

Semantic Web Applications for the Social Sciences

by Thomas Bosch¹ and Benjamin Zopilko²

Abstract

In recent years, Semantic Web technologies have matured and have made their way into various domains – for example, bioinformatics and eGovernment – where they are used in different applications to provide value-added services for users. In this paper, we present an overview of several representative applications that use Semantic Web technologies and show the potential of these technologies in the Linked Data world. We then present existing Semantic Web applications specifically for the social sciences and highlight their impact on this domain. Integrating the Semantic Web vision into the social sciences results in clear benefits, which we identify and discuss.

Keywords

Semantic Web, Linked Data, Semantic Web Applications, Social Sciences

Introduction

The corporate landscape is moving to the Semantic Web in a big way. Major companies like Adobe, Oracle, IBM, HP, Software AG, GE, Northrop Grumman, Altova, and Microsoft offer (or plan to offer) Semantic Web tools or systems. Others like Novartis, Pfizer, and Telefónica are using the Semantic Web (or are considering using it) as part of their own operations. Active participants in W3C Semantic Web related groups include HP, Agfa, SRI International, Fair Isaac Corp., Oracle, Boeing, IBM, Chevron, Siemens, Nokia, Pfizer, and Eli Lilly. In addition to the corporate sector, we see major communities such as digital libraries, defense sector, eGovernment, the energy sector, financial services, health care, the oil and gas industry, and the life sciences adopting the technologies.

For social science researchers, the Semantic Web and Linked Data hold great promise as Gregory and

Vardigan (2010) illustrate in detail. The adoption of semantic technologies has the potential to make the discovery of data and metadata in the Web more efficient, to enhance the reuse of social science metadata, and to decrease the technical barriers to employing data in research. Through Linked Data for the social sciences, users can easily discover the existence of data and can determine how the metadata are structured and whether the data are suitable for their interest. Therefore, it is necessary that the data is well-documented and that quality and provenance are explicit. This will enable the identification of complex relationships such as whether datasets are comparable or whether there are relationships to other versions of the same dataset. Since data collections published as Linked Data can easily be linked with each other, the integration and merging of heterogeneous datasets is also facilitated.

In this paper, we provide an overview of tools and applications that apply Semantic Web technologies and Linked Data. We also present social science projects and applications that make use of semantic technologies, and we discuss the benefits of

The number of applications utilizing Semantic Web technologies and Linked Data is increasing and illustrates the potential of these technologies for various domains

applying these technologies to the domain of the social sciences.

Semantic Web Applications

The number of applications utilizing Semantic Web technologies and Linked Data is increasing and illustrates the potential of these technologies for

various domains and heterogeneous communities. In this section of the paper we provide illustrative examples.

HCLS Demo

The W3C Health Care and Life Sciences Interest Group (HCLS) have developed demonstrations of the usage of Semantic Web technologies (W3C 2008a), which Herman (2011) has described. The demonstrations are designed to show the HCLS community how the Semantic Web can be used and to show the Semantic Web community how this Semantic Web technology can be useful in this specific application area. The core of the HCLS demo is the access and the integration of public datasets via the Semantic Web. One of the HCLS demonstrations is the Allen Brain Atlas, which is currently available only through an HTML interface. Mouse brains are cut in slices and stained for the presence of gene expression: 20,000 genes, 400,000 images at high resolution. The Allen Brain Atlas is 'mashed up' with Google maps. Google maps allows the user to upload his or her own 'maps', i.e., the URLs of bitmap images. The user then gets the navigation of large bitmaps for free. The goal of the demonstration is to find the right images in the Atlas data to really provide navigable data.

How is it done? What happens is that the query on the Atlas data is based on scraping the HTML structures, extracting the URI, and using SPARQL to combine these into more complex queries that result in the right image URIs being fed into the Google service. Via RDF, one gets a standard query SPARQL interface for free, providing much power to the end user. If the original authors of the brain Atlas had stored the data in a public MySQL database, one could have achieved the same user interface, but the fact is that they did not. RDF and SPARQL allow one to produce the relevant images easily without interfering with the original data in any way, using standard, off-the-shelf tools. The demo shows that sometimes RDF and SPARQL help in a very simple way – Semantic Web applications are not necessarily very complex.

NASA

Grone and Schain (2008) describe how to find the right experts at NASA using Semantic Web technologies. The NASA tool is an expertise locator for nearly 70,000 NASA civil servants. The tool uses RDF integration techniques across geographically distributed databases, data sources, and web services.

The authors use internal ontologies/vocabularies to describe the knowledge areas, and a combination of the RDF data and these ontologies to search through the (integrated) databases for specific knowledge expertise. The dump results from a faceted browser developed by the company to view the data on experts.

Semantic MediaWiki

Semantic MediaWiki (2012) is a module of the MediaWiki (2012) software and it extends wikis with ideas from the Semantic Web discipline. Semantic MediaWiki enables the user to make facts available for machines, thus making it easier for humans to search and reuse information. Articles, such as an article about the IASSIST conference in 2014, can be annotated semantically using the RDF format. In this way, RDF triples can be built and stored in the wiki code directly. Relations between subjects and objects can be defined semantically. The IASSIST conference in 2014 would be the subject of the RDF triples stated on the article. It could be specified that the IASSIST conference has participants, presentations, publications, key note speeches (e.g., relation 'hasKeyNoteSpeech'). Authors can also write articles for each relation and for each object to explain the semantics. In the context of this example, authors can write articles about the meaning of the relations to key note

speeches and also about each key note speech. On the site about a specific key note speech, the key note speaker could write an article about the content of the presentation.

Humans as well as machines can query all the information included in the wiki sites. People can query information about other articles when they edit new articles. Then the query results appear directly in the wiki site. These articles are updated automatically when the dependent sites are updated. As a consequence, overview articles are always up-to-date and consistent with the detail sites. Visitors of Semantic MediaWikis can download semantically enriched information directly in RDF. A large number of general-purpose RDF tools and specialized external programs can now reuse and process the RDF data in an easy and standardized way. The data and metadata exchange in RDF enables the combination of information from different sources like wikis.

Semantic MediaWiki is free software under the GPL license. More than 150 public wikis use the Semantic MediaWiki extension. In particular, semantic annotations in wikis (e.g., LexWiki, Concept-Hub-Wiki) are adopted in medical and biology sciences to create biomedical terminologies and ontologies collaboratively.

Web Repositories and Server Systems

In this section we present two popular systems for accessing, maintaining, and publishing data on the Web that utilize Semantic Web technologies.

