A Workbench for Thesaurus Mapping in the Cultural Heritage domain

Marjolein van Gendt

June 20, 2006

Contents

Abstract

Access to heterogeneous data sources asks for interoperability between controlled vocabularies. Cultural Heritage collections can be integrated by linking their respective controlled vocabularies. Such linking resembles ontology matching, but is different due to thesauri (the most usual form of controlled vocabulary used in Cultural Heritage) characteristics and application specifics. We propose to investigate these domain specifics and based on this analysis provide a modular, tunable Cultural Heritage thesaurus mapping workbench.

1 Introduction

Knowledge organisation structures, for example controlled vocabularies, are very important for information retrieval and become increasingly important in the light of current collection integration efforts. Thesaurus mapping within the Cultural Heritage (CH) domain facilitates collection integration. Currently such mapping efforts are mainly conducted manually, which consumes lots of effort and time. And one-shot mapping does not eternally solve the integration issue; every thesaurus update asks for an updated alignment. Automation within this field would be favorable.

The CH domain poses a challenge for automatically creating thesauri alignments; the controlled vocabularies often are huge (65.000 terms [8] is no exception), have their own semantics [10] [4], and frequently use glosses instead of simple terms, just to name a few specific properties. One of the research questions we address is which mapping methods are appropriate for thesauri in the CH domain. For each mapping task the possible alignment is influenced by the way a specific thesaurus instantiates those characteristics, e.g. if a thesaurus is almost flat, structure based mapping will be of no use. So the *possible output* depends on the input. The *preferred output*, however depends on the application in which the alignment is used. So a mapping task is case dependent, and should thus have case based configuration.

Case based configuration requires a workbench with an adequate set of implemented mapping methods. Our goal is to develop a workbench for mapping CH thesauri. Our requirements for the system are the following. The system should:

- contain a large set of mapping components relevant for CH applications and thesauri,
- be able to select and configure mapping components to fit input and output based requirements,
- be able to integrate multiple alignments.

This workbench will obtain characteristics of its input (thesauri, instances, reference alignment) and desired output (application), select the appropriate mapping modules, perform the matching and output an alignment between the thesauri, all in a (semi-)automatic fashion.

The main research question is:
How can we construct a workbench which will supply situation dependent alignments between CH thesauri to assist CH collection integrators?
Subquestions are:

1. Which input and application characteristics have which consequences on mapping?
2. Which mapping components and which configurations are appropriate for which combinations of requirements?
3. Which evaluation measures are applicable for thesaurus alignment evaluation and how can evaluation outcome be used to tune mapping components?

This paper sketches our approach for creating the workbench, based on a thorough analysis of relevant factors (thesauri characteristics, use cases, instance data, a reference alignment, evaluation measures) for this complex mapping problem. First, related research efforts are discussed, followed by a chapter describing our research steps. A planning for the coming three years, a conclusion section and a list of definitions of relevant concepts can be found at the end of this paper.

2 Context and related work

Below we describe research efforts which are relevant concerning our aim to develop a thesaurus mapping workbench.

2.1 Manual mapping

The MACS project [2] and the ECHO-project[1], note the importance of thesauri integration, but work via the labor-intensive, time-consuming route of manual mapping. MACS especially has shown the relevance of CH thesaurus alignment. In their interface documents can be retrieved that are annotated using heterogeneous controlled vocabularies, in several different languages. [2] admits that link establishment is very time-consuming.

2.2 Automatic mapping

The Semantic Web, on the other hand, shows numerous attempts at automatic alignment [6], but these focus mainly on full-fledged ontologies. Some efforts in the direction of thesaurus mapping are made [4] [3], but no application-oriented, tunable thesaurus mapping tools can be found.

[1] `http://www.mpi.nl/echo/`

Within previous Semantic Web alignment efforts extensive effort has been put in developing mapping algorithms: lexical, structure based, instance based and oracle based techniques [14] exist. All can be applied for thesaurus mapping. The outcome of STITCH's pilot project (which is described in section 2.4) motivates adaptations. Some mappers, such as S-Match, implement all mapping techniques in a modular fashion, which seems a first step towards flexible, tunable mapping. [15] and [16] provide ontology mapping toolkits, which can be extended to fit more mapping techniques, but do not supply a situation dependent selection of mapping methods.

