

Distributional Concept Analysis

A Computational Model for History of Concepts

PETER DE BOLLA, EWAN JONES, PAUL NULTY,
GABRIEL RECCHIA, AND JOHN REGAN

ABSTRACT

This article proposes a novel computational method for discerning the structure and history of concepts. Based on the analysis of co-occurrence data in large data sets, the method creates a measure of “binding” that enables the construction of verbal constellations that comprise the larger units, “concepts,” that change over time. In contrast to investigation into semantic networks, our method seeks to uncover structures of conceptual operation that are not simply semantic. These larger units of lexical operation that are visualized as interconnected networks may have underlying rules of formation and operation that have as yet unexamined—perhaps tangential—connection to meaning as such. The article is thus exploratory and intended to open the history of concepts to some new avenues of investigation.

KEYWORDS

computational linguistics, conceptual functions, diachronic analysis of word distributions, semantic fields, statistical methods for conceptual analysis

This journal, *Contributions to the History of Concepts*, has, since its inception, been commendably open to different approaches to the history of concepts and, since its first volume, in 2005, has significantly contributed to the heterogeneous development of the field across national and disciplinary boundaries. As it moves into its second decade, calls for a “transdisciplinary history of concepts” and “comparative transnational”¹ approaches feature more prominently in its pages, and this might be taken to imply that some of the earlier head-scratching around the methodological grounding of the

1. See Julian Bauer, “From ‘Organisms to World Society’: Steps toward a Conceptual History of Systems Theory, 1880–1980,” *Contributions to the History of Concepts* 9, no. 2 (2014): 51–72, here 54; José María Rosales, “Liberalism’s Historical Diversity: A Comparative Conceptual Exploration,” *Contributions to the History of Concepts* 8, no. 2 (2013): 67–82, here 73.



field is no longer appropriate, or even very common. Despite this trend, this article goes back to some grounding questions to provide an example of how some experimental methodologies that adapt computational and statistical techniques developed in both information extraction and computational linguistics might contribute to the heterogeneity of the field. Perhaps the most basic question we mean to invoke is this: Are concepts the same as words? We suspect many readers of this journal will find such a question naive or inconsequential, as Willibald Steinmetz does in his 2012 contribution when he indicates in a note, “I leave out the tedious discussion about the concept of concept,” preferring instead to work with the self-declared imprecision of “concept”: “Thus I speak of concept throughout the article although, strictly speaking, it would be correct to distinguish carefully between the terms (or words) referring to concepts and the concepts themselves.”²

Our purpose here is not to criticize a significant contribution to the ongoing reception of Reinhart Koselleck’s work but rather to underline the awkward and consistent fact that most research into conceptual history, insofar as it has settled into a discipline (or inter-discipline) with clear outlines, assumes what is most commonly taken to be the case, namely that there is a difference between a label for a concept—say, “justice” in English—and what Steinmetz calls the “concept itself.” Intuitively, this seems right, and it prompts observations of the following kind, taken from a review of a book by Brent Nongbri, *Before Religion: A History of a Modern Concept*, written by Helge Årsheim, who notes that tracking the history of the concept of religion from the ancient world to our own time “ultimately comes down to the somewhat basic and blunt observation that the ancient world had no word or corresponding concept to our modern-day conception of religion.”³ The bump in the road here is “no word *or* corresponding concept.” Hedging whether it would be enough to identify the lack of a *term* for religion, Årsheim does what we all do: invoke the assumption that words and concepts are not the same thing and leave us hanging as to whether the “or” marks a real or hard distinction between word and concept as opposed to a soft or inconsequential one. This hedging is common in the literature.

This contribution to the ongoing development of the field of conceptual history emerges from a four-year project, the Cambridge Concept Lab, which took as its starting point the axiom that words and concepts are not the same kind of thing: they have different ontologies.⁴ “Ontology” is used in

2. Willibald Steinmetz, “Some Thoughts on a History of Twentieth Century German Basic Concepts,” *Contributions to the History of Concepts* 7, no. 2 (2012): 87–100, here 88.

3. Helge Årsheim, “Critique of the Modern Concept of Religion: A Review of Brent Nongbri, *Before Religion: A History of a Modern Concept*,” *Contributions to the History of Concepts* 9, no. 2 (2014): 90–92.

4. The lab was funded by the Foundation for the Future and directed by Peter de Bolla,

a variety of disciplines and discourses, and this richness of use can impede understanding. We use the term here to refer to something quite basic, the difference in category or kind of thing. One might note, for example, that apples do not share the same ontology as motion pictures. In computation and information science, the word is used in a technical sense to refer to an artifact designed for a specific purpose, “which is to enable the modelling of knowledge about some domain.”⁵ Although we do not mean to invoke this strictly technical sense, we find the word helps us focus our own inquiry into how concepts model knowledge in ways that are different to how words do. Noting this does not preclude the possibility that the two entities, words and concepts, might share properties, one of which is meaning.

If one starts here, one immediately must consider how a putative history of concepts is to be distinguished from a history of the meanings of words.⁶ This article sketches out a possible answer that is intended to be exploratory. Unsurprisingly, this issue has been addressed in many of the contributions to this journal since its inception, and we find that what has emerged is a (weak) consensus around the following headline proposal: “Conceptual history entails assessing semantic changes and exploring the diverse historical settings where concepts are semantically recast, while comparative conceptual history maps those transnational paths.”⁷ Although there is surely much to be learned from tracking the history of semantic changes (both across natural languages and within them), we find there is an obvious question here: Does this assume an identity between types of things, in our terms, a shared ontology between words and concepts? No one would argue that “semantic change” must be based in the altering uses of words—Reinhart Koselleck, a principal guiding figure for some versions of the history of concepts, explicitly recognizes this in his introduction to the *Geschichtliche Grundbe-*

whose *The Architecture of Concepts: The Historical Formation of Human Rights* (New York: Fordham University Press, 2013) set the agenda for its work.

5. See Tom Gruber, “Definition of Ontology 07,” in *Encyclopedia of Database Systems*, ed. Ling Liu and M. Tamer Özsu (Berlin: Springer, 2009), <http://tomgruber.org/writing/ontology-definition-2007.htm>.

6. Contributors to this journal have not been slow to note that Quentin Skinner, for example, is in print indicating that he does not think a history of concepts possible—though he later recast this view. See, among other places: “My almost paradoxical contention is that the various transformations we can hope to chart will not strictly speaking be changes in concepts at all. They will be transformations in the applications of the terms by which our concepts are expressed.” Quentin Skinner, “Retrospect: Studying Rhetoric and Conceptual Change,” in *Visions of Politics: Regarding Method*, vol. 1 (Cambridge: Cambridge University Press, 2002), 175–187, here 179; Jan-Werner Müller, “On Conceptual History,” in *Rethinking Modern European Intellectual History*, ed. Darrin M. McMahon and Samuel Moyn (Oxford: Oxford University Press, 2014), 74–93, here 75.

