King's Research Portal

DOI:

[10.1109/DEST.2012.6227932](https://doi.org/10.1109/DEST.2012.6227932)

Document Version

Early version, also known as pre-print

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Blanke, T., Bodard, G., Bryant, M., Dunn, S., Hedges, M., Jackson, M., & Scott, D. (2012). Linked data for humanities research—The SPQR experiment: 6th IEEE International Conference on Digital Ecosystems Technologies (DEST), 2012. In *2012 6th IEEE International Conference on Digital Ecosystems Technologies (DEST)* (pp. 1-6) <https://doi.org/10.1109/DEST.2012.6227932>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Linked Data for Humanities Research — The SPQR experiment

Tobias Blanke, Gabriel Bodard, Michael Bryant,
Stuart Dunn, Mark Hedges
Centre for e-Research
King's College
London
firstname.secondname@kcl.ac.uk

Michael Jackson, David Scott
EPCC
University of Edinburgh
Edinburgh
michaelj@epcc.ed.ac.uk, d.scott@ed.ac.uk

Abstract—Ancient texts represent a primary source for research in the classics. A substantial body of digital material has evolved enriching these texts. Unfortunately these data are often distributed across myriad locations, stored in diverse and incompatible formats and are either not available online or are made available only in isolation. This paper describes an investigation into using linked data principles and technologies to build bridges between these islands of data to deliver an integrated data landscape through which researchers can explore and so seek to understand this data. The evaluation revealed that researchers were of the opinion that the linked data representation, and its visualisation as graphs, offers an intuitive and usable means of exploring and understanding the data, exceeding the capabilities offered by current online portals to classics data.

Keywords—Digital Humanities; Linked Data; Data Integration;

I. INTRODUCTION

For many years, researchers in the Classics and ancient history (as well as the rest of the Humanities) have been producing digital outputs of various forms, a significant proportion of them being databases or semi-structured corpora of marked-up texts. In many cases, these digital resources address related topics, communities and time periods, and would be of much greater utility to researchers if they could be linked up in such a way that allowed them to be explored as if they were part of a single rich data landscape. While the use of standards for representing these datasets provide one route to such unification, in the less than perfect world of heterogeneous resources there exists a great deal of legacy data in diverse formats. More importantly, standards are generally developed within particular domains, whereas research is often inter-disciplinary, making use of varied materials, and incorporating data conforming to different standards. There will inevitably be diversity of representation when information is gathered together from different domains and for different purposes, and consequently there will always be a need to integrate these diverse representations.

The LaQuAT (Linking and Querying of Ancient Texts) project [1], explored these issues within the domain of ancient Greek inscriptions and papyri. LaQuAT investigated

linking distributed heterogeneous datasets using the OGSA-DAI¹ software to provide an integrated view across these datasets, making them appear to researchers as a single virtual resource. In essence, this approach took a relational database-centric view of the data resources.

While this integration was technologically achievable, evaluation by humanities researchers raised significant issues. Running queries across datasets required a great deal of understanding about the semantics of the data at a fine-grained level. These semantics were for the most part left implicit in the underlying relational databases, and complicated further by the variety of conventions used in representing data. For example, dates may be given in very different forms, and may be expressed in very different precisions and with different levels of confidence. The fuzzy, uncertain and interpretative nature of the available data — an inevitable factor in the primary research data for antiquity — complicated the semantics of integration, and made it difficult to describe the semantics of the relationships between data sets. It was not always clear whether similarly named columns in independent databases really represented the same sort of information and could validly be linked.

LaQuAT concluded that this relational approach to integrating such material does not help researchers to explore the datasets and to understand and exploit the connections between them, nor, in any case, is the approach scalable beyond closely related datasets. The datasets remain as isolated silos, albeit accessible from a single place. The key is the connections themselves, how they are expressed and how researchers can understand and use them to explore the data more effectively.

It is in precisely this regard that Semantic Web (SW) and Linked Data (LD) approaches [2] have great potential, as they allow researchers to formalise resources and the links between them more flexibly, and to create, explore and query these linked resources. Closely allied to LD has been work on ontologies for providing agreed meanings for both links and the resources they connect. Thus ontologies can act as the semantic mediator between heterogeneous databases,

¹<http://www.ogsadai.org.uk/>

enabling researchers to explore, understand and extend these datasets more productively and so improve the contributions that the data can make to their research.