Fedora Repository Project

The Fedora (Flexible Extensible Digital Object Repository Architecture) Repository Project (2013) is an open source software system originally developed by researchers at Cornell University. The underlying architecture enables the storage, management, and access of digital content. Fedora allows for expressing digital objects and relationships among them by assigning so-called 'behaviors' (i.e., services) to them. In addition to a core repository service with well-defined APIs, Fedora includes services for searching, OAI-PMH, messaging, and administrative clients, to name a few.

Regarding Semantic Web technologies Fedora supports interaction with RDF data, since the repository software can be connected with RDF triple stores. Data stored in a triple store can be accessed and used by every service of Fedora.

There are various scenarios and domains dealing with digital content where Fedora is applied. It can be used for "digital collections, e-research, digital libraries, archives, digital preservation, institutional repositories, open access publishing, document management, digital asset management, and more" (Fedora, 2013).

Virtuoso

Virtuoso (2013) is a data server that can be applied to various ways of storing, maintaining, and accessing different kinds of data. The core of Virtuoso is an object-relational SQL database. For serving dynamic web pages, Virtuoso provides a flexible built-in web server, which can process pages written in Virtuoso's own web language (VSP) and other standard languages like PHP or ASP. Virtuoso also provides functionalities for maintaining and managing the published web pages like versioning, automatic metadata extraction, and full text searching. It is also available as an open source edition at <http://www.openlinksw.com/wiki/main/>.

With respect to semantic technologies Virtuoso currently enables the storing and querying of RDF data in its database. Since this is currently a SQL database, a translation of SPARQL queries into SQL is supported in order to query the RDF data and to provide RDF as output format as well. There are plans to extend the access and storage capabilities of connected databases, which would also enable particular technologies like inferencing on RDF data.

Thesaurus Management Tools

Thesauri and classifications are commonly used instruments for describing and annotating metadata about various kinds of documents. They have a long tradition in libraries and archives. In the following paragraphs we describe two thesaurus management tools that use Semantic Web technologies and datasets.

PoolParty Thesaurus Server

PoolParty Thesaurus Server (PPT) (2013) is a software platform that enables the management and maintenance of complex knowledge models such as taxonomies, thesauri, controlled vocabularies, and similar data. The metadata are fully organized and modeled using W3C’s Semantic Web standards RDF and SKOS, since SKOS focuses particularly on knowledge organization systems. Managing this data inside PPT is enhanced by text mining functionalities and linked data mapping technologies. In addition to processing RDF data, the APIs provided by PPT are also based on semantic technologies, the SPARQL standard of W3C. This allows also an integration of the maintained knowledge models in other systems, e.g., CMS, ERP-Systems, or Wikis. By using complex Semantic Web based approaches like text corpus analysis, entity extraction, linked data enrichment, and SKOS thesaurus management, it is possible to build, maintain, and publish large and complex knowledge models based on RDF data.

VocBench

VocBench (2013) is a vocabulary editing and workflow tool developed by the Food and Agriculture Organization (FAO). The web-based application enables the transformation of multilingual knowledge organization systems like thesauri, authority lists, and glossaries into SKOS/RDF concept schemes. Thus, traditionally

maintained thesauri can easily be used in Semantic Web applications. Besides the transformation, VocBench also allows for managing, maintaining, and editing the data. This includes collaborative editing as well as validation and quality assurance tasks. VocBench is an open source project and based on Protegé.

Currently, VocBench is used to manage several datasets held at FAO like the AGROVOC thesaurus (2013), the Biotechnology Glossary (2013), and bibliographic metadata used in FAO. For future releases FAO plans to include a native interface for SKOS and SKOS-XL and configurable support for hosting of different triple store technologies. Also, they plan to support generic OWL ontologies.

Vivo Project

The VIVO project (2013) is an open source Semantic Web application, which was originally developed and implemented at Cornell University. The application maintains profiles of researchers and organizations. These profiles can be populated with additional information like activities or interests. Through extensive search and browse capabilities, it is possible to discover information across institutions and disciplines. Although the VIVO software is installed locally, the different installations worldwide are connected with each other in a network, which also enables integrated searching and browsing across the information of all connected installations. The information that can be discovered can be used in different contexts, e.g., in visualizations or in applications like VIVO Searchlight, which allows to search for VIVO profiles based on textual information from any web page. The open source project is available at <http://vivo.sourceforge.net>.

The structured data in VIVO is represented in RDF using the VIVO ontology (Mitchell et al., 2011). This ontology focuses on describing researchers and networks of researchers across organizations and disciplines. It also covers researchers’ teaching activities, their expertise, their research, and which service activities they provide. The ontology has been developed inside the VIVO project.

Semantic Web Applications for the Social Sciences



Figure 1. SOFISwiki – Search for research institutions

So far, there are just a few Semantic Web applications for the social sciences. We present them in this section in detail. For each Semantic Web application for the social sciences, we show individual benefits for users in the social sciences community regarding Semantic Web supported functionalities.

A preview of these social science applications:

- SofisWiki provides research institutions the possibility to publish information about their research institution, their research activities, and their research projects
- The Microdata Information System (MISSY) is an information system to document German and European studies at the variable and the study level.
- NESSTAR enables publishing of a huge amount of statistical data and metadata using Semantic Web technologies.
- Colectica is a fast way to design, document, and publish survey research using open data standards.

SOFISwiki

Are you conducting a social science research project? Are you writing a thesis or doing advanced work on a social science topic? If so, then you may want to create your project in SOFISwiki and make your research work transparent for the scientific community.

GESIS has developed SOFISwiki (2012), the first Semantic MediaWiki in the social science community.

SOFISwiki informs researchers about research activities and projects as well as research institutions in the German-speaking social sciences. SOFISwiki contains all entries from the last ten years of SOFIS (2012), a central database hosted by GESIS that delivers information about over 50,500 research projects. Using SOFISwiki, research institutions like universities can enter and represent information about their research projects and their institutions in convenient forms. Other researchers can get an overview of the wiki content and can also search for and find this information. The collected project information is also available using the social science portal Sowiport (2012). As part of future work, SOFISwiki could be extended to support the documentation of research activities, projects, and institutions all over the world. Figure 1 shows the graphical user interface of SOFISwiki to enable a search for research institutions by institution name, country and location, content orientation, and institution type.