7. Rosales, “Liberalism’s Historical Diversity,” 73.

griffe: “This work explores how modern and old words begin to overlap and shift their meanings.”⁸ One might observe that saying words have histories of use, and therefore meaning, does not preclude the possibility that concepts may also alter over time. But how are they altered? If we say—or, more commonly, assume—the meanings of concepts change over time, in what sense are these meanings to be distinguished from the meanings of words?⁹ If they are not so distinguished, why make a case for a history of concepts as opposed to semantic history in general? Here, the distinction between a history of ideas and a history of concepts—claimed by Elias Jose Palti to be the major difference between Ernst Cassirer’s project and Koselleck’s—does not settle this issue. It may well be that concepts have the outline Koselleck claims for them (the semantic shifts that occur from alterations in the context of use are inscribed in the concept), but the means for identifying such change remains the same: meaning.¹⁰

Of course, Koselleck himself famously insisted that concepts exceed words: “The concept is bound to a word, but is at the same time more than a word,”¹¹ and he and his followers seek to further this observation by insisting on what might be called parallel semantic fields, one holding for words and the other for concepts. This allows one to mount the hypothesis that the meanings of words and concepts may sometimes overlap or converge and at other times be discrete or distant from each other. So, to return to our opening axiom, that words and concepts have distinctive ontologies, we now need to ask if the meanings of each are ontologically distinct as well, which, to be absolutely clear, is to ask if the meaning of a word is the same kind of thing as the meaning of a concept. Some work in the history of concepts addresses this question by expanding the field within which concepts

8. See Reinhart Koselleck, “Introduction to GG,” trans. Michaela Richter, *Contributions to the History of Concepts* 6, no. 1 (2011): 1–37, here 8. For one of many descriptions of Koselleck’s method in the same vein, see: “There is no doubt that for Koselleck, doing conceptual history entailed a word history, or rather a historical semantics, based on the study of language in the sources (*Quellensprache*).” Jan Ifversen, “About Key Concepts and How to Study Them,” *Contributions to the History of Concepts* 6, no. 1 (2011): 65–88, here 72.

9. One can see the issue clearly in this contribution: “Conceptual history (here understood as the description and analysis of concrete historical semantics, origins, derivation and alterations of concepts).” O. Hidalgo, “Conceptual History and Politics: Is the Concept of Democracy Essentially Contested,” *Contributions to the History of Concepts* 4, no. 2 (2008): 176–201, here 177. How is “concrete historical semantics”—a history of the meanings of words—distinguished from “alterations of concepts”?

10. See Elias Jose Palti, “Reinhart Koselleck: His Concept of the Concept and Neo-Kantianism,” *Contributions to the History of Concepts* 6, no. 2 (2011): 1–20, here 4.

11. Reinhart Koselleck, “Begriffsgeschichte and Social History,” in *Futures Past: On the Semantics of Historical Times* (Cambridge, MA: MIT Press, 1985), 84.

are assumed to operate. Such an expanded field was always central to Koselleck's project, since he was interested not solely in plotting the history of the changing meaning of words or concepts but also in the social and political uses to which such concepts were put. Here, we agree that Skinner's attention to "the uses to which words were put" at different times in history and by different actors in the past is more closely aligned with Koselleck's project than is sometimes claimed.¹² For others such as the historian of epistemology Hans-Jörg Rheinberger, this expanded field must include not only words but also deeds and things.¹³ But even when the discursive in its largest sense is moved to the center of attention a residual assumption persists, namely that the study of concepts must always entail inquiry into their meaning howsoever that meaning may be constructed, or what kind of thing it may be. The history of concepts as it has emerged is, therefore, at base a semantically motivated field of enquiry.

Although there are very good reasons for this investment in the diachronic study of meaning, we believe it hampers attempts to model concepts as distinct entities from words. Difficult as this is, this article sets out some ways in which one might begin to dampen down our interest in semantics as the primary motivation for studying concepts diachronically so as to open up for inspection something that heretofore has been impossible to see: the functions by which concepts operate.¹⁴ In this article, which is intended to be exploratory and not definitive, we sketch some ways in which a putative conceptual functionality may be identified from the distributions of lexis in large data sets of language use. We acknowledge that our method, based as it is on lexical co-occurrence data, does not decisively break the tie

12. See Quentin Skinner and Javier Fernández Sebastián, "Intellectual History, Liberty and Republicanism: An Interview with Quentin Skinner," *Contributions to the History of Concepts* 3, no. 2 (2007): 103–123, here 115. For Koselleck's account of his own method, see Javier Fernández Sebastián and Juan Francisco Fuentes, "Conceptual History, Memory, and Identity: An Interview with Reinhart Koselleck," *Contributions to the History of Concepts* 1 (2006): 99–127. Müller argues something similar, although his emphasis suggests that the history of concepts is slightly misleading: noting of Skinner and Koselleck "both in a sense agree that concepts do not actually change at all; what changes is the usage of words." Müller, "On Conceptual History," 87.

13. See Hans-Jörg Rheinberger, *Toward a History of Epistemic Things: Synthesizing Proteins in the Test Tube* (Stanford, CA: Stanford University Press, 1997).

14. Here our approach shares some similarity to recent work in cultural theory and history in which function moves to the front of attention. See, for example, Stengers's notion of "operations of propagation" and "operations of passage," which are said to be properties of "nomadic concepts." These are very close to some of the functions we believe we have identified in this article. Isabelle Stengers, ed., *D'une science à l'autre: Des concepts nomades* [From one science to another: Nomadic concepts] (Paris: Editions du Seuil, 1978), 17–19, 24.

between words and concepts because, as Jan Ifversen notes, “since concepts are expressed in words, this means that they are always tied to words.”¹⁵ We also see great merit in an approach that identifies larger units of lexical operation, in the literature referred to as semantic fields, as more amenable to specifically conceptual rather than word history. But what if these larger units of lexical operation are not only semantically motivated but also determined by an underlying set of rules of formation and operation that have as yet unexamined—perhaps tangential—connection to meaning as such? These objects of study—semantic fields or networks—are constructions based on meaning relations. Now, however, we have the ability to gather information on large constellated lexical behavior through computational methods that do not necessarily assume a model of semantic equivalence, even if some computational linguists use these methods to discover the basis for such equivalence.¹⁶ Our hypothesis, then, is that we may be able to inspect conceptual function at a level that is quasi-semantic or even perhaps non-semantic. If such functions change over time, then we may then be able to identify one of the causes for change in the meanings of concepts. The project from which the article emerges has been testing such a hypothesis, and our method has been first to establish functional descriptions and analyses of concepts based on data derived from word distributions in massive data sets of printed text. In this article, we restrict our few examples to Eighteenth-Century Collections Online, since the aim is to establish the outlines of a data-driven approach as exemplary for further work that might connect to more standard conceptual history.

These distributions enable us to map the ways in which concepts are constructed through “constellations” or “bundles” of individual words—hence, the name we give to our enquiry: distributional concept analysis. To some extent, the approach we take here can be compared to the project “Linguistic DNA” based in Sheffield, which also use methods developed in corpus and computational linguistics. This project focuses on the early modern period, using a transcribed subset of Early English Books Online in combination with a thesaurus categorization of word senses from the period to examine the change over time of raw word association frequencies and pointwise mutual information scores between pairs of terms of interest. The main distinction between that project and ours is its focus primarily on semantic change over time.¹⁷ Our second step, to be explored in a subsequent

15. Ifversen, “About Key Concepts,” 69.

16. We are encouraged to note that Sarasin proposes a similar turn to digital methods, even if his model keeps close to the semantic. Philipp Sarasin, “Is a ‘History of Basic Concepts of the Twentieth Century’ Possible? A Polemic,” *Contributions to the History of Concepts* 7, no. 2 (2012): 101–110, here 108–110.

17. Susan Fitzmaurice, Justyna A. Robinson, Marc Alexander, Iona C. Hine, Seth Mehl,

article, is to investigate such data-driven descriptions diachronically and then to map any identified changes onto diachronic investigations of changing lexical “bundles” or “constellations.”¹⁸ The emphasis we wish to make, then, is on the structure or shape of these entities, the “bundles” or “constellations,” in contrast to the semantic values of the words that comprise them.

Background

We begin with some background on the computational methods we have adapted from disciplines that rarely connect to the history of concepts: cognitive science and computational linguistics. Such methods are also commonly employed in research into natural language processing where projects are directed at tasks such as machine translation, analogy solving, question answering, and natural language inference systems. Researchers have achieved considerable success in these domains using distributional semantic models derived from large text corpora.¹⁹ Furthermore, some work in digital humanities is beginning to investigate the utility of vector semantics or “vector space models” for understanding concepts.²⁰ Our own modeling of concepts follows a similar direction and the tools the Concept Lab has created allow one to transparently operate arithmetical procedures (addition, subtraction, multiplication, and division) with the vectors we derive from co-association.²¹ Irrespective of theoretical motivations, the computational implementations of these methods have much in common.

and Fraser Dallachy, “Linguistic DNA: Investigating Conceptual Change in Early Modern English Discourse,” *Studia Neophilologica* 89, no. S1 (2017): 21–38.

