

Aligning Anatomy Ontologies in the Ontology Alignment Evaluation Initiative

Patrick Lambrix, Qiang Liu, He Tan
Department of Computer and Information Science
Linköpings universitet
581 83 Linköping, Sweden

Abstract

In recent years many ontologies have been developed and many of these ontologies contain overlapping information. To be able to use multiple ontologies they have to be aligned. In this paper we present and discuss results from aligning ontologies in a real case, the anatomy case in the 2008 Ontology Alignment Evaluation Initiative. We do this by briefly describing a base system for ontology alignment, SAMBO, and an extension, SAMBOdtf, and present and discuss their results for the anatomy case. SAMBO and SAMBOdtf performed best and second best among the 9 participating systems. SAMBO uses a combination of string matching and the use of domain knowledge. SAMBOdtf uses the same strategies but, in addition, uses an advanced filtering technique that augments recall while maintaining a high precision. Further, we describe the first results on ontology alignment using a partial reference alignment.¹

1 Introduction

In recent years many ontologies have been developed. Intuitively, ontologies (e.g. [6]) can be seen as defining the basic terms and relations of a domain of interest, as well as the rules for combining these terms and relations. They are considered to be an important technology for the Semantic Web. Ontologies are used for communication between people and organizations by providing a common terminology over a domain.

¹This paper is partly a revised and updated version of the paper [12] focusing on the anatomy ontology alignment task, and partly an extended version. The former paper contains brief descriptions of the systems, but, following the tradition of the Ontology Alignment Evaluation Initiative, was written before the final results were available.

They provide the basis for interoperability between systems. They can be used for making the content in information sources explicit and serve as an index to a repository of information. Further, they can be used as a basis for integration of information sources and as a query model for information sources. They also support clearly separating domain knowledge from application-based knowledge as well as validation of data sources. The benefits of using ontologies include reuse, sharing and portability of knowledge across platforms, and improved documentation, maintenance, and reliability (e.g. [7]). Ontologies lead to a better understanding of a field and to more effective and efficient handling of information in that field.

Many of the currently developed ontologies contain overlapping information. For instance, Open Biomedical Ontologies (OBO, <http://www.obofoundry.org/>) lists 26 different anatomy ontologies (January 2009). Often we would want to be able to use multiple ontologies. For instance, companies may want to use community standard ontologies and use them together with company-specific ontologies. Applications may need to use ontologies from different areas or from different views on one area. Ontology builders may want to use already existing ontologies as the basis for the creation of new ontologies by extending the existing ontologies or by combining knowledge from different smaller ontologies. In each of these cases it is important to know the relationships between the terms in the different ontologies. Further, the data in different data sources in the same domain may have been annotated with different but similar ontologies. Knowledge of the inter-ontology relationships would in this case lead to improvements in search, integration and analysis of data. It has been realized that this is a major issue and some organizations have started to deal with it.

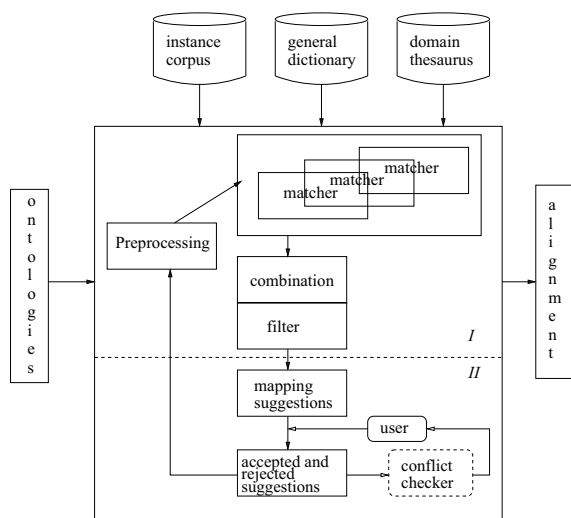


Figure 1: Alignment framework [11, 10].

