

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/320195474>

Finding Talk About the Past in the Discourse of Non-Historians

Article · October 2017

CITATIONS

0

READ

1

3 authors, including:



Alex Olieman

University of Amsterdam

14 PUBLICATIONS 31 CITATIONS

SEE PROFILE



Jaap Kamps

University of Amsterdam

325 PUBLICATIONS 4,125 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



Social Book Search [View project](#)



TREC Reports [View project](#)

All content following this page was uploaded by [Alex Olieman](#) on 31 October 2017.

The user has requested enhancement of the downloaded file.

Finding Talk About the Past in the Discourse of Non-Historians

Alex Olieman
University of Amsterdam
olieman@uva.nl
Stamkracht BV
alex@stamkracht.com

Kaspar Beelen
University of Amsterdam
k.beelen@uva.nl

Jaap Kamps
University of Amsterdam
kamps@uva.nl

ABSTRACT

A heightened interest in the presence of the past has given rise to the new field of memory studies, but there is a lack of search and research tools to support studying how and why the past is evoked in diachronic discourses. Searching for temporal references is not straightforward. It entails bridging the gap between conceptually-based information needs on one side, and term-based inverted indexes on the other. Our approach enables the search for references to (intersubjective) historical periods in diachronic corpora. It consists of a semantically-enhanced search engine that is able to find references to many entities at a time, which is combined with a novel interface that invites its user to actively sculpt the search result set. Until now we have been concerned mostly with user-friendly retrieval and selection of sources, but our tool can also contribute to existing efforts to create reusable linked data from and for research in the humanities.

Keywords: Colligatory Concepts, Semantically-Enhanced Search, Interactive Information Retrieval, Corpus Selection, Digital Humanities

1 INTRODUCTION

There has been a monumental shift from the future to the past in the cultural orientation of Western societies, starting in the 1980s. In a sense, the past has increasingly gained in presence: in the literary and artistic expressions of (traumatic) memories that cannot be contained by the evidence that forms the basis of historical studies, the proliferation of museums and archives, and the commodification of the past as marked by docudramas, historically-themed amusement parks, and memorabilia of pasts that never existed [3]. Many of the resulting representations of segments of the past serve a deeply different purpose than the representations created by professional historians. Under the influences of positivist science and literary realism in the 19th century, history was detached from “its former habitation in rhetoric” [10, p. 10] and has developed a rigorous manner of dealing with sources and evidence that leads to the production of distinctly historical accounts of past events. While changes in how historians have interpreted past events has been widely studied (e.g. under the banner of historiography), we are still in the early days of developing methods to analyze how

other groups, such as journalists and politicians, have incorporated the past into their narratives.

Identifying and interpreting the presence of the past in the discourse of non-historians is worthwhile, not because we expect to establish new facts about past events, but rather because laypersons and practitioners of other disciplines draw upon the past to make all kinds of judgments and decisions in daily life [10]. The availability of large diachronic corpora has opened new avenues to study how and why people make reference to the past in their discourse (e.g. to convince others or to express emotion), and to analyze differences within particular time intervals as well as across time. The size of such corpora, however, combined with the scatteredness of references to the past, can make these corpora daunting to explore without the right tools.

In response to the “spatial turn” in Digital Humanities, substantial effort has been put into the development of tools that allow for spatial navigation through text collections. In the Pelagios project, for example, the Pleiades gazetteer serves to anchor locations mentioned in text to machine-readable representations of these locations, which can be combined with linked data to form rich map-based visualizations and allows for spatial access to the texts through the Peripleo search interface [5]. The recognition that “space and time are no more separate in human cognition than they are in theoretical physics” [5, p. 43] now motivates the development of tools that provide access to texts by the historical entities that they reference, to complement the evolving spatial approaches.

In this paper, we propose an approach to support researchers who aim to identify and interpret (indirect) references to historical periods in a particular discourse. It consists of a semantically-enhanced search engine that is able to find references to many entities at a time in diachronic corpora, which is combined with an interface that invites its user to actively sculpt the search result set. As several pilot studies have indicated [4], this approach is useful for those studying how groups of people pragmatically feature historical entities in discourse, remember or commemorate them, and understand their own identity in relation to these entities. Moreover, we are able to capture in linked data how definitions of historical periods are operationalized by researchers, and generate semantic annotations for the search results that users consider to be relevant.