As a variety of tasks exists, we want our workbench to be application driven. Optimization or parameterization is scarcely addressed in current ontology mapping efforts. Apfel [11], which uses Machine Learning for configuration, does parameterize. It is goal driven in the sense that users define a training set of correct mappings, for which the algorithm optimizes. A drawback is that this tool combines all available mapping algorithms and data features and thus does not seem efficient for a large set of complex mapping algorithms.

2.3 Evaluation

Apfel is an example that shows that, for an optimization task, an evaluation method is important. Within the ontology mapping field usually precision, recall and f-measure [12] are used. During our pilot, however, it became apparent how ambiguous determining a golden standard, which is needed for all three measures, is. [11] mentions only 60% agreement between experts on mappings. Evaluation thus is an intrinsic part of a workbench approach and needs special attention.

Numerous evaluation approaches exist [12] [17] [18] [19], from which we will select appropriate ones for CH thesaurus mapping. Measures based on comparison with a reference alignment or golden standard (compliance measures) are [18]:

- precision
- recall
- fallout
- f-measure
- overall
- hamming distance
- weighted hamming distance

More subjective measures let the user evaluate the alignment method based on: the effort it takes to use it, whether it delivers an answer fast enough, whether the results are correct enough, whether all expected results are shown, and whether she finds the results understandable. Performance measures, which address scalability, memory consumption and speed of an algorithm, can also be used. [17] proposes a more advanced version for measuring precision and recall. Their method provides a proximity measure instead of a boolean function on whether a mapping co-occurs in the target alignment and the reference alignment. This way a distinction between a near miss and a complete failure regarding mapping can be quantified. Many research efforts address such ameliorations of evaluation techniques. [19] for example argues that certainty factors provided by mapping algorithms can give an indication of the performance of the algorithm itself. It is investigated whether the determination of certainty factors is monotonic: will this value increase if the mapping becomes better. If you have an algorithm using certainty factors and you can determine it is monotonic, then this algorithm indeed performs better if the factor is higher.

2.4 STITCH

In the STITCH pilot project [13] we evaluated two off-the-shelf, state-of-the-art mappers (Falcon [7] and S-Match [1] - the first an ontology mapper, the second a tree-like structure mapper) for mapping two CH thesauri (ARIA[2] and Iconclass[3]). ARIA has 500 concepts, multiple inheritance and is maximally two levels deep. Iconclass has 25.000 concepts, is up to ten levels deep and has some auxiliaries: keywords, cross references, keys, bracketed text and structural digits (for definitions see the Iconclass site). Here we describe the way Falcon and S-match performed on this task. The S-Match description will be more elaborate as my work focussed on this mapper.

2.4.1 Falcon

Falcon [7] is one of the best performing tools[4] for aligning complex RDFS/OWL ontologies. It relies on a combination of lexical comparison and graph-matching techniques. First, it compares concepts based on the set of weighted terms derived from their lexical "environment": their own identifiers, labels, comments, but also the ones of their immediate neighbors – parents, children – in the ontology. These similarities are used as input for the second step, which exploits a graph representation of the semantic information and matrix computation processes to finally return equivalence links between the concepts and relations of the compared ontologies.

[2] http://www.rijksmuseum.nl/aria/aria_catalogs/index?lang=en
[3] http://www.iconclass.nl
[4] See the 2005 OAEI campaign, http://oaei.ontologymatching.org/

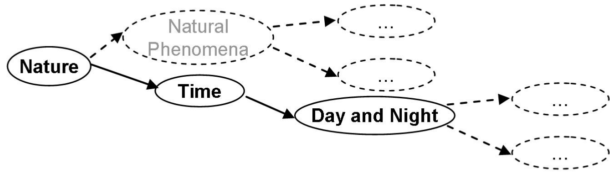

Figure 1: Iconclass concept 23Q ”Day and Night” and its parents

Results

Falcon performed very poorly against our vocabularies. First, as it uses a complex algorithm, it was practically very difficult to have it run with full IconClass as input. Some division of the huge classification scheme had to be done beforehand, which could prove harmful for the mapping process, since some associative links across general Iconclass branches would then become meaningless.

