Semantic Web Opportunities for Digital Libraries

Laura Hollink, Antoine Isaac, Véronique Malaisé, Guus Schreiber

Department of Computer Science
Vrije Universiteit Amsterdam
de Boelelaan 1081 HV
The Netherlands

April 11, 2008

1 Introduction

The World Wide Web has enabled easy access to a myriad of information resources. People can access the objects and metadata in digital collections of libraries, museum and archives from anywhere in the world. This opens the possibility to reach a broader public.

This also creates the opportunity to connect the archives of different institutes, so users could access multiple collections from one location, with one vocabulary, using one interface. In this paper, we discuss what knowledge and techniques from the Semantic Web community can contribute towards the realization of cross-archive access. The archive and library fields are particularly suitable for the Semantic Web approach, since rich and well-structured knowledge sources are often available. There is a long standing tradition of developing knowledge organization schemes (KOSs) such as controlled vocabularies, thesauri, classification schemes, subject heading systems and taxonomies, to index large collections of objects like books or museum artifacts. Portals such as the Memory of The Netherlands[1] suggest a trend towards cross-collection access.

One of the problems that have to be overcome in order to realize cross-archive access, is the issue of interoperability between different KOSs (Isaac *et al.*, 2008). Consider, for example, two collections, each indexed by its own KOS. A search term from the vocabulary of the first collection will not give any results from the second collection, and vice versa. A query for ‘Charlemagne’ or ‘Charles the Great’ will not match any objects indexed with ‘Charles I’, which describes the same person but stems from another vocabulary. Similarly, a query for ‘ballet’ will not retrieve documents indexed as ‘Swan Lake’.

The Linking Open Data project[2] promotes the exchange of data and metadata vocabularies. The project publishes various open data sets on the web in a

[1] `http://www.geheugenvannederland.nl/`

[2] `http://esw.w3.org/topic/TaskForces/CommunityProjects/LinkingOpenData/`

standard Semantic Web format, including Wikipedia[3], Geonames[4], WordNet[5], the World Factbook[6], DBLP[7] and many more. Links are created that connect these data sets. In October 2007, the participating data sets had a total size of over two billion statements, and they were connected by around 3 million links. This has a huge potential for interlinking the collections of different libraries.

With this paper we intend to provide the reader with some insight into recent developments in the Semantic Web field regarding interoperability and access. It is not meant as an exhaustive summarization of Semantic Web activities. In the next section, we will give a brief overview of what the Semantic Web is. Readers familiar with the Semantic Web may skip this section. Section 3 will discuss interoperability issues. In section 4 we discuss how one can improve access to data using semantics. In section 5 we present opportunities for enrichment of metadata and vocabularies. In section 6 we conclude with a discussion of the maturity of Semantic Web techniques and the feasibility of using them in real life scenarios.

2 The Semantic Web

Today's Web content is mostly meant for human consumption. The Semantic Web (Berners-Lee *et al.*, 2001) is an extension of the current Web, augmented with a machine-processable representation of the information. At the lowest level, the Semantic Web consists of *resources*, which denote any element that can be identified on (or even outside) the Web. Files and documents can be resources, but also concepts or terms is a controlled vocabulary, people, places, etc. Resources are identified (or named) by Unified Resource Identifiers (URI's). They are linked together by three-part statements consisting of a *subject resource*, a *property resource* and an *object resource*. An example of three resources that form such a triple is `ELAG2008 - location - Wageningen` stating that ELAG 2008 is located in Wageningen, or `Wageningen - type - inhabited_place` stating that Wageningen is an inhabited place. The latter is an actual triple from the Thesaurus of Geographical Names[8]. Together, such triples form a directed, labeled graph, such as the example in Figure 1.

One of the cornerstones of the Semantic Web is the Resource Description Framework (RDF)[9], which is used to describe the graphs in the above examples. On top of RDF, RDF-Schema[10] provides the facilities needed to define new resources, for example to represent the classes and properties of a controlled vocabulary.

