

Hacking History: Automatic Historical Event Extraction for Enriching Cultural Heritage Multimedia Collections*

Roxane Segers¹, Marieke van Erp¹, Lourens van der Meij¹, Lora Aroyo¹, Guus Schreiber¹, Bob Wielinga¹, Jacco van Ossenbruggen^{1,2}, Johan Oomen³
and Geertje Jacobs⁴

¹ VU University Amsterdam

² Centre for Mathematics and Computer Sciences (CWI)

³ Netherlands Institute for Sound and Vision

⁴ Rijksmuseum Amsterdam

Abstract. Within cultural heritage collections, objects are often grounded in a particular historical setting. This setting can currently not be made explicit, as structured descriptions of events are either missing or not marked up explicitly. This poster reports a study on automatic extraction of an historical event thesaurus from unstructured texts. We also present a demo in which relations between events and museum objects are visualised to accommodate event- and object-driven search and browsing of two cultural heritage collections.

1 Introduction

Events have recently gained attention in the knowledge representation community as valuable constructs [4, 7, 8] that can help tie together relevant but yet unrelated elements of information. In the cultural heritage domain, knowledge about historical events is often concealed in textual descriptions that can only be accessed via keyword search. As such, the available knowledge can not be reused across collections as it is not part of the shared metadata and controlled vocabularies.

In this study, we investigate how historical events in unstructured text collections can be captured and modeled to create an event thesaurus for enriching metadata in cultural heritage collections. We adopt the SEM event model [8] to distinguish event types, actors, locations, and dates. We experiment with natural language processing (NLP) techniques to extract event names and their associated actors, dates and locations. Additionally, we show how this resulting preliminary event thesaurus is employed in a new platform for event- and object driven search and browsing of the collections of the Rijksmuseum Amsterdam (RMA) and the Netherlands Institute for Sound and Vision (S&V).

* This work was previously presented as a poster at The Sixth International Conference on Knowledge Capture (K-CAP 2011)

Slachtoffers gemaakt door de Nederlandse troepen op weg naar Jogjakarta¹ (Object)



Slachtoffers gemaakt door de Nederlandse troepen op weg naar Jogjakarta. Kinderschilderij van de inname van Jogjakarta tijdens de tweede politionele actie, december 1948.
NG-1998-7-10

Associated Events
DepictsEvent: [Tweede politionele actie¹](#)

biographical aspects
Creator: Toha Adimidjojo, Mohammed (4) **Date:** 1948-12-18 (3) 1949-06-30 (3) 20e eeuw (18) tweede kwart 20e eeuw (17)

material aspects
Type: aquarel (3) tekening (3)
Technique: aquarelleren (3)
Material: hardboard (4)

semiotic aspects
Subject: Jogjakarta (4) Tweede politionele actie (7)
1948-12-18 (4) 1949-06-31 (1)
militaire geschiedenis (12)

Associated Objects (25) [< prev](#) [1](#) [2](#) [3](#) [4](#) [5](#) [next >](#)


[President Soekarno o.
Associated Press](#)


[Sintin nanjano met s...
Anonymous](#)


[Indonesië vrij
Hatta, Mohammed](#)


[Schild van een Alahier
Anonymous](#)


[Aankomst van Van Soe...
Anonymous](#)


[Het kasteel van Bata...
Reuchman, Andras](#)

Fig. 1. Screenshot of object page in the Agora Event Browsing Demonstrator

2 Event Extraction from Text

As no annotated historical document collections exist in Dutch, our approach is focused on extracting named events with minimal manual effort. For this study we selected 3,724 historical Wikipedia articles as a test set. The event extraction process consists of three steps: in the **first step**, we recognize *actor names* and *locations* using the Stanford Named Entity Recognition system [2] adapted for Dutch historical texts. Dates were recognized via regular expressions. This step resulted in 18,623 candidates for actors (F-measure of 0.77), 7,023 locations (F-measure of 0.66) and 7,981 dates. In the **second step**, we use a pattern-based method for recognizing *event names* such as *French Revolution*. We harvest patterns from the Web (e.g., *destroyed during the, before the*) using the Yahoo! search API ⁵ and a seed set of one hundred historical events. Patterns are ranked by frequency of co-occurrence with two or more seed events [6]. To retrieve event candidates, we applied the patterns to the Wikipedia corpus. The event candidates are then filtered, based on a threshold on the pattern score, resulting in a set of 2,444 unique events. The precision score of this set is 56.3%.