Falcon's resulting output was not satisfactory. Even with maximal involvement we reached only 16% precision for the notation 25 subset of Iconclass (1500 concepts) vs. complete ARIA (500 concepts).

2.4.2 S-Match

S-Match [1] has been developed for mapping classification schemes. It has a modular approach where a lexical, string based matching component, an oracle or background knowledge component and a structure based mapping module are most prominent. The input for S-Match is tree-like structures. The path from the root to a specific concept denotes the meaning of that concept. For each concept the path is rewritten as a propositional formula. Then matching is subsequently solved as a propositional unsatisfiability problem.

A node in a classification scheme has a local and a global semantics. For the example in figure 1 the local meaning of concept 23Q is all the Iconographical objects with the subject ”Day and Night”. Due to the structure of the tree, however, 23Q also has a global meaning which takes the parent nodes of the concept into account. For 23Q this means that Iconographical objects annotated with 23Q are also about ”Time” and ”Nature”. Therefore S-Match uses two steps in the mapping of concepts: local semantics matching and global semantics matching.

23Q's local semantics are constructed as follows: first the concept is tokenized into <day, and, night>, wherafter the tokens will be lemmatized if needed e.g. photos → photo. Then all senses for the lemmatized tokens are extracted from Wordnet and a logical formula is built e.g. <day, U(WNday)> ∨ <night, U(WNnight)>, where U is the union of all wordnet senses for the lemma.

Now our concepts' local semantics are mapped to that of concepts from the other tree using lexical methods and domain knowledge. A concept can either

IC notation	Iconclass textual correlate	Relation	ARIA label
23L	‘the twelve months represented by landscapes’	Less General	‘Landscapes’
25A271	‘(map of) the North Pole’	Less General	‘Charts, maps’
23U1	‘calendar, almanac’	Less General	‘Publications’

Table 1: Some good S-Match mapping results

have no relation to another concept, be equivalent, be disjoint or less or more general. All these relations can be rewritten in propositional formulas e.g. “A” less general “B” is written as “A” → “B”. This is the first step in the mapping process.

The global semantics for ”Day and Night” follows from its place in the scheme: it is the intersection of all concepts from the present concept to the root of the graph. In a formula for our example concept (from now called C1): (<day, U(WNday)> ∨ <night, U(WNnight)>) ∧ <time, U(WNtime)> ∧ <Nature, U(WNNature)>.

Matching the global semantics of two concepts can then be reduced to logical relations between their global semantics. The two concepts are equivalent if the logical formula “C1 ↔ C2” holds, they are disjoint if “¬ (C1 ∧ C2)”, and C1 is less general than C2 if “C1 → C2” (and vice versa for more general). To produce such a proof the results of the local semantics matching are taken as axioms and fed into a Boolean Satisfiability solver together with the logical formula describing the relation we are trying to prove.

Results

In table 1 some nice mappings produced by S-Match are shown, where the first mapping was produced based on pure lexical mapping, the second using stemming, and the third making use of background knowledge.

Notice that all Iconclass concepts are less general than the ARIA concepts. This is caused by the difference in depth of the hierarchy of both thesauri. An example of an Iconclass concept with its path to root: “/Top/Nature/earth, world as celestial body/animals”. The matching concept in ARIA would be: “/Top/Animals”. The SAT method leads to such subsumption links if only one unrelated term is used, e.g. “/house/table” vs. “/house/room/table” would correctly map, but “/house/table” vs. “/house/rabbit/table” would not. So the difference in the depth of the two hierarchies in combination with vague terms, such as ‘celestial’, makes S-Match conclude on subsumption, which is debatable if discussed with CH-experts.

For S-Match a precision of 46% is measured on the same subset as we tested Falcon on.

