DIVE into the Event-Based Browsing of Linked Historical Media

Victor de Boer, Johan Oomen[1], Oana Inel[2], Lora Aroyo[2], Elco van Staveren[3], Werner Helmich[4], and Dennis de Beurs[4]

[1] Netherlands Institute for Sound and Vision, Hilversum, the Netherlands
vdboer@beeldengeluid.nl, joomen@beeldengeluid.nl
[2] Dept. of Computer Science, VU University Amsterdam, the Netherlands
lora.aroyo@vu.nl, oana.inel@vu.nl
[3] Dutch National Library, the Hague, the Netherlands
elco.vanStaveren@kb.nl
[4] Frontwise, Utrecht, the Netherlands
werner@frontwise.com, dennis@frontwise.com

Abstract. DIVE is a linked-data digital cultural heritage collection browser. It was developed to provide innovative access to heritage objects from heterogeneous collections, using historical events and narratives as the context for searching, browsing and presentating of individual and group of objects. This paper describes the DIVE Web Demonstrator[5]. This demonstrator uses semantics from existing collection vocabularies and linked data vocabularies to establish connections between the collection media objects and the events, people, locations and concepts that are depicted or associated with those objects. The innovative interface combines Web technology and theory of interpretation to allow for browsing this network of data in an intuitive "infinite" fashion. DIVE focuses to support digital humanties scholars in their online explorations and research questions.

Keywords: Digital History, Maritime Data, Heterogeneous Data Cloud, Digital Hermeneutics, Historical Events, Crowdsourcing

1 Background

The Web has offered cultural heritage institutions and their public a medium, changing their traditional task from information interpreters to that of information providers [3] and collections are being made digitally available in increasing numbers. Public repositories such as Europeana and the Digital Public Library of America, for instance, offer access to tens of millions of digital artifacs from museums, archives and libraries. As a result, however, the users can access much more, but understand much less, because of the limited curration. This urges cultural heritage institutions to rethink the access provision strategies to their collections to allow the public to interpret and contribute to their collections.

[5] The DIVE demonstrator is available at http://dive.frontwise.com

Search and browsing interfaces provide access to both professionals as well as the general public, searching for cultural heritage objects in either a single collection or in multiple collections at the same time. The traditional information access to cultural heritage assumes that experts interpret and curate their collections in such a way that the users of their information systems perform simple or complex keyword search to come to a selection of items matching the query. However, research has shown that many users seek more *exploratory* forms of browsing [1]. Recent Linked Data innovations make it possible to create interactive access to heterogeneous cultural heritage collections not only through direct textual keyword search, but also through structured links between cultural heritage objects and related events, persons, places and concepts.

In the Agora project, browsing of cultural heritage collections through events and their links to collection object has proved succesful for supporting the interpretation of end users, and thus realising the so-called 'digital hermeneutics'[4]. Building on the event modelling of the Agora project, the DIVE demonstrator presented in this paper implements the event-based browsing of cultural heritage objects from two heterogeneous historical collections. Within DIVE, new interaction concepts for events and event-based narratives have been explored and developed. We explicitly support a diversity of user groups, including Digital Humanities researchers, professional (commercial) users and the general public.

The heterogeneous collections are made available through the demonstrator and are interlinked in a common linked data network. This interconnected network of events, persons, places, and concepts, provide context to the cultural heritage objects, which are represented in the same networks. Thus, the objects are contextualized with events and narratives, which is crucial for the findability and hermeneutics. Core contribution is the innovative user interface supporting information interpretation in multimedia collections through dynamic browsing experience with linked data and explicit event representation.

2 The Data

2.1 Data Sources

Within the project, content from the two cultural heritage institutions are enriched, linked and made available:

- The Netherlands Institute for Sound and Vision (NISV)[6] archives Dutch broadcasting content, including television and radio content. Within the project, a subset of the NISV collection was made available using the OAI-PMH protocol. These are videos of news broadcasts. For these videos descriptive metadata is available including free-text content description.
- The Dutch National Library (KB)[7] provides access to a number of historical datasets one of which we use in the DIVE - the ANP Radio News Bulletin dataset[8]. This dataset is made up of digitized typoscripts (radio news

[6] http://www.beeldengeluid.nl
[7] http://www.kb.nl
[8] http://radiobulletins.delpher.nl/

scripts, to be read during news broadcasts) from the period 1937-1984. These have been made public through a Web interface and API. Here, the scanned images, OCRed content and descriptive metadata is available. The original data and metadata are available in Dutch.