In the **third step**, we associate events with actors, locations and dates. We experiment with both redundancy and co-occurrence of data on the Web, inspired by the work of Geleijnse et al. [3] and Cilibrasi & Vitanyi [1]. Each combination of an event name and actor/location/date is sent to Yahoo! and for each pair a score is computed. We discovered 392 event names that were paired with an actor, a location and a date. Through manual evaluation we conclude the following: 71.9% (323) are correct event names, 45.6% (179) are correct actors, 41.1% (161) are correct locations and 51.5% (202) are correct dates.

⁵ <http://developer.yahoo.com/search>

3 Enrichment by Events

The extracted events are linked to the RMA and S&V collections. In total 35 unique events provide direct relations from 435 S&V objects to 675 RMA objects. An additional 34 unique events provide links from 391 S&V objects to 362 RMA objects, but this link exists indirectly through the event instance (e.g., S&V object - Actor - RMA object). We hypothesize that these links are potentially useful for navigating cultural heritage collections.

4 The Agora demonstrator

The automatically generated event thesaurus is applied in a new historical event browser called Agora⁶ which provides an integrated access route to museum objects and audio-visual material from RMA and S&V respectively. It is a first step towards a platform to investigate the added value of historical events and narratives for the exploration of integrated collections. For each event and object there is an automatically generated page that shows (1) all associated objects, e.g., museum and audio-visual objects; (2) all associated events and the type of their relationship, e.g., previous-in-time event, sub-event; (3a) the event descriptive metadata, e.g., actors, place, period; or (3b) object descriptive metadata organized in three groups, e.g., biographical, material and semiotic dimensions – see figure 1 for a screenshot –and finally (4) the navigation path. The current version of the event thesaurus will be extended further to accommodate searching for relations between events such as temporal inclusion, causality and meronymy.

5 Discussion

In this paper, we presented a modular pipeline for capturing knowledge about historical events from Dutch texts. Compared with previous approaches (i.e., [5]), it relies on a minimum of manual annotation and can be repurposed for other languages. To the best of our knowledge, this is the first work to extract events from unstructured Dutch text. Although our results are promising, more sophisticated techniques are necessary to obtain more fine-grained extractions and define measures for the historic relevance of the extracted events. Additionally, we also aim to find and represent relations between events such as causality, meronymy and correlation.

6 Acknowledgements

This research was funded by the CAMeRA Institute of the VU University Amsterdam and by the CATCH programme, NWO grant 640.004.801.

⁶ <http://agora.cs.vu.nl/demo>

References

1. R. Cilibrasi and P. Vitanyi. The google similarity distance. *IEEE Trans. Knowledge and Data Engineering*, 19(3):370–383, 2007.
2. J. R. Finkel, T. Grenager, and C. Manning. Incorporating non-local information into information extraction systems by gibbs sampling. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL 2005)*, 2005.
3. G. Geleijnse, J. Korst, and V. de Boer. Instance classification using co-occurrences on the web. In *Proceedings of the ISWC 2006 workshop on Web Content Mining (WebConMine)*, Athens, GA, USA, November 2006.
4. N. Gkalelis, V. Mezaris, and I. Kompatsiaris. Automatic event-based indexing of multimedia content using a joint content-event model. In *ACM Events in MultiMedia Workshop (EiMM10)*, Oct 2010.
5. N. Ide and D. Woolner. Exploiting semantic web technologies for intelligent access to historical documents. In *Proceedings of the Fourth Language Resources and Evaluation Conference (LREC)*, pages 2177–2180, Lisbon, Portugal, 2004.
6. E. Riloff and R. Jones. Learning dictionaries for information extraction by multi-level bootstrapping. In *Proceedings of AAAI '99*, pages 474–479, 1999.
7. R. Shaw, R. Troncy, and L. Hardman. Lode: Linking open descriptions of events. In *4th Annual Asian Semantic Web Conference (ASWC'09)*, 2009.
8. W. R. van Hage, V. Malaisé, G. de Vries, G. Schreiber, and M. van Someren. Abstracting and reasoning over ship trajectories and web data with the Simple Event Model (SEM). *Multimedia Tools and Applications*, 2011.