IC notation	Iconclass textual correlate	Relation	ARIA label
23H	‘seasons of the year represented by concepts other than [...] landscapes [...]’	Less General	‘Landscapes’
29D	‘natural forms in stones, wood, clouds’	Less General	‘Jewellery’

Table 2: Some bad S-Match results

Applying S-Match to our use case was not as easy as wished for, because it cannot deal with SKOS or OWL/RDF(S) input. We had to transform our thesauri to TAB indented trees, which is a format with less semantics than the SKOS we already had it in. Cross references between concepts are lost, for instance.

Table 2 shows the lack of interpretation of S-Match of “other than” in the first example and the lack of disambiguation of the link between “stones” and “jewellery” by the remainder of the concept (“natural forms in...”) in the second example. These results clearly show misinterpretation of thesaurus concepts.

This exercise showed major differences between thesauri and full-fledged ontologies. For example the gloss-like lexical size of concepts; a typical CH concept is “Noah’s sacrifice; various animals are offered, possibly a lamb, a dove and a ram (often combined with the rainbow of the covenant)”. Concepts in an ontology normally consist of just one word or term. This is one explanation for ontology mappers’ failure to interpret the complete meaning of a CH concept.

	S-Match	Falcon
Precision	46%	16%

Considering the results from both S-Match and Falcon makes us conclude that the state-of-the-art ontology mappers are not really suitable for Cultural Heritage thesaurus mapping.

In the following chapter we describe our ideas on the research steps needed for creating a thesaurus mapping workbench.

3 Research Steps

To build a workbench as we envision, we need an architecture streamlining the alignment process, but we also need thorough understanding of the domain. Figure 2 shows an initial architecture where our main research questions underlie construction of the groups of components (1, 2 and 3):

1. Conditions (requirements) analysis: distills requirements or conditions for alignment from input and application features. Our related research question is “Which input and application characteristics have which consequences on mapping?”

2. Mapping component(s) selection and configuration: selects relevant mapping methods for the given requirements and parameterizes these components. Useful combination of the mapping techniques is essential. Related question: “Which mapping components and which configurations are appropriate for which combinations of requirements?”

3. Evaluation and repair: evaluates alignment quality and decides whether improvements are needed, which will cause a loop back to the “Mapping component selection and configuration” phase. Related question: “Which evaluation measures are applicable for thesaurus alignment evaluation and how can evaluation outcome be used to tune mapping components?”

Below we elaborate on our ideas on each of these steps.

3.1 Requirements analysis

Mappers heavily rely on their input assumptions. The STITCH pilot project, for instance, showed that thesaurus size has significant impact on mapping output, see section 2.4.1. Without a thorough understanding of expected input features, mapping output is highly unlikely to be of good quality. Figure 2 shows that mapping requirements depend on availability of instance data, a reference alignment, and on characteristics of both thesauri and applications. These characteristics are discussed below, as they are CH specific.

Thesaurus characteristics In order to decide which thesaurus characteristics influence mapping, a full list of characteristics must first be constructed. Examples of thesaurus characteristics are thesaurus size, glosses (see section 2.4), auxiliaries, and thesaurus relations. Auxiliaries are additional semantics in the notation of thesaurus concepts, e.g. *qualifiers*: (...), which serve to disambiguate homonyms: cranes (lifting equipment) and cranes (birds)[9][10]. In Iconclass such ”contexts” are given by (∼), instead of (...). Such symbolic variances must be noted, to ensure correct interpretation of thesaurus concepts. If thesauri are ISO compliant the following relations might be present: BT (broader than), NT (narrower than), RT (related to), USE (points at the preferred term), UF (points at the non-preferred term)[10]. Thesaurus guidelines for original employment (e.g. subject indexing) also determine thesaurus characteristics. For instance some thesauri are meant for precoordination, some for postcoordination, which would argue respectively for compound or for one-to-one mapping.