2 SEARCHING FOR COLLIGATORY CONCEPTS

The search for references to historical periods entails bridging the gap between conceptually-based information needs on one side, and term-based inverted indexes on the other. When a researcher is looking for the fragments of text that refer to, e.g., the French

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

Drift-a-LOD'17, September 2017, Amsterdam, Netherlands

© 2017 Copyright held by the owner/author(s).

ACM ISBN 123-4567-24-567/08/06...\$2/3

https://doi.org/10.475/123_4

Revolution in a specific collection, they are not served well by providing only the documents that contain the literal phrase “French Revolution.” This is the case because such periods do not pre-exist in reality, waiting to be discovered and named, but rather come into being when the disparate observable elements of a phenomenon are “seen together” as a synthetic whole [1]. Philosophers of science and of history have called this process colligation—a binding together of selected historical facts, culminating in the proposal of a “colligatory concept” which represents the historian’s understanding of the facts as an ‘entwined whole’ in a form that can be communicated to others [7].

Periods are but one kind of colligatory concept that is commonly constructed by historians; the others being characters, such as ‘Louis XVI’ and ‘the French people,’ and ideal types, for instance ‘capitalism’ and ‘revolution’ [7]. Whereas characters are localized in time and space, and ideal types bind together certain aspects of the events in which these characters participate across time and space, periods are bound together by narratives that feature selected events which are distributed within (flexible) temporal and spatial bounds [7]. The PeriodO project has already produced a linked data gazetteer of periods as they are represented in the published works of historians and historiographers, using a nanopublication approach that connects the qualitative definition of a period by an individual author to its name, spatial extent, and a time interval, expressed together in RDF [1, 5].

This model captures accurately the multivocal aspect of period definitions, but is not sufficient to retrieve references to these periods in discourse that is about another subject. Here we are confronted with the difference between colligatory concepts which are necessary constructs in the writing of history, and the *subjects* that are intended to group together multiple histories that exhibit common “patterns of colligation” [7, p. 1098]. As a colligatory concept, a period is a particular representation of the past that binds together historical entities in a unique narrative. When this individual representation leads to further discourse, the discourse as a whole is not about the original colligation, but rather about a homonymous subject that allows us to, e.g. ask a librarian for ‘novels set during the French revolution.’ In the practice of information organization, we establish common referents and shared structure between colligations in order to group a multiplicity of perspectives under a single label [7].

We have no way of searching directly for the period that the researcher has in mind (it exists only as a cognitive representation), but by starting from its associated subject we can make an educated guess about the elements of the period. Named entities—events, people, artifacts—are central in our search approach, because they serve as the points of consensus that enable the search for as of yet unknown perspectives on segments of the past. The aim is to collect (and represent) an intermittent discourse, consisting of sentences that make (indirect) reference to the target period of the search, with as much context as is needed to understand them. To be sure, we do not equate the colligatory concepts that are put forward by historical work with the scattered references to the past that are found in non-historical discourse. Rather, we expect the (re)searcher to have a particular period (i.e. colligatory concept) in

mind, and aim for the system to find references to the entities that are bound together by this period.

3 APPROACH

The task that we aim to facilitate—selecting a research corpus of text fragments that refer to a particular period—is not supported well by either subject indexing [7] or full-text search [2], which nowadays provide the primary means of access to much of the text in archives and libraries. Subjects are traditionally assigned to whole documents. Our task, however, depends on the retrieval of individual sentences (and their context), given *their subjects*. It seems infeasible to provide such granular access with manually assigned subjects, but recent advances in Entity Linking have enabled the automated identification of the referents in individual sentences. Even though entity linking systems are prone to errors that human annotators would not make, the entity links that they produce can still be useful to search for many entities simultaneously [4]. Semantically-enhanced search is made possible by incorporating entity links or similar semantic annotations into search indexes [2]. Our approach extends this practice by employing linked data to bridge the semantic gap between the search target (i.e. a period) and references to the entities that are bound together by this concept.