In the remainder of this paper we say that we align two ontologies when we define the mapping relationships between terms in the different ontologies. We discuss results from aligning ontologies in a real case, the anatomy case in the 2008 Ontology Alignment Evaluation Initiative, one of the best known benchmark cases. We do this by presenting the two state-of-the-art ontology alignment systems that performed best and second best, and present and discuss their results for the anatomy case.

2 Background

2.1 Framework

A large number of ontology alignment systems have been developed. For an overview of most of these systems, we refer to review papers (e.g. [11, 16, 15, 8]), the ontology matching book [4], and the ontology matching web site at <http://www.ontologymatching.org/>.

Many ontology alignment systems are based on the computation of similarity values between terms in different ontologies and can be described as instantiations of the general framework defined in [11, 10] (figure 1). The framework consists of two parts. The first part (*I* in figure 1) computes mapping suggestions. The second part (*II* in figure 1) interacts with the user to decide on the final mappings.

An alignment system receives as input two source ontologies. The ontologies can be preprocessed, for instance, to select pieces of the ontologies that are likely to contain matching terms. The alignment algorithm includes one or several matchers, which calculate similarity values between the terms from the different source ontologies and can be based on knowledge about the linguistic elements, structure, constraints and instances of the ontology. Also auxiliary information can be used. Mapping suggestions are then determined by combining and filtering the results generated by one or more matchers. By using different matchers and combining and filtering the results in different ways we obtain different alignment strategies. The suggestions are then presented to the user who accepts or rejects them. The acceptance and rejection of a suggestion may influence further suggestions. Further, a conflict checker is used to avoid conflicts introduced by the mappings. The output of the ontology alignment system is an alignment which is a set of mappings between terms from the source ontologies.

2.2 SAMBO and SAMBOdtf

SAMBO and SAMBOdtf are based on the framework described in section 2.1. They do not have a preprocessing step. SAMBO and SAMBOdtf contain currently five basic matchers [11]: two terminological matchers (a basic matcher and an extension using WordNet; extension described below), a structure-based matcher (which uses the is-a and part-of hierarchies of the source ontologies), a matcher based on domain knowledge (described below), and a learning matcher (which uses life science literature that is related to the concepts in the ontologies to define a similarity value between the concepts). In addition to these techniques we have also experimented with other matchers [13, 18, 21].

The user is given the choice to employ one or several matchers during the alignment process. We have two strategies to combine the results from different matchers. One strategy is to give weights to the different matchers and the similarity values are then computed as a weighted sum of the similarity values computed by the different matchers. The other strategy defines the similarity of a pair of terms as the maximum value of the similarity values for the pair computed by the different matchers.

The filtering method in SAMBO is single threshold filtering. Pairs of terms with a similarity value higher than or equal to a given threshold value are returned as

mapping suggestions to the user.



Figure 2: Combination and filtering.

Figure 2 shows a screenshot from the SAMBO system with the five matchers, a weighted sum combination and the single threshold filtering.

SAMBODtf implements the double threshold filtering method developed in [2]. The double threshold filtering approach uses the structure of the ontologies. It is based on the observation that (for the different approaches in the evaluation in [11]) for single threshold filtering the precision of the results decreases and the recall increases when the threshold decreases. Therefore, we propose to use two thresholds. Pairs with similarity value equal to or higher than the upper threshold are retained as suggestions. The intuition is that this gives suggestions with a high precision. Further, pairs with similarity values between the lower and the upper threshold are filtered using structural information and the rest is discarded. We require that the pairs with similarity values between the two thresholds are 'reasonable' from a structural point of view.² The intuition here is that the recall is augmented by adding new suggestions, while at the same time the precision stays high because only structurally reasonable suggestions are added. The double threshold filtering approach contains the following three steps. (i) Find a consistent suggestion group from the pairs with similarity value equal to or higher than the upper threshold. We say that a set of suggestions is a consistent suggestion group if each concept occurs at most once as first argument in a pair, at most once as second argument in a pair and for each pair of suggestions (A, A') and (B, B') where A and B are concepts in the first ontology and A' and B' are concepts in the second ontology: $A \subset B$ iff $A' \subset B'$. (ii) Use the consistent suggestion group to partition the original ontologies. (iii) Filter the pairs with similarity values between the lower and upper thresholds using the partitions. Only pairs of which the elements belong

²In our implementation we have focused on the is-a relation.