The Web Ontology Language OWL is more expressive and more formal than

[3] http://www.wikipedia.org/
[4] http://www.geonames.org/
[5] http://wordnet.princeton.edu/
[6] https://www.cia.gov/library/publications/the-world-factbook/
[7] http://dblp.uni-trier.de/
[8] http:://www.getty.edu/research/conducting_research/vocabularies/tgn/
[9] http://www.w3.org/TR/rdf-primer/
[10] http://www.w3.org/TR/rdf-schema/

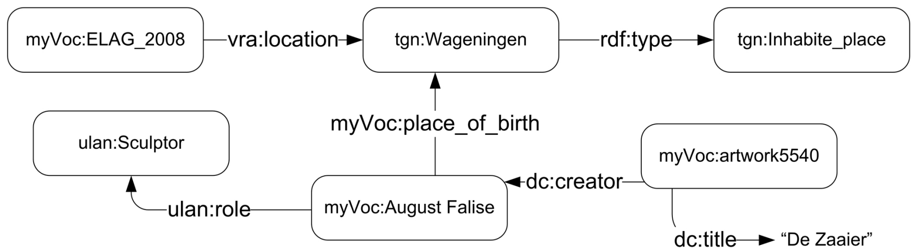

Figure 1: Example of an RDF graph with resources from ULAN, TGN, Dublin Core (dc), VRA and an example vocabulary called myVoc.

RDF and RDF-S (McGuinness and van Harmelen, 2004). It allows a user to model things such as inverse relations (e.g. `creates` versus `is created by`) and disjoint concepts (e.g. one cannot be both `male` and `female`).

As World Wide Web Consortium (W3C) recommendations, RDF, RDF-S and OWL provide a common format that can be used to exchange information across a diverse set of applications. For a more detailed discussion of RDF, RDF-S and OWL we refer the reader to the respective websites.

3 Interoperability

Publishing collections on the web invites the creation of links between archives. It creates the opportunity to provide access across different collections and to create links to other web resources. This, however, requires interoperability of the KOSs. In general, two types of interoperability are discerned, both of which are necessary for unified access to more than one archive:

Syntactic interoperability is about the formats of KOSs: some libraries encode their KOS in XML, while others use plain text or databases. More fundamental differences might also occur: one library might use a hierarchy of terms, while another uses a hierarchy of concepts where each concept is described by one of more synonymous terms.

Semantic interoperability deals with the meaning of concepts. When two KOSs are about similar topics, they will most likely contain concepts that have the same or similar 'meaning'. The concepts "Zeus" in one vocabulary and "Jupiter" in another vocabulary can be said to have the same meaning. The concepts "Christianity" and "Religion", for example, have a related meaning. If this type of links can be established, two KOSs are said to be semantically interoperable.

3.1 Simple Knowledge Organisation System

The Simple Knowledge Organization System (SKOS) is a common data model for knowledge organization systems (Miles and Bechhofer, 2008). SKOS provides a standard way to represent KOSs using the Resource Description Framework (RDF). Encoding this information in RDF allows it to be passed between computer applications in an interoperable way.

Isaac *et al.* (2008) explain the basic model: "SKOS has chosen a concept-based approach for the representation of controlled vocabularies. As opposed to a term-based approach, where terms from natural language are the first-order elements of a KOS, SKOS describes abstract concepts that may have a different materialization in language (lexicalizations)." SKOS concepts are identified by URIs, enabling anyone to refer to them unambiguously, and making them a part of the World Wide Web (Miles and Bechhofer, 2008). Isaac *et al.* (2008) summarize the features that SKOS provides to characterize the concepts:

Labeling properties The properties skos:prefLabel and skos:altLabel link a concept to the terms that represent it in language. The prefLabel value shall be a non-ambiguous term that uniquely identifies the concept, and can be used as a descriptor in an indexing system. The term altLabel is used to introduce alternative entries - synonyms, abbreviations etc. SKOS allows concepts to be linked to prefLabels and altLabels in different languages. SKOS concepts can thus be used seamlessly in multilingual environments.

Semantic properties They are used to represent the structural relationships between concepts, which are usually at the core of controlled vocabularies like thesauri. The construct skos:broader denotes the generalization link (BT in standard thesauri), while skos:narrower denotes its reciprocal link (NT), and skos:related the associative relationship (RT).