18. See Peter de Bolla, Ewan Jones, Gabriel Recchia, John Regan, and Paul Nulty, “The Idea of Liberty 1600–1800: A Distributional Concept Analysis,” *Journal of the History of Ideas* (forthcoming), <https://doi.org/10.17863/CAM.35230>.

19. See Colin Kelly, Barry Devereux, and Anna Korhonen, “Acquiring Human-Like Feature-Based Conceptual Representations from Corpora,” *Proceedings of the NAACL HLT 2010 First Workshop on Computational Neurolinguistics* (Association for Computational Linguistics, 2010): 61–69; Francisco Pereira, Matthew Botvinick, and Greg Detre, “Using Wikipedia to learn semantic feature representations of concrete concepts in neuroimaging experiments,” *Artificial Intelligence* 194 (2013): 240–252.

20. See Michael Gavin, “Intellectual History and the Computational Turn,” *The Eighteenth Century* 58, no. 2 (2017): 249–253; Michael Gavin, “Vector Semantics, William Empson, and the Study of Ambiguity,” *Critical Inquiry* 44, no. 4 (2018): 641–673; Michael Gavin, Colin Jennings, Lauren Kersey, and Brad Pasanek, “Spaces of Meaning: Conceptual History, Vector Semantics, and Close Reading,” in *Debates in the Digital Humanities 2018*, ed. Mathew Gold and Lauren Klein (Cambridge, MA: MIT Press, 2019), 243–267.

21. The tools are available at The Concept Lab, “Tools and Resources,” <https://concept-lab.lib.cam.ac.uk> (accessed 13 April 2019).

Word, phrase, or document meanings are approximated by deciding on a word-distance window within which to count word co-occurrences or compare the paradigmatic context, and counts of word or context co-occurrence are tabulated into a vector. The vectors can be compared directly to measure word associations, or combined into a matrix to measure the similarity of documents.

One of the most widely applied document-based methods today is a generative probabilistic model of the relationship between words and documents known as topic modeling.²² Topic modeling discovers groups of associated words by modeling a process in which documents are generated by selecting from subsets of words with varying probabilities. The resulting clusters of associated words are intended to capture topics in the corpus and allow comparison of documents by topic and measurement of diachronic change in topic emphasis. The method has been widely applied outside of computer science, for example, to measure political attention and literary style.²³

Applications that focus on measuring word rather than document association often use paradigmatic similarity of contexts shorter than document length.²⁴ The meaning of a word in such models is represented as a point embedded in a high-dimensional space defined by the contexts in which its occurrences are counted. The context may be a short window of words around the target word or defined by a grammatical relation discovered by automatic syntactic parsing of the text.²⁵ Since the number of possible con-

22. See David M. Blei, Andrew Y. Ng, and Michael I. Jordan, "Latent Dirichlet Allocation," *Journal of Machine Learning Research* 3 (2003): 993–1022.

23. See Kevin M. Quinn, Burt L. Monroe, Michael Colaresi, Michael H. Crespin, and Dragomir R. Radev, "How to Analyze Political Attention with Minimal Assumptions and Costs," *American Journal of Political Science* 54, no. 1 (2001): 209–228; James M. Hughes, Nicholas J. Foti, David C. Krakauer, and Daniel N. Rockmore, "Quantitative Patterns of Stylistic Influence in the Evolution of Literature," *Proceedings of the National Academy of Sciences* 109, no. 20 (2012): 7682–7686.

24. This kind of similarity as a method for word translation was first suggested by Warren Weaver, "Translation," *Machine Translation of Languages* 14 (1975): 15–23. See also Will Lowe, "Towards a Theory of Semantic Space," *Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society* (Mahwah, NJ: Lawrence Erlbaum Associates, 2001), 576–581.

25. See Kevin Lund and Curt Burgess, "Producing High-Dimensional Semantic Spaces from Lexical Co-occurrence," *Behavior Research Methods, Instruments, & Computers* 28, no. 2 (1996): 203–208; Dekang Lin, "Automatic Retrieval and Clustering of Similar Words," *Proceedings of the 17th International Conference on Computational Linguistics* 2 (1998): 768–774; Omer Levy and Yoav Goldberg, "Dependency-Based Word Embeddings," *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics* 2 (2014): 302–308.

texts may be many times the size of the vocabulary, dimension-reduction methods such as singular value decomposition are often used to reduce the complexity of the model, with the drawback that the resulting dimensions are no longer directly interpretable. Word-context models with dimensionality reduction result in a “word vector,” usually of the order of a few hundred dimensions; these are commonly referred to as “word-embedding” models. An efficient and widely used implementation of this kind of model is the Word2Vec package, which uses a neural network trained to predict word contexts and encodes the meaning of the word in its parameters rather than explicitly counting word contexts.²⁶

Some work in computational linguistics describes types of word co-occurrence in terms of Saussure’s syntagmatic and paradigmatic relations: words that tend to occur close to one another in a passage of text have a syntagmatic relation, while words that tend to have similar contexts—one word may be easily substituted for another—have a paradigmatic relation.²⁷ These different kinds of contextual similarity reveal different conceptual relations: syntagmatic similarity captures meronymy, phrasal association, and general conceptual association, while paradigmatic similarity captures synonymy, antonymy, and hyponymy.

An advantage of syntagmatic counts is that the type of semantic association required can be precisely defined. Fine-grained relations such as that between “dog” and “bark” are more clearly instantiated in syntagmatic co-occurrence patterns. Researchers have used specific co-occurrence patterns to learn precise kinds of semantic relation.²⁸ For example, a meronymic relation between word x and word y may be determined by computational analysis of counts for phrases like “ x s consisting of y ” in a large corpus, whereas syntagmatic patterns for arbitrary semantic relations can be learned from example seed pairs.²⁹ Descriptions of distributional semantic methods focus on measuring and evaluating word similarity, but although it is not always explic-

26. See Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean, “Efficient Estimation of Word Representations in Vector Space,” *arXiv* (2013): arXiv:1301.3781; David E. Rumelhart, Ronald J. Williams, and Geoffrey E. Hinton, “Learning Representations by Back Propagating Errors,” *Nature* 323, no. 6088 (1986): 533–536. For a more digital humanities focus, see Ben Schmidt, “Vector Space Models for the Digital Humanities,” *Ben’s Bookworm Blog*, 25 October 2015, <http://bookworm.benschmidt.org/posts/2015-10-25-Word-Embeddings.html>.

27. E.g., Magnus Sahlgren, “The Word-Space Model: Using Distributional Analysis to Represent Syntagmatic and Paradigmatic Relations between Words in High-Dimensional Vector Spaces” (PhD diss., Stockholm University, 2006).

28. See Marti A. Hearst, “Automatic Acquisition of Hyponyms from Large Text Corpora,” *Proceedings of the 14th Conference on Computational Linguistics* 2 (1992): 539–545.

29. See Roxana Girju, Adriana Badulescu, and Dan Moldovan, “Automatic Discovery of Part-Whole Relations,” *Computational Linguistics* 32, no. 1 (2006): 83–135.

itly stated, the literature recognizes the possibility that these models encode information that could be considered conceptual rather than simply lexical.

