

SEMANTIC TIMELINE INTERFACES FOR ANNOTATED MULTIMEDIA ASSETS

Jacco van Ossenbruggen, Lynda Hardman

Centrum voor Wiskunde en Informatica

Kruislaan 413, P.O. Box 94079

1090 GB, Amsterdam, The Netherlands

email: Firstname.Lastname@cwi.nl

fax: (+31) 20 592 4312

Abstract

One of the major advantages of the Semantic Web is that resources can be annotated with arbitrary properties. It is, however, this flexibility of expression which makes it hard to develop general interface components that can be used to interact with the resources and their annotations.

For many resources, however, their location is a key aspect to understand their context. The recent success of semantics maps and “virtual earth” applications proves this. Initially only used for “way finding” tasks, these tools now provide a generic access method for any type of information that has longitude/latitude information attached to it. Typical examples range from location based services (e.g. find nearby restaurants based on a current location) to “geotagged” personal images and bookmarks.

We believe that a similar general access mechanism can be developed that conveys the temporal aspects of a selected set of resources to the user. Where a semantic map or satellite image allows a user to interact with the underlying data based on its location, a *semantic timeline* would allow interaction based on the temporal aspects of the underlying resources.

Understanding, combining and reasoning with temporal information is, however, not trivial. Many schemes to deal with time have been developed. These include the use of different calendar systems, different names for the same period in disciplines such as art history and geology, the use of time points versus time intervals, qualitative versus quantitative notions of time, absolute versus relative time, etc.

The proposed interface component can therefore only be realized if it has an ontological basis that describes and relates the notions of time used in the annotations and in the interface.

1 Introduction

Traditionally, the publication date of books and journals has been a key property used for indexing and cataloging in libraries. The same holds for art objects — displayed in musea or described in text books — where the year of creation plays an important role in understanding the art historical period the object was made in. Time also plays a role in ordering our daily information, such as diary entries, meeting schedules and structuring our personal digital photo and video albums. Finally, time is also of crucial importance when describing change, that is, resources with properties whose values evolve over time.

Interactive timelines are well known and well described in the literature, see [3] for a survey. Because such a wide variety of our data has temporal properties, timelines have been proposed as a user interface to databases and digital libraries [2]. A wide spectrum of user interface techniques, display metaphors, interaction styles, etc., has been described. Selecting the right set of techniques for a specific application is important, but not the topic of this paper. The paper focuses on the practical aspects of combining interactive timeline interfaces with Semantic Web technology:

- Semantic Web technology makes temporal (meta) data more easily accessible in commonly agreed upon formats. For example, a timeline-based application can now make use of often used properties such as `dc:date`, `vra:date`, `foaf:birthdate` etc. to position events on a timeline. Widespread use of previous timeline applications have been hindered by the need to convert data to the particular format the application could handle. The Semantic Web should make conversion unnecessary, or at least much easier.
- Semantic Web languages allow explicit, share-