3.1 Bootstrapping with DBpedia

Some groupings of entities that people might want to search for already exist as Linked Open Data. The category network of Wikipedia serves an analog function to that of the subject access systems of librarians and archivists. It is a knowledge organization system (KOS) that relates more specific subjects to more general ones by *broader-than* relations [8]. As the product of the contributions of a diverse community, it encodes multiple perspectives on the world in a single structure which allows for multiple paths to exist between any two subjects. We used DBpedia’s RDF representation of this category network for our proof-of-concept, because we were working with a corpus that we enriched with entity links that point to DBpedia URIs. Its simple structure, represented in the SKOS ontology, may not be ideal for all research purposes, but it is sufficient to start a search process with a period-as-subject and invite the researcher to operationalize the particular period he/she wants to search for.

In order to obtain possible mappings between periods of interest and the entities that are bound together by such periods, we extracted a subgraph of DBpedia, corresponding to Wikipedia’s category network and its related entities, into a property graph database (see [6]). The category network is used at runtime to select potentially relevant entities given a root category, by traversing `skos:broader`¹ and `dct:subject` relations in reverse direction. Starting, for example, from the root category `dbc:French_Revolution`, the traversal would proceed through subcategories such as `dbc:Montagnards` and `dbc:French_First_Republic` to collect entities including `dbr:Reign_of_Terror`, `dbr:Maximilien_Robespierre`, `dbr:Bastille`, and `dbr:Drownings_at_Nantes`.

Our proof-of-concept makes use of DBpedia, but any knowledge graph that conforms to the SKOS ontology can be loaded easily.

¹For namespace prefixes, see <https://dbpedia.org/sparql?nsdecl>.

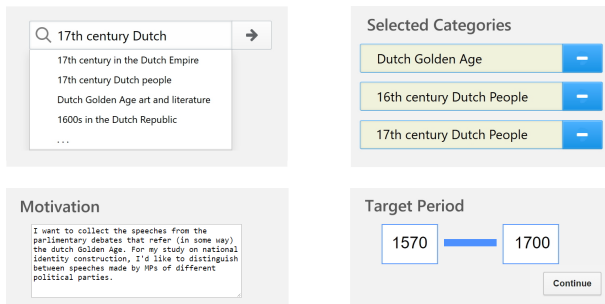


Figure 1: Initial query specification in WideNet.

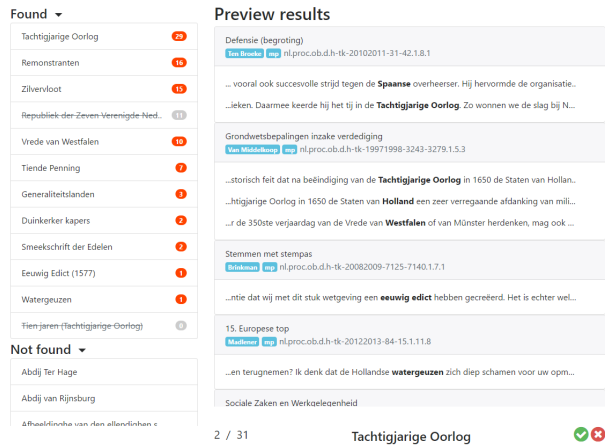


Figure 2: Assessing the relevance of categories and entities.

Linked data that is structured differently can also be used, as long as a grouping or categorization that is familiar to the intended users can be derived from it. It is important that the representations of entities in this data are identifiable by the same URIs as those used in the entity links in the corpus, to be able to connect periods, via colligated entities, to the text fragments that refer to these entities.

Finally, the system needs access to coarse temporal clues about entities. Because DBpedia does not provide this data reliably across entity types, we extract mentioned years from the `rdfs:comment` values of DBpedia resources with a simple regular expression, and add them to the graph. The same technique may be successful for other linked data sources that provide textual descriptions which often include temporal expressions. It would be preferable, however, to use representations that incorporate structured temporal relations in the form of RDF literals for all entities.

