Anchoring Dutch Cultural Heritage Thesauri to WordNet: two case studies

Véronique Malaisé and **Antoine Isaac**
Vrije Universiteit
Amsterdam
The Netherlands
{vmalaise, aisaac}@few.vu.nl

Luit Gazendam
Telematica Instituut
Enschede
The Netherlands
Luit.Gazendam@telin.nl

Hennie Brugman
Max Planck Institute
for Psycholinguistics, Nijmegen
The Netherlands
Hennie.Brugman@mpi.nl

Abstract

In this paper, we argue on the interest of anchoring Dutch Cultural Heritage controlled vocabularies to WordNet, and demonstrate a reusable methodology for achieving this anchoring. We test it on two controlled vocabularies, namely the GTAA thesaurus, used at the Netherlands Institute for Sound and Vision (the Dutch radio and television archives), and the GTT thesaurus, used to index books of the Dutch National Library. We evaluate the two anchorings having in mind a concrete use case, namely generic alignment scenarios where concepts from one thesaurus must be aligned to concepts from the other.

1 Introduction

Cultural Heritage Institutions are the keepers of large collections of data. To optimize the core tasks of indexing and searching through these collections, controlled vocabularies like thesauri are often used. These vocabularies are structured concept networks[1] and help indexers to select proper subjects for description, and users to formulate queries or to browse collections using the concepts that appear in the metadata.

The Netherlands Institute for Sound and Vision[2], for example, uses the GTAA thesaurus for indexing public radio and TV programs – GTAA is a Dutch abbreviation for "Common Thesaurus [for] Audiovisual Archives". Its hierarchy of subjects contains about 3800 Preferred Terms and 2000 Non Preferred terms. A second example is the GTT thesaurus, which contains 35000 concepts, gathering 50000 preferred and non-preferred Dutch terms. This thesaurus is used to index and retrieve books from the Dutch National Library[3] – GTT is a Dutch abbreviation for "GOO keyword thesaurus", GOO referring to the Joint Subject Indexing system used by many Dutch libraries.

Besides this classic scenario, thesauri can also allow for (semi-)automatic optimization of search processes, like query expansion exploiting their hierarchical structure. But the available structure might not be rich and regular enough for such purposes. In fact, it has been shown that a mapping to a richer and sounder terminology, like the English WordNet (Fellbaum, 1998), would enable more sophisticated query expansion or other inferencing possibilities (Voorhees, 1994; Hollink, 2006). This will become especially true now that WordNet exists in the form of an RDF ontology (van Assem et al., 2006).

Mapping Cultural Heritage controlled vocabular-

[1]The typical semantic relationships found between elements from thesauri are **Broader Term** linking a specialized concept to a more general one, **Narrower Term**, its inverse relationship, and **Related Term**, which denotes a general associative link. Thesauri also contain lexical information, where the *preferred terms* used for description are given *synonyms* or *non-preferred terms* (**Use** and **Used for** links), as well as general **scope notes** giving indexers instructions regarding the use of a term.

[2]http://www.beeldengeluid.nl

[3]http://www.kb.nl

Proceedings of the Workshop on Language Technology for Cultural Heritage Data (LaTeCH 2007), pages 57–64,
Prague, 28 June 2007.

ies in Dutch to WordNet can also be beneficial for sharing information across institutions, which is difficult when the metadata attached to the different documents come from different thesauri. This issue can be solved by building equivalence links between the elements from these different vocabularies, as in (van Gendt et al., 2006). This *vocabulary alignment* problem is comparable to the *ontology matching* one, and techniques similar to the ones developed by the Semantic Web research community can be applied here. As found e.g. in (Euzenat, 2004), the existing methods are quite diverse, and proposed strategies often mix several individual techniques:

- lexical techniques, trying to compare the labels found in vocabularies;
- structural techniques, assessing similarities between concepts from the structure of vocabularies (e.g. hierarchical links);
- instance-based techniques, looking at the objects that are actually populating the ontologies to infer from their similarities correspondences between the concepts they instantiate.
- techniques making use of some background knowledge source, by trying to derive from the information found there relations between the elements from the original vocabularies.