Advantages and Disadvantages of Existing Methods for Detecting Conceptual Coherence

There are numerous well-established approaches in computational linguistics for identifying statistical associations among words based on their patterns of co-occurrence in text corpora.³⁰ However, a measure that would serve as a foundation for discovering conceptual structure within large textual data sets needs to satisfy additional desiderata. We have therefore developed a measure to satisfy three particularly important requirements for our purposes—namely transparency, sensitivity to frequency, and sensitivity to distance—while at the same time acknowledging that approaches that embed co-occurrence-based measures to discriminate between paradigmatic and syntagmatic relations are possible and could be a worthy future direction to pursue. The reasoning behind our methodology can be summed up as follows:

- (1) *Transparency*: An ideal measure of association should not be a “black box.” If such a measure identifies two words in a corpus as having a strong association to each other, it should be possible to understand why.
- (2) *Sensitivity to frequency and data sparsity*: A word (e.g., “democracy”) will likely co-occur more frequently with words that are very frequent in the corpus as a whole (“was,” “of”) than with less frequent words (“government,” “aristocracy”), even though “democracy” and “government” serve similar functions. An ideal measure of association should therefore consider frequency in a way that also considers data sparsity. That is to say, no matter how large one’s corpus is, it will always contain a large number of words whose counts are relatively infrequent and less statistically reliable as a result. This difference in statistical reliability should somehow be accounted for as well.
- (3) *Sensitivity to distance*: Because we wish to capture associations that lay outside narrowly semantic operations, it is important that our measure be able to distinguish how strongly associated two words are at one distance (e.g., one hundred words away) from another (e.g., five words away).

30. Note that we do not lemmatize the corpus in order to retain important distinctions in abstraction often indicated by different word forms, for example, chose versus chosen.

The standard procedures for deciding whether two words x and y “co-occur” entail determining (a) whether y appears in a “window” of text that extends some specific number of words to the left and/or right of x , or (b) whether x and y appear in the same document. The most common way to implement (a) is for the window of text to begin at x and extend outward in both directions. As such, a classic “window” of one hundred words will include all words that appear just one, two, three, or more words away, which is not desirable if we wish to exclude words that primarily only appear with x in adjectival phrases, idiomatic expressions, and other relations that have more to do with syntax than conceptual relatedness. One reason we use “co-association” rather than “co-occurrence” for these “long-range” co-occurrences is that “co-occurrence” frequently implies the two words appear in the same phrase or collocation. Another reason is its distinctness from “association,” which is not necessarily symmetric.³¹ In contrast, if A co-associates with B, then the reverse must necessarily be true as well. We believe such standard procedures do not meet all the aforementioned desiderata. Neural networks, word embedding, and often topic models can produce nontransparent mathematical representations, although there has been some research aimed at improving their interpretability.³²

Simple measures of co-occurrence are much more transparent. For example, pointwise mutual information (PMI) is frequently used either as a measure of lexical association on its own or as a starting point for the construction of more complex measures.³³ PMI is given as

$$dpf(x, y) = \frac{P(x, y)}{(P(x)P(y))^\alpha}$$

31. See Lukas Michelbacher, Stefan Evert, and Hinrich Schütze, “Asymmetry in Corpus-Derived and Human Word Associations,” *Corpus Linguistics and Linguistic Theory* 7, no. 2 (2011): 245–276.

32. See M. Setnes, R. Babuska, and H. B. Verbruggen, “Rule-Based Modeling: Precision and Transparency,” *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 28, no. 1 (1998): 165–169; Julian D. Olden and Donald A. Jackson, “Illuminating the ‘Black Box’: A Randomization Approach for Understanding Variable Contributions in Artificial Neural Networks,” *Ecological Modelling* 154, nos. 1–2 (2002): 135–150; Hongyin Luo, Zhiyuan Liu, Huanbo Luan, and Maosong Sun, “Online Learning of Interpretable Word Embedding,” *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing* (Association for Computational Linguistics, 2015), 1687–1692.

33. See Kenneth W. Church and Patrick Hanks, “Word Association Norms, Mutual Information, and Lexicography,” *Computational Linguistics* 16, no. 1 (1990): 22–29; Jeffrey Pennington, Richard Socher, and Christopher D. Manning, “GloVe: Global Vectors for Word Representation,” *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing* (Association for Computational Linguistics, 2014), 1532–1543.

where $P(x)$ and $P(y)$ can each be approximated as the frequency of x and y (respectively) divided by the total number of words in the corpus, and $P(x, y)$ as the number of times x and y co-occur divided by the total number of words in the corpus.³⁴ $P(x, y)/P(x)P(y)$ has an intuitive interpretation: it simply expresses the ratio between the number of times x and y co-occur, divided by the number of times one would expect them to appear if their appearances throughout the corpus were randomly distributed. A ratio of 2 would indicate that one is twice as likely to see x and y together than one would expect if these words were randomly and uniformly distributed throughout the corpus—that is, if we assume x and y were to retain the same corpus-wide frequencies, but their positions relative to other words in the corpus were completely random. What qualifies as a “co-occurrence” is up to the user of the measure, and the most popular choice is appearance in the same window or document.

PMI is highly transparent—two words will rank as highly associated if and only if they appear together in the same contexts more frequently than would be expected by chance, and the contexts in which they appear together in the corpus can be directly inspected. Although PMI is sensitive to the global frequency with which words appear in a corpus, it is insensitive to data sparsity—the fact that the most infrequent words in the corpus also provide the least reliable information. If x and y each appear a hundred times in a corpus, and co-occur with each other on ninety-nine of those instances, this should be treated as more important than if x and y appear only once in a corpus and happen to co-occur in that singular instance. However, PMI does not penalize rare, statistically unreliable events in any way. As a result, if one were to rank the words in a corpus that co-occurred with some fixed word x according to its PMI with y , one would find an inverse relationship between rank and frequency: the higher the rank on the list, the lower the frequency.³⁵ A common bandage is to impose a minimum frequency or co-occurrence count, an approach that mitigates the problem but does not remove it.³⁶ In sum, most simple measures of co-occurrence tend to be transparent but have varying degrees of success with respect to their approach to frequency and data sparsity.³⁷ Approaches such as topic models and neural networks tend to be less transparent.

34. See Chris Manning and Hinrich Schütze, *Foundations of Statistical Natural Language Processing* (Cambridge, MA: MIT Press, 1999).

35. See Patrick Pantel and Dekang Lin, “Discovering Word Senses from Text,” *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Association for Computational Linguistics, 2002), 613–619.

36. See Manning and Schütze, *Foundations*.

37. For how the measure described in this article compares with similar measures in this regard, see Gabriel Recchia and Paul Nulty, “Improving a Fundamental Measure of

Finally, although most practitioners define co-occurrence in terms of appearance in the same window or document, this is more a matter of convention than necessity. Window-based measures can be readily adapted to “distance-based” measures: one can specify that in order to count as a co-association with x , a word y must appear some distance d words away from x , plus or minus a word or two. We have taken this approach because we wish to capture data on patterns of lexical distribution that moves from close up, where semantic or syntactic coherence is strongest, to far away, where it is weakest. By taking this wider view, we mean to discover relations of binding that go beyond or underpin strictly local semantic ties. Thus, we use what could be called a “sliding window,” as it involves a window of a fixed size that, rather than being centered on x itself, is centered on a word d words away. The investigations we will present on constructing conceptual profiles employ a sliding window of five words (e.g., y will co-associate with x at distance d if it appears anywhere between $d - 2$ and $d + 2$ words away from x).

Constructing a Measure for the Discovery of Conceptual Coherence

We have suggested that an ideal measure would combine the transparency of PMI with the ability of more complex methods, such as word embeddings, to exhibit sensitivity to lexical frequency and data sparsity. One possible method is suggested by Omer Levy, Yoav Goldberg, and Ido Dagan, who reported an intriguing connection between PMI and one of the most popular word-embedding models, Word2Vec SGNS (skip-grams with negative sampling). They note that one of the features of Word2Vec corresponds to a “smoothed” version of PMI in which “all context counts are raised to the power of α ,” and they apply this modification, along with many others, to a relatively simple and transparent model of lexical association.³⁸ They report that among all modifications applied, this was one of the most effective in increasing performance on a range of benchmark tasks, reporting their performance at $\alpha = 0.75$. They did not hypothesize why this value worked well, except to point out that similar work had found success with this value.