3.2 Search Interface Design

Our proof-of-concept, named WideNet [4], provides access to a semantically enriched version of the Dutch parliamentary proceedings—the “verbatim” records of the debates that take place in the Houses of Parliament (the *Staten Generaal*). These discussions touch almost on every issue that moved Dutch public opinion over more than two centuries.




Figure 3: A closer look at the retrieved documents.

The interface guides users through three research phases: (1) selection of root category, (2) assessment of the categories’ and entities’ relevance, and (3) close-reading. In the first step the user selects one or several root categories from a typeahead search box (see Figure 1), and demarcates the query by selecting a time period, which is used to prune the underlying entities of the selected categories. WideNet subsequently retrieves the network of narrower categories for each selected root category, and collects the contained entities as potentially relevant query components. Behind the scenes, each entity is compared with the target period, and is considered to be outright relevant to the period, or not, or a borderline case, or as lacking temporal clues altogether. In the current implementation this classification is achieved with simple rules, based on the features: ‘fraction of years within period,’ ‘fraction of intervals that overlap with the period,’ and ‘has at least one year in period.’ The system uses this information to deselect (sub)categories where more than half of the dated member entities are out-of-period, i.e., those categories are excluded from the query.

The next step for the WideNet user is to assess which of the retrieved subcategories actually contain entities that lead to relevant results. By doing this, researchers can operationalize their own definition of their target period, at least for the purpose of retrieval. The interface facilitates this task by showing, per subcategory, which entities are mentioned in the corpus, and how frequently, as well as which entities did not occur (see Figure 2). It also displays a list of preview results, showing limited context, to offer quick clues about the relevance of the category. This preview is also useful to identify individual entities that are not relevant after all, which can be deselected by the user. At the end of this step, the researcher’s decisions amount to a motivated and organized representation of ‘relevant’ or ‘possibly present’ elements of the target period.

After inspecting and selecting relevant categories of entities—thereby sculpting the final query, the WideNet interface allows further scrutinizing of the sources by providing an environment in which the retrieved documents can be studied up-close, as shown in Figure 3. By situating the close reading activity within the same interface, the user is able to compile a corpus of relevant documents which may be saved and exported. Moreover, the user can examine

the results in relation to the document metadata, e.g. to look for saliency by plotting the annotations over time, or to study bias by comparing how often different political parties refer to the entities of interest. The selected corpus, representing a fragmented discourse, may also be used to analyze changes in how and why the demarcated segment of the past was evoked, establishing contrasting perspectives and trends over time, provided that enough references could be found.

3.3 Capturing Reusable Data

As a product of the search, semantic annotations are created which link the source documents to the entities that are referenced in the corpus. These expert-approved annotations have much more value than the automatically generated entity links. For one, the identity of the referents is established with a greater confidence when a user chooses to include a particular document fragment in his/her research corpus. We cannot interpret this as a direct assessment of the entity links, but when a user has confirmed that the document fragment (indirectly) refers to the target period, it is tempting to assume that the document was retrieved because the entity links are correct, and not by some mistake. To model the provenance of such relevance decisions, it is necessary to produce a representation of the broader context of the search process to which individual assertions can be explicitly related (see [1]).

In the first screen that users encounter when starting a new search process in WideNet (depicted in Figure 1), we are able to capture the motivation and a rough demarcation of the search. This information forms the backbone of a representation to which subsequent assertions can refer. The selection of categories and entities (in Figure 2) provides assertions about their relevance for the search target, according to the researcher and dependent on the corpus. We currently store each decision that is made by the researcher, so when a previously made decision is revisited, we create representations of e.g. both the act of deselecting an entity, and reselecting it after more preview results had been inspected. Capturing this process data, rather than only the final selection of entities, benefits the richness of the provenance of the subsequent assertions about the relevance of text fragments.