The primary aim of ‘Supporting Productive Queries (SPQR)’ was to build on this previous work and investigate the potential of an SW/LD approach for linking and integrating datasets related to classical antiquity, focusing on certain targeted datasets as test cases. We investigated ways to represent the information in our targeted datasets using RDF or equivalent formalism, taking particular account to address the semantic issues — incompleteness, uncertainty, fuzziness, etc. — identified by LaQuAT and described above. In parallel, we investigated mechanisms for breaking this information out of its current silos, and transforming it from its legacy formats into our chosen representation, and exposing it as LD. There are two broad approaches to this: wrappers for on-the-fly conversion and converting data before exposing it. We assessed the pros and cons of each. Moreover, the research community for classical antiquity is already developing a number of data resources with SW/LD in mind, and resources could provide additional ‘glue’ for linking the datasets into a wider network of knowledge. We focused here on geospatial resources, providing links to ancient (and modern) placenames.

The ultimate objective would be to bring the transformed information into a common corpus or ‘RDF warehouse’ where it can be explored and searched in an integrated way, and where new connections (corresponding to new RDF or similar statements) can be made by the researcher and added to the corpus of information. This was outside the scope of the project, however.

This paper presents the results of our SPQR experiments. It identifies advantages and disadvantages of the LD approach for current Humanities data integration. In ??, we first describe the datasets, we have experimented with, and then present how we integrated these datasets into a joint set of triples. The evaluation in Section III finally details the results and open questions of the SPQR approach.

II. IMPLEMENTATION: CONVERTING EPIGRAPHIC DATA TO LINKED DATA

SPQR used as a starting point the following datasets, which overlap in terms of both time and location and thus are suitable, realistic test cases for integrated exploration:

- The *Heidelberger Gesamtverzeichnis der griechischen Papyrusurkunden Aegyptens* (HGV)², a collection of metadata records for 65,000 Greek papyri from Egypt, from the period 300 BC to 700 AD.
- The *Inscriptions of Aphrodisias* (InsAph)³, a corpus of 2,000 ancient Greek inscriptions from the Roman city of Aphrodisias in modern Turkey, from the period 200 BC to 700 AD.

²<http://www.rzuser.uni-heidelberg.de/~gv0/>

³<http://insaph.kcl.ac.uk/index.html>

- The *Inscriptions of Tripolitania* (IRT)⁴, a corpus of over 1,000 inscriptions from modern Libya, from the period 100 BC to 700 AD.

The first two datasets were the main ones used by LaQuAT, facilitating comparison of results with the earlier project. Each dataset is available as a collection of XML documents, one per inscription or papyrus, that conform to the EpiDoc⁵ XML recommendations. EpiDoc follows the guidelines of the Text Encoding Initiative (TEI)⁶, which promotes adoption of a set of flexible but rigorous standards and tools for the digital encoding and interchange of texts.

The datasets have particular characteristics that have made them intractable to other integration approaches:

- These are ‘hand-crafted’ datasets, resulting from a very great deal of arduous research and intellectual effort by individual researchers, and raise quite different issues from (e.g.) datasets generated by experiments and other automated processes. Such hand-crafted datasets occur across disciplines, in the humanities but also in the sciences when research activities are focused around the activities of individual scientists.
- Due to the fragmentary nature of the evidence on which they are based, the datasets are often incomplete and ambiguous, and indeed may contain errors and contradictions — this was one of the key obstacles to integrating such material identified by LaQuAT.
- Some of the datasets (in particular databases) contain much embedded, implicit semantics. This information is very difficult for the researcher to make use of. SW/LD approaches offer a way of exposing it explicitly.

To convert InsAph, IRT and HGV to LD required parsing EpiDoc XML documents into sets of triples. Once a set of triples is obtained it is straightforward to convert between these and analogous formats, such as RDF/XML. As RDF can be expressed in XML form (i.e. as RDF/XML), and the EpiDoc documents are also XML, a straightforward approach would have been to apply an XSLT transformation to transform each EpiDoc document directly into RDF/XML, to produce results such as in listing 1, which represents information relating to the inscription iAph010002:

```
<?xml version="1.0" encoding="utf-8"?>
<rdf:RDF xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:j.0="http://spqr.epcc.ed.ac.uk/"
  <rdf:Description rdf:about="http://insaph.kcl.ac.uk/iaph2007/iAph010002/">
    <j.0:material>WHITE MARBLE</j.0:material>
    <j.0:wasLocatedAt rdf:resource="http://insaph.kcl.ac.uk/iaph2007/iAph010002/location"/>
```

⁴<http://irt.kcl.ac.uk/irt2009/>

⁵ EpiDoc: Epigraphic Documents in TEI XML, <http://epidoc.sourceforge.net>.

