexical semantic change can be defined as ”changes in ’sense,’ the concepts associated with expressions” [35]. This change in sense may either stem from the context in which a word is used over time (diachronic or temporal shift) or across different domains (synchronic or domain shift). This work focuses on the former. Change in sense may have different origins; amongst the different types are for example pejoration and amelioration (association of a term with a negative or positive meaning, respectively) or the narrowing/restriction of a term as opposed to the broadening/generalization of a concept [35].

The automatic detection of semantic shifts (also referred to as lexical semantic change detection or LSCD) has gained significant attention in NLP in recent years [33, 28, 9]. To detect diachronic change, a large text corpus for each time period of interest is assembled. There exist two types of computational approaches: type-based (word embeddings) and token-based (language models or contextualized embeddings). Using type-based methods usually implies

[1]https://www.oed.com/dictionary/privacy_n?tab=meaning_and_use

training one word embedding model (e.g. fastText, [23]) for each corpus slice, aligning the vector spaces (e.g. orthogonal Procrustes) and then comparing the word vectors across time slices. For the token-based approach, embeddings are extracted from a single language model (e.g. BERT, [6]) through the sentences the words of interest are used in. As an intermediate step, clustering may be used to identify different word usages before the comparison. The distances of the token vectors of one word for adjacent time slices then indicate change, for example through average pairwise distances (APD) or prototype/centroid distances [18].

Disadvantages of the token-based method are therefore currently the computational scalability [10] and the fact that token clusters are formed by word forms rather than semantics [18]. On the other hand, token-based methods enable the contextualization of each token through additional information from model pre-training and joint processing of the word and its sentence. This enables a deeper look into individual word senses and their shifts.

There are multiple ways in which change may be detected in vector space. Most commonly, words are analyzed in isolation: The distance of one word to itself in another time period or domain shows the stability of that word. The lower the distance, the higher is its stability, and in turn, less change is detected. In previous studies, distances were typically translated to a binary or graded scale for evaluation [28]. There are multiple evaluation datasets that capture this notion of word change through the creation of word usage graphs (WUG) [12, 29, 30]. However, there are also other ways in which change could be analyzed. While some studies working with word embeddings have analyzed change in terms of changing word associations in the past [11, 38], the trend towards using language models has streamlined the task towards single word analysis. Kutuzov, Øvrelid, Szymanski, and Velldal [16] argue that "most current studies stop after stating the simple fact that a semantic shift has occurred". Similarly, Hengchen, Tahmasebi, Schlechtweg, and Dubossarsky [13] criticize that "an emerging or evolving concept, almost by definition, will not be constrained to a single word. Rather, methods will probably have to be adapted to study a cluster of words".

The notion of words shifting semantically in groups was hypothesized by a handful of pre-computational linguists. In a 1931 treatise, Stern appealed to "find out to what degree [synonyms'] development runs in parallel lines and is conditioned by identical factors" [32]. In 1985, Lehrer expanded this hypothesis to different word relations: "semantically related words are more likely to undergo parallel semantic changes because of their semantic relationships. [...] If one word changes meaning, it will drag along other words in the domain" [19]. Computation now allows us to explore these notions in greater depth.

In this paper, we propose to pick up on these previous intuitions on change detection, that is analyzing the relationship of word pairs across time on a larger scale, using contextualized word embeddings.

3. Experiments

3.1. Corpus and Dictionary

We selected The Corpus of Late Modern English Texts, version 3.1, a genre- and time-balanced text corpus of late modern English (1710-1920) of about 34M word tokens [5]. Originally, the corpus consists of three balanced slices of 10-12M tokens each. To create more subcorpora, we split each of these slices in half to create six subcorpora based on date: 1710-1745, 1745-1780, 1780-1815, 1815-1850, 1850-1885, 1885-1920. The splits had an impact on the overall genre balance, although the proportions of genres for all slices between 1745 and 1920 stayed mostly consistent (fiction being about 45-65% of each sub-corpus, while the other genres have a share of 5-25% each). The 1710-1745 slice is a bit of an outlier in genre mix to keep in mind in the experiments below, being about 45% letters.

For a list of words relating to privacy, we selected Vasalou et al.'s Privacy Dictionary [37], a dictionary of 616 words and phrases for automated content analysis on privacy-related texts. Some benefits of this dictionary include its creation through a rigorous methodology, as well as its application to studies in a variety of fields, e.g. [1, 2, 36, 3], including humanistic research [24, 15].

3.2. General Approach

To obtain the contextualized word embeddings, we turn to the MacBERTh model implemented by Manjavacas and Fonteyn [20], which is entirely pre-trained on historical English. Their training corpus consists of texts of varying genres (literary works, articles, etc.) the size of about 3.9B tokens, covering a time span from 1473 to 1950.