Figure 2 shows the result of the above query for the department ‘Monitoring Society and Social Change’ of the research institution

GESIS - Leibniz-Institut für Sozialwissenschaften Dauerbeobachtung der Gesellschaft (Mannheim)

Institutionsname: [GESIS - Leibniz-Institut für Sozialwissenschaften Dauerbeobachtung der Gesellschaft](#)

Institutions-Nr.: 069258

Inhaltliche Ausrichtung: [Sozialwissenschaften](#)

Organisationstyp: [andere außeruniversitäre Forschungseinrichtung](#)

Homepage: <http://www.gesis.org/das-institut/wissenschaftliche-abteilungen/dauerbeobachtung-der-gesellschaft/>

Straße: B2,1

Ort: [Mannheim](#)

PLZ: 68072

Postfach: 122155

Land: [Bundesrepublik Deutschland](#)

Email: christof.wolf@gesis.org

Telefon: 0621 1246-0

Projekte

Zur Zeit sind 12 Projekt(e) für diese Institution vorhanden.

↕

[Data without Boundaries \(DwB\)](#)

[Erwerbs- und Betreuungspotenziale von Paaren mit Kindern: Realisierungschancen einer gleichmäßigen Arbeitsteilung](#)

[Externe Managementunterstützung zur Erleichterung von Ausgründungsvorhaben \(Good Practice\) mit dem Ausgründungsvorhaben "Bodymonitor" aus dem GESIS Leibniz-Institut für die Sozialwissenschaften](#)

[German Longitudinal Election Study \(GLES\)](#)

Figure 2. SOFISwiki – Research institution and associated research projects

'GESIS'. Department metadata such as address, homepage, and contact person as well as the associated research projects are delivered.

Benefits for the social sciences community

Research institutions like universities can manage and publish information about their research institutions and their research projects using this tool, and researchers can get an overview of research projects, institutions, and activities. Internally, the metadata are represented using RDF to ease the integration of the metadata items. RDF is also used to integrate SOFISwiki with diverse other portals like Sowiport and SOFIS.

The Microdata Information System (MISSY)

General Description of MISSY

The Microdata Information System (MISSY) (GESIS 2012) maintains the largest household survey in Europe – the German microcensus, which provides statistics about the general population in Germany, including the employment market (occupation, professional education, income, legal insurance). MISSY consists of approximately 500 variables and questions and captures data for 25 years, since 1973. Figure 3 shows the graphical user interface of MISSY. In the detailed view of the variable, the user gets details about the variable "gender" including associated question, values, value labels, and absolute and relative frequencies.

MISSY has two parts:

- Missy Web, the end-user front-end
- Missy Editor for the metadata documentation, which is the back-end

Several use cases are covered by MISSY:

- Thematic classification: variables by thematic classification and year
- Variables by year
- Generated variables by year
- Details of variables with statistics
- Variable-Time Matrix: Variables by thematic classification and year (selectable)
- Questionnaire Catalogue

In the third generation of MISSY, additional surveys such as EU-SILC (European Union Statistics on Income and Living Conditions), EU-LFS (European Union Labour Force Survey), and EVS (European Values Study) will be integrated. The MISSY Editor will be implemented as a web application. In future, it should also be possible to browse variables by survey and by country.

MISSY and Linked Data

The MISSY-specific data model is based on the DDI-RDF Discovery Vocabulary – Disco (Bosch et al. 2012) for the following reasons:

- Disco contains the most salient components of both DDI-Codebook and DDI-Lifecycle for data discovery (as a