⁶ TEI: Text Encoding Initiative, <http://www.tei-c.org>.

```

</rdf:Description>
<rdf:Description
  rdf:about="http://insaph.kcl.ac.uk/iaph2007/iAph010002/location">
  <j.0:name>Aphrodisias</j.0:name>
  <j.0:type>ANCIENTFINDSPOT</j.0:type>
  <j.0:externalURL>http://sws.geonames.org/323763/</j.0:externalURL>
</rdf:Description>
</rdf:RDF>

```

Listing 1. RDF representation

However, we concluded that XSLT is insufficiently flexible to enable the automation during conversion of certain actions that we considered highly desirable. The foremost among these was the creation of triples that link out to existing datasets. Instead, we developed Clojure⁷ scripts to convert EpiDoc documents to sets of N-triples. Clojure is a functional programming language that has a close relationship to Java. It is sufficiently flexible for implementing automated traversal of EpiDoc documents, and also allowed the implementation of functions to parse EpiDoc elements relating to places and to output triples that link to URIs corresponding to ancient and modern place names, as recorded in Pleiades and GeoNames respectively. The scripts can be customised and extended to handle the conversion of other EpiDoc datasets⁸.

A significant decision was which ontologies would be used when representing the types of resources, their properties, and relationships between resources. We held discussions with participants in the Europeana project — SPQR’s Humboldt colleagues are also partners in Europeana — as to make our model compatible with the Europeana Data Model (EDM) [3], a model for representing cultural and historical artifacts. Moreover, generic standards such as Dublin Core⁹ and OAI-ORE¹⁰ were also taken into account. In the end, we decided to adopt a bottom-up approach, creating a simple SPQR-specific ontology, for the following reasons:

- The project had limited resources and wanted to direct them to producing data that we could give to researchers, for evaluation, as soon as possible. If their views on a LD representation proved to be positive then we could return to ontologies as an important aspect of future work.
- It demonstrates a pragmatic model of adoption to data providers. Publishing data conforming to a custom ontology and linking that to existing ontologies — e.g. using the OWL ‘sameAs’ relationship [4] to assert that concepts, resources or relations mean, or relate to, the same thing — offers a lower barrier to publication than insisting that data providers must conform to existing ontologies prior to publication.

⁷ Clojure, <http://clojure.org>.

⁸ The conversion scripts are available at <http://code.google.com/p/spqr>.

⁹ <http://dublincore.org/>

¹⁰ <http://www.openarchives.org/ore/>

The SPQR ontology consisted of URIs for a range of epigraphic concepts. The subject of a triple can be an epigraph, a date of origin, a person or a location. Examples of the predicates that are used are shown in Table I.

An additional reason for creating our own ontology was that, at the beginning of SPQR, we could not identify an existing one that covered epigraphy. In the latter stages of the project, we extended our scripts to generate additional triples that conformed to an ontology defined by the Pelagios community project¹¹. Pelagios provides a digital resource of references to ancient places, and conformance with their ontology allowed SPQR to serve as a data provider for Pelagios. The mapping has proven to be straightforward. Pelagios functions as an international collective of projects that were connected by a shared vision of the interconnected and interactive geography of the ancient world. Each project brought a different perspective on the ancient world to the project — some based on texts, others on maps or archaeological records — with the aim of combining our diverse contributions to produce a resource that is much more than the sum of its parts.

The scripts produce data that has links both within and between datasets:

- Links between the resources in the datasets: InsAph, IRT and HGV resources are linked in terms of object types, e.g. white marble.
- Links out to other resources: the datasets are linked to Pleiades, GeoNames and text-specific web pages within InsAph and IRT.
- Links in from other resources: conformance to the PELAGIOS ontology allows their datasets to link to the SPQR datasets.

SPQR’s RDF conversion scripts were used to transform data from the existing datasets into LD. These linked datasets were deposited into an OpenLink Virtuoso linked data server¹². OpenLink has been at the forefront of linked data technology development and Virtuoso was adopted as it has been proven as a stable linked data server technology, e.g. it is used by both DBPedia and Bio2RDF.