As proposed by Hengchen, Tahmasebi, Schlechtweg, and Dubossarsky [13], to study a cluster of concepts, keywords may be chosen either manually or in a data-driven way. To construct a list of words around the concept of *privacy*, we followed both suggestions. We turned to 1) the Privacy Dictionary created by Vasalou, Gill, Mazanderani, Papoutsi, and Joinson [37], selecting only the unigrams and discarding phrases, and 2) we expanded this word list with 70 nearest neighbors of *privacy* from fastText embeddings trained on our first and last corpus slices each, to create a new dictionary of 454 words in total.

We adopt the method of using large language models (in our case MacBERTh) for detecting semantic change, i.e. retrieving token embeddings for all of our words of interest from the corpus, resulting in one embedding matrix for each of the six time slices. To obtain a stable representation, we only consider words in their respective time slice when they have more than 10 occurrences in that slice. For computational scaling purposes, we sample 100 occurrences for each word per slice.

For the embedding extraction, we use the average of the last 4 layers (9-12) in the MacBERTh model. To determine the degree of change for one word for adjacent time slices, we calculate the cosine distance of the mean of both matrices (centroid distance) with a sliding window over all slices:

$$d_{t_i} = \cos\left(\frac{1}{N_w^{t_i}} \sum_{j=1}^{N_w^{t_i}} \mathbf{w_j^{t_i}}, \frac{1}{N_w^{t_{i+1}}} \sum_{j=1}^{N_w^{t_{i+1}}} \mathbf{w_j^{t_{i+1}}}\right)$$

where $w_j^{t_i}, w_j^{t_i+1}$ are individual word vectors of two adjacent time slices, and N represent the counts of word vectors in the two time slices. The results are five distance measurements $d_{t_1}, ..., d_{t_5}$ per word.

We begin our analysis with the standard procedure of binary/graded word level semantic shift detection. First, we calculate mean and standard deviations for all five points in time

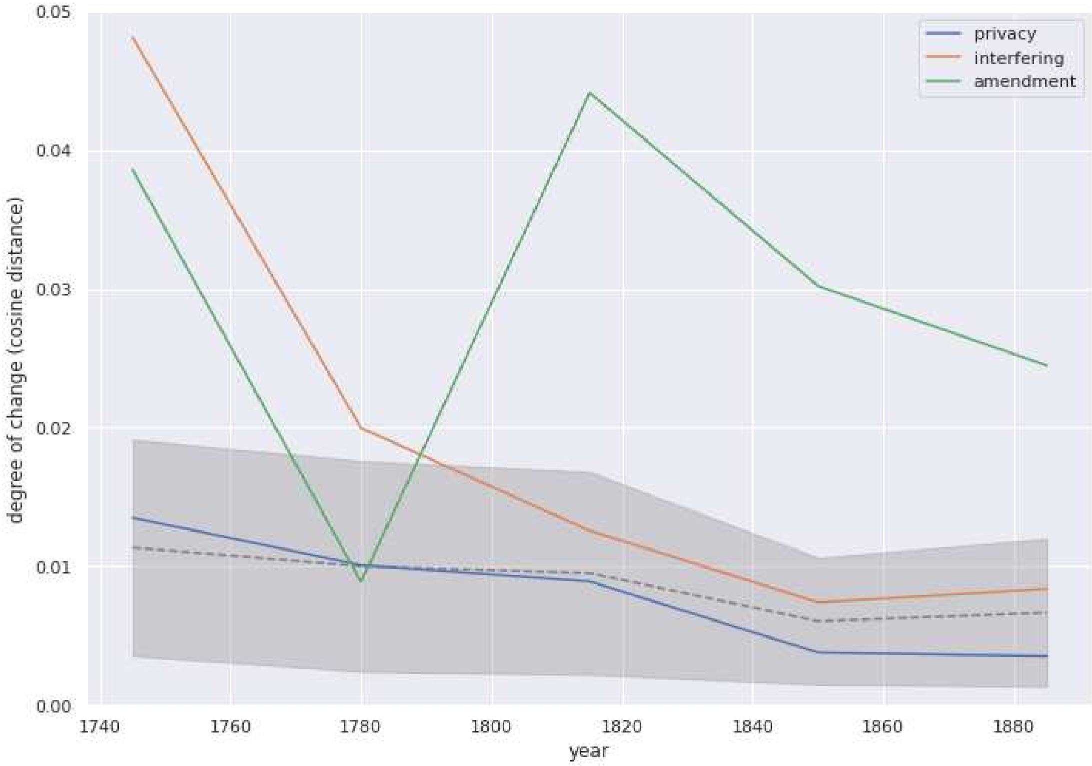

Figure 1: Semantic shift of *privacy* in comparison to two other, highly unstable words. The grey dotted line indicates mean cosine distances across all 454 words. The grey area indicates the standard deviation.

across all words as our point of reference:

$$\text{Mean}_{t_i} = \frac{1}{n}\sum_{j=1}^{n} d_j^{t_i}$$

$$\text{Standard Deviation}_{t_i} = \sqrt{\frac{1}{n}\sum_{j=1}^{n}(d_j^{t_i} - \text{Mean}_{t_i})^2}$$

where n represents the number of words in our dictionary (454). We then extract the words that show the highest distances across all measurements, additionally to *privacy*.