Figure 3: Mapping suggestion.

to corresponding pieces in the partitions are retained as suggestions. For details we refer to [2].

Based on the results from the matchers, combination and filtering algorithms, mapping suggestions are provided to the user. Figure 3 shows such a suggestion. SAMBO displays information (definition/identifier, synonyms, relations) about the source ontology terms in the suggestion. For each mapping suggestion the user can decide whether the terms are equivalent, whether there is an is-a relation between the terms, or whether the suggestion should be rejected. If the user decides that the terms are equivalent, a new name for the term can be given as well. Upon an action of the user, the suggestion list is updated. If the user rejects a suggestion where two different terms have the same name, she is required to rename at least one of the terms. The user can also add comments on a mapping relationship. At each point in time during the alignment process the user can view the ontologies represented in trees with the information on which actions have been performed, and she can check how many suggestions still need to be processed. A similar list can be obtained to view the previously accepted mapping suggestions. In addition to the suggestion mode, the system also has a manual mode in which the user can view the ontologies and manually map terms.

2.3 Ontology Alignment Evaluation Initiative - Anatomy case

The Ontology Alignment Evaluation Initiative (OAEI, <http://oaei.ontologymatching.org/>) is a yearly initiative that was started in 2004. The goals are, among others, to assess the strengths and weaknesses of alignment systems, to compare different techniques and to improve evaluation techniques. This is to be achieved through controlled experimental evaluation. For this purpose OAEI publishes different cases of ontology alignment problems, some of which are open (reference alignment

is known beforehand), but most are blind (reference alignment is not known - participants send their mapping suggestions to organizers who evaluate the performance).

In the anatomy case (version 2008) participants are required to align the Adult Mouse Anatomy (2744 concepts) and the NCI Thesaurus - anatomy (3304 concepts). The case is divided into 4 tasks (of which task 4 was new for 2008). The anatomy case is a blind case. The reference alignment (the correct solution according to the organizers) contains 1523 equivalence mappings of which 934 are deemed trivial (i.e. they can be found by a relatively basic string-based matcher). Only equivalence correspondences between concepts are considered.

In all tasks the two ontologies should be aligned. The results of the experiments are given in terms of the quality of the mapping suggestions. The evaluation measures are precision, recall, recall+ and f-measure. *Precision* measures how many of the mapping suggestions were correct. It is defined as the number of correct suggestions divided by the number of suggestions. *Recall* measures how many of the correct mappings are found by the alignment algorithm. It is defined as the number of correct suggestions divided by the number of correct mappings. *Recall+* is the recall computed with respect to non-trivial mappings. *F-measure* is the weighted harmonic mean of precision and recall.

In task 1 the system should be tuned to optimize the f-measure. This means that both precision and recall are important. The systems are compared with respect to precision, recall, f-measure and recall+. For the f-measure in task 1, precision and recall are evenly weighted. Nine systems participated in this task.

In tasks 2 and 3, in which four systems participated, the system should be optimized with respect to precision and recall, respectively. The f-measure is computed with an unevenly weighted precision and recall (factor 5).

In task 4, in which four systems participated, a partial reference alignment is given which can be used during the computation of mapping suggestions. It contains all trivial and 54 non-trivial mappings in the reference alignment. In this case precision, recall and f-measure are computed with respect to the non-given part of the reference alignment.

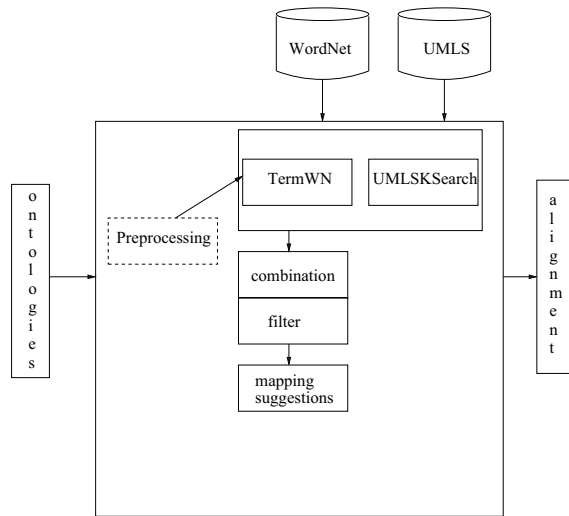


Figure 4: SAMBO and SAMBOdtf for OAEI.

