Semantically Augmented Annotations in Digitized Map Collections

Rainer Simon
AIT Austrian Institute of Technology
Vienna, Austria
rainer.simon@ait.ac.at

Bernhard Haslhofer
University of Vienna
Liebiggasse 4/3-4
Vienna, Austria
haslhofer@cs.univie.ac.at

Werner Robitza
University of Vienna
Liebiggasse 4/3-4
Vienna, Austria
robitza@cs.univie.ac.at

Elaheh Momeni
University of Vienna
Liebiggasse 4/3-4
Vienna, Austria
momeni@cs.univie.ac.at

ABSTRACT

Historic maps are valuable scholarly resources that record information often retained by no other written source. With the YUMA Map Annotation Tool we want to facilitate collaborative annotation for scholars studying historic maps, and allow for semantic augmentation of annotations with structured, contextually relevant information retrieved from Linked Open Data sources. We believe that the integration of Web resource linkage into the scholarly annotation process is not only relevant for collaborative research, but can also be exploited to improve search and retrieval. In this paper, we introduce the COMPASS Experiment, an ongoing crowdsourcing effort in which we are collecting data that can serve as a basis for evaluating our assumption. We discuss the scope and setup of the experiment framework and report on lessons learned from the data collected so far.

Categories and Subject Descriptors

H.3.3 [**Information Storage and Retrieval**]: Information Search and Retrieval—*Retrieval models*; H.3.7 [**Information Storage and Retrieval**]: Digital Libraries—*User issues*

General Terms

Experimentation, Verification

Keywords

Maps, Annotation, Linked Data, Tagging, Crowdsourcing

1. INTRODUCTION

Historic maps are increasingly being recognized as a valuable scholarly resource. They record historical geographical information often retained by no other written source [17], and are thus relevant to the study of a range of environmental, ecological or socio-economic phenomena: from the development of land use [16] to the effects of river channel changes [4] or floods [20], to the reconstruction of past urban environments [11]. At the same time, they capture more than mere geographic facts: they also draw a fascinating picture of the cultural, political and religious context in which they were created. Their degree of accuracy tells much about the state of technology and scientific understanding at the time of their creation [17]. Consequently, historic maps are cultural heritage artifacts in their own right, part of the artistic heritage as much as of the history of science and technology as a whole [3].

Annotations are a fundamental scholarly practice common across disciplines [19]. They enable scholars to organize, share and exchange knowledge, and work collaboratively in the interpretation and analysis of source material. At the same time, annotations offer additional context: they supplement the item under investigation with information that may better reflect a user's setting [5]. With the YUMA Map Annotation Tool[1] we want to provide social annotation functionality to scholars studying digitized historic maps. A central feature of this tool is that it integrates *semantic linking* into the annotation process. This way, map annotations are augmented with relevant semantic context from the *Linked Data* Web [1]. For example, if a user annotates the region of Ottawa on a map of Canada, or mentions the city of Ottawa in the annotation text, the system suggests links to possibly relevant resources - e.g. on Geonames or DBpedia - and prompts the user to verify these links.

A recent survey of existing scholarly annotation systems is provided by Hunter [10]. Prior efforts particularly relevant to our work are found in the field of multimedia annotation and tagging in general, and in the area of semantic annotation systems in particular: examples in the former field include MADCOW [2] and Zotero[2], two collaborative annotation tools for (multimedia) Web documents; Fotonotes[3]

JCDL'11, June 13–17, 2011, Ottawa, Ontario, Canada.

[1] Online demonstration at: http://dme.ait.ac.at/annotation, source code at: http://github.com/yuma-annotation
[2] http://www.zotero.org
[3] http://www.fotonotes.net

an open source online image annotation tool; or ImaNote[4] as well as the commerical Zoomify[5] zoomable Web image viewer, both of which allow users to annotate high-resolution zoomable images. Examples from the latter field include e.g. the One Click Annotator [8], a Web-based semantic annotation system for HTML content or LODr [15], a tool for enriching existing tagged Web content with semantic links.