Our top five most unstable words are *interfering*, *amendment*, *suppressing*, *licence*, and *sensitive*. We find that *privacy* actually shows a stabilizing trend in our corpus across the 18th and 19th century (see Figure 1). In the following experiments, we will take a closer look especially at some of these more unstable words in addition to the focus word of our dictionary, *privacy*.

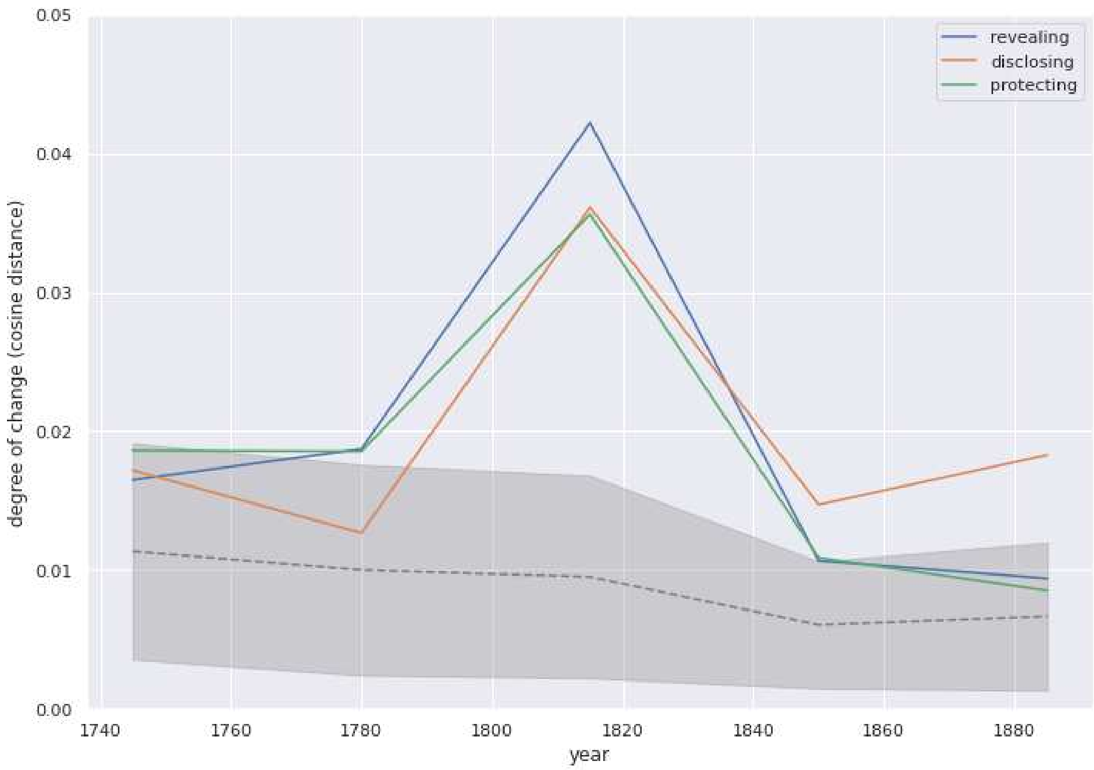

Figure 2: Correlating shifts in stability scores for three keywords. The grey dotted line indicates mean cosine distances across all 454 words. The grey area indicates the standard deviation.

3.3. Shift Correlation

For our first novel approach, we take up Kutuzov et al.'s proposal to identify "groups of words that shift together in correlated ways" [16]. Correlating shifts may reveal what we dub *traveling word pairs*: pairs of words which shift in meaning over time with a similar correlative pattern. We see computational change correlation as an opportunity to dive deeper into why meanings may have changed.

For this experiment, we investigate each pair of words in our list of 454, by calculating shift detecting correlations from the stability scores. This reveals which word pairs show a similar shift in stability scores across the 18th and 19th century. To calculate correlations, we use Pearson's ρ.[2] The top 10 pairs are reported in Table 1.

If we look for *traveling pairs*, we would need to identify words which, first of all, semantically shift significantly; otherwise, two words with no shift over time would correlate. We define high shift as a cosine distance of at least 0.02 (our highest mean plus standard deviation) in our list of measurements.

[2]One reviewer brought to our attention that Pearson's ρ might not be the ideal measurement for correlation in this case, as it cannot detect similarities between shifted time-series, and advised to use dynamic time-warping or similar instead. We would like to thank them and will certainly take up their suggestion for future work on this matter.