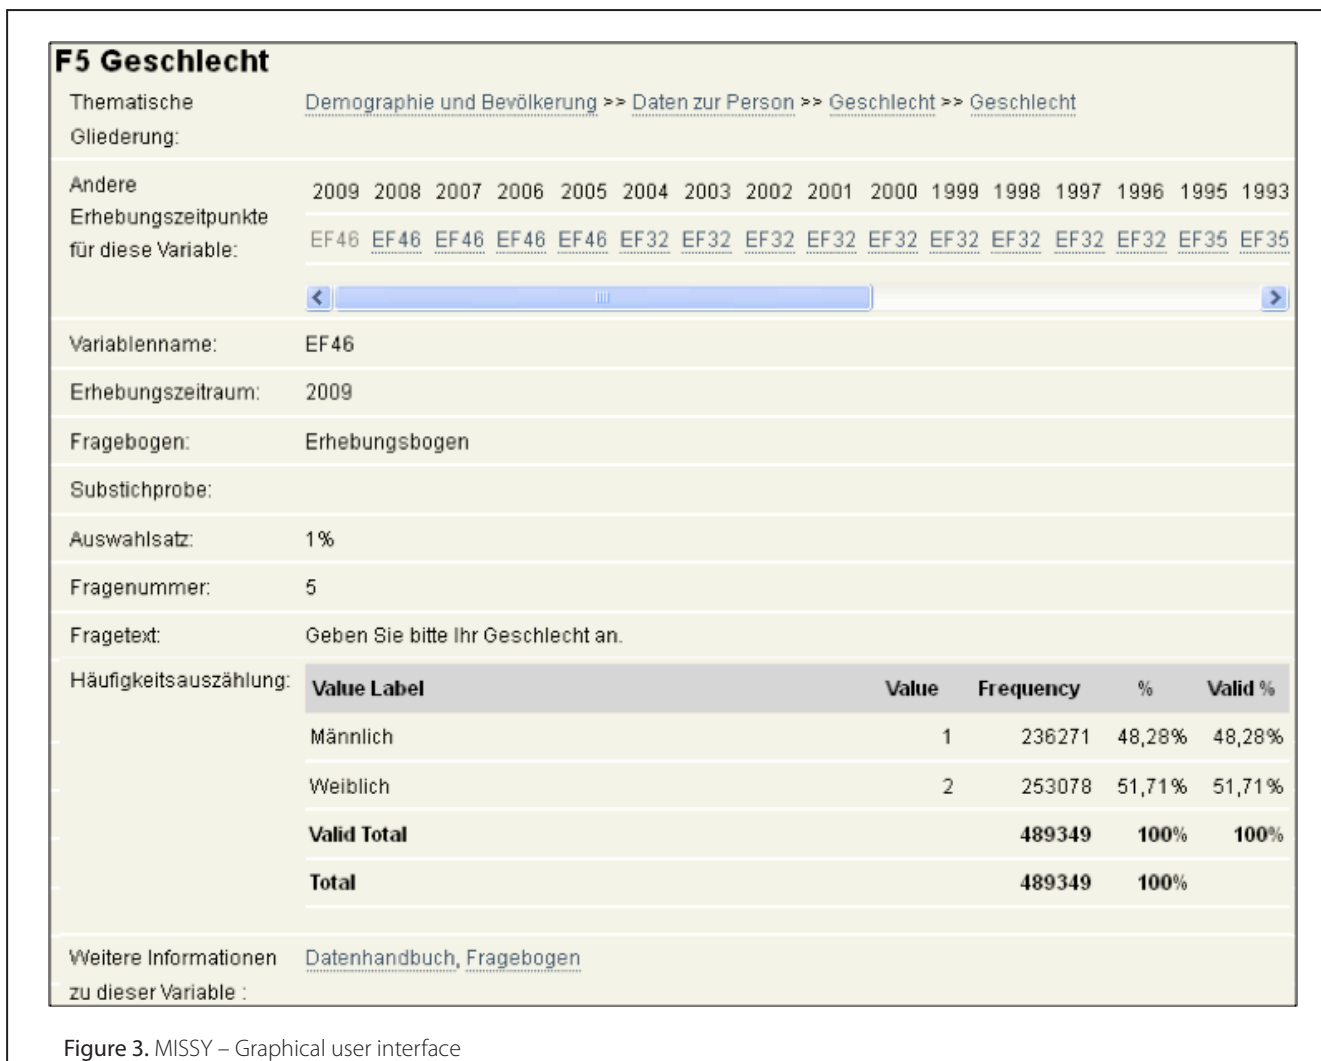


Figure 3. MISSY – Graphical user interface

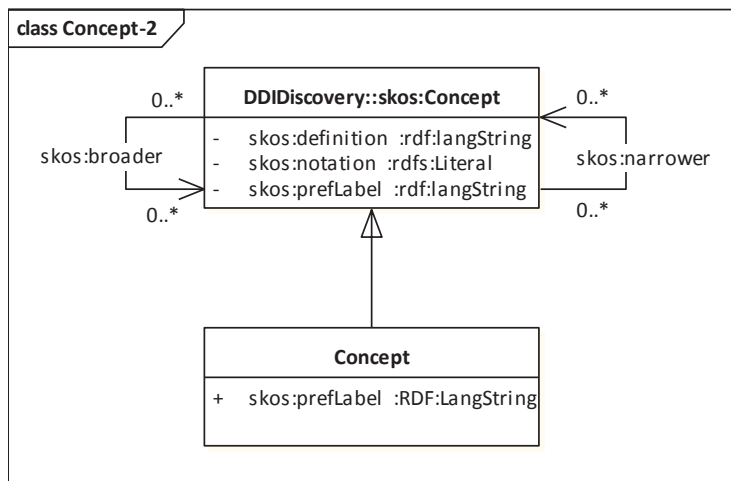


Figure 4. MISSY – Abstract and individual data model

consequence, not all of the over 830 XML elements of DDI 3.1 are covered)

- Disco serves as a first step in developing a model-driven DDI specification to document microdata within the social, behavioral, and economic sciences
- Disco will be officially published and publicly available
- More than 20 experts from the statistics and the Linked Data community of eight different countries have contributed to the development of Disco in three workshops and additional working groups

As not every requirement within the MISSY context is covered by the DDI Ontology, an individual data model is defined on top of the common abstract data model. Other software projects intended to document studies at the study and variable level using DDI-L can also be based on Disco and can reuse existing code, which is made available on a GitHub repository (MISSY 3 2012). To show how this works, an example is provided below. The SKOS (Simple Knowledge Organization System) (W3C 2009), whose purpose is to define hierarchies of concepts, is reused in the Disco ontology to a large extent (see Figure 4).

For instance, codes, categories, DDI concepts, and study subjects are represented in Disco as `skos:Concepts`. The next figure visualizes `skos:Concepts` within MISSY. The `skos:Concept` used in the DDI Discovery Vocabulary is extended by the `Concept` defined in MISSY. Using the property `skos:prefLabel`, category labels can be stored in an RDF format. The datatype of the category labels is specified as a normal string. One requirement in MISSY is to store category labels in different languages such as English, German, and French. Thus, we have defined the type 'Multilingual' in MISSY. In order to represent category labels, the property `skos:prefLabel` of the datatype 'Multilingual' is used and therefore the initial common abstract data model is extended.

In a well-defined software architecture, the application itself does not need to know how the data are stored. The application just needs to know the API, i.e., methods that are provided to access and store objects. These methods may be abstracted away from the actual implementation. An actual implementation or strategy can just be a matter of configuration. A strategy is an implementation of the actual type of persistence or physical storage, e.g., DDI-L-XML, DDI-RDF, XML-DB, or Relational-DB. The persistence API defines the persistence functionality for model

components regardless of the actual type of physical persistence. Several components implement the persistence functionality defined in the persistence API with respect to the usage of relational DBs, DDI-XML, and DDI-RDF. One concrete implementation of the persistence API is DDI-RDF, the RDF representation of the developed DDI Ontology. MISSY will offer several export formats – one of them will be DDI-RDF.

We will implement two additional concepts providing RDF data. First, MISSY websites will be annotated semantically using RDFa (W3C 2012), which are generic annotations in XHTML documents. Thus, machines can crawl the MISSY websites in order to import exactly the information needed for further processing, since now machines 'know' the meanings of the provided information. RDFa metadata should and will be provided according to the DDI ontology Disco and according to the schema.org vocabulary (Schema.org 2012). Launched in 2011, Schema.org is an initiative from Bing, Google, and Yahoo to provide a vocabulary (a collection of concepts and their properties) to be used by web masters to markup web content in ways recognized by major search providers. Search engines will rely on this markup to improve the display of search results, making it easier for people to find the right pages they search for. The second way of exporting semantic information to the social science community is to build a SPARQL endpoint. RDF triples are stored in a triple store in parallel to the XML documents containing both the data and metadata of multiple studies offered by MISSY.

Benefits for the social sciences community

Other software projects documenting studies on the study and variable level using DDI-L can reuse existing GitHub repository code. MISSY will provide multiple export formats (e.g., DDI-RDF). DDI data as well as metadata can be published in the Linked Open Data cloud. MISSY websites will be annotated using RDFa according to DDI-RDF and schema.org. As a consequence, search engines will improve their query results. By writing SPARQL queries, DDI data and metadata can be accessed from SPARQL endpoints.

NESSTAR

NESSTAR (Norwegian Social Science Data Services 2012) is a Semantic Web application for documenting both statistical data and metadata. NESSTAR can be seen as an extraordinarily easy to understand SW application, as it does not specify sophisticated ontologies, does not use advanced RDF features such as reification, and does not use logical inference.

In 1998, the European Union funded the research and development project called 'Networked Social Science Tools and Resources,' abbreviated as NESSTAR. The EU project with the name 'FASTER' (Faster 2012) followed the goals associated with the NESSTAR project. Assini (2002) gives a rather detailed description of NESSTAR.