2.2 Data Conversion and Enrichment

The textual descriptions and descriptive metadata for both collections are retrieved and are converted to RDF. From the textual descriptions we extract *events* as well as places, persons and concepts linked to those events which in turn are depicted by the cultural heritage objects (videos or news bulletins). For this extraction, we employ an ensemble of methods. Named Entity Recognition (NER) and Event extraction tools for Dutch text are used including the xTas[9] and Opener[10] toolkits. In a second stage, crowdsourcing through the CrowdTruth platform[11] is employed to have human-recognized entities and to refine the results from Natural Language Processing. For the extraction of the News Bulletins, we also use the results of the NER employed by the KB, which is based on the Stanford parser, optimized for Dutch texts.

The results from the different tools and the crowdsourcing are consolidated to RDF. The data is modelled using the Simple Event Model (SEM) [5]. This model allows for the representation of events, actors, locations and temporal descriptions. We extend SEM with other Linked Data schemas, e.g. DC, SKOS, OpenAnnotation and FOAF to represent other types of resources linked to the media objects[12]. Links are established between the different datasets using the Amalgame alignment tool [6]. Links are also established to external sources, including Wikipedia and DBPedia. The resulting dataset is stored in an RDF Triple store, which provides a SPARQL endpoint[13].

3 User Interface

3.1 User Experience Design

The core of digital hermeneutics is formed by two components: *object-event relationships* and *event-event relationships*. By making explicit relationships between the objects and events, and between the events themselves we can facilitate users in their access and interpretation processes of objects in online cultural heritage collections. In DIVE we aim at implementing those relations in an intuitive event-centric browsing interface for browsing cultural heritage collections by means of underlying linked data graph.

[9] http://xtas.net
[10] http://www.opener-project.eu/
[11] http://crowdtruth.org/
[12] The DIVE datamodel is visualized at `https://github.com/biktorrr/dive/blob/master/imgs/dive_model_v3.pdf`
[13] The ClioPatria triple store can be accessed at http://eculture.cs.vu.nl:8877/dive/

Major effort was put in creating an interface with a clear identity and an engaging user experience that invite users to continue exploring at different levels of detail. Users become explorers diving deeper into the data, as a diver deeper and deeper into an ocean trench discovering new species. This metaphor makes the interface a "digital submarine", which provides navigation controls and supportive and manipulative tools. The design of the interface also forms an innovative "inifinite exploration path", which unlocks the potential of touchbased explorative user interfaces.

An extensive design phase in which multiple concepts have been tested resulted in the DIVE "infinity browsing" interface, a combination of two core interaction concepts that involve a *horizontal level* supporting the breadth and a *vertical level* supporting the depth of information exploration and interpretation.

The *horizontal level* displays the result set related to the seed keyword search in a dynamic presentation. At this level, user's exploration is supported by *event-centric filters* making explicit the relation of each object to either the depicted and associate events and their properties, e.g. people, locations and times involved in the events. Consistent *color coding* is used for each property type to allow for a quick discovery of desired type of objects. Large result sets are represented as a colored barcode to give an overview of the amount and composition of event properties in the search results. *Objects* are represented by type-color, type-icon, title and an image and associated with a set of *buttons* providing detailed information for each object, e.g. description, source and external links. To allow for active user engagement through sharing of personal perspectives and interpretations, users can add *comments* to each object and save them in private or shared *collections*. Additionally, we provide a set of related objects from the Europeana collections. Typical interactions at this level are:

- *Pinch or scroll the elements of the color barcode* zooms in on the objects to reveal more information, e.g. image, title, icon of the event-related property (for example, an icon for location indicates that this object depicts a location of the related event).
- *Drag right or left the row of related objects* reveals previous or next object on the horizontal level.
- *Arrows* are used for navigating to previous or next objects in the row.
- *Search option & event filters* allow to show sub-sets of related objects.

The *vertical level* is formed by the user exploration history, as a path of selections on the horizontal level - leading to the last selected object. Each selection of an object results in a new row with related entities loaded under the selected object. Users can scroll back to a previous step, zoom out, choose another object and build a new path from there. This allows for fluid dynamics in collection exploration, discovery of alternative paths, and ultimately supports deep interpretation of cultural heritage collections. Our intention is to allow *saving each exploration history* as a collection so that users can revisit or share their browsing experiences.