Table 1
Top 10 correlated word pairs with significant semantic shifts in our corpus.

word pair	Pearson's ρ
cabinet-predecessor	0.990
champion-prime	0.990
privileges-interfering	0.989
prior-unfair	0.989
protecting-revealing	0.986
sensitive-printer	0.986
privileges-shared	0.985
community-interfering	0.985
ensure-shared	0.984
privileges-embarrassment	0.984

Among our most correlated words that also show high semantic shift, we find *disclosing*, *protecting*, and *revealing* (see Figure 2). The three words seem to undergo a significant semantic change shortly past 1800. By looking at the sentences in each of the corpus slices, we find that *revealing* and *disclosing* show a similar trend. Both go through a shift from the abstract to the concrete, e.g. ”revealing meanwhile a flannel shirt”[3] or ”disclosing spaces of faint yet clearest blue”[4] as opposed to ”revealing his affairs to him”[5] and ”disclosing my Love and Esteem”.[6]

That same sentiment is mirrored in *protecting*: While in earlier text passages, it is values or acts that need protection (”protecting and securing the trade”,[7]) while the later texts mostly use *protecting* in a physical sense (”protecting that building from the fury of the populace”,[8] ”stretch out a protecting arm”.[9])

3.4. Shifting Word Relations

Previous studies on LSCD working with word embeddings have analyzed nearest neighbors across time, either for qualitative interpretation or evaluation [11] or to calculate another dimension of meaning shift (word level change in k nearest neighbors [38]). In contrast, we suggest the close tracking of a multitude of neighbor relations as another way to discover and quantitatively interpret change, which may create opportunities to identify factors that condition similar development, as Stern hypothesized [32]. We thus propose a second method that further explores the idea of using nearest neighbors: to investigate which words in a predefined dictionary become semantically closer or more distant to a keyword over time, to allow a richer interpretation a single word's semantic history over time.

In these experiments, we selected two keywords: *privacy*, as this is the primary topic of our investigation and dictionary, and *revealing*, because this word was salient in the results

[3]George Gissing, 1891
[4]Mary St Leger Kingsley (Lucas Malet), 1901
[5]Henry Fielding, 1751
[6]Gilbert Langley, 1745
[7]Samuel Johnson, 1740
[8]Henry Hunt, 1820
[9]Percy James Brebner, 1910

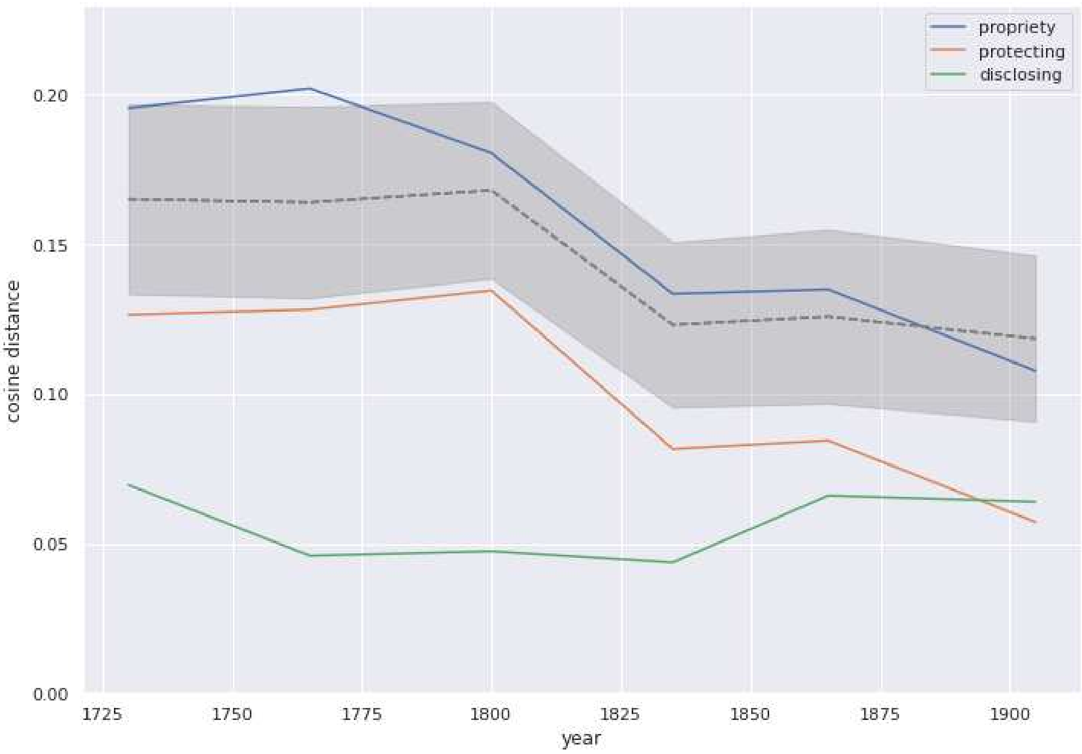

Figure 3: Cosine distances of selected dictionary words for keyword *revealing* over time. Grey dotted line indicates the mean cosine distance of all words for *revealing*. Grey area indicates the standard deviation.

of our experiment above. We then calculate cosine distances between each keyword and each dictionary word, in each time slice, and detect which of these pairs display the highest divergence across time. Note that in this experiment, we trace the shift of word relations over time as opposed to contrasting single words. This is why in Figure 1 and 2, the cosine distances are comparatively small (~0.01), because we compare the word to itself across different time periods. In Figure 3 and 4, we compare the distances of two different words over time, which means cosines distances are generally higher (~0.2).