Here, we are interested in the last kind of techniques. In these approaches, concepts from the vocabularies to be aligned are first attached – "anchored" – to the concepts from a third vocabulary (Aleksovski, 2006). Then, these anchors in the background vocabulary are compared together. When a relation is found between them[4], a similar relation can be inferred between the elements from the vocabularies to be aligned. This is especially interesting when the lexical overlap between the vocabularies is low or when the vocabularies are quite poorly structured: it is expected then that the background knowledge will alleviate these shortcomings. The choice of this knowledge is therefore crucial, and WordNet, which has a rich structure and a broad coverage, has been exploited in many existing alignment methods (Giunchiglia et al., 2005; Castano et al., 2005).

For these reasons – even if this paper will only focus on the alignment scenario – we wanted to experiment the anchoring of two aforementioned Dutch thesauri to WordNet. Unlike literature about linking English thesauri to WordNet, we propose in this paper an anchoring method for vocabularies in other languages, and experiment it on these two thesauri, testing its usefulness in terms of possibilities for vocabulary alignment. The remainder of the paper is organized as follows: in section 2, we present the general anchoring methodology. The anchoring experiment is described in section 3: first the GTAA case (section 3.1) and then the GTT one (section 3.2), as a reusability test. We evaluate the two anchoring processes in section 3.3 and conclude on general reflexions about this method. Then, we show examples of such anchorings in the context of a possible alignment between GTAA and GTT in section 4. We conclude on perspectives to this research in section 5.

2 Anchoring methodology

The anchoring experiment presented in this paper is based on a comparison of lexical descriptions of the thesaurus terms with the ones of WordNet synsets, the *glosses*: WordNet is a lexical database of English, which entries are grouped "into sets of cognitive synonyms (synsets), each expressing a distinct concept"[5]. In contrast to many anchoring methods, like the one in (Khan and Hovy, 1997), we do not compare the terms from our thesauri to the labels of synsets, but measure the lexical overlap of their descriptions. The same approach has already been followed, for example, by (Knight and Luk, 1994).

As the thesauri we focus on in this paper are in Dutch, we first need to map their terms to English descriptions, and possibly translations, to make a comparison with the English glosses. Given the fact that these thesauri cover a broad range of topics, we hypothesize that using a general language bilingual dic-

[4] The reader can turn to (Budanitsky and Hirst, 2006) for an overview of the different methods that have been proposed in this field.

[5] `http://wordnet.princeton.edu/`

tionary will lead to a good coverage of their content. Additionally, it might give on top of the definitions – *i.e.* the natural language descriptions of a term's meaning – useful information such as term translations and Part Of Speech (POS) tags – their grammatical category: noun, verb, etc. For each thesaurus term which has been associated to an English definition, the rest of the anchoring procedure consists in checking the overlap between the lexical content of the definitions and the one of the different WordNet glosses, considered as bags of words. The hypothesis is that the closest gloss should give us a pointer to a synset semantically equivalent to the intended meaning of a thesaurus term.

3 Anchoring feasibility experiments and evaluations

3.1 Anchoring GTAA concepts

First step: Finding English definitions for GTAA terms The first step in mapping Dutch terms from the GTAA to WordNet was to select an online dictionary that would cover a significant part of the thesaurus entries and that would allow automatic queries for these terms. We have tested the bilingual dictionary LookWAYup[6], which returned a 2222 results – definitions and translations – on our query set.

This query set consisted in the list of GTAA Preferred terms (3800), Non preferred terms (2000) and their singular forms[7] (3200). These singular forms were computed in the context of a MultimediaN project[8], on the basis of linguistic derivational rules and a manual correction.