3 SAMBO and SAMBOdtf for OAEI

The OAEI only evaluates the non-interactive part of the ontology alignment systems. Therefore, we used a variant of the systems without the user interface (see figure 4). Further, it would not make sense to have mapping suggestions where a concept appears more than once as the user would not be able to make a choice. Therefore, we decided to filter our systems' mapping suggestion lists such that only suggestions are retained where the similarity between the concepts in the mapping suggestion is higher than or equal to the similarity of these concepts to any other concept according to the mapping suggestion list. (In the case there are different possibilities, one is randomly chosen. In the implementation the first in the list is chosen.)

For the OAEI we used the following matchers. The matcher *TermWN* contains matching algorithms based on the textual descriptions (names and synonyms) of concepts and relations. In the current implementation, the matcher includes two approximate string matching algorithms (n-gram and edit distance), and a linguistic algorithm that also uses WordNet (<http://wordnet.princeton.edu/>) to find synonyms and is-a relations. Our matcher *UMLSKSearch* uses the Metathesaurus in the Unified Medical Language System (UMLS, <http://www.nlm.nih.gov/research/umls/>). The similarity of two terms in the source ontologies is

determined by their relationship in UMLS. In our experiments we used the UMLS Knowledge Source Server to query the UMLS Metathesaurus with source ontology terms. The querying is based on searching the normalized string index and normalized word index provided by the UMLS Knowledge Source Server. We used version 2008AA of UMLS. As a result we obtain concepts that have the source ontology term as their synonym. We assign a similarity value of 0.99 if the source ontology terms are synonyms of the same concept and 0 otherwise.

The combination algorithm used for OAEI 2008 is a maximum-based algorithm. The similarity value for a pair of concepts is the maximum value obtained from TermWN and UMLSKSearch for this pair of concepts.

As in the full SAMBO and SAMBOdtf systems, SAMBO uses single threshold filtering and SAMBOdtf double threshold filtering.

4 Results for the Anatomy Case

The results for the participating systems and discussions are available from <http://oaei.ontologymatching.org/2008/> and the paper [1].

For SAMBO and SAMBOdtf tests were performed on a IBM R61i Laptop, WinXP Intel(R) Pentium(R) Dual T2370 @ 1.73GHz, 1.73GHz, 1.99G RAM.

Task 1. We used the matchers, combinations and filtering described in section 3. SAMBO used threshold 0.6. SAMBOdtf used upper threshold 0.8 and lower threshold 0.4. These thresholds were chosen based on our experience with previous experiments with biomedical ontologies. SAMBO generated 1465 mapping suggestions and reached a precision of 0.869, a recall of 0.836 and an f-value of 0.852. Further, it reached a recall+ of 0.586. This was the best result for all 9 participating systems in OAEI 2008.³ In 2007 we used a version of SAMBO that used Term instead of TermWN and a previous version of UMLS. The 2007 version obtained a better recall for non-trivial mappings, but at the cost of an overall decrease in precision and recall. A possible explanation for this is our strategy for choosing maximum one mapping suggestion per concept. In 2008 exact matching strings were preferred, while in

2007 there was no preference between pairs that had exact matching strings or pairs that were proposed based on domain knowledge.

SAMBOdtf generates 1527 mapping suggestions. Of these suggestions, 1440 have a similarity value between 0.6 and 0.8. This means that SAMBOdtf filtered out 25 of the suggestions obtained by SAMBO with threshold 0.6. (A manual check seems to suggest that most of these are correctly filtered out, but some are wrongly filtered out. One reason for removing correct suggestions is that the source ontologies have missing is-a links.) Further, SAMBOdtf also filtered out 19 suggestions with similarity values between 0.4 and 0.6. (A manual check seems to suggest that these were correctly filtered out.) SAMBOdtf obtained a precision of 0.831, a recall of 0.833, an f-value of 0.832 and a recall+ of 0.579. This was the second best result for all 9 participating systems in OAEI 2008.