A potential reason for the success of this value is that it may hit the sweet spot in the trade-off between frequency and data sparsity, which are naturally in opposition. An alpha of 1 corresponds to standard, unsmoothed

Lexical Association,” *Proceedings of the 39th Annual Conference of the Cognitive Science Society* (2017), <https://doi.org/10.17863/CAM.30302>.

38. Omer Levy, Yoav Goldberg, and Ido Dagan, “Improving Distributional Similarity with Lessons Learned from Word Embeddings,” *Transactions of the Association for Computational Linguistics* 3 (2015): 211–225, here 215.

PMI, which ignores data sparsity and gives too much weight to the specious evidence provided by highly infrequent words. However, an alpha of 0 ignores a word’s corpus-wide frequency. We have adopted a simple variant of PMI, which drops the logarithm and introduces a smoothing exponent in the denominator, because we are primarily interested in the rank order of words as they are co-associated with a focal word x . Since the logarithm does not affect this rank, ordering it can be dropped without loss of generality. Doing so highlights the fact that our measure is essentially a small modification to the measure sometimes referred to simply as “observed over expected”: the number of times two words are observed in conjunction, divided by the number of times one would expect to see them together by chance. This measure, which we refer to as a distributional probability factor (DPF) since it expresses the extent to which a co-association at distance x is predicted to be distributed across the data set, is as follows:

$$I(x,y)\log_2 = \frac{P(x,y)}{P(x)P(y)}$$

When we calculated the value of α that eliminated the aforementioned inverse correlation between rank and frequency, we found the optimal value for our corpus to be 0.78, very close to 0.75. Higher values corresponded to a negative correlation between rank and frequency, while lower values corresponded to a positive correlation. In other words, with an appropriate value of alpha, DPF is sensitive to data sparsity as well as frequency, and retains the simplicity and transparency of PMI. Subsequent references in this work to the DPF between two words therefore refer to formula (2) with an alpha of 0.78. Recall that we estimate $P(x,y)$ in this formula by counting “co-associations”—cases in which y appears some particular distance d away from x . By combining this distance-based approach with DPF, we have constructed a measure that meets all three of our fundamental desiderata. Next, we will describe how this measure is employed to build a “profile” of a lexical co-association with respect to a particular time period and distance, and how these profiles can be used in our investigation of concepts.

Constructing a Profile of Lexical Co-associations

The following examples are derived from the data set Eighteenth-Century Collections Online (ECCO), which comprises some 180,000 titles, 200,000 volumes, and more than 33 million pages of text. This data set is well known as the world’s largest digital archive of books from the eighteenth century, containing “every significant English-language and foreign-language title printed

in the United Kingdom between the years 1701 and 1800.”³⁹ To begin the work of creating these profiles, we first calculated 280,000 constellations of co-associated lexis from our data set ECCO. We immediately discovered that computing DPF scores between all possible pairs of words in a corpus (between a focal and bound word) yields an enormous range of values. Thousands of pairs appear high to the eye—but how high is high? In this kind of environment, mainstays of parametric statistics, such as the use of *p*-values to determine statistical significance, are not appropriate.⁴⁰ How should we decide whether a co-association between two words is strong enough to merit our attention?

The approach we selected is similar to the common technique of determining which components of a factor analysis to preserve by looking for the “elbow” or “bend in the curve” in a scree plot.⁴¹ Consider the list created when one computes the DPF between a particular word of interest (the “focal word”) and a large number of other words (Figure 1). If these DPF val-

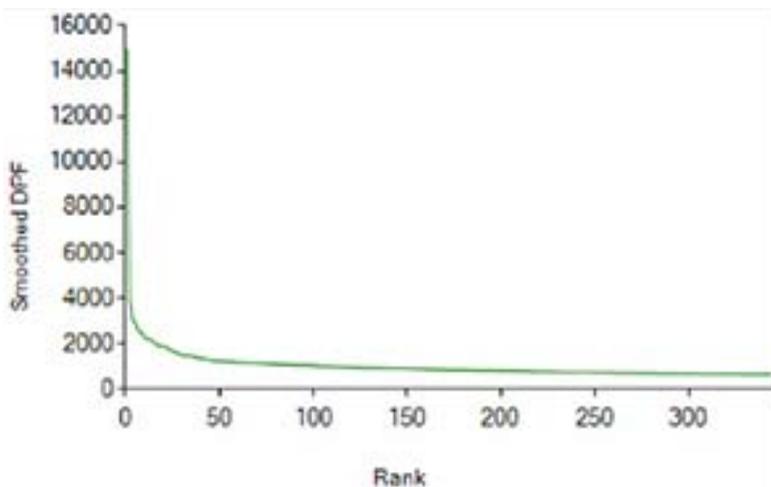


Figure 1. DPF values plotted against rank for the focal word “despotism,” distance 50, from all works published from 1780 to 1800 in Eighteenth-Century Collections Online. The height of the curve at any given point represents the DPF between “despotism” and some other word in the lexicon. The smooth, long-tailed curve seen here is typical.

39. Gale, “Eighteenth-Century Collections Online,” *Gale: A Cengage Company*, <https://www.gale.com/primary-sources/eighteenth-century-collections-online> (accessed 13 April 2019).

40. See Anders Søgaard, Anders Johannsen, Barbara Plank, Dirk Hovy, and Hector Martinez, “What’s in a P-Value in NLP?” *Proceedings of the Eighteenth Conference on Computational Language Learning* (Association for Computation Linguistics), 1–10.

41. See Donald A. Jackson, “Stopping Rules in Principal Components Analysis: A Comparison of Heuristical and Statistical Approaches,” *Ecology* 74, no. 8 (1993): 2204–2214.

Table 1. The Top Twelve Co-associated Words, Ranked by Their DPF, with “Despotism” at a Distance of Fifty in a Corpus of All Works Published from 1780 to 1800 in Eighteenth-Century Collections Online

| Bound word | DPF |
|----------------|----------|
| despotism | 14,978.4 |
| despotic | 3,928.8 |
| monarchy | 3,572.3 |
| arbitrary | 3,012.2 |
| discontinuance | 2,947.1 |
| governments | 2,726.7 |
| aristocracy | 2,557.7 |
| decaying | 2,548.0 |
| statesmen | 2,389.1 |
| revolution | 2,291.6 |
| repelling | 2,236.0 |
| tyranny | 2,222.7 |

ues are plotted against their rank on the list, one typically obtains a smooth curve (Figure 2) that is well fit by a power function. By fitting a power function to this curve and solving for where the derivative is equal to -1 , the DPF at which the “bend in the curve” occurs can be identified. This value serves as a threshold that separates the “short head” of the curve from the “long tail.” Next, we will refer to words occurring above the threshold as “bound” to the focal word and their specific DPF values as “strength of binding.” Table 1 presents an example of the data our method produces, here truncated at twelve co-associations, to give a sense of how we can begin to construct profiles of words and their bound lexis.

DPF is similar to PMI and many other measures commonly employed by computational linguists to quantify lexical asso-

ciation in that the specific identities or lengths of the documents in which the lexical co-associations occur do not figure in the calculation. Like most computational linguists, we see this as an advantage: the measure is sensitive to the number of co-associations, and not to the number of documents or contexts in which co-associations appear (which would require semi-arbitrary decisions to be made about what counts as a “document” or “context”). We should also note that we make no claim as to the specific importance of a distance or span: we investigate DPFs at large distances not to suggest there is anything deeply meaningful about a peak in DPF at some specific distance but rather to distinguish between co-associations that are more likely to be derived from syntax or the “syntax-semantics interface” from those where such derivation is likely to be weak. If binding is strong in such cases (at long spans), then we must posit a reason for this—hence, our interest in this as a possible index to a distinctive conceptual ontology.