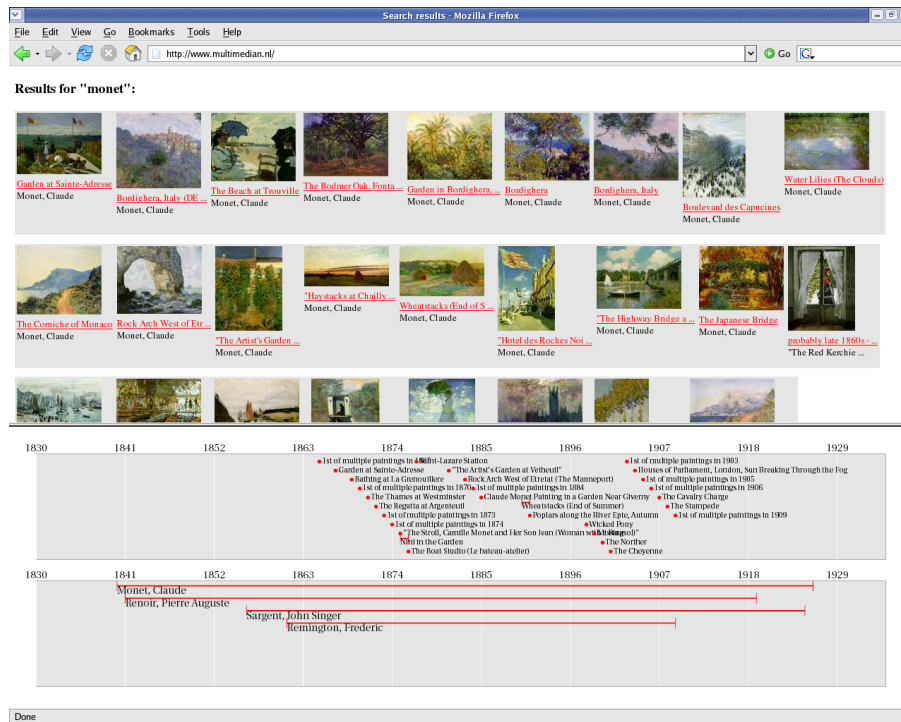


Figure 1: Timeline view of RDF literal keyword search results

able and reusable formalizations of a model's underlying timeline. This allows the development of flexible, generic timeline applications into which the specifics of the applicable temporal models can be plugged. This would enable these applications to use different calendar systems, or to plug in domain specific knowledge such as the names of art periods and movements in a cultural heritage application or the names of the periods of the geological timescale in an archaeological application.

- In addition to temporal metadata, other semantic relationships that are encoded in the metadata can be used to enrich the timeline interface. For example, an application in the art domain could, in addition to positioning the life spans of artists on the timeline, also visualize other relationships among the artists, such as “influenced-by” or “teacher-of”. Pre-Semantic Web applications had to rely on application-specific formats to encode such relationships.

2 Example application: Semantic Search in Cultural Heritage

We are currently exploring the use of semantic timelines in the context of an e-culture project that is part of the larger Dutch national MultimediaN¹ project. The goal is to use Semantic Web technology to improve search, reasoning and presentation of media assets of our national digital cultural heritage portal². The portal provides access to several extremely heterogeneous collections, including paintings, sculptures, archaeological artifacts, lyrics and samples of Dutch folksongs, etc. The vocabularies used to annotate the artifacts vary significantly, as does the granularity and quality of the annotations. Search, reasoning and navigation strategies based on a single, fixed metadata vocabulary is thus not an option.

While the project is just starting up, we have been working on a first prototype, based on an on-line painting collection³. The prototype is built using the SWI-Prolog SeRQL engine⁴. For the prototype, we first uploaded the metadata of the painting collection (after conversion to RDF, using the VRA vocabulary). Second, to be able to use some semantic

¹<http://www.multimedien.nl>

²<http://www.cultuurwijzer.nl>

³<http://www.artchive.com/>

⁴<http://www.swi-prolog.org/packages/SeRQL/>

background knowledge, we uploaded RDF versions of the Getty thesauri (AAT, ULAN and TGN)⁵. Finally, to find search results referring to synonyms or other closely related words, we also uploaded an RDF-version of Wordnet.

Figure 1 shows a screen-dump of the current version of the prototype. It displays the results of a (fuzzy) keyword search on the string “monet”. The engine has indexed all literals in the RDF triple store, which allows us to quickly find all RDF literals matching the search term. For every match, we look in the neighborhood in the RDF graph for resources of a particular type. In the figure, the user searched for paintings, so we selected resources of `rdf:type vra:visualResource` that were within a certain distance (in this case $n = 2$) of the matching literal. Thumbnails of all resulting resources are generated on the fly, along with the `vra:title` and `vra:creator` are shown on the upper half of the screen. The lower half shows the same titles in an interactive timeline view, using the `vra:date` property. A second timeline shows the lifespan of the creators of the results. Note that the latter timeline uses background knowledge that originates from the RDF version of Getty’s ULAN (Universal List of Artist Names). In addition to conveying the temporal relations among the paintings and their creators, the timelines are also used for semantic navigation and search. In the current prototype, for example, all painting titles are clickable and can be used to navigate to a web-based annotation form where (privileged) users can add, delete and edit the RDF metadata of the selected painting. The names of the creators in the bottom timeline are also clickable and can be used to refine the search. For example, in contrast with the fuzzy keyword search on ‘monet’ in the initial query, clicking on the name “Monet, Claude” in the timeline, will result in an exact search matching all `vra:visualResources` with a `vra:creator` property that is equal to the URI that identifies Claude Monet in the Getty ULAN database.