The aim of NESSTAR is to make a huge quantity of statistical data and metadata accessible using Semantic Web technologies. Before the implementation of NESSTAR, statistical data as well as metadata was typically only available in a human-readable and understandable form and not additionally in a machine-understandable form that could be further processed by

computer programs. NESSTAR was intended to revolutionize the way people access statistical information, bringing the advantages of instant access to the world of statistical data dissemination. On the Nesstar website (Nesstar 2013), a list of Nesstar catalogues (e.g., surveys, tables) is provided by the Nesstar's Demo Server. Figure 5, for example, shows information such as the associated question, the values and the categories, summary statistics, interviewer instructions, and the total responses about the variable gender of a demo survey.

Conceptual model of NESSTAR

The NESSTAR object model is defined in RDFS. About 15 classes represent the key domain-specific concepts within the statistical domain like studies, data files, variables, indicators, and tables. The conceptual model also includes relationships between domain-specific concepts: studies, for example, may contain cubes having one or more dimensions. Additionally, 10 domain independent

support classes are part of the object model of statistical data and metadata. The class Server, for instance, represents the server where the metadata objects are hosted. It provides basic administrative functionality such as file transfer, server reboot, and server shutdown. Starting from a server and by recursively traversing the objects' relationships, applications can reach all the server's objects. The domain independent support class Catalog groups metadata objects. Instances of Catalogs can be browsed and you can get a list of all the metadata objects which are included in the Catalog. There is also the possibility to search for particular objects.

Many research studies contain sensitive information that cannot be made available without restrictions. Within NESSTAR, access control policies can be defined in order to follow a security model. To implement this, classes such as User, Role (e.g., administrator, final user, data publisher), and Agreements (e.g., 'I agree to

use this data only for non-commercial research purposes') are specified.

NESSTAR is based on lightweight, object-oriented web middleware named NEOOM - NESstar Object Oriented Middleware. NEOOM is based on Web and Semantic Web standards like HTML, HTTP, RDF, and RDFS. NEOOM is a set of guidelines on how to use web as well as Semantic Web technologies to build distributed object-oriented systems. The NEOOM guidelines are described extensively in Assini (2001b) and very briefly in Assini (2001a). RDFS does not provide a way to describe the behavior, the operations of statistical objects (e.g., queries, statistical operations, file transfers, tabulate, and frequency). How to specify the operations formally? In the NEOOM Object Model, specific methods (e.g., Login) are defined as sub-classes of the Method class. Concrete method invocations are then instances of the Method class.

Behavioral view on NESSTAR

According to the NESSTAR conceptual model, data publishers make their statistical data and metadata available on the web as objects. These objects are represented by RDF resources according to the NESSTAR object model. Each data publisher runs its own server, which is an instance of the class Server. NESSTAR