Given the fact that most of the thesaurus terms are in plural form, but not all of them[9], and knowing that the dictionary entries are only standard lemma forms (most of the time in singular), we first assumed that queries on the dictionary with a plural form would not generate a result, and simply added the singular forms to the singular ones in the query set. It turned out that the dictionary gave result for some plural forms, creating noise: some plural forms corresponded to lemmas of verbs, and a spelling correction facility provided definitions for some plural forms.

Removing doubles We cleaned manually the first set of errors, and automatically the last one, based on POS tag information. In the future, we will avoid introducing duplicate lemmas in our the query set.

After cleaning, 1748 terms had one or more translation in English together with their associated POS tag(s) and definition(s)[10]. This low number, compared with the original set of 5800 distinct thesaurus terms can be explained by the fact that our vocabulary contains numerous multi-words terms and also compound entries, both of which are rarely dictionary entries. We discuss possible solutions to this shortcoming in section 3.3.

POS tag-based cleaning We did then a rough manual evaluation of these candidate definitions. The evaluation was conducted by three people and took about one day each. It turned out that some of the definitions were irrelevant for our task: the Dutch *Bij* was associated with the English *Bee* and *Honey bee*, but also with the preposition *by*. We used again the information given by the POS tag to remove these irrelevant definitions: we kept only definitions of Nouns and (relevant) Verbs. After this last cleaning, some terms still had more then one definition.

Cleaning based on thesaurus relationships We used the hierarchical relationship in the thesaurus to check the intended meaning of these terms: for example, *Universiteit* (*University*) had a Broader Term relationship with *Wetenschappelijk onderwijs* (*Scientific education*), so its meaning is restricted to the "Educational aspect", and it should not be used to describe TV programs about University buildings for instance. We used this information to restrict the

[6]Built by RES Inc., Canada, online at the URL: `http://lookwayup.com/free/`.

[7]Following the recommendations of the ISO standard, most of GTAA terms are in plural form.

[8]MultimediaN Project 5 – Semantic Multimedia Access, `http://monetdb.cwi.nl/projects/trecvid/MN5/index.php/Main_Page`, transformation done by Gijs Geleijnse, from the Philips Research group.

[9]For example, the term corresponding to *Baptism* is in singular form.

[10]1299 terms have more than one definition.

Step	Result
Gathering query set	3800 + 2000 + 3200 terms
Querying dictionary	2222 defined terms
Removing doubles	1748 different defined terms
POS tag-based cleaning	1655 def. terms, 7530 definitions
Thesaurus-based cleaning	
Anchoring to WordNet	1060 anchored concepts

Table 1: GTAA term anchoring experiment

number of valid candidate definitions associated with every GTAA term. But in some cases the distinction was hard to make between the different definitions, or no clue was provided by the thesaurus to dismabiguate the senses of the term: sometimes it did not have any relationship to other concepts nor explanatory text (Scope Note).

Conclusion of the first step As a final result, as summarized in table 1, 1655 GTAA terms had one or more English equivalent and their related candidate definitions (7530). We decided to postpone a more in-depth validation to the evaluation of anchoring results with WordNet: we kept all candidate definitions and translations that were not obviously incorrect, and checked the WordNet anchoring result to see if some further refinement had to be done. The idea was that the anchoring process would only work for parts of the definitions, so we wanted to keep as many data as possible.

Second step: Anchoring to WordNet synsets We stemmed the candidate definitions of GTAA terms and the glosses from WordNet with the Porter stemmer to augment mapping possibilities. Stemming is the operation of reducing words to a root, for example by removing the "s' character at the end of an (English) word in plural form. This process can reduce different unrelated words to a same root, and hence should be handled with care, but it requires less resources then a full fledged lemmatizing and helps comparing a larger number of words then on the basis of the graphical forms only. As announced, in order to map synset to GTAA terms, we compared their lexical descriptions: we compared the different sets of stems in a simple bag-of-words approach. We actually found out that the definitions of the online dictionary were exact matches with WordNet glosses, thus all defined terms could be straightforwardly anchored to one or more synsets. In the end, 1060 concepts from GTAA are successfully anchored to a synset, which represents 28% of the total number of concepts.