The results for keyword *revealing* and selected words from the dictionary list are in Figure 3, which displays the progression of cosine distances between *revealing* and dictionary words. Amongst the words displaying the most semantic divergence is *propriety*, which we interpret as evidence that *revealing* and *propriety* were semantically distant in the 1700's and gradually became more close by the 1920's. Our method thus yields any number of comparisons in shift to a keyword, with a calculated "divergence" score. As a brief observation of *revealing* and *propriety* becoming semantically closer over time, one hypothesis is that both of these words shifted from more abstract to more concrete meanings (similar to the example of *revealing* and *disclosing* in our method above). As another observation, we report a stable relation between synonyms *revealing* and *disclosing* over time, whereas the fuzzy antonyms *revealing* and *protecting* become closer. This shows that even though stability scores may correlate, the distances

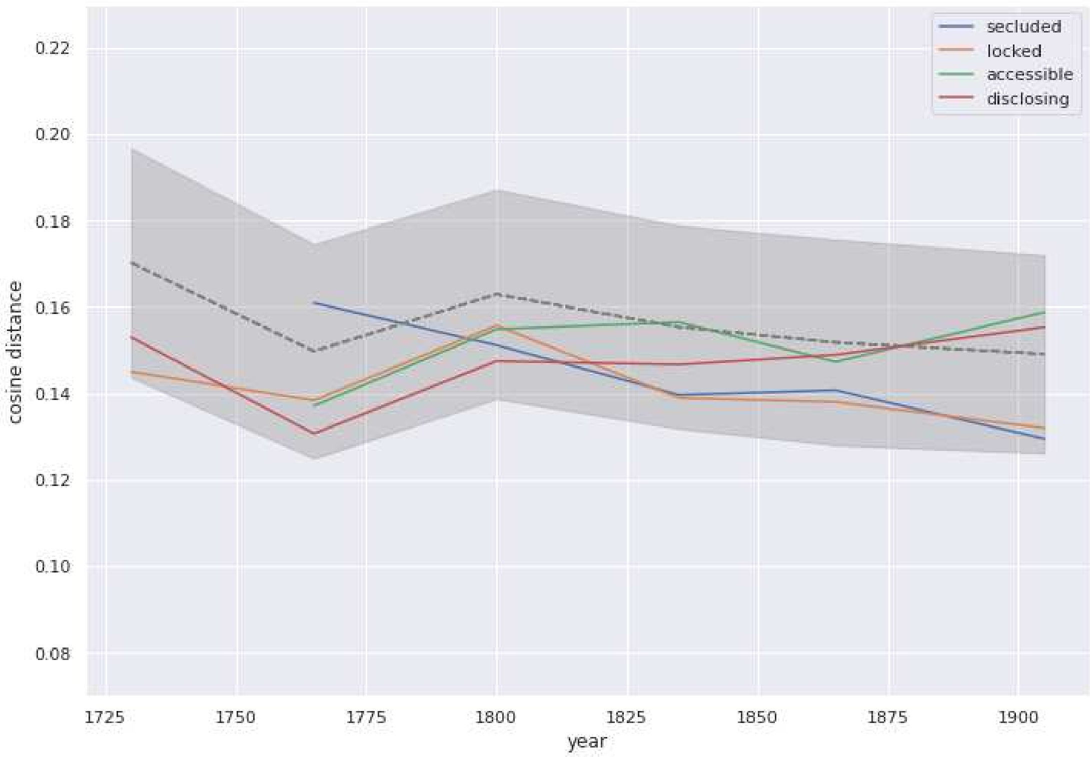

Figure 4: Cosine distances of selected dictionary words for keyword *privacy* over time. Grey dotted line indicates the mean cosine distance of all words for *privacy*. Grey area indicates the standard deviation.

between those words might not change in the same way. Further observations on these specific words would require more dedicated studies, but our method will provide such studies with new information about comparison of semantic shift.

We then apply our method to keyword *privacy* (see Figure 4). Although *privacy* was the starting point of our use case exploration, our method actually reveals only subtle information for the relation of our dictionary words to privacy, as *privacy* remains fairly stable in semantics, by our metrics, throughout the time periods (as shown in the General Approach). This alone is a tantalizing result that warrants further investigation, as *privacy* is widely considered to be extremely rich semantically and subject to a complex evolving history. A consistent, uniform theory of privacy has proven elusive [37], and taxonomies exist which group privacy by theory, e.g. Tavani [34], who suggests nonintrusion ("being let alone"), seclusion ("one's being secluded from others"), control ("one has privacy if and only if one has control over information about oneself"), and limitation ("one has privacy when information about oneself is limited or restricted in certain contexts"). Returning to our experiments, while the changes detected are very subtle, we observe that two of the most semantically distanced words from *privacy* — *accessible* and *disclosing* — refer to concepts of access, whereas, for example, *secluded* and *locked* draw a bit closer to the meaning of *privacy*.