When the triples were loaded, these datasets from distributed sources together formed a single federated graph within Virtuoso. Virtuoso allows the LD to be browsed online using RDF facets and also queried via a web service that executes SPARQL queries [5]. It also provides an interface supporting free-text searches across the content of literals in triples. These services can both be accessed by a researcher using a web browser or called from third-party applications.

A. Data Issues

The InsAph and IRT data are available online and each inscription has a URI and an associated web

¹¹ <http://pelagios-project.blogspot.com/>

¹² Virtuoso Universal Server, <http://virtuoso.openlinksw.com>.

Table I
EXAMPLES FROM THE SPQR ONTOLOGY

Predicate	Subject	Object
http://spqr.epcc.ed.ac.uk/material	epigraph	Literal (e.g. "WHITE MARBLE")
http://spqr.epcc.ed.ac.uk/wasLocatedAt	epigraph	location
http://spqr.epcc.ed.ac.uk/notBefore	date of origin	Literal (e.g. "-0050")
http://spqr.epcc.ed.ac.uk/notAfter	date of origin	Literal (e.g. "1000")
http://spqr.epcc.ed.ac.uk/personType	person	Literal (e.g. "EMPEROR")
http://spqr.epcc.ed.ac.uk/region	location	Literal (e.g. "???????")

page. It was decided to use these URIs to refer to the corresponding inscriptions within the LD set created by SPQR. However, these URIs do not support dereferencing — browsing to the resource URI <http://insaph.kcl.ac.uk/iaph2007/iAph010002> gives a page-not-found error, despite the fact that there is a web page at <http://insaph.kcl.ac.uk/iaph2007/iAph010002.html> that provides a representation of that resource. Nevertheless, the scripts were written to use this existing URI set so as to conform to the linked data recommendations. The InsAph and IRT data providers were informed of this issue.

Some EpiDoc documents were found to have typographical errors that would result in undesirable resources being created. For example, one text had the material listed as 'White marbl' rather than 'White marble'. These were corrected manually as they were encountered, and the data providers were notified of the errors.

Different versions of EpiDoc are used by the datasets. InsAph data is encoded in a version of EpiDoc based on TEI.2, whereas the IRT and HGV data is encoded in EpiDoc v6, which is based on TEI P5. Consequently, similar data is encoded in different ways. For example, the InsAph EpiDoc represents material (i.e. the material on which the inscription is written) using an XML element 'rs' with attribute 'material', whereas the IRT EpiDoc represents it as an XML element 'material'. As a consequence, scripts specific to each dataset were required, with common functionality abstracted out into a set of general purpose utility scripts.

The original, untranslated ancient texts of InsAph are in Greek. There was a desire to retain the Greek in resource URIs and literals for people and places. The question arises as to whether URIs containing Greek characters are valid. The W3C recommendation 'RDF Test Cases'¹³ states that for N-triples the character encoding is '7-bit US-ASCII text'. In contrast, RDF/XML allows a UTF-8 character encoding. Therefore, strictly speaking, Greek appearing in the URIs or literals of N-triples must be encoded using Unicode escape sequences. This can be a problem if the interface to a triple store holding N-triples does not recognise these escape sequences and cannot render them appropriately.

III. EVALUATION

Once we had created an epigraphic linked data set, we were able to evaluate the effectiveness of a linked data approach from the perspective of the targeted researchers. The key question that we wanted to answer was: does using a LD representation enable the data to be explored and understood more readily, especially when considering relationships between discrete datasets originating from different data providers?

The researchers that we involved in the evaluation were all classicists or related specialists, and were all 'digitally aware'. We targeted specifically researchers involved in the Digital Classicist community¹⁴. For the evaluation we decided to use off-the-shelf tools, as there are already a number of tools that are suitable for this purpose, and in any case the project did not have the staff resources to develop a new one. Virtuoso provides a faceted browsing extension that allows a researcher to follow RDF links and filter results based on defined RDF facets. A researcher can start with the results from a full-text search and refine the search by applying filters using classes, properties and other qualities of the LD. This functionality is also accessible via a web service that returns results as XML. Initially we considered using these extensions as a basis for evaluation, but we encountered issues with specifying Unicode queries, e.g. searches for Greek words, and, similarly, in rendering Unicode in triples. OntoWiki [6] was another option that has been successfully used in other humanities LD projects, but prior ontology engineering is required to exploit its capabilities to the best effect.