Evaluation of the results We evaluated the number of semantically relevant anchorings for a random representative part of the the 1655 GTAA terms that had one or more WordNet anchor: we evaluated 1789 mappings out of 7530. On these 1789 mappings, 85 were not equivalence links: 5 out of these 85 links were relating Related Terms (like *zeerov* anchored to *corsair*, the first being in GTAA a profession and the second a ship in Wordnet), 17 pointed to Broader Terms, and the others were mapping a term with a correct translation that was correct *per se* but did not correspond to the intended meaning of the term in GTAA. For example, two anchorings were proposed for *Vrouwen*: *married_woman* and *female_person*, the latter one being the only valid for our thesaurus. The first cases (RT and BT relationships between the original term and its anchoring) still provide useful information for aligning vocabularies, but we took only equivalence relationships into account in this experiment.

An additional evaluation that was also performed on a sample set was to check that non-preferred terms that were given a definition were pointing to the same synset as their related preferred terms. It turned to be correct for the evaluated pairs.

On a qualitative perspective, we found different types of mappings:

- some GTAA terms had more then one translation, all of them pointing to the same synset: this was the confirmation that the mapping from the term to the synset was correct;
- some GTAA terms had more then one translation, pointing to different but close synsets: nothing in the thesaurus content could help us distinguish between the different synsets, thus we kept the different possibilities;

- some different GTAA terms pointed to a same synset and, although they were not linked in the thesaurus, they had a semantic relationship. This information can be used to enrich the structure of the GTAA.

We can conclude that the anchoring was quite successful: only 4.7% of the anchorings were incorrect in the test sample. And this was due to cases where multiple senses were linked to a same term, which would not cause a big problem in a semi-automated anchoring process. Moreover, this process can bring an additional value to the thesaurus structure itself, on top of the possible applications mentioned in the introduction.

3.2 Anchoring GTT concepts

Setting We carried out for GTT the same experiment as for GTAA, but did not compute singular forms, although GTT terms are generally in plural form. Also, because GTT had 70% of its concepts already translated to English by human experts, we decided that we would measure the global performance of our method based on this translation gold standard, additionally to manually assess the relevance of the produced anchorings from GTT to WordNet.

Results Out of the 35194 GTT general subjects, only 2458 were given some English definition and translation by the dictionary service we used. For the set of 25775 concepts for which there was already a translation, the figure drops down to 2279, slightly less than 9%.

As said, we tested the validity of these definitions and translations by comparing them to the expert translations. Our assumption was that an English definition for a concept would prove to be correct if its associated term matched one of the expert translations of the concept[11]. We found that 1479 of the 2279 concepts being given both expert and automatic translations had the expert translation confirming one of the automatically found ones, *i.e.* a precision rate of 65% in terms of defined concepts.

When measuring accuracy of the found English definitions for the 2279 defined concepts, we saw that out of a total 3813 English definitions associated to a concept, 2626 – 69% – had an associated term confirmed by the expert translation.

We also tried to assess the quality of the translations associated to the concepts of this set by our method: out of 5747 terms proposed as translations, 1479 matched the expert translation. This precision rate is low (25.7%) but it actually highlights one of the problem of the expert translations found in the thesaurus: the manual translation had a very low lexical coverage, having provided with very few synonyms for the "preferred" translations. The set of 25775 translated GTT concepts only brings 26954 English terms in total…

The evaluation by comparison to the expert translation brings useful information, but it has some drawbacks, especially the limited coverage of the translation work and a correctness assumption bringing a (small) error rate. To complete it, we carried out a manual investigation, inspired by what had been done for the GTAA thesaurus.