Toward Conceptual Functionality

Here, we will introduce a second computational tool we have developed that we call the grapher, which supplements the aforementioned DPF ranker. Our methodology here follows what is emerging as “good practice” in digital humanities-based research: tool construction is driven in the first instance by research questions.⁴² The grapher also calculates the relative probability of the co-association of a focal word with one or more bound words, but it reports the binding coefficient (DPF) at every separate lexical distance in a given range: where the ranker offers a rapid survey of the whole environment in which a focal word is embedded (at a given distance), the grapher enables us to trace in more detail the force that the same word exerts on a finite number of specified bound words by calculating the relative probability of a bound word occurring.⁴³ This relative probability is plotted on the x-axis, while the y-axis plots each separate lexical distance within a specified range (in this case, two hundred words before and after the word). The grapher therefore presents a more fine-grained pattern of co-association across large lexical distances; this precision allows us to identify what we call preferential order distribution, where a bound word appears more frequently either before or after the focal word.

Figure 2 charts the co-association of “freedom” (the focal word) and “liberty.” Here we find a pronounced spike close to 0, that is, within five words either preceding or following the focal word, but also note that while the tendency for pronounced co-association declines as we move further out, it does not get close to reverting to random distribution (as indicated on the graph by the black line above the baseline). Thus, at a distance of one hundred words, “liberty” is likely to occur both before and after freedom almost seven times more likely than in random distribution. When we inspect the close-up distance, however, we find a clear preference: “liberty” is more likely to occur immediately before than immediately after “freedom.” Is this preference in the eighteenth-century data set indicative of a habit of thinking as well as a habit of writing for the period?

If we consider for a moment the semantic values of these two words, we might say “liberty” prepares the ground for “freedom,” where the latter term is to be taken not as a pure synonym but rather as some kind of adjunct or inflection. When we move from these data returns to the individual texts that cumulatively produce them, we see precisely such inflections. “By this

42. See Mark J. Hill, “Invisible Interpretations: Reflections on the Digital Humanities and Intellectual History,” *Global Intellectual History* 1, no. 2 (2016): 130–150; Jennifer London, “Re-imagining the Cambridge School in the Age of Digital Humanities,” *Annual Review of Political Science* 19 (2016): 351–373.

43. To preserve the data on raw frequency, we report them in brackets following the specified word.

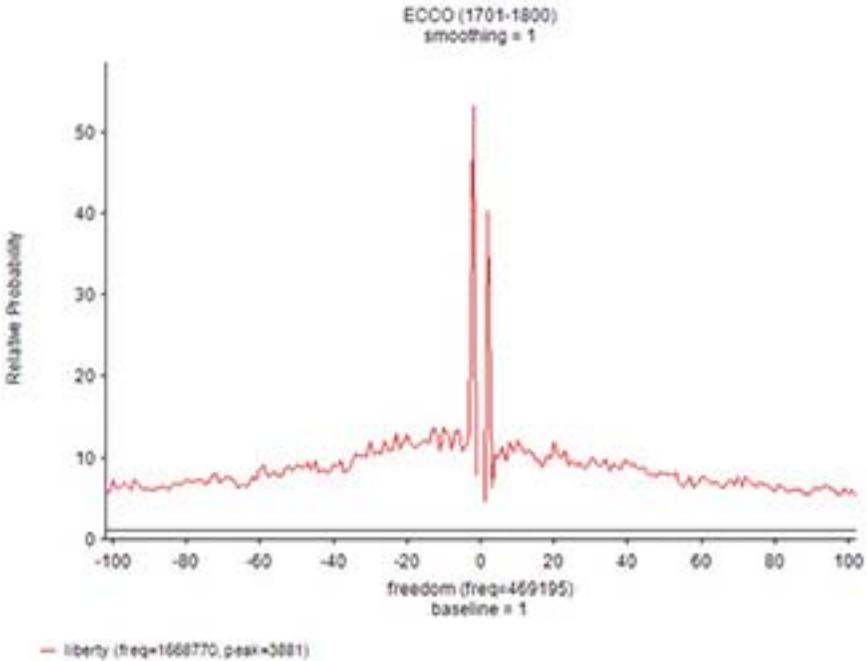


Figure 2. Relative probability of “freedom” and “liberty” co-associating plus or minus 100 words from all works published from 1700 to 1800 in Eighteenth-Century Collections Online, graph smoothed.

means our liberty becomes a noble freedom,” states Edmund Burke in his *Reflections on the Revolution in France*.⁴⁴ Alternatively, such a process may fail to take place, in which case the two terms become desynonymized: “This whole state of commercial servitude and civil liberty taken together is certainly not perfect freedom,” Burke similarly states in a 1774 speech on American taxation.⁴⁵ We posit that such semantic effects can be separated from the underlying distributional pattern, the preferred order distribution, which in this case functions as a kind of “seeding,” whereby one term prepares for its bound partner without prejudice to its semantic coloration (i.e., the function of preferred order can “seed” for both synonymic and antonymic freighting). This can be thought of as a conceptual analogue to what the linguist Michael Hoey calls “lexical priming.”⁴⁶

44. Edmund Burke, *Reflections on the Revolution in France*, ed. Conor Cruise O’Brien (London: Penguin, 1968), 86.

45. Edmund Burke, *The Works of Edmund Burke with a Memoir*, 3 vols. (New York: George Dearborn, 1835), 1:201.

46. See Michael Hoey, *Lexical Priming: A New Theory of Words and Language* (London: Taylor & Francis, 2005).

Let us take a slightly more complex example. Figure 3 charts the co-association of two words: “sensibility” (the focal word) and “irritability.” Here, we can compare the example of “liberty” and “freedom,” where the co-association is around seven times more frequent than in a random distribution at a distance of one hundred words before the focal word. In the case of “irritability” and “sensibility,” that frequency is one hundred times above random. The co-association is still more striking at close distance: we have cropped Figure 3 to make it legible, but the peak on the x-axis indicates a frequency 4,200 times more than what would be predicted in a random distribution. In this case, there is no evidence of preferential order distribution: the relative frequency of the bound word is comparable both ante- and post-focal word. Here, we can identify a different underlying pattern and function by noting the extent to which both words co-associate frequently not only with one another but also with a wide range of further common lexical material.

To capture this, we have collated successive searches using the ranker to reveal which bound words appear prominently at the lexical distances of five, ten, twenty, thirty, and a hundred words away. In collating bound terms across these different distances, we surely capture close-up syntactic and

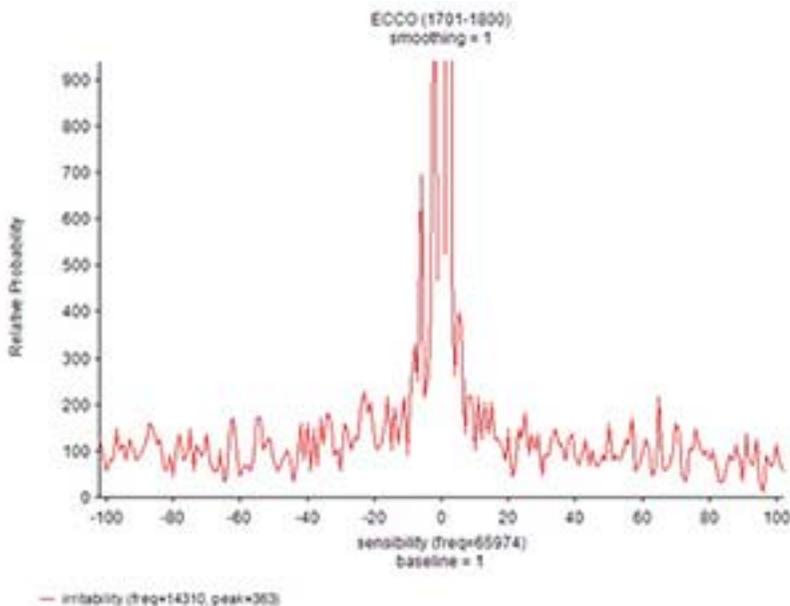


Figure 3. Relative probability of irritability and sensibility co-associating plus or minus 100 words from all works published from 1700 to 1800 in Eighteenth-Century Collections Online, graph smoothed.