Finally, it was decided to use the Gruff-browser for the evaluation. Gruff renders LD as graphs, with subjects and objects as nodes and predicates as links. Gruff is not interoperable with Virtuoso, but this was not an issue since the purpose of the evaluation was to assess the LD approach from the researcher's perspective, so the underlying triple store was not relevant. The SPQR triples were uploaded into an AllegroGraph triple store (version 3.3), which is compatible with Gruff. To facilitate the evaluation, a tutorial was written for researchers on how to download and install

¹³<http://www.w3.org/TR/rdf-testcases/>

¹⁴<http://www.digitalclassicist.org/>

AllegroGraph and Gruff and upload triples into AllegroGraph.

Evaluation proceeded by a series of sessions with researchers, during which detailed observations were recorded about how researchers might use a LD browser to support or enhance their research activities. On the basis of these observations we developed a set of usage scenarios.

One observation was that researchers find it useful to be able to contextualise facts by using graph visualisation to browse related content; two scenarios were based on this observation. The first addressed how on-site research could be supported by the ability to contextualise particular words found in inscriptions. In the second scenario, researchers want to identify inscriptions which refer to a particular person (e.g. a particular Emperor), and filter these by certain attributes, such as time period or location. The linked data corpus — as explored through the Gruff browser — was considered to be potentially far more useful than a simple gazetteer or other reference resource, and reflected an advance beyond existing resources such as the online InsAph or IRT web portals, which allow keyword searches.

As well as the ability to browse related content, researchers found graph visualisation useful as a means of general exploration of the data. Gruff can also visualise the data as tables, displaying a resource URI and the relations of that resource. In future work, the integration of graph and table views should be improved, allowing properties for multiple nodes to be viewed simultaneously, to allow resources to be compared.

As well as the opportunities provided by the ability to visualise epigraphic information as graphs, researchers also discussed the issues and risks this raises. One risk is that, as more source information is incorporated and the datasets grow larger, the user becomes overwhelmed and lost within the data. The users need ways of orientating themselves within the data, beyond just browsing a graph. We considered a number of ways of addressing this issue:

- Returning to Virtuoso, we implemented filters using Virtuoso's faceted browsing features, to filter nodes according to specific criteria, for example by including just those from InsAph, or by date, referenced people or location.
- Allowing users to search for and display texts with particular attributes, e.g. all of those whose material is 'white marble', or which were found at a certain place, or date from a certain time. Gruff actually supports this functionality already, but this was not clear to the users during the evaluation.
- Keeping track of previously-visited nodes and branches using some form of 'breadcrumb trail'. Maintaining such a trail, and adding annotations recording why certain resources were visited and branches chosen, allows researchers to capture their browsing rationale and thus allow their reasoning and insights to be shared

with their peers.

Researchers identified functionality that they are used to from other services that should be replicated to allow for effective browsing of the SPQR datasets. For example, with GeoNames or Pleiades URIs can be entered into a web browser to access maps of these locations. Similarly, the IRT web portal supports pops-up maps for locations. Researchers expect a similar feature from a LD browser. For these features, we will evaluate more existing LD tools to enhance the SPQR offerings.

IV. CONCLUSIONS AND FUTURE WORK

SPQR sought to assess whether LD offered an appropriate representation for epigraphic data that would be more easily explorable, understandable and usable. Three epigraphic data sets were transformed to LD using EpiDoc-to-linked data conversion scripts. The LD was explored using an off-the-shelf browser by digital classicists. Their experiences revealed that a LD representation does have the potential to facilitate usable navigation of distributed data sets and offers an approach to exploring the data beyond that offered by existing online resources, which primarily support navigation via hyperlinks through indices or keyword searching. This applies to navigation both within a single dataset and across datasets from multiple providers.

These findings contrast positively with the outcomes of LaQuAT, which revealed that providing a single access point for multiple, related, datasets was of limited utility if the data could not readily be explored and understood. Queries in LaQuAT required a significant understanding of how the data was structured, whereas the universal subject-predicate-object schema of LD allows more intuitive data exploration and querying.