We believe that, beyond providing valuable related information to the user, semantically augmented annotations can also be exploited to improve search and retrieval in larger map collections. To test this assumption, we have started a crowdsourcing experiment in which we are collecting data to serve as basis for a deeper information retrieval evaluation. The planned outcome of the experiment will be the *COMPASS* data set (*Collection Of MaPs, Annotations and Semantic linkS*), which we plan to make publicly available. It will comprise (i) digitized maps provided by the Library of Congress together with their original metadata, (ii) a corpus of real-world user queries collected in the the Library of Congress' map search portal, (iii) binary relevance judgements between maps and search queries produced manually by volunteer human judges and (iv) free-text annotations augmented with semantic links that users have added to a subset of the maps. In this paper, we present some early results from this experiment, and an outlook on upcoming work.

2. THE YUMA MAP ANNOTATION TOOL

The YUMA Map Annotation Tool (Figure 1) is a browser-based application that displays a full-screen Google-Maps-like drag- and zoom-able representation of the digitized map. A floating window shows a list of all annotations that exist for the map. The window also contains the necessary GUI elements for creating, editing and replying to annotations, and a basic moderation feature that allows users to report inappropriate annotations to the system administrator. When creating annotations, users also have the option to draw polygon shapes on the map to denote the area to which an annotation pertains. Using the tool's integrated geo-referencing functionality [18], the shapes can be translated to geographical coordinates, and overlaid on top of a present-day map, shown in a separate window.

To enrich annotations with structured semantic information, the tool employs a *semi-automatic linking approach* [7]. While the user creates an annotation, the tool suggests links to potentially relevant resources on the Linked Data Web. In the user interface, these suggestions are represented in the form of a tag cloud: the user can accept a suggestion by clicking on the tag, which will then add the link to the annotation. Two sources of information are used to derive the suggestions: if the map has been geo-referenced, the system suggests geographical entities (cities, countries) that intersect with the annotated map region. For example, if the user annotates the area of Yosemite National Park on a map of California (see Figure 1), the system will suggest the country (United States) and relevant cities in the area (such as Yosemite Lakes). Furthermore, the text of the annotation is analyzed using named entity recognition (NER). Recognized entities are queried against an encyclopedic Linked Data set to obtain dereferencable URIs. The current prototype makes

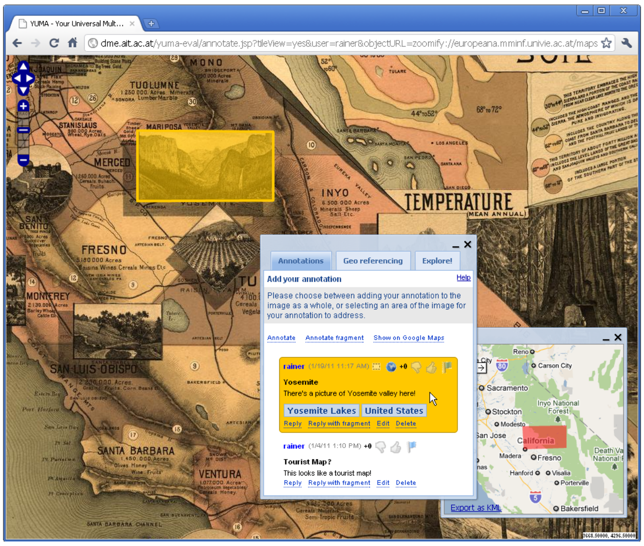

Figure 1: YUMA Map Annotation Tool Screenshot.

use of the spatial query API of the Geonames[6] online geographical database to obtain geographical resources for the annotated map region. NER is performed through the public OpenCalais[7] Web service. URIs for identified entities are retrieved via the DBpedia Lookup service[8].

Annotations are also exposed as Linked Data: each annotation is assigned a unique URI, which returns an RDF representation when resolved. At the time of writing, the tool relies on the LEMO[9] multimedia annotation vocabulary [6], which reuses and refines terms from the W3C Annotea [12] vocabulary. However, the tool is currently being re-designed, and new annotation models - in particular the OAC model[10] - are being investigated for future use.