For this, we selected the 179 concepts that were translated by our method but had not previously been assigned English labels by experts. For this subset, 441 glosses had been assigned. Of these, 172 were correct, concerning 138 concepts. We therefore obtain a 77% precision rate in terms of anchored concepts. However, if we aim at assessing the quality of the method and its potential to be used in a semi-automatic anchoring process, we have to consider the obtained glosses themselves. And here precision falls to 39%, which is a far less satisfactory figure.

Feasibility of the proposed method in GTT case Some of the previously mentioned anchorings to wrong glosses could have been successfully found by applying the heuristics mentioned in section 3.1. The use of POS tags and the checking of the singular form of terms allowed to manually spot 41 obviously wrong results. The other irrelevant glosses were mainly found using the thesaurus information:

[11] A manual checking of this assumption on the first 150 concepts matching the criterion demonstrated an error rate of 4%: 4% of the concepts had no correct definition in their associated glosses while there was a match between the expert translation and one of the terms linked to the definitions.

Comparison with expert Gold Standard	
Concepts with expert translation	25775
Concepts with a definition	2279
Concepts with def. confirmed by GS	1479
Total definitions given	3813
Definitions confirmed by GS	2626
Total translations given	5747
Translations confirmed by GS	1479
Manual evaluation	
Concepts	179
Concepts with correct definition	138
Total definitions given	441
Correct definitions	172
Global results	
Total GTT concepts	35194
Concepts with a definition	2458
Concepts with correct definition	1617
Total definitions given	4254
Correct definitions	2798

Table 2: GTT term anchoring evaluation

the Broader Term information helped to discriminate 68 cases, compared with 6 for Related Term, 6 for synonyms and 15 for scope notes.

It is however still uncertain whether these different kinds of information can be used in a more automatised setting. If we could count on translation of broader and related terms to be done by the process we have applied, taking into account scope notes would require more effort. And the poor structure of thesauri such as GTT – some 20000 concepts have no parents at all – makes such validations by semantic links difficult. It is also important to notice that in 14 cases, it was necessary to check the books which have been indexed by a concept to find out its precise meaning.

This could yet be compensated by an interesting result we have observed: the anchoring method gave us material for inferring new semantic links, as in the GTAA case. Amongst the translated GTT concepts, 689 concepts are sharing at least one synset and are not connected by a thesaurus link. We found interesting matches, such as *gratie* (pardon) and *absolutie* (absolution) or between *honger* (hunger) and *dorst* (thirst). This potential for enriching thesauri could actually be used to spark some positive feedback loop for the anchoring process itself: a richer vocabulary enables for example to use with greater profit the selection strategies based on thesaurus structure.

An important problem for the implementation of such strategies remains to deal with disambiguation (when several English definitions are found, which one shall be selected?) in a context of fine-grained vocabularies. Both GTT and WordNet have a high level of precision, but they are focused on different matters. Especially, for a same GTT term the dictionary pointed at several meanings that were very close, but considered as different synsets in WordNet. A typical example is the distinction made between the gloss attached to moderation and temperance, "the trait of avoiding excesses", and the one attached to moderateness and moderation, "quality of being moderate and avoiding extremes". Looking at the books indexed by the concepts which these glosses were attached to, it was not clear whether the indexers systematically considered such a distinction.

Finally, we made rough estimattions of recall – the number of concepts that were correctly anchored compared to the number of concepts anchored in the ideal case. If we compare the 1479 correctly defined concepts to the 25775 concepts being given an expert translation, we find a very disappointing recall rate of 5.7%. This very low performance is in fact largely due to three recurrent situations in which the online dictionary could not give any translation:

- terms containing some special Dutch characters – especially the so-called Dutch *ij*, where *i* and *j* make a single character – and which occurs for more than 2000 concepts;
- specialized scientific terms, like *kwantumhalleffect*;
- complex notions, rendered in Dutch by compound words (e.g. *gebruikersinterfaces* for *user interfaces*), multi words (*Algemene kosten* for *general costs*) or a mixture of the two (*Grafische gebruikersinterfaces* for *graphic user interfaces*).