Documentation properties Often, informal documentation plays an important role in a KOS. SKOS introduces explanatory notes - skos:scopeNote, skos:definition, skos:example - and management notes - skos:changeNote, skos:historyNote etc.

Concept Scheme properties A KOS as a whole also has to be represented and described. SKOS coins a skos:ConceptScheme construct for this. It also introduces specific properties to represent the links between different KOSs and the concepts they contain. The term skos:inScheme asserts that a given concept is part of a given concept scheme, while skos:hasTopConcept states that a KOS contains a concept as the root of (one of) its constituent hierarchical tree(s), i.e., a concept without a broader concept.

In the same paper, Isaac *et al.* give an example of a SKOS representation: a subject 11F coming from the Iconclass concept scheme[11], "the Virgin Mary",

[11] `http://www.iconclass.nl/`

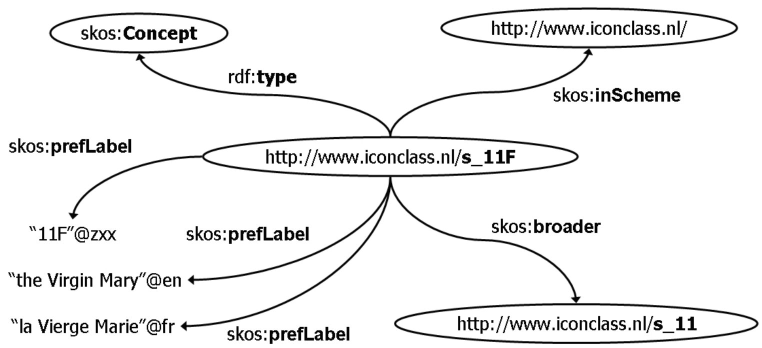

Figure 2: A SKOS graph partly representing the Iconclass subject 11F. Quoted strings are plain literals. "@" specifies the language of a literal: "en" is the tag for "English", "fr" for "French" and "zxx" stands for any "artificial language".

identified by the (as yet fictive) resource $http://www.iconclass.nl/s_1 1F$, could be partly represented by the graph in Figure 2.

The SKOS specifications are currently published as W3C Working Drafts, which means they are work in progress but on the way to become a W3C recommendation. In this respect, SKOS is interesting as a model for interoperability on the Web (Malaisé *et al.*, 2007c).

3.2 Conversion to Semantic Web formats

Semantic interoperability between the metadata of different collections can be achieved by converting or mapping the different data to a common generic (and preferably standard) representation language and model (Malaisé *et al.*, 2007c).

Conversion of metadata vocabularies to Semantic Web languages has been studied by Assem, van *et al.* (2004). They propose a method to translate a KOS to RDF/OWL while preserving the KOS' original semantics. The method consists of four steps: (1) preparation; (2) syntactic conversion; (3) semantic conversion; and (4) standardization. In the first step, an analysis is made of the KOS and its format. This is used in step two to convert to very basic RDF. For example, the data elements in the KOS are converted into classes (concepts) and properties (relations). In step three, the properties of the data model are made explicit. For example, a relatedTerm property would be defined as a an OWL `symmetric property`, explicating that if A is related to B, than B is related to A. In step four, links are created between the KOS and SKOS.

In a further paper, Assem, van *et al.* (2006) take the standardization one step further. They describe how a KOS can be converted to SKOS, thus adding to its interoperability.

3.3 Ontology Alignment

Having the metadata vocabularies of different collections in a standardized format solves the syntactic interoperability problem. In order to also acquire semantic interoperability, semantic links are required between the elements of the different vocabularies; concepts of one KOS need to be linked to their semantic counterparts in the other KOS. In the Semantic Web field, this task is referred to as alignment, matching or mapping of ontologies.

Alignments between two vocabularies can be used for various tasks, such as merging two vocabularies into one, reformulating a user query from one vocabulary to another, translation of metadata, or browsing from one collection to the other.