3. THE COMPASS EXPERIMENT

We believe that a map retrieval system that indexes annotations with user-verified links to Linked Data resources can be more effective w.r.t. retrieval than systems that index only metadata or purely textual annotations. To test this assumption, we compiled a data set from real world data we received from the Library of Congress Geography and Map Division for research purposes. It consists of 130,935 user *search queries* extracted from query logs collected over a period of two years, 6,306 high-resolution *digitized map images* in JPEG 2000 and TIFF file format, and, for each map, its descriptive MODS metadata harvested via OAI-PMH.

In the first step, we want to evaluate the effectiveness of existing, metadata-based map retrieval approaches. We indexed the metadata using Apache Lucene and executed a randomly chosen subset of the user queries against the metadata index. From the results, we constructed a collection consisting of the queries, and a list of the top ranked results returned for each query.

In order to quantify the effectiveness of the retrieval sys-

[4]http://imanote.uiah.fi
[5]http://zoomify.com
[6]http://www.geonames.org/ontology
[7]http://www.opencalais.com/calaisAPI
[8]http://lookup.dbpedia.org
[9]http:// lemo.mminf.univie.ac.at/annotation-core
[10]http://www.openannotation.org/spec

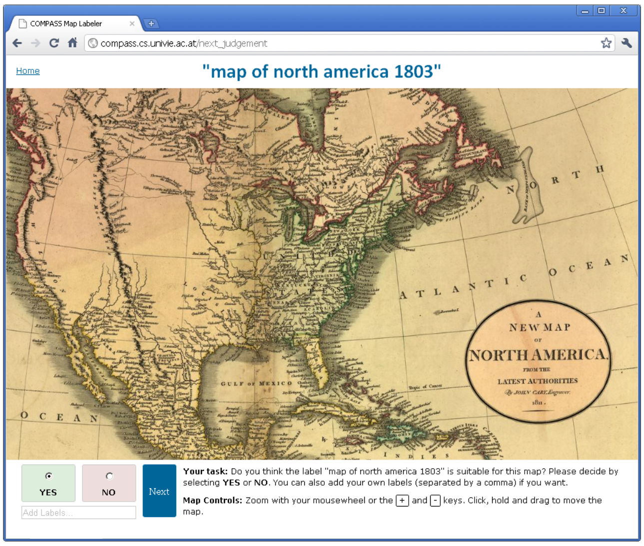

Figure 2: Relevance Judgement Application GUI.

tem, however, a *ground truth* is needed: a set of judgements as to whether a map is relevant to a particular query or not. To establish such a foundation for our evaluation, we created a Web-based crowdsourcing application[11] with which we are currently collecting binary *relevance judgements* from invited volunteer users. We hope that the collected data can serve as a starting point for what might ultimately constitute an accepted *gold standard* for map retrieval evaluation. Volunteers were recruited from appropriate mailing lists in the digital library and map history domain (IFLA DIGLIB, INETBIB and MapHist lists, respectively). Participants were provided with basic information about the experiment background and goals, as well as with instructions on how to use the experiment application, and a contact e-mail address to provide feedback or ask for assistance. Upon registration, participants were furthermore asked to provide a brief statement to explain whether they have expert (professional or educational) background in the field. In order to reduce the total amount of maps to be assessed by users, we applied a basic *pooling* technique by limiting the test set to the top-ten ranked results for each query. A screenshot of the experiment application is shown in Figure 2: the title bar displays a single search query, selected randomly from the collection. The main area of the screen shows one random map which was among the top-ten ranked search results for this query. The map is displayed as a zoomable Web image, so that users are able to inspect the map in full detail, at original resolution. At the bottom of the screen, YES/NO buttons allow users to submit a relevance judgement for this map/query pair.

In the next step, we would like to analyze the effect of user-contributed annotations and semantic linkage on the effectiveness of the map retrieval system. We will invite domain experts to annotate historic maps using the YUMA Map Annotation Tool, and index the collected annotations along with selected properties of linked resources, such as labels, descriptions, alternative names in different languages, etc. We can then measure the effects of including certain kinds and combinations of data (metadata, annotations, linked resources) in the retrieval process.