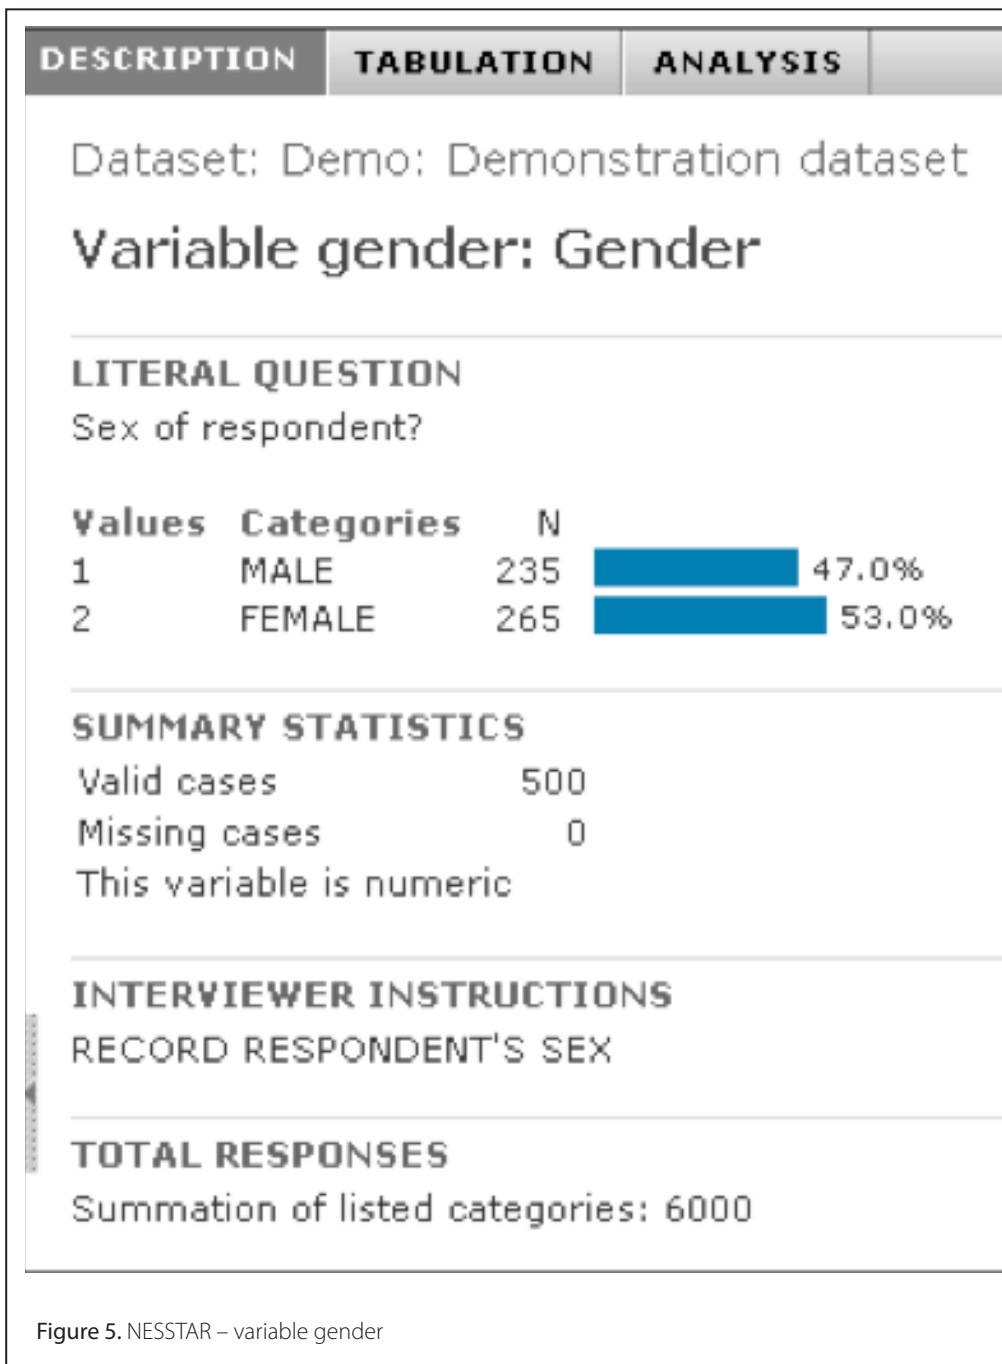


Figure 5. NESSTAR – variable gender

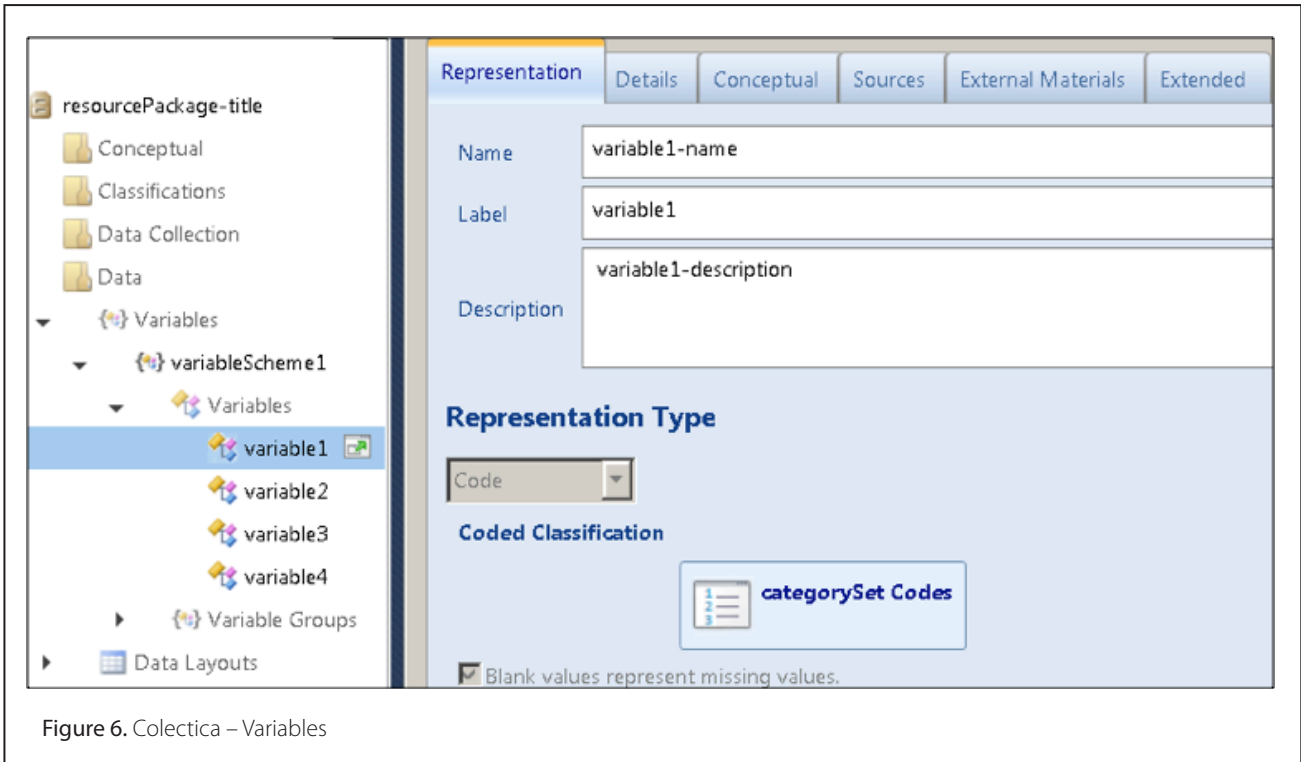


Figure 6. Colectica – Variables

servers host the maintained objects. NESSTAR servers provide WWW resources such as HTML pages and images as well as statistical objects. NESSTAR is fully distributed and each server is totally independent and integrated. Users have the possibility to access statistical objects remotely by simply typing objects' URLs.

SOAP (W3C 2007) is used for remote object-oriented calls. Similar to using search engines like Google, users can search for remote statistical objects: they could for example type the search term 'find all variables about political orientation'. In NESSTAR, there are different kinds of user access possibilities: NESSTAR Explorer,