3 Future Work

As we mentioned previously, the work described in this paper is part of a project that is currently in the start-up phase, and in this sense the “real” work still needs to be done. We plan to use the prototype described in this paper as a starting point for exploring semantic user interfaces to RDF-annotated media repositories. We are currently working on several extensions that are relevant for this workshop.

⁵http://www.getty.edu/research/conducting_research/vocabularies/

First, we are improving both the background knowledge in the RDF triple store and the GUI, so we can transform the GUI in a true “semantic timeline” interface. For example, we are developing a reasoning module that can make informed guesses about the link between the various historic art styles as described in the AAT, the paintings as described in the Artchive collection and the absolute and relative time periods associated with the relevant art styles. With this information made explicit in the RDF, it is easy to see how the timelines can be semantically enriched, both as a presentation and as an interaction device. Secondly, we plan to develop a (Prolog) library that supports temporal queries on the RDF-triple store (such as semantic searches limited to events in a particular interval, or after a particular point in time). To be able to reuse such a library in other collections and domains, we need to define its API in terms of commonly used temporal datatypes. We are currently investigating to what extent OWL-Time can provide the basis for such an API.

Second, we are planning to extend the current interface with a multi-facet interface (see [1] for a good example) that supports navigation along the many hierarchical ordered property values that are relevant for this domain. Relevant properties that give rise to an intuitive hierarchical navigation structure include `ulan:teacherOf` and several properties with values defined by their position in the AAT hierarchy, such as (sub)style, material, geographical region, etc. That last property brings us to another planned alternative to the timeline interface: an extra facility to support search, presentation and interaction based on a geographic map or satellite-image GUI. Similar to the temporal information, we have several types of explicit geographical information in our RDF triple store. Examples include places depicted by paintings, the place where a painting has been painted, its current location (e.g. the museum or gallery), the relevant places from the creator’s biography, etc. In particular, we plan to integrate the interfaces for time, space and other semantics. This would support combined search and navigation strategies, such as: “show me artifacts and painters from the same period from other countries”, or “show me artifacts that belong to the same period, but are classified as a different art style”. Note that relative time periods often differ per location. For example, the interval traditionally associated with the Renaissance in Italy differs from the interval of the Renaissance in the Netherlands and other North West European countries. Also note that the combination of time and space also allows the use of historic maps that are part of many collections as a intuitive basis for the GUI.

Third, we will extend the current collection of paint-

ings from the Artchive to several other collections that are part of the Dutch national e-culture portal. An important research aspect here is to investigate how we can reduce the amount of manual semantic integration that is needed to link the information in those new collections to the background knowledge used. Note that many repositories are not annotated using structured and standardized vocabularies such as the Getty AAT, and many only use unqualified Dublin Core attributes filled with free text.

While this project is focused on cultural heritage, in other projects we would like to explore the use of interfaces developed for museum artifacts for other domains that deploy spatio-temporal metadata. Potential examples include “geotagged” personal images and video, geotagged bookmarks, etc.

Acknowledgments

The authors would like to thank Guus Schreiber (VU) and the other MultimediaN project partners for their input. Jan Wielemaker (UvA) developed the Semantic Web and SWI-Prolog infrastructure and a major part of the demo. The screen-dump shows scans of paintings made by Mark Harden from <http://www.artchive.com/>.

References

- [1] E. Hyvönen, M. Junnila, S. Kettula, E. Mäkelä, S. Saarela, M. Salminen, A. Syreeni, A. Valo, and K. Viljanen. Finnish Museums on the Semantic Web. User’s Perspective on MuseumFinland. In *Proceedings of Museums and the Web 2004 (MW2004)*, Arlington, Virginia, USA, March 31 – April 3, 2004. Selected papers.
- [2] Vijay Kumar, Richard Furuta, and Robert B. Allen. Metadata Visualization for Digital Libraries: Interactive Timeline Editing and Review. In *Proceedings of the third ACM conference on Digital Libraries*, pages 126–133, Pittsburgh, Pennsylvania, United States, 1998. Available from World Wide Web: http://portal.acm.org/ft_gateway.cfm?id=276689&type=pdf.
- [3] Sonia Fernandes Silva and Tiziana Catarci. Visualization of Linear Time-Oriented Data: A Survey. In *First International Conference on Web Information Systems Engineering (WISE’00)*, volume I, pages 310–319, 2000. Available from World Wide Web: <http://csdl.computer.org/dl/proceedings/wise/2000/0577/01/05771310.pdf>.