3.2 Implementation details

The interface is developed in HTML5, Javascript and CSS3. A number of libraries are used to provide specific functionality: jQuery handles the major part of the functionalities like DOM interaction and manipulation, event handling and AJAX. Velocity.js is used to improve the performance of animations. Hammer.js supports the handling of touch events. Moment.js makes dates managable. As the amount of entities to display can be near the feasible limits of the web-browsers effort was put in the optimization of the javascript code while maintaining readibility of the code. Examples of this include the gradual buildup of DOM elements, use of Prototypes and limiting features like animation and CSS3 filters on large collections.

The interface acquires data from the data layer using the triple store's SPARQL API. Several queries are used to search entities by keyword, get related entities (events, persons, etc.) and get entity details. The returning data is handled by a client-side datamapper in Javascript which maps the datafields to an internal format which is used to build the interface representations. This approach relieves the server of unneccesary data parsing and contributes to compatibility with other data sets.

A smart image cache has been implemented to provide a visual representation for other entities. Based on keywords from entity titles, images are retrieved from the five most relevant Wikipedia searches using the Wikipedia API[14]. If no images are found, another query is made to the OpenCultuurData API[15] which covers an extensive set of Dutch open heritage- and cultural data. The quality and availabity of images through this system is acceptable and provides a powerful way of filling in the (visual) data gaps. These images increase the user experience by supporting the visual navigation throught the interface and rememberability and recognizability of individual entities. Figure 1 shows the current version of the interface, optimized for tablets and modern web browsers.

4 Current status and future work

The DIVE demonstrator has been published at `http://dive.frontwise.com`. We are continually updating both the data as well as the user interface. Currently, a subset of data from the two institutions is available in the datastore and through the interface. We are scaling this up to more data from the two institutions as it becomes available, as well as (Dutch) historical data from other sources. More links to datasets (cf. Europeana) and background knowledge on the Web of Data (cf. GeoNames) will provide more context for the items in the DIVE datasets. Future work includes follow-up on the initial Agora evaluations with DIVE demo evaluation with media professionals and humanities researchers. This will include information-retrieval-type user evaluations, with a specific aim at measuring the success of explorative search. We will also log usage

[14] http://www.mediawiki.org/wiki/API:Main_page
[15] http://www.opencultuurdata.nl/api/

Fig. 1. DIVE Screenshot: shows a person (Berlage), with two related events and KB media objects associated with them (transcripts from two ANP radio programs), related location (Bilt) and related person (Dutch Queen). In the bottom, objects are projected on a timeline as well (if timestamp is available).

by members of the general public as the demonstrator will be made available as part of the websites of NISV and KB.

Acknowledgements

This work is partially supported by NWO CATCH programm (http://www.nwo.nl/catch) and the Dutch Ministry of Culture. The authors would like to thank Theo van Veen and Olaf Janssen for their help with the KB data.

References

1. M. van Erp, J. Oomen, R. Segers, C. van de Akker, L. Aroyo, G. Jacobs, S Legne, L. van der Meij, J.R. van Ossenbruggen, and G. Schreiber. Automatic heritage metadata enrichment with historic events. In *Museums and the Web 2011*, 2011.
2. Oana Inel, Khalid Khamkham, Tatiana Cristea, Arne Rutjes, Jelle van der Ploeg, Lora Aroyo, Robert-Jan Sips, Anca Dumitrache, and Lukasz Romaszko. Crowd truth: Machine-human computation framework for harnessing disagreement in gathering annotated data. In *To appear in: Proceedings of the International Semantic Web Conference 2014 Replication, Benchmark, Data and Software track*, 2014.
3. Klaus Mueller. Museums and virtuality. In *In R. Parry, editor, Museums in a Digital Age, chapter 30, pages 295-305 Routledge*, 2007.
4. Chiel van den Akker, Susan Legne, Marieke van Erp, Lora Aroyo, Roxane Segers, Lourens van der Meij, Jacco van Ossenbruggen, Guus Schreiber, Bob Wielinga, Johan Oomen, , and Geertje Jacobs. Digital hermeneutics: Agora and the online understanding of cultural heritage. In *Proc. of the 3rd International Web Science Conference (WebSci '11). ACM, New York, NY, USA, 7 pages.*, 2011.