Additionally, we suggest that we may be able to discover, through average changes in word

Table 2
10 most diverging dictionary words (in order) of *privacy* and *freedom*, calculated via subtracting the minimum from the maximum cosine distance across all points in time for all words. Color indicates direction (red for distancing, green for converging).

freedom	legitimacy, duchy, troubling, suppressing, incivility, lawsuit, posted, gossip, apologize, scrutinised
privacy	interdict, suppressing, apologize, ruling, comment, confided, interim, sensitive, community, indignity

relations, whether a term may have undergone semantic *broadening* or *narrowing*. If, on average, other concepts have smaller distances to a target word, we might discover *broadening*. Otherwise, if concepts grow more distant on average, we might discover *narrowing*. This can be done either by the mean changes (as indicated in all figures so far) or majority vote across all words of interest, or the top k most volatile dictionary words. Through the latter for example, among the top ten most volatile dictionary words relations of *privacy*, we find that all but one word have become more close (see Table 2), which could imply a generalization of the concept. On the other hand, the concept of *freedom* in our corpus for example has equal amounts of terms that grow closer or more distant. The same trends are mirrored in the mean distances. Through a targeted selection of sentiment keywords as dictionary words, this method could also be used to identify changes in connotation. Finally, both approaches could potentially be combined to discover correlated neighbor shifts.

4. Discussion and Future Work

We have presented two experimental approaches to expanding the investigation of historical semantic change by not only limiting the analysis to words in isolation, but instead taking changing word relationships into account. The word *privacy* has undergone relatively little change, but rather key issues surrounding the concept have. Through an analysis of stability correlations, we for example were able to show the tie between *revealing/disclosing* and *protecting*, as these concepts shift to a more concrete representation in our corpus. Through a closer analysis of the word relationships of *revealing*, we found that a shift towards *propriety* aligns with this hypothesis. Our interpretation also aligns with previous findings concerning concreteness shifts; Hills and Adelmann provide evidence that concrete words were gradually used more often over time [14], and Snefjella et al. add that distinct word types also undergo a shift towards concreteness [31]. While both of these examined word types and tokens at large, this work focused on a select group of words around *privacy*. Future work could explore this shift in more depth, for example using abstract and concrete seed words [31] for the neighbor shift analysis to get a larger perspective, or employ additional methods to acquire concreteness ratings of the specific word types from the corpora.

More work will be necessary to refine and expand on the proposed method. Most importantly, there are several limitations to our work that are also ongoing issues in the field of LSCD. The first issue is orthographic variation. Especially at the start of the 18th century, we face a high degree of spelling variation that we did not normalize, which limits the recall of embedding extraction. For future work, orthographic normalization as well as lemmatization should be considered as pre-processing steps.

The results of our study are highly dependent on the choice of corpus. Even though our corpus has balanced slices in terms of genre, the observed variation can still be dependent on a specific context of a particular slice. This means that rather than finding semantic shifts, we find unintended domain shifts because words were used in specific contexts in certain time slices. This is a well-known underlying issue of using the data-driven, distributional approach that is subject of a larger discussion about what semantic shifts actually are; per Kutuzov et al., "If one does not employ external data sources [...], there is no reliable way to discern 'semantic changes' from 'differences in the underlying textual data': they are simply the same thing" [17].

Additionally, by selecting our dictionary words, we limit our result space and another list of words would have produced different results. As the method is computationally expensive, this pre-selection is still necessary at this point in time. To study a different concept where a curated list of words is not available, researchers could either turn to resources such as knowledge graphs, where related words for most concepts can be retrieved up to a desired path depth (e.g. WordNet or FrameNet), or select words based on distributional approaches such as fastText entirely.

A final limitation concerns word frequency and polysemy. Both highly polysemous words and words with instable frequency produce a higher semantic change score with this method [28]. We set a minimum and maximum frequency limit for our experiments. However high fluctuation is still possible, especially due to orthographic variation as mentioned before, which we did not control for.

At the moment, we provided suggestions for further insights into semantic change detection and presented a demonstration of that idea. In future work, evaluation data will be necessary to support the results and methods discussed so far. We propose to construct a new type of evaluation dataset based on the idea of changing word pairs. Current evaluation data consists of words with assigned stability scores [29, 12]. Instead, new datasets may be created by tracing the relationship of e.g. synonyms or associated words through a similar methodology. By grading sentence pairs that contain synonyms based on their similarity instead of the same word for example, the degree of synonymy may be determined across different time periods. Existing word similarity datasets such as SimLex333 or WordSim353 could provide the basis for such word pair selection. Such word pair ratings could also be established to evaluate concreteness shifts for example.