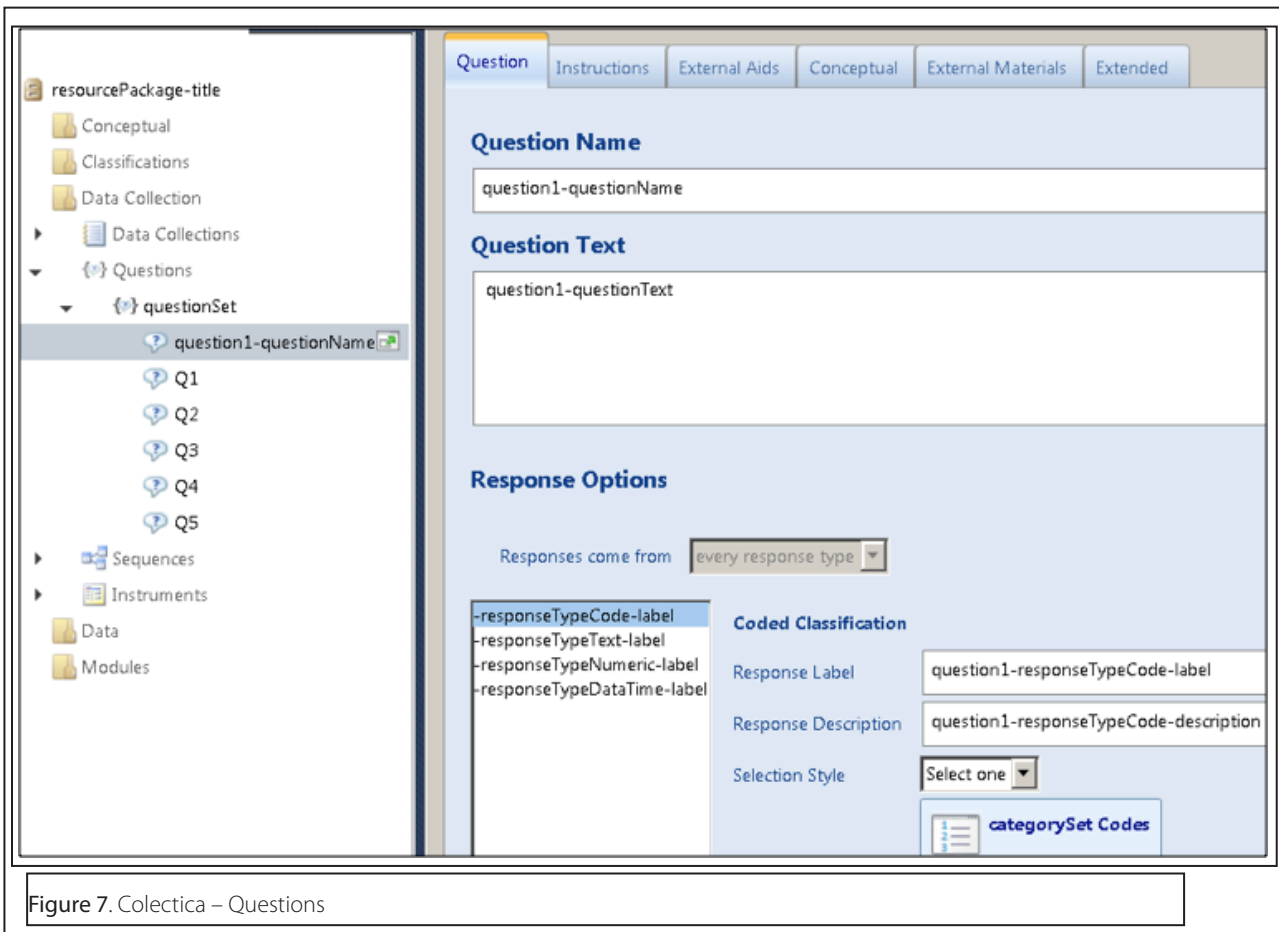


Figure 7. Colectica – Questions

NESSTAR Light Explorer, NESSTAR Publisher, and Object Browser. NESSTAR Explorer is similar to a common web browser. Users can enter objects' URLs and WWW resources are displayed as they would be displayed in a web browser. NESSTAR Publisher is a tool for editing metadata, for validating, and for publishing. The Object Browser's purpose is to test and to administer statistical objects.

Benefits for the social sciences community

NESSTAR enables users to publish a huge amount of statistical data and metadata using Semantic Web technologies. Now statistical data and metadata is not only available in a human-readable and understandable form but also in a machine-understandable form which can be further processed.

Colectica RDF Services

Colectica is a fast way to design, document, and publish survey research using open data standards. The Colectica Platform provides features for statistical agencies, survey research groups, public opinion researchers, data archivists, and other data intensive operations.

Colectica can increase the expressiveness and longevity of the data collected through standards-based metadata documentation (Colectica 2012c). DDI-L allows for reuse and harmonization of metadata items through the use of referencing. With the Colectica 4.0 Repository Addin, the relationships between metadata items are indexed (Colectica 2012b), which makes it possible to execute queries on these relationships. Figure 6 displays the documentation of variables using Colectica Designer. You can specify variable names, variables labels, variable descriptions, the response unit, associated concepts and universes, and the representation of the variable (e.g., numeric, textual, or coded representation).

Figure 7 shows how to document questions using Colectica Designer. For questions you can specify question names, question text, the question scheme, and the response domain (e.g., numeric, textual, or coded).

The Colectica RDF model is created by hand based on the Colectica DDI-L model. Each description of a DDI metadata item is stored as a named graph. The RDF Services Architecture can be deployed with Colectica Repository. All DDI metadata items, which are stored and versioned in the Repository, are also stored in RDF (Colectica 2012a). Several external vocabularies such as RDF, RDFS, simple Dublin Core (DC), the DCMI Metadata Terms (DCTERMS), OWL, XSD, and FOAF are reused (Colectica 2012b).

SPARQL (W3C 2008b) is a query language created for searching RDF data and is standardized by the W3C. It allows for searching based on the relationships and literal data stored in an RDF graph or store. SPARQL can be used to construct very precise questions about DDI metadata items referencing multiple metadata items. There are two deployment scenarios which can be distinguished in Colectica: internal RDF stores and external RDF stores. Using Colectica Repository's internal RDF store, SPARQL 1.0 as well as the draft version 1.1 of SPARQL are supported. The SPARQL Update functionality is disabled in order to maintain consistency with the versioned DDI metadata items in the repository. For Colectica Repository, it is also possible to replicate the RDF to external already existing RDF stores, which is the second deployment scenario (Colectica 2012a).

One can query DDI-L as RDF either using a web service from Colectica Repository or using a SPARQL endpoint on Colectica Web. In addition, each DDI-L metadata item, which is stored in the Colectica Repository, can be downloaded as an RDF dump

(Colectica 2012a). One example of such a SPARQL query in the statistical domain could be: Which studies has Dan Smith - the software developer of Colectica - authored since the beginning of 2010 (Colectica 2012a)?

```
PREFIX ddi: <urn:ddirdf:>
PREFIX ddit: <urn:ddirdf:type:>
PREFIX dc: <http://purl.org/dc/elements/1.1/>
SELECT ?study
WHERE {
    ?study a ddit:StudyUnit;
    dc:date ?creation_date;
    dc:creator <http://dan.smith.name/who#dan>.
    FILTER (xsd:dateTime(?creation_date) > "2010-01-01
    00:00:00"^^xsd:dateTime) . }
ORDER BY ?study
```

Another example of a SPARQL query would be: How many times has a variable been reused across multiple datasets (Colectica 2012a)?

```
PREFIX ddi: <urn:ddirdf:>
PREFIX ddit: <urn:ddirdf:type:>
PREFIX dc: <http://purl.org/dc/elements/1.1/>
SELECT ?variable COUNT (?parent) AS c
WHERE {
    ?variable a ddit:Variable ;
    ?parent ddi:HasVariable ?variable .
    ?parent a ddit:Dataset. }
GROUP BY ?variable
```

Dan Smith's website (Colectica 2012b) also offers several examples of DDI-RDF serializations. A further SPARQL example from the DDI-L US 2010 Census sample file is also provided. As part of future work, predicates will be updated when the official and community adopted DDI Discovery Vocabulary is available (Colectica 2012a).

Benefits for the social sciences community

DDI-L data and metadata can be queried as RDF using the Colectica Repository web service or the Colectica Web SPARQL endpoint. DDI-L metadata items, stored in the Colectica Repository, can also be downloaded as RDF dumps.