LaQuAT considered the potential for researchers to annotate existing datasets with their own observations, providing their own data that could be linked to existing read-only datasets. The relational model made this difficult to achieve, whereas it is straightforward in a LD model. This is demonstrated by the ease with which the corpus created by SPQR could be linked with GeoNames, Pleiades and the other datasets created and enhanced by Pelagios, as well as with existing InsAph and IRT data already online. These results suggest that further explorations of linked data as a means of representing, exploring and understanding humanities data would be fruitful.

Our immediate future work concerns the generation of larger amounts of linkable data. More work needs to be done on the conversion of EpiDoc documents to LD. A particular issue to be addressed is how to handle ambiguous data, such as place names that may refer to one of several different places. Such ambiguities represent decision points in the conversion workflow researcher input is required to direct the conversion. This raises the question as to whether conversion should be an interactive process or an entirely

automated process in which the conversion scripts make a selection but record their decision so a researcher can change it subsequently if required.

A second improvement would be to try and keep the data more up to date. Currently, data sets are converted statically into LD and then uploaded to the triple store, which in effect holds a snapshot of the data at a point in time. A more flexible alternative would be on-demand conversion and uploading, allowing existing data sets to be used as is and allowing any updates to be immediately accessible in LD format. Alternatively, the LD could be regenerated and cached at regular intervals. In addition, to ensure proper provenance and reproducibility, versioning of such compounded, and rapidly evolving, datasets would need careful consideration.

Virtuoso, finally, can integrate heterogeneous data sources into a federated RDF graph that allows non-RDF data sources, such as documents, to be parsed in real-time. LD from ancient texts could be combined with linked data derived from online data sources to provide more information about people or places cited in the texts. For example, GoogleBooks¹⁵ and the online journal resource JSTOR¹⁶ both provide APIs¹⁷ that can be used to query those resources, and the results can be converted to linked data [2]. These would need custom convertors to parse the data into triples and, again, some form of local caching of these triples.

REFERENCES

- [1] M. Jackson, M. Antonioletti, A. Hume, T. Blanke, G. Bodard, M. Hedges, and S. Rajbhandari, "Building bridges between islands of data - an investigation into distributed data management in the humanities," in *Proceedings of the 2009 Fifth IEEE International Conference on e-Science*, ser. E-SCIENCE '09. Washington, DC, USA: IEEE Computer Society, 2009, pp. 33–39. [Online]. Available: <http://dx.doi.org/10.1109/e-Science.2009.13>
- [2] C. Bizer, T. Heath, and T. Berners-Lee, "Linked data - the story so far," *International Journal on Semantic Web and Information Systems*, 2009. [Online]. Available: <http://tomheath.com/papers/bizer-heath-berners-lee-ijswis-linked-data.pdf>
- [3] M. Doerr, S. Gradmann, S. Henniecke, A. Isaac, C. Meghini, and H. van de Sompel, "The Europeana Data Model (EDM)," in *WORLD LIBRARY AND INFORMATION CONGRESS: 76TH IFLA GENERAL CONFERENCE AND ASSEMBLY 10-15 August 2010, Gothenburg, Sweden, Aug. 2010*. [Online]. Available: <http://www.ifla.org/files/hq/papers/ifla76/149-doerr-en.pdf>
- [4] B. Motik, P. F. Patel-Schneider, and B. C. Grau, "Owl 2 web ontology language direct semantics," *Director*, no. October, pp. 1–17, 2009. [Online]. Available: <http://www.w3.org/TR/2009/REC-owl2-direct-semantics-20091027/>
- [5] E. Prud'hommeaux and A. Seaborne, "SPARQL query language for RDF," W3C Recommendation, Tech. Rep., 1 2008. [Online]. Available: <http://www.w3.org/TR/2008/REC-rdf-sparql-query-20080115/>
- [6] S. Auer, S. Dietzold, and T. Riechert, "OntoWiki - A Tool for Social, Semantic Collaboration," in *International Semantic Web Conference*, ser. Lecture Notes in Computer Science, I. F. Cruz, S. Decker, D. Allemang, C. Preist, D. Schwabe, P. Mika, M. Uschold, and L. Aroyo, Eds., vol. 4273. Springer, 2006, pp. 736–749. [Online]. Available: <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.61.6346>

¹⁵ GoogleBooks, <http://books.google.com>.

¹⁶ JSTOR, <http://www.jstor.org>.

¹⁷ JSTOR Data for Research, <http://about.jstor.org/node/19881>.