[11]http://compass.cs.univie.ac.at

4. PRELIMINARY RESULTS

At the time of writing, we have received registrations from more than 75 users from at least 12 countries worldwide. In total, more than 1,600 judgements have been collected so far, which is approximately 1/3 of the judgements we need to create a corpus for 400 sample queries. The distribution of the number of judgements per user is relatively skewed: approximately 60% of all judgments were produced by the top 10 contributors, whereas about a quarter of registered users have so far not submitted a single judgement at all. This however, is not necessarily surprising. Rather, it corresponds with the phenomenon of *participation inequality*, a pattern of user participation in online communities according to which only a small fraction of users contribute at a large scale [14]; and which has been observed in similar crowdsourcing efforts [9].

Out of the active participants, about 1/3 provided a statement regarding their expert background, either in the general library domain (approx. 15% of active participants); more specifically in the map library field (approx. 15%); or in fields related to the scope of the experiment, such as GIS or cartography (approx. 4%). The remaining active participants did either not provide explicit information on expert background or indicated that they consider themselves as "amateur" users, rather than experts. In general, the level of involvement was higher for expert users: among the top 10 contributors, 6 participants had declared themselves experts. The same distribution could be observed in the overall experiment, with approx. 59% of all submitted judgments being provided by expert users.

Using the data collected so far, we also performed a first analysis of the effectiveness of the purely metadata based map retrieval approach: when disregarding all map/query pairs that have not yet received relevance assessments, we calculated a Mean Average Precision of 0,41. It has to be noted that the omission of map/query pairs, as well as the fact that only a few map/query pairs have received more than one relevance assessment so far both distort this result. Nonetheless, the low value indicates that the performance of a metadata-only based retrieval approach is clearly limited.

5. FUTURE WORK

As future work we will address two areas: first, we will complete the ongoing crowdsourcing effort. We will continue to collect relevance judgements to meet the required goals for a 400-query corpus; and we will collect annotations from invited domain experts through the YUMA Map Annotation Tool. Based on this data, we can repeat our analysis on precision and recall for a retrieval approach that takes into account semantically augmented annotations. Second, we will repeat the crowdsourcing effort in order to compile a bigger corpus of judgements, and obtain more judgements for each map/query pair. Recruiting more users by disseminating information about the experiment more widely in appropriate communities, and improving the motivation to contribute will be crucial success factors in this effort.

Informal feedback we have received from users on the overall experience with our experiment application so far was mostly positive. However, there were several issues that were raised independently by multiple participants, and which we aim to address in the next phase: for example, several users critizied the lack of a "Skip this Map" button in the judgement interface. We deliberately decided against the possibility to skip maps initially, as we thought we would otherwise risk getting no judgements at all for some maps. As more detailed feedback from some users revealed, however, lack of this feature may have caused frustration and prevented some users from spending a lot of time with the application alltogether. Consequently, despite a risk of skewing the distribution of judgements, a "Skip" button may have helped sustain long-term motivation. Furthermore, some users voiced concern as to whether "they were doing the right thing", i.e. whether they were applying appropriate criteria when deciding on the relevance of a map to a query, how they were supposed to react in case of obvious typing errors in search queries, etc. For a successful continuation, we will therefore need to provide improved instructions which are more explicit about the goals and expected outcomes of the experiment; lower the entry barrier for first time users by providing inituitive examples, e.g. through a video screencast or "guided tour" (a similar approach has been applied by [13]); and to "raise the bar" [9] for dedicated users, i.e. sustain motivation by offering more complex tasks, by integrating the YUMA Map Annotation tool more closely with the crowdsourcing application used to collect relevance judgements.

6. ACKNOWLEDGMENTS

We want to thank the Library of Congress for providing the base data for this work. This paper presents work done for the best practice network Europeana*Connect*. Europeana*Connect* is funded by the European Commission within the area of Digital Libraries of the eContent*plus* Programme and is lead by the Austrian National Library.

7. REFERENCES

[1] C. Bizer, T. Heath, and T. Berners-Lee. Linked Data - The Story So Far. *International Journal on Semantic Web and Information Systems*, 5(3):1–22, 2009.

[2] P. Bottoni, S. Levialdi, A. Labella, E. Panizzi, R. Trinchese, and L. Gigli. MADCOW: a visual interface for annotating web pages. In *Proceedings of the Working Conference on Advanced Visual Interfaces*, AVI '06, pages 314–317. ACM, 2006.