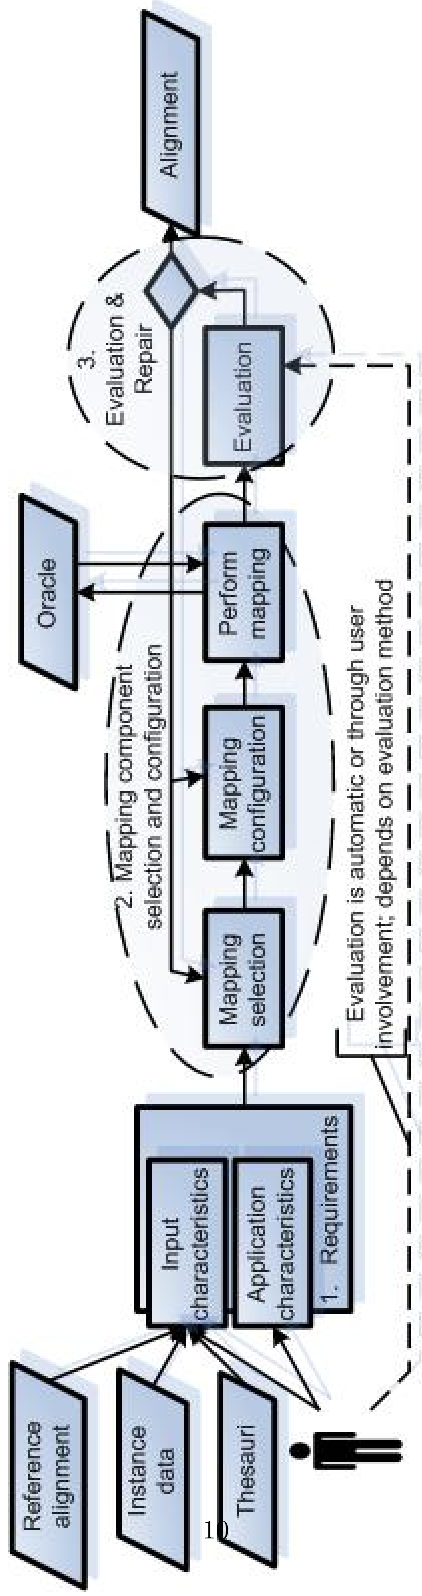

Figure 2: Workbench architecture

Subsequently an analysis of the relevance of each of these features for mapping and an analysis of the semantics of all thesaurus relations is needed. For example glosses could argue for compound mapping. We also have to investigate the impact of combinations of such characteristics, e.g. what happens if you try to map a very large thesaurus to a very small one?

Possibly relevant thesaurus characteristics are:

- size
- glosses as concepts
- growth; no. per timeunit
- facet or thesaurus subject, e.g. shape, subject, geography, personsname, corporation
- usage of lexical standard notation concerning abbreviations etc. (can for instance come from "het groene boekje")
- usage of updates and checks of outdated terms; checks for internal consistency and correctness by checking for orphans, multiple occurrences of terms, etc.
- singular vs. plural; normalisation or lemmatization
- auxiliaries; additional semantic constructs such as keywords
- coding of relations (mathematical symbols, words, ...)
- ISO compliant
  - BT/NT; BTG/NTG, BTP/NTP
  - RT
  - antonymy
  - homonymy
  - synonymy
  - see/see also, Use/UF
- language
- multilinguality (between or within thesauri)
- percentage topterms on complete set of terms; gives an indication of depth
- multiple inheritance
- depth

- combined with other thesaurus/classification scheme
- chance for misspelling (e.g. URI's?); for instances, but also thesaurus itself
- annotation rules: pre- or postcoordination (compound terms), no. of terms
- instances: kind of?, doubly indexed instances?
- size of the instances database

Application characteristics Alignment use poses mapping requirements, e.g. a mapping must be executed in real-time, must have excellent recall, etc. Such requirements can be distilled from common scenarios, so our second step is to determine those use cases. An example of a *migration* scenario is the following: at the Dutch National Library (KB) currently two thesauri are used for partially different, but largely overlapping collections. A lot of time and effort is invested in annotating all incoming documents using both systems and keeping both up to date. For efficiency reasons one of the two might have to go, but all documents that are already indexed against that thesaurus should still be retrievable. So a mapping between these thesauri is needed. An important mapping criterion determined by this application is that it should be complete: none of the terms in the to-be-discarded thesaurus should be left out. This example shows that scenarios pose requirements for mappings, in this case a focus on recall, so that experts can manually filter out bad results but can be sure not to miss any correct mappings.