Whereas the encoding problem appears fairly simple, the last ones are more serious – they were indeed also encountered in the GTAA case – and shall be discussed further.

3.3 Conclusion on the anchoring methodology

As just mentioned, a drawback of our anchoring method is the fact that there are very few multi-word entries in dictionaries but they compose a large part of thesauri, and particularly thesauri in Dutch. Previous work about assigning a semantic relationship between a multi-word term and its components (see (Ibekwe, 2005)) could be used in order to give elements of solution to this problem. Using this pre-processing, we could apply our method to the single-word part that corresponds to the generic meaning of the original multi-word term, and try to anchor the single-word corresponding to the semantic root of the thesaurus' multi-word term (*Kosten* for *Algemene kosten* – *Cost* for *General cost* – for instance).

From a more conceptual point of view, however, further effort would be needed to adapt our anchoring method – and the subsequent alignment of one vocabulary with the other – to the cases where a concept from one vocabulary should be anchored to more than one element from WordNet. More complex heuristics come closer to traditional anchoring problems cases – without translation – and could be solved using existing solutions, as proposed by (Giunchiglia et al., 2005; Castano et al., 2005).

The last problem encountered in the anchoring process was the fact that specialized notions, that also appear in general purpose thesauri, have usually no definition in a general language dictionary. Specialized dictionaries should be used as a complementary resource.

These different shortcomings reduced the coverage of the anchoring, but our method has still positive points: the number of obviously wrong anchors was rather low for the found pairs and additional links could be provided for both of the source thesauri. This method also provides a starting point for anchoring complex and large vocabularies to WordNet, which is also a large lexical resource, and both are hard to grasp completely by a human expert.

4 GTAA and GTT alignment using WordNet anchoring: a qualitative evaluation

Once the anchoring is performed, the synsets corresponding to the terms from the different thesauri can be compared, in order to infer from them equivalences between the original concepts, as is done in classical alignment techniques using background knowledge. In this section, we present some examples illustrating the kind of alignment results one can expect from a proper anchoring of our Dutch controlled vocabularies.

First, we can confirm alignments of equal Dutch labels: *gtaa:arbeiders* is aligned to *gtt:arbeiders* since they are both anchored to the synset "someone who works with their hand, someone engaged in manual labor". In some cases, though, a first stemming or lemmatizing process would have been needed to achieve alignment, as in the example of *gtaa:bekeringen* and *gtt:bekering* (*Conversion*, respectively in plural and singular form), or *gtaa:biljart* and *gtt:biljartspel*[12] (*Billiard* and *Billiard game*).

Nevertheless, the more interesting cases are the ones involving concepts with large semantic overlap but a small lexical one, as in the case of *gtaa:plant* (*Plant*) and *gtt:begroeiing* (*Excessive growth of vegetation*) via the WordNet *flora* synset. *Begroeiing* is actually semantically related in the GTT to the concept *Planting*. Here, the translation process compensates for the lack of lexical coverage in the respective vocabularies, which precisely corresponds to one of the traditional features background knowledge-based techniques boast. We can also derive general conceptual similarity relationship based on the overlap between glosses, such as the one between *gtaa:drank* and *gtt:alcohol*, which are not direct matches but for which our method has found some common glosses like "an alcoholic beverage that is distilled rather than fermented".

[12] Notice that substring-based matching could also give these results, but this method is usually very noisy for alignment processes and therefore must be used cautiously.