grammatical ties between words, as well as looser ties that can be understood partly in terms of a common topic. But for our purposes of discerning putative functional attributes to concepts, we are less interested in the semantic values of individual words or the larger semantic field they comprise than in the data profile of such common sets of lexis. Once again, we are asking if we can discern an underlying pattern to lexical distributions and their accompanying binding profiles, which may point toward structures or

Table 2. Common set of Co-associated Terms with “Irritability” at Distances post or ante Five, Ten, Twenty, Thirty, and a Hundred Words Away in a Corpus of All Works Published from 1700 to 1800 in Eighteenth-Century Collections Online

| | 1701–1720 | 1721–1740 | 1741–1760 | 1761–1780 | 1781–1800 |
|--------------|-----------|-----------|-------------|--------------|---------------|
| irritability | 0 | 0 | animal | accumulation | abates |
| | | | animals | acrid | accumulated |
| | | | dr | aliment | accumulation |
| | | | fibres | animal | aion |
| | | | intestines | animals | animal |
| | | | motion | arteries | animals |
| | | | muscles | bark | applications |
| | | | muscular | bladder | arteries |
| | | | nerve | blisters | bark |
| | | | nerves | brain | bladder |
| | | | property | causes | blisters |
| | | | sensibility | ceases | blood |
| | | | system | chronic | bowels |
| | | | | circulation | brain |
| | | | | contraction | causes |
| | | | | convulsions | chronic |
| | | | | cough | circulation |
| | | | | debility | constitutions |
| | | | | degree | contraction |
| | | | | delirium | convulsions |
| | | | | disease | cutaneous |
| | | | | +53 | +78 |

functions that are not entirely determined by meaning. Table 2 charts those bound words with “irritability” across the eighteenth century (note that the lists for 1761–1780 and 1781–1800 have been cut off; both continue for another fifty-three and seventy-eight words, respectively) subdivided into specific historical periods (1700–1720, 1721–1740, 1741–1760, 1761–1780, and 1781–1800). The data for sensibility has a significantly different tenor (in this case, the lists are complete).

Such a protocol enables us to identify which focal words consistently bind a large amount of lexical material, and to track this binding profile diachronically. At the start of the century, neither has a common bound set across distance; by the century’s close, however, “sensibility” has a common bound set of fifteen words across each of the specified distances. The data for “irritability,” by contrast, indicate a still more pronounced transformation: having no common bound set up to 1740, we find, by the final date slice 1781–1800, a set of ninety-nine bound words. Once again, we note this certainly provides evidence for a semantic history, but we wish to look away

Table 3. Common Set of Co-associated Terms with Sensibility at Distances post or ante Five, Ten, Twenty, Thirty, and a Hundred Words Away in Corpus of All Works Published from 1700 to 1800 in Eighteenth-Century Collections Online

| | 1701–1720 | 1721–1740 | 1741–1760 | 1761–1780 | 1781–1800 |
|-------------|-----------|-----------|--------------|--------------|--------------|
| sensibility | 0 | delicate | feeling | delicacy | affection |
| | | | irritability | delicate | feeling |
| | | | sensibility | emotion | feelings |
| | | | | feeling | her |
| | | | | impressions | impressions |
| | | | | irritability | irritability |
| | | | | irritation | morbid |
| | | | | muscles | nerves |
| | | | | nerves | nervous |
| | | | | nervous | pathetic |
| | | | | organs | sensation |
| | | | | sensation | sensibility |
| | | | | sensations | sympathy |
| | | | | sympathy | temperament |
| | | | | temperament | tenderness |

from the meanings of these words and their co-associates to ask if we can discern a particular way of behaving, a functionality, that sits, as it were, under or beside the level of semantics.

Looking at Tables 2 and 3, we see “sensibility” has begun consistently to bind only three words (variants on itself and irritability) by 1741–1760, whereas “irritability” binds several terms that comprise a discourse of physiology—nerves, muscles, organs, and so forth. Looking ahead to 1761–1780, one can see “sensibility” has accrued much of the same lexical material that had constellated around “irritability” much earlier; by this date, both words circulate in a recognizable discourse of affect that includes “sympathy,” “emotion,” and “feeling.” One way of understanding these shifts in semantic fields is to consider the wider contexts within which this language operated that would include, for example, a historical account of the development of psychology as a scientific discipline. We believe these data could certainly contribute to current conceptual histories that raise such research questions. But we mean to propose something more radical: what if the coalescence of such lexical units—call them the “bundles” or “constellations” of words that build concepts—are determined not only by sense, meaning, or topic but also by habits of cultural thinking, conceptualization, based on a small suite of functional operations? If we consider concepts as tools to do something with—as Wittgenstein did—what functions do these tools have, and do they change over time?

Toward a History of Conceptual Functioning

Finally, we will outline some ways in which our experimental methodology may contribute to an enriched history of concepts that tracks both their changing meanings and their altering functionalities. The following figures are derived from the same calculations we have made to discern the DPF of terms in our data set (ECCO) with the distance set at one hundred. In this case, however, we have exported the data into a network environment that helps us identify complex patterns of similarity and connection between data points, here represented as words in our data set. The figures are based on screenshots from these network representations displayed in an interactive web application. These are “neighborhood” or “ego” graphs of order two; that is, they show the graph containing nodes within at most two steps from a specific focal node or nodes entered by the user. An edge exists between two nodes if their DPF association is above a user-specified value.

The network is drawn using the Three.js R package, with a force-directed layout algorithm that models the network mechanically as repelling

particles connected by a spring.⁴⁷ The result is that, in a graph of suitable density and degree, nodes are spaced apart enough to be distinguished, but the edges pull together nodes into clusters that share many relations. Community structure is detected with an unsupervised modularity optimization algorithm and indicated by color. For visual clarity, nodes with only a single connection (degree one nodes) are not displayed. As we learn to interpret these graphs, we wish to draw attention to their shape and structure as much as to the lexical labels that are mapped in relation to each other. Our working hypothesis is that “concepts” are most usefully understood as larger units than words (they are like semantic fields): in this case, the networks and clusters presented by Figures 4–7. Furthermore, these plots enable us to see how these larger units operate or function with respect to each other and over time.

As we can see from Table 2, the common set of bound words across distance for “irritability” increases in size as the century lengthens and the semantic freighting of the growing network—a semantic field—is clear enough: a language of nerves, fibers, organs, and so forth. When we plot our data in the format of a network graph, we should find precisely this displayed, as indeed we do. But we believe the following figures also provide historical evidence for an alteration in function whereby “irritability” and its associated network of closely bound terms grows in density, generating subnetworks or clusters within its close orbit and, at the same time, holds off at a distance another network within which “sensibility” operates.

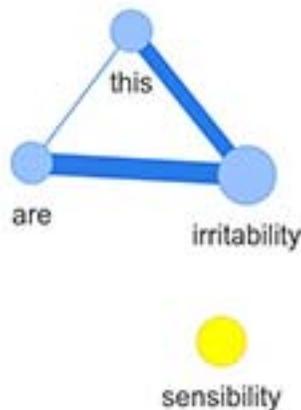


Figure 4. “Irritability” and “sensibility,” DPF at distance 100, from all works published from 1750 to 1760 in Eighteenth-Century Collections Online. Network plot, top twenty ranked terms.

47. See Thomas M. J. Fruchterman and Edward M. Reingold, “Graph Drawing by Force-Directed Placement,” *Software: Practice and Experience* 21, no. 11 (1991): 1129–1164.