Possible use cases are:

- migration
- integrated search and browse
  - on the fly mapping during search
- enrichment of an existing thesaurus
- update of an existing alignment

The requirements above form preconditions for the mapping components; the useable parts of the input constrain mapping methods and the application requirements constitute the goal. Given these preconditions the real mapping can start.

3.2 Mapping component selection and configuration

Four main mapping methods are prominent; lexical, structure based, oracle based and instance based matching [14]. All these approaches can be applied to Cultural Heritage thesaurus mapping. We will decompose existing black-box tools into small mapping components. Extending or adapting these mapping techniques to the CH thesaurus domain forms a significant part of our research. Therefore hypotheses for CH thesauri specific interpretation of all four mapping categories are presented in the following paragraphs. S-Match and Falcon are often used as examples here, due to their role in STITCH's pilot project. For each mapping method category an example is given of a thesaurus characteristic or use case that has an impact on the method.

Lexical or term based mapping Section 2 suggests glosses are Cultural Heritage specific thesaurus features. Glosses ask for compound mapping, but also for much more elaborate Natural Language Processing techniques than simple terms do. S-Match considers every word or term in a concept as equally important (except for stopwords). Extended NLP techniques are needed to interpret constructs such as "other than", see table 2. So the thesaurus characteristic "glosses as concepts" votes for a lexical mapping component with extensive NLP.

Symbolic variances as mentioned in section 3.1 should also be accounted for in a lexical mapping component. Important information concerning concept semantics can be extracted from such parts of notations. In our research we want to provide ways to use them for mapping.

Structure based mapping What structure components are present in the thesauri determines which structure based mapping approaches should be used. For example, all thesaurus relations (see subsection 3.1) can be used for mapping. S-Match just uses the hierarchy (BT and NT). Falcon, as most ontology mappers, assumes a stricter defined structure in the form of properties, property inheritance, value restrictions, etc. Such constructs are not present in CH thesauri though, and the semantics found in thesauri are not used by Falcon. Our goal is to investigate the level of structure within CH thesauri (first research objective) and analyse those constructs to determine the appropriate mapping algorithms. If the thesaurus characteristic "depth" has the value "one", then hierarchical matching techniques should remain unemployed.

Oracle based mapping Currently S-Match uses Wordnet as background knowledge for improving matching. From Wordnet information as synonymy, hyper- and hyponymy is used to clarify the meaning of a concept. Other mappers use other oracles and thus it is relevant to decide on appropriate ones, for instance AAT (domain-specific) or Wordnet (general), for mapping within the CH domain. The impact such a choice has on the mapping process must be investigated. “Multilinguality” between thesauri determines that dictionaries should be used as oracle, maybe in addition to normal oracles which can define synsets.

Instance based mapping We assume instance based mappers to operate quite similar in both ontology or thesaurus mapping circumstances, as will be tested in this research. This depends on the semantics of “instance” e.g. ontologies argue “leo” to be an instance of “Lion”, but thesauri claim a book annotated with the term “Lion” to be an instance of the thesaurus-concept “Lion”. We will evaluate the effects these semantic differences have on mapping performance. This mapping method is of course only relevant if doubly indexed instances exist.

Our research will further focus on combinations of and dependencies between the techniques above. During the research we will take automatic configuration, for instance using agents or machine learning, into consideration. In the above many times the impact of a certain mapping algorithm on output quality is mentioned. The next section describes our ideas on evaluation methods.