On the historical semantics of *privacy* specifically, we note that our findings do not align with an analog study on the historical semantics of *privacy* in American historical jurisprudence by Prestidge, who, based on close reading, reported *privacy* "shifting in meaning from the literal to the figurative" [26]. In our experiments, *privacy* remained quite semantically stable, and we present evidence that some words relating to *privacy*, to the contrary, shifted from figurative to literal. These findings are certainly worth exploring further.

We hope that the ideas on automatic semantic change analysis presented here can provide a step towards the pairwise analysis of shifts in word meaning as a new direction within the field.

References

[1] A. Caliskan Islam, J. Walsh, and R. Greenstadt. “Privacy detective: Detecting private information and collective privacy behavior in a large social network”. In: *Proceedings of the 13th Workshop on Privacy in the Electronic Society*. Scottsdale Arizona USA, 2014, pp. 35–46.

[2] F. Casillo, V. Deufemia, and C. Gravino. “Detecting privacy requirements from User Stories with NLP transfer learning models”. In: *Information and Software Technology* 146 (2022), pp. 1–12. DOI: 10.1016/j.infsof.2022.106853.

[3] D. D’Acunto, S. Volo, and R. Filieri. ““Most Americans like their privacy.” Exploring privacy concerns through US guests’ reviews”. In: *International Journal of Contemporary Hospitality Management* 33.8 (2021).

[4] S. De Pascale. *Meaning change in token space: a token-based computational approach to diachronic prototype semantics*. Kick-off Conference Day of the RJ Program. 2022.

[5] H. De Smet, S. Flach, H.-J. Diller, and J. Tyrkkö. *The Corpus of Late Modern English Texts, version 3.1*. Online. 2011.

[6] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding”. In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Minneapolis, Minnesota, 2019, pp. 4171–4186.

[7] C. Dwork. “Differential privacy”. In: *International colloquium on automata, languages, and programming*. Springer. Venice, Italy, 2006, pp. 1–12.

[8] K. Franco, M. Montes, and K. Heylen. “Deconstructing destruction: A Cognitive Linguistics perspective on a computational analysis of diachronic change”. In: *Proceedings of the 3rd Workshop on Computational Approaches to Historical Language Change*. Dublin, Ireland, 2022, pp. 23–32.

[9] M. Giulianelli, M. Del Tredici, and R. Fernández. “Analysing Lexical Semantic Change with Contextualised Word Representations”. In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Online, 2020, pp. 3960–3973. DOI: 10.18653/v1/2020.acl-main.365.

[10] M. Gruppi, P.-Y. Chen, and S. Adali. “Fake it till you make it: Self-supervised semantic shifts for monolingual word embedding tasks”. In: *Proceedings of the 35th AAAI Conference on Artificial Intelligence*. Vol. 35. Online, 2021, pp. 12893–12901.

[11] W. L. Hamilton, J. Leskovec, and D. Jurafsky. “Diachronic Word Embeddings Reveal Statistical Laws of Semantic Change”. In: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Berlin, Germany, 2016, pp. 1489–1501. DOI: 10.18653/v1/P16-1141.

[12] A. Hätty, D. Schlechtweg, and S. Schulte im Walde. “SURel: A Gold Standard for Incorporating Meaning Shifts into Term Extraction”. In: *Proceedings of *SEM 2019*. Minneapolis, Minnesota: Association for Computational Linguistics, 2019, pp. 1–8. DOI: 10.18653/v1/S19-1001.

[13] S. Hengchen, N. Tahmasebi, D. Schlechtweg, and H. Dubossarsky. “Challenges for computational lexical semantic change”. In: *Computational approaches to semantic change*. Berlin: Language Science Press, 2021, pp. 341–372.

[14] T. T. Hills and J. S. Adelman. “Recent evolution of learnability in American English from 1800 to 2000”. In: *Cognition* 143 (2015), pp. 87–92.

[15] E. Ketzan, J. Edmond, and C. Vogel. “Need a Good Book about Privacy? Evaluating Dictionary-Based Corpus Query for Detecting the Topic of Privacy in Literary Texts”. In: *Journal of Computational Literary Studies* (in press).

[16] A. Kutuzov, L. Øvrelid, T. Szymanski, and E. Velldal. “Diachronic word embeddings and semantic shifts: a survey”. In: *Proceedings of the 27th International Conference on Computational Linguistics*. Santa Fe, New Mexico, USA, 2018, pp. 1384–1397.

[17] A. Kutuzov, E. Velldal, and L. Øvrelid. “Contextualized embeddings for semantic change detection: Lessons learned”. In: *Northern European Journal of Language Technology* 8 (2022), pp. 1–14. DOI: 10.3384/nejlt.2000-1533.2022.3478.