[3] C. Boutoura and E. Livieratos. Some fundamentals for the study of the geometry of early maps by comparative methods. *e-Perimetron*, 1(1):60–70, 2006.

[4] G. Braga and S. Gervasoni. Evolution of the Po River: an Example of the Application of Historic Maps. In H. M. G.E. Petts and A. Roux, editors, *Historical Change of Large Alluvial Rivers: Western Europe*, pages 113–126, New York, 1989. John Wiley.

[5] M. E. Frisse. Searching for Information in a Hypertext Medical Handbook. In *Proceedings of the ACM Conference on Hypertext (HYPERTEXT '87)*, pages 57–66, Chapel Hill, North Carolina, 1987. ACM.

[6] B. Haslhofer, W. Jochum, R. King, C. Sadilek, and K. Schellner. The LEMO Annotation Framework: Weaving Multimedia Annotations with the Web. *International Journal on Digital Libraries*, 10(1):15–32, 2009.

[7] B. Haslhofer, E. Momeni, M. Gay, and R. Simon. Augmenting Europeana Content with Linked Data Resources. In *6th international Conference on Semantic Systems*, Graz, Austria, 2010. ACM.

[8] R. Heese, M. Luczak-Rösch, A. Paschke, R. Oldakowski, and O. Streibel. One Click Annotation. In *6th Workshop on Scripting and Development for the Semantic Web, colocated with ESWC 2010*, Crete, Greece, May 2010.

[9] R. Holley. *Many Hands Make Light Work : Public Collaborative OCR Text Correction in Australian Historic Newspapers*. National Library of Australia Staff Papers, 2009.

[10] J. Hunter. Collaborative semantic tagging and annotation systems. *Annual Review of Information Science and Technology*, 43(1):187–239, 2009.

[11] Y. Isoda, A. Tsukamoto, Y. Kosaka, T. Okumura, M. Sawai, K. Yano, S. Nakata, and S. Tanaka. Reconstruction of Kyoto of the Edo era based on arts and historical documents: 3D urban model based on historical GIS data. *International Journal of Humanities and Arts Computing*, 3:21–38, 2010.

[12] J. Kahan and M.-R. Koivunen. Annotea: an open RDF infrastructure for shared Web annotations. In *Proceedings of the 10th international conference on World Wide Web*, WWW '01, pages 623–632, New York, NY, USA, 2001. ACM.

[13] C. J. Lintott, K. Schawinski, A. Slosar, K. Land, S. Bamford, D. Thomas, M. J. Raddick, R. C. Nichol, A. Szalay, D. Andreescu, P. Murray, and J. Vandenberg. Galaxy Zoo: morphologies derived from visual inspection of galaxies from the Sloan Digital Sky Survey. *Monthly Notices of the Royal Astronomical Society*, 389(3):1179–1189, 2008.

[14] J. Nielsen. Participation inequality: Encouraging more users to contribute, 2006. http://www.useit.com/alertbox/participation_inequality.html (accessed January 26, 2011).

[15] A. Passant. Linked Data tagging with LODr. In *Semantic Web Challenge (International Semantic Web Conferrence)*, Crete, Greece, 2008.

[16] A. W. Pearson. Digitizing and analyzing historical maps to provide new perspectives on the development of the agricultural landscape of England and Wales. *e-Perimetron*, 1(3):178–193, 2006.

[17] D. Rumsey and M. Williams. Historical Maps in GIS. In A. Knowles, editor, *Past Time, Past Place: GIS for history*, pages 1–18, Redlands, CA, 2002. ESRI Press.

[18] R. Simon, J. Korb, C. Sadilek, and M. Baldauf. Explorative User Interfaces for Browsing Historical Maps on the Web. *e-Perimetron*, 5(3):132–143, 2010.

[19] J. Unsworth. Scholarly Primitives: what methods do humanities researchers have in common, and how might our tools reflect this? In *Humanities Computing: formal methods, experimental practice.* King's College, London, May 2000. Available at: http://www3.isrl.illinois.edu/ unsworth/Kings.5-00/primitives.html.

[20] S. Witschas. Landscape dynamics and historic map series of Saxony in recent spatial research projects. In *21st International Cartographic Conference*, number August, pages 59–60, Durban, South Africa, 2003.