3.3.1 Semi-automatic alignment

Tools and algorithms have emerged that automate the task of matching two ontologies (see e.g. Shvaiko and Euzenat (2005) or Kalfoglou and Schorlemmer (2003) for an overview). Fully automatic alignment gives satisfactory results only if the two thesauri are alike and well structured. Hollink *et al.* (2008b) have shown that The Art and Architecture thesaurus can be fully automatically matched to the SVCN (Stichting Volkenkundige Collectie Nederland) thesaurus with a precision of around 90% and recall of around 80%. Semi-automatic alignment has been shown to be feasible in many domains (e.g. Cruz *et al.* (2004)). Ehrig and Euzenat (2005) consider the semi-automatic matching process by measuring the quality of an alignment by the effort it will take an expert to correct it.

The Ontology Alignment Evaluation Initiative (OAEI)[12] organizes a yearly evaluation event in which several alignment tools are compared. Although they are focused on evaluating tools, they also target the semi-automatic alignment scenario that includes a domain expert. In the 'Anatomy task' of OAEI 2007, for example, participating groups were asked to generate a high-recall alignment so that incorrect correspondences can be filtered out manually.

3.3.2 Representation of alignments

Although much of the current research is focused on finding equivalent concepts in two ontologies, there is a trend towards also taking into account concepts that are sub-concepts of each other, or have another relationship with each other. In SKOS, five properties exist to represent the range of alignment properties (Miles and Bechhofer, 2008):

- skos:mappingRelation
- skos:exactMatch
- skos:broadMatch

[12] `http://oaei.ontologymatching.org/`

- skos:narrowMatch
- skos:relatedMatch

The property skos:mappingRelation is a super-property of (i.e. more general than) the other four properties. The skos:exactMatch property is used when two concepts in different concept schemes are so similar that they can be used interchangeably. The properties skos:broadMatch and skos:narrowMatch denote a hierarchical related between two concepts in different KOSs; skos:relatedMatch is used for associative links between two concepts in different KOSs. Note that the last three mirror the SKOS semantic relations that are meant for use within a KOS: skos:broader, skos:narrower and skos:related[13].

4 Access

Semantic relations and standardized formats provide opportunities for retrieval within a collection as well as across multiple collections. Using semantic relations, we can now match documents to queries based on semantic similarity, even if there is no textual match between the query and the index. A query for 'eating', for example, could result in documents annotated with 'banquet', since in WordNet 'banquet' (or 'feast') is related to 'feasting', which is related to 'eating'. Hollink *et al.* (2007) have investigated which types of semantic relations between a query and a document are likely to improve search results. Experiments performed on WordNet have shown that next to the hyponym/hypernym relation (similar to the broaderTerm/narrowerTerm relations), the meronym/holonym (part-of) relation is particularly useful for search.

In a recent study, Hollink *et al.* (2008a) have performed a similar experiment on the thesaurus of the Netherlands Institute for Sound and Vision. The findings confirmed that the narrowerTerm relation is beneficial to search results; a query for flower is also satisfied with a document about tulips. In addition, it was shown that the use of relatedTerm works well in situations were a large number of results is more important than the relevance of each individual returned document. Also sequences of different types of relations was found to be beneficial to retrieval.

Tools have emerged that demonstrate this type of semantic search. The Multimedian E-Culture project[14] explores how Semantic Web technologies can be deployed to provide better indexing and search support within large virtual collections of cultural-heritage resources (Ossenbruggen, van *et al.*, 2007). The E-Culture demonstrator, the winner of the Semantic Web Challenge 2006 (Schreiber *et al.*, 2006), hosts four thesauri, namely the three Getty vocabularies[15] (the Art and Architecture Thesaurus (AAT), the Union List of Artists

[13] There is still discussion within the Semantic Web Deployment Working Group as to whether the same semantic relations could be used both to denote links within and to denote links between KOSs. See `http://www.w3.org/TR/skos-reference/`

[14] `http://e-culture.multimedian.nl/`

[15] `http://www.getty.edu/research/conducting_research/vocabularies/`

Names (ULAN) and the Thesaurus of Geographical Names (TGN)), as well as the lexical resource WordNet. The Getty thesauri were converted from their original XML format into a Semantic Web representation using the conversion methods of Assem, van *et al.* (2004) as described in section 3.2.