[18] S. Laicher, S. Kurtyigit, D. Schlechtweg, J. Kuhn, and S. Schulte im Walde. “Explaining and Improving BERT Performance on Lexical Semantic Change Detection”. In: *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Student Research Workshop*. Online, 2021, pp. 192–202. DOI: 10.18653/v1/2021.eacl-srw.25.

[19] A. Lehrer. “The influence of semantic fields on semantic change”. In: *Historical Semantics - Historical Word-Formation*. Ed. by J. Fisiak. Berlin, New York, Amsterdam: Mouton de Gruyter, 1985, pp. 283–296.

[20] E. Manjavacas and L. Fonteyn. “Adapting vs. Pre-training Language Models for Historical Languages”. In: *Journal of Data Mining & Digital Humanities* Nlp4dh (2022). DOI: 10.46298/jdmdh.9152.

[21] B. McGillivray. “Computational methods for semantic analysis of historical texts”. In: *Routledge International Handbook of Research Methods in Digital Humanities*. London: Routledge, 2020, pp. 261–274.

[22] B. McGillivray, S. Hengchen, V. Lähteenoja, M. Palma, and A. Vatri. “A computational approach to lexical polysemy in Ancient Greek”. In: *Digital Scholarship in the Humanities* 34.4 (2019), pp. 893–907.

[23] T. Mikolov, K. Chen, G. Corrado, and J. Dean. “Efficient estimation of word representations in vector space”. In: *arXiv preprint arXiv:1301.3781* (2013).

[24] G. R. Milne, B. Kaplan, K. L. Walker, and L. Zacharias. “Connecting with the future: The role of science fiction movies in helping consumers understand privacy-technology trade-offs”. In: *Journal of Consumer Affairs* 55.3 (2021), pp. 737–762.

[25] V. Perrone, M. Palma, S. Hengchen, A. Vatri, J. Q. Smith, and B. McGillivray. "GASC: Genre-Aware Semantic Change for Ancient Greek". In: *Proceedings of the 1st International Workshop on Computational Approaches to Historical Language Change*. Florence, Italy, 2019, pp. 56–66.

[26] A. Prestidge. "Semantic Change in Supreme Context: Semantics in the Privacy Line and Originalist Interpretation". In: *Brigham Young University Prelaw Review* 24 (2019), pp. 117–137.

[27] M. A. Rodda, M. S. Senaldi, and A. Lenci. "Panta rei: Tracking semantic change with distributional semantics in ancient Greek". In: *IJCoL. Italian Journal of Computational Linguistics* 3.3-1 (2017), pp. 11–24.

[28] D. Schlechtweg, B. McGillivray, S. Hengchen, H. Dubossarsky, and N. Tahmasebi. "SemEval-2020 Task 1: Unsupervised Lexical Semantic Change Detection". In: *Proceedings of the 14th Workshop on Semantic Evaluation*. Barcelona (online), 2020, pp. 1–23.

[29] D. Schlechtweg, S. Schulte im Walde, and S. Eckmann. "Diachronic Usage Relatedness (DURel): A Framework for the Annotation of Lexical Semantic Change". In: *Proceedings of NAACL-HLT*. New Orleans, Louisiana, 2018, pp. 169–174.

[30] D. Schlechtweg, N. Tahmasebi, S. Hengchen, H. Dubossarsky, and B. McGillivray. "DWUG: A large Resource of Diachronic Word Usage Graphs in Four Languages". In: *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics, 2021, pp. 7079–7091. DOI: 10.18653/v1/2021.emnlp-main.567.

[31] B. Snefjella, M. Généreux, and V. Kuperman. "Historical evolution of concrete and abstract language revisited". In: *Behavior Research Methods* 51 (2019).

[32] G. Stern. *Meaning and change of meaning; with special reference to the English language*. Bloomington: Indiana University Press, 1964.

[33] N. Tahmasebi, L. Borin, A. Jatowt, Y. Xu, and S. Hengchen. *Computational approaches to semantic change*. Berlin: Language Science Press, 2021. DOI: 10.5281/zenodo.5040241.

[34] H. T. Tavani. "Philosophical theories of privacy: Implications for an adequate online privacy policy". In: *Metaphilosophy* 38.1 (2007), pp. 1–22.

[35] E. C. Traugott. *Semantic change*. Oxford research encyclopedia of linguistics. 2017.

[36] K. A. Vakeel, S. Das, G. J. Udo, and K. Bagchi. "Do security and privacy policies in B2B and B2C e-commerce differ? A comparative study using content analysis". In: *Behaviour & Information Technology* 36.4 (2017), pp. 390–403.

[37] A. Vasalou, A. J. Gill, F. Mazanderani, C. Papoutsi, and A. Joinson. "Privacy dictionary: A new resource for the automated content analysis of privacy". In: *Journal of the American Society for Information Science and Technology* 62.11 (2011), pp. 2095–2105.

[38] M. Wevers and M. Koolen. "Digital begriffsgeschichte: Tracing semantic change using word embeddings". In: *Historical Methods: A Journal of Quantitative and Interdisciplinary History* 53.4 (2020).