The semantic annotations that are derived from relevance assertions on text fragments can provide richly indexed access to statements about the past. This notion is similar to Ryan Shaw's proposal for "deep gazetteers," in which multiple descriptions of the same named entity are linked to fragments of discourse in which its name is used [8]. In our case, however, we use representations of periods-as-subjects to connect particular conceptions of these periods to discourse that refers to these periods from a different perspective, rather than use of the same name per se. Our approach can produce "projections" of the parts of the past that were present in a particular discourse, which can be organized according to the KOS that was used to perform the search, but also according to any other KOS that includes representations of the same entities, provided that equivalence of identity (e.g. owl:sameAs) can be established between the knowledge organization systems.

While we capture the described representations of search processes and semantic annotations in our current proof-of-concept, we do not yet publish this data. Our approach is related to ongoing efforts to produce reusable data for research in the humanities

(e.g. [1, 8, 9]) and we are still investigating how we can best link to the existing models. We have prioritized the development of a useful search tool over the production of reusable data in order to investigate which data can be captured during actual research in the humanities, rather than designing our models first and finding out later that they are not as usable as we had hoped [9].

4 CONCLUSION AND OUTLOOK

In searching diachronic corpora for fragmented discourses about historical periods, we need to come to a fundamental understanding of conceptual drift over time. Our approach for finding references to periods consists of a semantically-enhanced search engine that is able to find references to many entities at a time in diachronic corpora, which is combined with an interface that invites its user to actively sculpt the search result set. Besides yielding sources that are useful for the searcher, the search tool also produces an operationalization of the search target by the (re)searcher as well as semantic annotations that are much richer than those that we can generate algorithmically.

Although a pragmatic treatment of concepts is sufficient to search for multifaceted subjects, we envision how the products of such search processes can collectively provide a shared source for more elaborate knowledge representations. In designing a tool that implements this approach, we were faced with some trade-offs between usability of the tool and reusability of the data it captures. We prioritize supporting the present-day researcher well, and facilitate publishing the search process and its results as linked open data, so that subsequent refinement of the captured data may be a community effort.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their suggestions and remarks. This work was supported by the Netherlands Organization for Scientific Research (ExPoSe project, NWO CI # 314.99.108).

REFERENCES

- [1] Patrick Golden and Ryan Shaw. 2016. Nanopublication Beyond the Sciences: the PeriodO period gazetteer. *PeerJ Computer Science* 2:e44 (2016). <https://doi.org/10.7717/peerj-cs.44>
- [2] Annika Hinze, David Bainbridge, Sally Jo Cunningham, and J. Stephen Downie. 2016. Low-cost Semantic Enhancement to Digital Library Metadata and Indexing. *Proceedings of JCDL '16* (2016), 93–102. <https://doi.org/10.1145/2910896.2910910>
- [3] Andreas Huyssen. 2000. Present Past: Media, Politics, Amnesia. *Public Culture* 12, 1 (2000), 21–38.
- [4] Alex Olieman, Kaspar Beelen, Milan van Lange, Jaap Kamps, and Maarten Marx. 2017. Good Applications for Crummy Entity Linkers? The Case of Corpus Selection in Digital Humanities. In *Proceedings of SEMANTICS 2017*. arXiv:1708.01162
- [5] Adam Rabinowitz, Ryan Shaw, Sarah Buchanan, Patrick Golden, and Eric Kansa. 2016. Making sense of the ways we make sense of the past: The PeriodO project. *Bulletin of the Institute of Classical Studies* 59, 2 (2016), 42–55.
- [6] Marko Rodriguez. 2015. The Gremlin Graph Traversal Machine and Language. In *Proc. 15th Symposium on Database Programming Languages*. 1–10. arXiv:1508.03843
- [7] Ryan Shaw. 2013. Information Organization and the Philosophy of History. *Journal of the American Society for Information Science and Technology* 64, 6 (jun 2013), 1092–1103. <https://doi.org/10.1002/asi.22843>
- [8] Ryan Shaw. 2016. Gazetteers Enriched: A Conceptual Basis for Linking Gazetteers with Other Kinds of Information. In *Placing Names: Enriching and Integrating Gazetteers*. 51–63.
- [9] Ryan Shaw, Patrick Golden, and Michael Buckland. 2015. Using Linked Library Data in Working Research Notes. In *Linked Data and User Interaction*. 48–65.
- [10] Hayden White. 2014. *The Practical Past*. Northwestern University Press.