In the previous section we have argued that for semantic interoperability, links are required between the elements of the different vocabularies. The Getty institute maintains such links between its vocabularies: places in ULAN (such as the city in which a painter is born) are linked to places in TGN. The E-Culture project has added another type of links: artists in ULAN are linked to art styles in AAT (Boer, de *et al.*, 2006). The artist 'Picasso' is, for example, linked to the style 'cubism'.

In the E-Culture project, semantic relations are used for two purposes. First, they are used to expand search results. Ossenbruggen, van *et al.* (2007) give the following example: a query for 'Art Nouveau' retrieves not only images annotated as Art Nouveau, but also images that are 'created by an artist with a matching style', making use of the links between artists and art styles. Although the results may be imperfect, it is a good strategy if there are no (or only few) images directly annotated with Art Nouveau.

Second, semantic relations are used to present search result in a meaningful way. The list of returned images is clustered based on the path from the query to the result. A query for 'Paris' would return a cluster of works of art depicting Paris, a cluster of works created in Paris, a cluster of works that are made by artists born in Paris, etc.

5 Enrichment of Metadata and Vocabularies

In the CHOICE project[16], semantic relations in combination with natural-language processing (NLP) techniques are employed to support documentalists in their indexing tasks. Gazendam *et al.* (2006) suggest a candidate set of thesaurus concepts to the documentalist who is indexing a television program. The candidate concepts are derived from textual resources associated to the program, such as TV guides. The suggested concepts are ranked not only based on how often they occur in the texts, but also based on the semantic relations they have with other concepts in the suggestion list. The underlying assumption is that terms which semantically relate to a lot of other terms found in the text are more representative of the core topics of the TV program than terms without any semantic relations to other terms.

Malaisé *et al.* (2007b) use the same principle for disambiguation of textual metadata. In their paper they give the following example: in the text "Snacks do not contain a lot of minerals", the term minerals can have three meanings (i.e is a label of three thesaurus concepts): food, fertilizer and ore. It is not clear which is the intended meaning in this case. Due to the fact that the thesaurus contains a direct semantic relation between the terms 'food' and 'snacks', food

[16] `http://ems01.mpi.nl/CHOICE/`

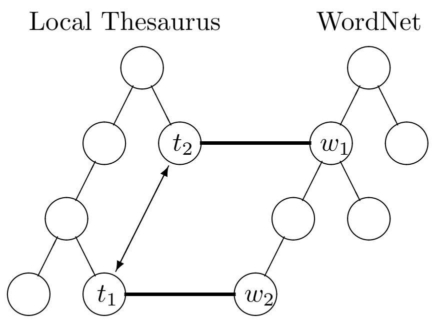

Figure 3: Links between a local thesaurus and WordNet to infer relations within the local thesaurus.

is a more likely interpretation than the other two meanings. This means that their algorithm here interprets 'minerals' as referring to 'food'.

Applications that use semantic relations between concepts, such as the E-Culture search demonstrator or the CHOICE system to support documentalists, rely on a rich thesaurus structure. However, local thesauri are often limited in breadth and depth. In Hollink *et al.* (2008a), a local thesaurus is enriched with additional structure in order to improve retrieval of television programs indexed with this thesaurus. To that end, the thesaurus was linked to an external, semantically richer resource, WordNet. The anchoring to WordNet was used to infer new relations within the local thesaurus; pairs of thesaurus terms were related that were not previously related. Figure 3 illustrates how a relation between two terms in the local thesaurus, t_1 and t_2, is inferred from their mapping to WordNet concepts w_1 and w_2. If t_1 is mapped to w_1 and t_2 is mapped to w_2, and w_1 and w_2 are closely related, a relation between t_1 and t_2 is inferred.