5. Willem Robert van Hage, Vronique Malais, Roxane H Segers, Laura Hollink, and Guus Schreiber. Design and use of the simple event model (sem). *Web Semantics: Science, Services and Agents on the World Wide Web*, 9(2), 2011.
6. Jacco van Ossenbruggen, Michiel Hildebrand, and Viktor de Boer. Interactive vocabulary alignment. In *Research and Advanced Technology for Digital Libraries - International Conference on Theory and Practice of Digital Libraries, TPDL 2011, Berlin, Germany, September 26-28, 2011. Proceedings*, pages 296–307, 2011.

Appendix: Challenge requirements

In this appendix, we address how this submission meets the evaluation criteria for the Semantic Web Challenge.

4.1 Mandatory criteria

1. *The application has to be an end-user application, i.e. an application that provides a practical value to general Web users or, if this is not the case, at least to domain experts. It should show-case functionalities that the use of semantic web technologies can bring to an application.* The DIVE demonstrator is indeed an end-user application, available for all. The demonstrator's UI is designed for two user groups. The primary user group is that of cultural heritage domain researchers, to browse the integrated collections in an explorative fashion. However, the demonstrator is set up in an intuitive way, allowing the general public to navigate the integrated graph in without the need of academic prior knowledge.
2. *The information sources used should be under diverse ownership or control; should be heterogeneous (syntactically, structurally, and semantically); and should contain substantial quantities of real world data (i.e. not toy examples).* We use two primary information sources that are under diverse ownership: That of the Netherlands Institute for Sound and Vision and that of the Dutch Royal Library. The data have been enriched within the DIVE project and these enrichments are gathered in a central triple store for computational reasons. However, this extracted information refers back to landingpages as well as media objects (videos and images), hosted by the individual institutions. Furthermore, the data is enriched with information from external datasources (DBPedia). At the time of writing, the demonstrator deals with a considerable amount of data, resulting in [VIC: numbers here(do this last)] triples describing X and Y videos and images. In the near future, we will extend this to all OpenImages videos as well as all KB news bulletins.
3. *The meaning of data has to play a central role. Meaning must be represented using Semantic Web technologies; Data must be manipulated/processed in interesting ways to derive useful information; and this semantic information processing has to play a central role in achieving things that alternative technologies cannot do as well, or at all.* As described in this document, the original data is enriched using an ensemble of NLP and crowdsourcing techniques throught the CrowdTruth platform [2]. This results in a semantic

graph of extracted knowledge about these media objects. This knowledge is represented using Semantic Web standards and vocabularies (RDF(S), OWL, SKOS, SEM, OpenAnnotation). This allow us to represent connections between the content of these integrated datasets. The type of explorative browsing interface that is the front-end of this demonstrator is made possible by this explicit modelling of the knowledge as a semantic graph. For example, the interface exploits RDFS reasoning for search results and browsing.

4.2 Relevant desirable criteria

- *The application provides an attractive and functional Web interface (for human users).* The demonstrator's innovative UI is the result of a careful design process, in collaboration between Web and interaction design experts, data experts, humanities researchers and end users. We believe it is both attractive and functional and makes use of the visual nature of much of the media objects as well as the extracted knowledge. By making explicit relationships between the collection objects and events, and their related properties, e.g. people involved, locations and times we can facilitate users in their access and interpretation processes of objects in online cultural heritage collections.
- *Functionality is different from or goes beyond pure information retrieval.* The demonstrator is designed specifically for an innovative form of *explorative browsing*. It uses reasoning over the graph to display related items to the one that the user is focusing on. This goes well beyond standard information retrieval. The event-related filters allow to identify groups of objects depicting or associated to specific events, related people, locations and times. Additionally, we project all browsing results on an interactive timeline, providing alternative filtering and access to the related objects.
- *Multimedia documents are used in some way.* The DIVE demonstrator is centered around knowledge extracted from heterogenous multimedia objects, e.g. text, videos, images. It showcases how this knowledge and the media objects themselves can be linked through events and their properties, and presented in a single intuitive interface.
- *There is support for multiple languages and accessibility on a range of devices.* We currently support only browsing and search in the original language of the collections (Dutch), although work for exploiting mappings to multilingual vocabularies is underway. The demonstrator UI is designed specifically for multiple platforms and features a responsive design. In fact, the UI was designed with a"tablet-first" principle, where touch-actions and gestures are primarily considered for the browsing of the graph.