3.3 Evaluation and repair

Sound evaluation of the quality of a mapping is needed on two levels: for workbench construction and also for automatic adaptation during mapping execution. Constructing a workbench relies heavily on whether it is clear which mapping techniques are appropriate for certain sets of circumstances. Requirements for mappings based on user scenarios and thesaurus characteristics are highly relevant for determining the right evaluation measure: for instance a focus on recall or on precision. Section 2 describes current evaluation practices where, although the subjectivity of such standard is widely known, it still is common practice to let human experts construct a “golden standard”. We want to put current methods for constructing “golden standards” up for discussion, so our research will address algorithms or guidelines for expert mapping and also on alternative evaluation methods. The OAEI[5] mapping initiative, for instance, uses alternative measures, which we will study in detail.

Evaluation measures must also provide information on where a mapping method fails. Such information is needed to loop back in the workbench architecture and repair or tune the method used. More detailed measures will be needed than above, which might just be executed when the global measures detect low mapping quality. For instance, when recording low precision it might be checked whether concepts have multiple matches, from different branches in the target-thesaurus. If so, Word Sense Disambiguation should perform better, for instance by stricter structure based mapping.

[5] `http://oaei.inrialpes.fr/`

3.4 Workbench building

Possible users of the workbench would be knowledge experts, the general public, librarians or domain experts. Their acceptance of the tool would depend on several criteria: data availability and the ease with which random file formats can be imported into the system, their expectations and expected gains, quality of the results when they use the tool once, time (they will want concrete results, fast), reactiveness (will the system adapt based on its input and expert feedback), which hardware resources are needed, maintenance, ownership and how data is protected, and the ability of the tool to explain its output. These criteria can be partly used to evaluate the tool and partly to put up the tools requirements:

- no. of thesauri that can be aligned (just 2 or more?)
- output format; with certainty factor or without, kind of relations (BT/NT or just equality)
- option to manually adjust configuration
- 1-1 or compound mapping
- possibility to match just one concept instead of the entire ontology (a UI feature only)
- help function
- explanatory/verbose mode - for newbies

4 Planning

July 1st 2006 - June 30th 2007 This paragraph covers the tasks that need to be performed in the second year of my research.

The first half year will be dedicated to theory gathering on:

- Thesaurus features
- Mapping algorithms
- Plausible use-cases
- Evaluation

The second half of the year we will pursue these subjects, and draw theoretical conclusions combining the diverse subjects, e.g. the influence of certain use cases on evaluation metrics. This will provide a sound theoretical foundation for configuration and combination of mapping algorithms for the task at hand.

In addition to theoretical aspects the workbench architecture will also be written in this second half year and a first prototype will be build. This first prototype will incorporate every component (although restricted to some mappers only) as shown in figure 2, except for the repair loop.

In this first half of 2007, I will also spend a period of time abroad, preferably at Martin Doerr's group or in Tokyo.

July 1st 2007 - June 30th 2008 The third year consists of completing the workbench by implementing reconfiguration within an evaluate and repair loop and by implementing full functionality for the workbench. In addition to this a round of user evaluation regarding workbench functionalities will be executed, so that we can assess its quality from a practical perspective. This user evaluation will probably also provide insights for improvements and bug reports which can be incorporated in the system's final version.

Around January 2008 I will spend another period in a foreign country.

July 1st 2008 - June 30th 2009 A year of workbench finalisation and thesis writing.

5 Conclusion

Heterogeneous vocabularies mapping is a prominent issue in many domains, but is especially relevant for collection integration in the Cultural Heritage field. In this paper we have sketched a plan for solving this matching problem in an application based fashion, grounded in the present state-of-the-art theory on ontology mapping. The presented modular workbench construction will provide case based thesaurus mapping for the Cultural Heritage field. Using this workbench can heavily improve collection integration practices.

6 Definitions

Mapper / aligner = complete mapping tool

Mapping / match = one (1) correspondence between concepts

Mapping component / algorithm = one implemented method

Alignment = output of a mapping method; a complete set of mappings

References

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

[2] Clavel-Merrin, G.: MACS (Multilingual access to subjects): A Virtual Authority File across Languages. Cataloguing and Classification Quarterly 39 (1/2), 2004.

[3] Constantopoulos, P., Sintichakis, M.: A Method for Monolingual Thesauri Merging. ACM SIGIR Conference, Philadelphia, USA, 1997.