Two WordNet concepts w_1 and w_2 are considered to be 'closely related' if they are connected though either a direct (i.e. one-step) relation without any intermediate concepts or an indirect (i.e. two-step) relation with one intermediate concept. The latter situation is shown in Figure 3. From all WordNet relations, only meronym and hyponym relations are used, which roughly translate to part-of and subclass relations, since a previous study demonstrated that other types of WordNet relations do not improve retrieval results (Hollink *et al.*, 2007) (see also Section 4).

6 Conclusion

The Semantic Web is an active research area. Techniques are maturing and are finding their way into applications. The Linking Open Data project shows that there is a wealth of data in Semantic Web formats available on the Web. Conversion methods to transform data and vocabularies to Semantic Web formats have been developed and are being employed successfully. These developments increase interoperability and bring us closer to unified access across collections.

Alignment of vocabularies will be the second step towards interoperability. Although vocabularies are being aligned in the digital libraries domain (Isaac *et al.*, 2007, Malaisé *et al.*, 2007a, Wang *et al.*, 2007), this remains a difficult task. Fully automatic alignment is still an open and active research area. As Isaac *et al.* (2008) put it: "no single technique is universally applicable, or will return satisfactory results. In practice, different techniques have to be carefully selected and combined, depending on the characteristics of the case at hand, such as the richness of the semantic structures of vocabularies, their lexical coverage and the existence of collections simultaneously described by several vocabularies." The annual evaluation events organized by the OAEI monitor the development of alignment tools and techniques. A considerable advancement has been noticed over the years (Euzenat *et al.*, 2007), and further improvement is expected and necessary.

We have shown how vocabulary alignment can be used to enrich metadata and vocabularies, which in turn can lead to improved search results. Although results are promising, this is still in an experimental phase.

More and more Semantic Web research is aimed directly at cultural heritage data. The most obvious example is the acceptation of SKOS as a standard, but also the fact that the annual workshop for the evaluation of alignments (OAEI) now features a special 'track' devoted to the library domain[17] is a good sign. Projects such as MuseumFinland[18], E-Culture and the CATCH program[19] illustrate the benefits of using Semantic Web technologies in the digital libraries domain, such as cross-collection access, search using semantic relations and innovative presentation of search results.

Acknowledgements

The authors are supported by the NWO CATCH projects STITCH, MUNCH and CHOICE, the MultimediaN project E-Culture and the eContentPlus project TELplus.

References

Assem, van, M., Menken, M. R., Schreiber, A. Th., Wielemaker, J., and Wielinga, B. J. (2004). A method for converting thesauri to RDF/OWL. In *Proceedings of the Third International Semantic Web Conference*, pages 17–31.

Assem, van, M., Malaisé, V., Miles, A., and Schreiber, A. Th. (2006). A method to convert thesauri to skos. In *In Proceedings of the Third European Semantic Web Conference*, Budvar, Montenegro.

[17] `http://oaei.ontologymatching.org/2007/library/`
[18] `http://www.seco.tkk.fi/`
[19] `www.nwo.nl/catch/`

Berners-Lee, T., Hendler, J., and Lassila, O. (2001). The semantic web. *Scientific American*, **5**(284), 34–43.

Boer, de, V., Someren, van, M., and Wielinga, B. J. (2006). Extracting instances of relations from web documents using redundancy. In *Proceedings of the Third European Semantic Web Conference*.

Cruz, I. F., Sunna, W., and Chaudhry, A. (2004). Semi-automatic ontology alignment for geospatial data integration. In *Proceedings of the Third International Conference on Geographic Information Science*, Baltimore, MD, USA.

Ehrig, M. and Euzenat, J. (2005). Relaxed precision and recall for ontology matching. In B. Ashpole, M. Ehrig, J. Euzenat, and H. Stuckenschmidt, editors, *Integrating Ontologies*.

Euzenat, J., Isaac, A., Meilicke, C., Shvaiko, P., Stuckenschmidt, H., Šváb, O., Svátek, V., van Hage, W. R., , and Yatskevich, M. (2007). First results of the ontology alignment evaluation initiative 2007. In B. Ashpole, M. Ehrig, J. Euzenat, and H. Stuckenschmidt, editors, *Ontology Matching*, CEUR Workshop Proc.