Conclusion and Future Work

We have presented several representative applications that apply Semantic Web technologies to a high degree. While semantic technologies and Linked Data have yet not been widely used in the social sciences, we have identified initial applications exclusively developed for this domain. The impact of Semantic Web and Linked Data are exposed in these applications. Additional potentials and benefits for an adaption of semantic technologies for scientific purposes can easily be identified. We have shown individual benefits for users of the social sciences community regarding Semantic Web functionalities.

References

- AGROVOC 2013 AGROVOC Thesaurus, viewed 6 May 2015, < <http://aims.fao.org/vest-registry/vocabularies/agrovoc-multilingual-agricultural-thesaurus> >
- Assini, P 2001a 'Objectifying the Web the 'light' way: an RDF-based framework for the description of Web objects', Proceedings of the International World Wide Web Conference, Hong Kong, 01 Mai 2001, tenth International World Wide Web Conference.

- Assini, P 2001b NEOOM: A Web and Object Oriented Middleware System, [Online], Available: <http://www.nesstar.org/sdk/neoom.pdf> [6 May 2015].
- Assini, P 2002, 'A Semantic Web Application for Statistical Data and Metadata', Proceedings of the International World Wide Web Conference, Hawaii, 07 May 2002, 11th International World Wide Web Conference.
- Biotechnology Glossary 2013 Biotechnology Glossary, viewed 6 May 2015, <<http://www.fao.org/biotech/biotech-glossary/en/>>
- Bosch, T, Cyganiak, R, Wackerow J & Zapolko B 2012 'Leveraging the DDI Model for Linked Statistical Data in the Social, Behavioural, and Economic Sciences', Proceedings of the International Conference on Dublin Core and Metadata Applications, Kuching, 03 September 2012, International Conference on Dublin Core and Metadata Applications, pp46 - 55.
- Colectica 2012a Accessing DDI 3 as Linked Data: Colectica RDF Services, viewed 6 May 2015, <<http://www.iassistdata.org/conferences/2012/presentation/3326>>.
- Colectica 2012b DDI 3 meets RDF and SPARQL with Colectica Repository, viewed 6 May 2015, <<http://dan.smith.name/2011/10/ddi-3-meets-rdf-and-sparql-with-colectica-repository/>>.
- Colectica 2012c Colectica Website, viewed 6 May 2015, <<http://www.colectica.com/>>.
- Faster 2012 Faster, viewed 6 May 2015, <<http://fasterproject.eu/>>.
- Fedora 2013 Fedora Repository Project, viewed 6 May 2015, <<http://fedora-commons.org/>>.
- GESIS 2012 The Microdata Information System (MISSY), viewed 6 May 2015, <<http://www.gesis.org/missy/>>.
- Gregory, A & Vardigan, M 2010 The Web of Linked Data: Realizing the Potential for the Social Sciences, viewed 6 May 2015, <http://odaf.org/papers/201010_Gregory_Arofan_186.pdf>
- Grove, M & Schain, A 2008 Case Study: POPS — NASA's Expertise Location Service Powered by Semantic Web Technologies, viewed 6 May 2015, <<http://www.w3.org/2001/sw/sweo/public/UseCases/Nasa/>>.
- Herman, I 2011 Semantic Web Adoption and Applications, viewed 6 May 2015, <<http://www.w3.org/People/Ivan/CorePresentations/Applications/>>.
- MediaWiki 2012 MediaWiki.org, viewed 6 May 2015, <<http://www.mediawiki.org/wiki/MediaWiki>>.
- Mitchel, S, Chen, S, Ahmed, M, Lowe, B, Marks, P, Rejack, N, Corson-Rikert, J, He, B, Ding, Y 2011 'The VIVO Ontology: Enabling Networking of Scientists' ACM WebScience Conference, Koblenz
- MISSY 3 2012 MISSY 3 project, viewed 6 May 2015, <<https://github.com/missy-project>>.
- Nesstar 2013, Welcome to Nesstar's Demo Server, viewed 6 May 2015, <<http://nesstar-demo.nsd.uib.no/webview/>>.
- Norwegian Social Science Data Services 2012 Nesstar, viewed 6 May 2015, <<http://www.nesstar.com/>>.
- PPT 2013 PoolParty Thesaurus Server, viewed 6 May 2015, <<http://www.poolparty.biz/>>
- Schema.org 2012 Schema.org, viewed 6 May 2015, <<http://schema.org/>>.
- Semantic MediaWiki 2012 Semantic MediaWiki, viewed 6 May 2015, <<http://semantic-mediawiki.org>>.
- SOFIS 2012 SOFIS - Social Science Research Information System, viewed 6 May 2015, <<http://www.gesis.org/en/services/research/sofis-social-science-research-information-system/>>.
- SOFISWiki 2012 SOFISWiki, viewed 6 May 2015, <<http://www.gesis.org/sofiwiki/Hauptseite>>.
- Sowiport 2012 Sowiport, viewed 6 May 2015, <<http://sowiport.gesis.org/>>.
- Virtuoso 2013 Virtuoso Universal Server, viewed 6 May 2015, <<http://virtuoso.openlinksw.com/>>
- VIVO 2013 VIVO Project, viewed 6 May 2015, <<http://www.vivoweb.org/>>
- VocBench 2013 VocBench, viewed 6 May 2015, <<http://aims.fao.org/vest-registry/tools/vocbench-2>>
- W3C 2007 SOAP Version 1.2 Part 0: Primer (Second Edition) - W3C Recommendation 27 April 2007, viewed 6 May 2015, <<http://www.w3.org/TR/2007/REC-soap12-part0-20070427/>>.
- W3C 2008a HCLS/Banff2007Demo, viewed 6 May 2015, <<http://www.w3.org/wiki/HCLS/Banff2007Demo>>.
- W3C 2008b, SPARQL Query Language for RDF, viewed 6 May 2015, <<http://www.w3.org/TR/2008/REC-rdf-sparql-query-20080115/>>.
- W3C 2009 SKOS Simple Knowledge Organization System Namespace Document - HTML Variant - 18 August 2009 Recommendation Edition, viewed 6 May 2015, <<http://www.w3.org/2009/08/skos-reference/skos.html>>.
- W3C 2012 RDFa 1.1 Primer - Rich Structured Data Markup for Web Documents - W3C Working Group Note 07 June 2012, viewed 6 May 2015, <<http://www.w3.org/TR/2012/NOTE-rdfa-primer-20120607/>>.

Notes

1. Thomas Bosch GESIS - Leibniz Institute for the Social Sciences, Mannheim, Germany E-Mail: thomas.bosch@gesis.org
2. Benjamin Zapolko GESIS - Leibniz Institute for the Social Sciences, Köln, Germany E-Mail: benjamin.zapolko@gesis.org