As we would expect from the data in Tables 2 and 3, “irritability” and “sensibility” should have no overlap or connection up through the 1740s (Figure 4). During the following decade, we begin to observe a network becoming established; here, we can see the visualization software draws a thick line between “irritability,” “muscles,” “nerves,” and “sensibility,” which corresponds to a high degree of relation (i.e., high DPF value). We can think of this as the “core” of the conceptual network that includes all the terms in Figure 5. In the next decade, this strong connection between three of these terms (“irritability,” “sensibility,” and “nerves”) decreases as “fibre” comes to make the link or “bridge” the “irritability” cluster and the “sensibility” cluster, and “muscles” drops out entirely. It is also noteworthy that a new cluster including “nerve(s)” and “sensation” becomes established with the strength of connection between “fibre,” “nervous,” and “nerves” indicated by the thick line (Figure 6). By the last decade of the century, something very eye-catching happens: the network within which “irritability” operates

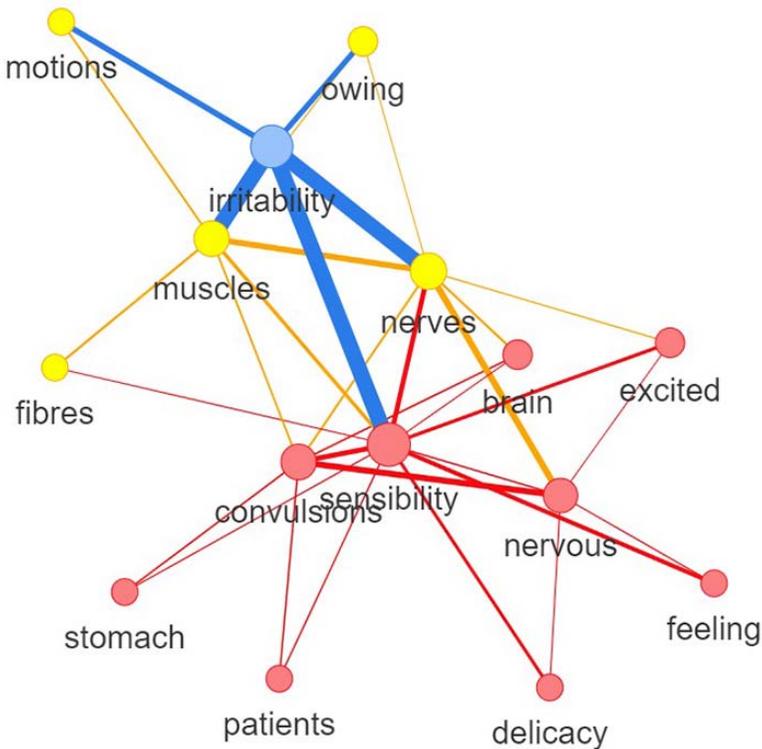


Figure 5. “Irritability” and “sensibility,” DPF at distance 100, from all works published from 1760 to 1770 in Eighteenth-Century Collections Online. Network plot, top twenty ranked terms.

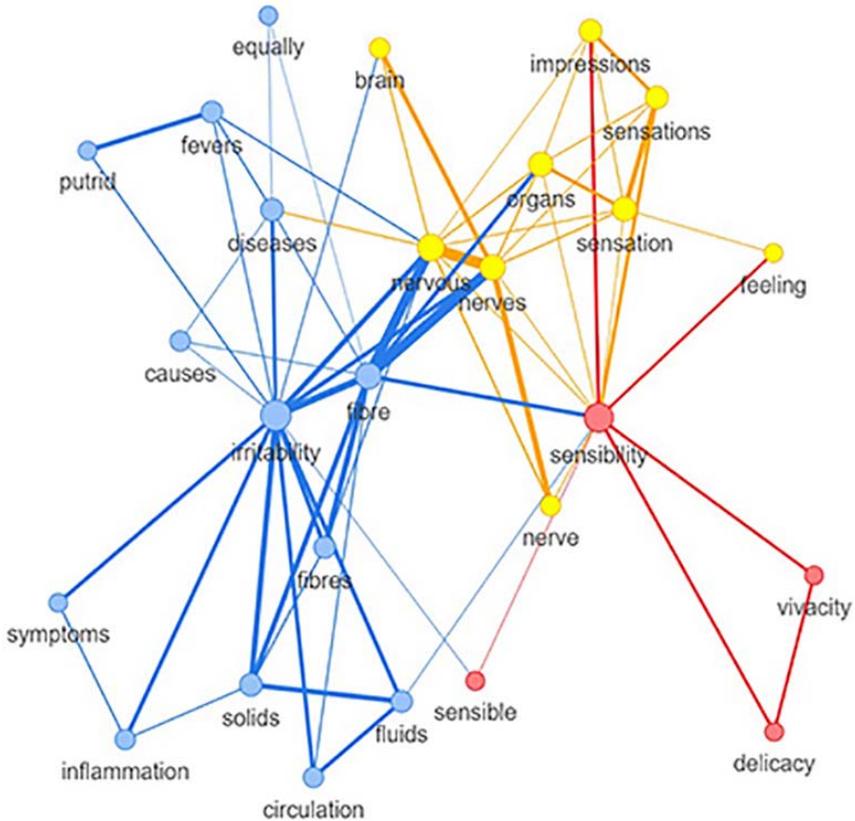


Figure 6. “Irritability” and “sensibility,” DPF at distance 100, from all works published from 1770 to 1780 in Eighteenth-Century Collections Online. Network plot, top twenty ranked terms.

becomes denser and more complex, and its connection to “sensibility” becomes etiolated and weaker.

It is unwise to read these network plots as directly representing changes in a kind of relation—essentially reading them as substantives rather than as conveniences according with the underlying mathematics that takes complex multidimensional vectors and fits them to a visually legible form. Notwithstanding this, the structure of these plots is based on invariant data: they are telling us something. Here, the elongation of the distance between the node “irritability” and “sensibility” corresponds to a weakening of DPF values and, we believe, indicates a particular force of repulsion or distancing is operating. We may think of irritability and its immediate cluster

Conclusion

This article is speculative and intended to be exploratory. The construction of the computational entities we have presented has been based on data derived from lexical distributions in a massive data set of printed materials from the Anglophone eighteenth century in order to test the hypothesis that “conceptual functions” such as “seeding” or “repulsion” become inspectable. The diachronic metadata of the data set have not been exploited in full in our examples, but in two additional articles, we have explored how our methodology can address more standard questions in the history of concepts.⁴⁸ The aim of this article is to make a theoretical intervention within the field of conceptual history to turn exclusive attention away from the meanings of concepts toward a more heterogeneous project that includes inquiry into their function and structure.

Peter de Bolla is Professor of Cultural History and Aesthetics at the University of Cambridge, and the Director of the Cambridge Concept Lab.
Email: pld20@cam.ac.uk

Ewan Jones is Lecturer in Nineteenth-Century Literature in the Faculty of English at the University of Cambridge.
Email: ejj25@cam.ac.uk

Paul Nulty is Marie Skłodowska-Curie Career-FIT Fellow at University College Dublin.
Email: paul.nulty@gmail.com

Gabriel Recchia is Research Associate at the Winton Centre for Risk and Evidence Communication in the Department of Pure Mathematics and Mathematical Statistics at the University of Cambridge.
Email: glr29@cam.ac.uk

John Regan is Lecturer in Literature and Digital and Creative Industries at Royal Holloway, University of London.
Email: john.regan@rhul.ac.uk

48. See de Bolla et al., “Idea of Liberty”; Peter de Bolla, Ewan Jones, Gabriel Recchia, John Regan, and Paul Nulty, “The Conceptual Foundations of the Modern Idea of Government in the British Eighteenth Century: A Distributional Concept Analysis,” *International Journal for History, Culture and Modernity* (forthcoming).