Gazendam, L., Malaisé, V., S. A. Th., and Brugman, H. (2006). Deriving semantic annotations of an audiovisual program from contextual texts. In *Proceedings of First International workshop on Semantic Web Annotations for Multimedia (SWAMM)*, Edinburgh, Scotland.

Hollink, L., Schreiber, G., and Wielinga, B. (2007). Patterns of semantic relations to improve image content search. *Journal of Web Semantics*, **5**, 195–203.

Hollink, L., Malaisé, V., and Schreiber, A. B. (2008a). Enriching a thesaurus to improve retrieval of audiovisual material. Submitted for publication.

Hollink, L., van Assem, M., Wang, S., Isaac, A., and Schreiber, G. (2008b). Two variations on ontology alignment evaluation: Methodological issues. In *Proceedings of the European Semantic Web Conference*, Tenerife, Spain.

Isaac, A., van der Meij, L., Schlobach, S., and Wang, S. (2007). An empirical study of instance-based ontology matching. In *Proceedings of the the 6th International Semantic Web Conference (ISWC 2007)*, Busan, Korea.

Isaac, A., Matthezing, H., Schlobach, S., and Zinn, C. (2008). Integrated access to cultural heritage resources through representation and alignment of controlled vocabularies. *Library Review, Special issue on digital libraries and the Semantic Web*. To appear.

Kalfoglou, Y. and Schorlemmer, M. (2003). Ontology mapping: the state of the art. *The Knowledge Engineering Review*, **18**(1), 1–31.

Malaisé, V., Isaac, A., Gazendam, L., and Brugmann, H. (2007a). Anchoring dutch cultural heritage thesauri to wordnet: two case studies. In *ACL 2007 Workshop on Language Technology for Cultural Heritage Data*.

Malaisé, V., Gazendam, L., and Brugmann, H. (2007b). Disambiguating automatic semantic annotation based on a thesaurus structure. In *in proceedings of the 14e conference sur le Traitement Automatique des Langues Naturelles (TALN-2007)*.

Malaisé, V., Gazendam, L., Brugman, H., Schreiber, G., Veenstra, M., de Jong, A., and Oomen, J. (2007c). Skos: a model for metadata representation and interoperability dutch cultural heritage institution thesaurus conversion use case. In *Proceedings of the DELOS and Multimatch Workshop*, Tirrenia, Italy.

McGuinness, D. L. and van Harmelen, F. (Editors). (2004). Owl web ontology language overview. Technical report, W3C Recommendation. Latest version available at http://www.w3.org/TR/owl-features/.

Miles, A. and Bechhofer, S. (Editors). (2008). SKOS simple knowledge organization system reference. W3C working draft. Electronic document. Accessed April 2008. Available from: http://www.w3.org/TR/skos-reference/.

Ossenbruggen, van, J., Amin, A., Hardman, L., Hildebrand, M., van Assem, M., Omelayenko, B., Schreiber, G., Tordai, A., de Boer, V., Wielinga, B., Wielemaker, J., de Niet, M., Taekema, J., van Orsouw, M.-F., , and Teesing, A. (2007). Searching and annotating virtual heritage collections with semantic-web techniques. In J. Trant and D. Bearman, editors, *Proceedings of Museums and the Web*, Toronto, Canada.

Schreiber, A. Th.., Amin, A., van Assem, M., de Boer, V., Hardman, L., Hildebrand, M., Hollink, L., Huang, Z., van Kersen, J., de Niet, M., Omelayenko, B., van Ossenbruggen, J., Siebes, R., Taekema, J., Wielemaker, J., and Wielinga, B. (2006). Multimedian e-culture demonstrator. In *the Semantic Web Challenge at the Fifth International Semantic Web Conference*, Athens, GA, USA.

Shvaiko, P. and Euzenat, J. (2005). A survey of schema-based matching approaches. *Journal on Data Semantics*, **3730**, 146–171.

Wang, S., Isaac, A., van der Meij, L., and Schlobach, S. (2007). Multi-concept alignment and evaluation. In *Proceedings of the Second International Workshop on Ontology Matching at ISWC*, Busan, Korea.