[4] Doerr, M.: Semantic Problems of Thesaurus Mapping. Journal of Digital Information, 1 (8), 2004.

[5] Hyvönen, E., Mäkelä, E., Salminen, M., Valo, A., Viljanen, K. et al.: MuseumFinland - Finnish Museums on the Semantic Web. Journal of Web Semantics, 3(2), 2005.

[6] Kalfoglou, Y., Schorlemmer, M.: Ontology Mapping: The State of the Art. The Knowledge Engineering Review Journal, 18(1), 2003.

[7] Jian, N., Hu, W., Cheng, G., and Y. Qu: Falcon-AO: Aligning Ontologies with Falcon. K-CAP Workshop on Integrating Ontologies, Banff, Canada, 2005.

[8] Handleiding voor de Gemeenschappelijke Onderwerpsontsluiting (GOO). Derde, herziene editie, april 2006.

[9] ISO 5964. Documentation - Guidelines for the establishment and development of multilingual thesauri, 1985.

[10] ISO 2788. Documentation - Guidelines for the establishment and development of monolingual thesauri, 1986.

[11] Ehrig, M., Staab, S. and Y. Sure: Bootstrapping Ontology Alignment Methods with APFEL. International Semantic Web Conference, 2005, pp. 186-200.

[12] van Rijsbergen, C.J.: Information Retrieval. Information Retrieval Group. University of Glasgow, 1979, London: Butterworths.

[13] van Gendt, M., Isaac, A., van der Meij, L., and S. Schlobach: Semantic web techniques for multiple views on heterogeneous collections: a case study. To appear in: European Conference on Research and Advanced Technology for Digital Libraries, 2006.

[14] van Harmelen, F.: Ontology Mapping: A Way Out of the Medical Tower of Babel? 10th Conference on Artificial Intelligence in Medicine, AIME 2005, Aberdeen, Proceedings, pp. 3-6.

[15] Ziegler, P., Kiefer, C., Sturm, C., Dittrich, K.R. and A. Bernstein: Detecting Similarities in Ontologies with the SOQA-SimPack Toolkit. To appear in: 10th International Conference on Extending Database Technology (EDBT 2006), Munich, Germany, 2006.

[16] Kalfoglou, Y., Hu, B., Reynolds D. and N. Shadbolt: CROSI project, final report. CUniversity of Southampton, Technical Report, E-Print No. 11717. October, 2005.

[17] Ehrig, M., Euzenat, J.: Relaxed Precision and Recall for Ontology Matching. In Benjamin Ashpole, Marc Ehrig, Jérôme Euzenat, Heiner Stuckenschmidt, Proceedings of the Workshop on Integrating Ontologies, volume 156, pp. 8. CEUR-WS.org, August 2005.

[18] Ehrig, M., Euzenat, J., and R. G. Castro: Specification of a benchmarking methodology for alignment techniques. Knowledge Web Deliverable 2.2.2, INRIA Rhone-Alpes. January 2005

[19] Anaby-Tavor, A., Gal, A., and Trombetta, A: Evaluating Matching Algorithms: The Monotonicity Principle. IJCAI-03 Workshop on Information Integration on the Web, Acapulco, Mexico, August 9-10, 2003.

Cultural Heritage institution often have a wealth of metadata to share and exploit. 24 Enabling bits & pieces Exploiting semantic links in CH vocabularies Rijksmuseum thesaurus: Concept â€œGizaâ€ narrower than concept â€œEgypteâ€ Mapping/alignment between CH vocabularies Louvreâ€™s â€œÃ‰gypteâ€ equivalent to Rijksmuseumâ€™s â€œEgypteâ€ Enrichment of existing metadata The string â€œEgyptâ€ in a metadata record indicates the concept of Egypt defined in Rijksmuseum thesaurus.Â A web-based repository service for vocabularies and alignments in the Cultural Heritage domain Lourens van der Meij Antoine Isaac Claus Zinn. Notes on ThoughtLab / Athena WP4 November 13, 2009 Antoine Isaac.