5 Conclusion and perspectives

Our experiments showed that the partial anchoring of large Dutch controlled vocabularies to WordNet can be done via a bilingual dictionary, even though there is an obvious loss in information: not every thesaurus concept can easily be found in a general language bilingual dictionary, and a preprocessing of multi-word and compound thesaurus entries has to be done. Yet, a significant part of the GTAA thesaurus could be anchored, and with some improvement to the method this could be true for GTT too. Besides multi-word and compound words processing, useful extensions should also take into account specialized dictionaries and have a closer look at methodologies for anchoring a thesaurus term to multiple WordNet synsets with close meanings. We plan to test such strategies in future experiments, and hope to obtain a better coverage of the thesauri.

In this paper, we have sketched a way to use of these anchorings in a vocabulary alignment scenario, and underlined the potential gains on test examples. Even if the number of results given by the current implementation of our method is quite low, the reader should notice that the process can already, as is, suggest new relationships between concepts of the source thesauri. Moreover, proposed strategies in the alignment field often advocate using combined methods: combined contributions can be used to proceed with some cross validation if they overlap, or to provide with larger number of candidate for further (semi-)automatic selection. In such a setting, every contribution of candidate links is welcome. In this respect, what is useful here is the ability of a WordNet-based method to provide with results that could not be obtained with other techniques because of the lack of explicit semantic information and hierarchical structure in the original vocabularies.

Finally, as mentioned in the introduction, there are other motivating use cases that we plan to experiment with. Especially interesting is the way a mapping with WordNet can enhance the existing access to document collections of the Dutch Cultural Heritage Institutes by providing with query refinement services and browsing possibilities.

Acknowledgements

This research was carried out in the context of the CATCH projects CHOICE and STITCH, funded by NWO, the Dutch organization for scientific research.

References

Aleksovski Z. 2006. *Matching Unstructured Vocabularies using a Background Ontology*. 15th International Conference on Knowledge Engineering and Knowledge Management (EKAW 2006).

van Assem M., Gangemi A. and Schreiber G. 2006. *RDF/OWL Representation ofWordNet*. W3C Working Draft, 19 June 2006. http://www.w3.org/TR/wordnet-rdf/

Budanitsky A. and Hirst G. 2006. *EvaluatingWordNet-based Measures of Lexical Semantic Relatedness*, volume 32(1). Computational Linguistics, 13–47.

Castano S., Ferrara A. and Montanelli S. 2005. *Matching Ontologies in Open Networked Systems: Techniques and Applications*, volume 5. Journal on Data Semantics (JoDS).

Euzenat J., coordinator. 2004. *State of the art on ontology alignment*. KnowledgeWeb Deliverable 2.2.3.

Fellbaum C. 1998. *WordNet An Electronic Lexical Database*. MIT Press.

van Gendt M., Isaac A., van der Meij L. and Schlobach S. 2006. *Semantic Web Techniques for Multiple Views on Heterogeneous Collections: a Case Study*. 10th European Conference on Research and Advanced Technology for Digital Libraries (ECDL 2006), 426–437.

Giunchiglia F., Shvaiko P., and Yatskevich M. 2005. *Semantic Schema Matching*. 13th International Conference on Cooperative Information Systems (CoopIS 2005).

Hollink L. 2006. *Semantic annotation for retrieval of visual resources*. PHD Thesis, Vrije Universiteit Amsterdam.

Ibekwe-SanJuan F. 2005. *Clustering semantic relations for constructing and maintaining knowledge organization tools.* volume 62 (2). Journal of Documentation, Emerald Publishing Group, 229–250.

Khan L. R. and Hovy E. 1997. *Improving the Precision of Lexicon-to-Ontology Alignment Algorithm*. AMTA/SIG-IL First Workshop on Interlinguas, San Diego, CA, October 28.

Knight K. and Luk S. 1994. *Building a Large-Scale Knowledge Base for Machine Translation*. In Proceedings of the AAAI-94 Conference.

Voorhees E. 1994. *Query expansion using lexical-semantic relations*. 17th International ACM/SIGIR Conference on Research and Development in Information Retrieval, 61–69.