The running time for SAMBO was ca 12 hours and for SAMBOdtf ca 17 hours. As discussed in [1] the best performing systems in 2007 and 2008 in terms of quality of the mapping suggestions heavily use domain knowledge. This comes at a cost of larger running time. However, another interesting observation discussed by the organizers of the anatomy track at OAEI 2008 was that circa 50% of the non-trivial mappings was found by at least one system using domain knowledge and at least one system that did not use domain knowledge. Circa 13% of the non-trivial mappings was found only by systems using domain knowledge. Circa 13% of the non-trivial mappings was found only by systems that did not use domain knowledge. The reason for this is that the used domain knowledge (most often UMLS) is not complete. Further, still circa 25% of the non-trivial mappings were not found at all. As [1] suggests, a combination of different strategies may improve the results. Taking the union of the SAMBO results with the results of the RiMOM [24] and Lily [22] systems would give a higher recall and recall+. RiMOM and Lily use linguistic and structure-based approaches, but no domain knowledge.

Tasks 2 and 3. We did not participate in tasks 2 and 3. As reported in [1] the best system for task 2 (RiMOM) obtained a precision of 0.964 (with a recall of 0.677). The best system for task 3 (RiMOM) obtained a recall of 0.808 (with a precision of 0.450 and a recall+ of 0.538). We note that the best recall is lower than the recall for SAMBO and SAMBOdtf in task 1. The best system for task 3 for non-trivial mappings (Lily) obtained a recall+ of 0.613 (with a recall of 0.790 and a

³The system with best f-measure in 2007 (AOAS [23]) obtained 0.928 precision, 0.815 recall, 0.523 recall+ and 0.868 f-measure. SAMBO (in its first participation) was second best system in 2007 regarding precision, recall and f-value, but best regarding recall+ [20].

precision of 0.490). As neither RiMOM nor Lily used domain knowledge, these can be considered to be good results.

Task 4. For task 4, we augmented SAMBO and SAMBOdtf in the following ways.

For SAMBO we added the mappings in the partial reference alignment to the list of mapping suggestions, but with a special status. These mappings could not be removed in any filtering step. SAMBO generated 1494 suggestions of which 988 are also in the partial reference alignment. SAMBO obtained the best results of the participating systems. With respect to the unknown part of the reference alignment, its precision increased with 0.024, its recall decreased with 0.002 and its f-value increased with 0.011. Our strategy for using the partial reference alignment helped remove wrong suggestions that conflicted with the partial reference alignment, although also some correct suggestions were removed.

For SAMBOdtf we also added the mappings in the partial reference alignment to the list of mapping suggestions with the special status. In addition, we used the partial reference alignment in the double threshold filtering step. We used a part of the partial reference alignment that satisfied the consistent group property as a consistent suggestion group. For upper threshold 0.8 and lower threshold 0.4 we obtained 1547 mapping suggestions. SAMBOdtf obtained the second best results of the participating systems. With respect to the unknown part of the reference alignment, its precision increased with 0.040, its recall with 0.008 and its f-value with 0.025. SAMBOdtf was the system with the highest increase in f-value and was the only system that used the partial reference alignment to increase both precision and recall. This result is most likely due to the fact that, in contrast to task 1 where the consistent suggestion group consists of suggestions, in this task the consistent suggestion group consists of true mappings. Therefore, the suggestions with similarity value between the two thresholds that are retained are structurally reasonable with respect to true mappings and not just (although with high confidence) suggestions.

We note that although the improvements seem small, as SAMBO and SAMBOdtf perform already well on their own, even small improvements are valuable. Further, due to the choice of the partial reference alignment all newly found mappings are non-trivial.

In a follow-up on task 4 we have started investigating the use of partial reference alignments in the different components of the framework in section 2.1 [10]. In ad-

dition to the techniques described above, we have used partial reference alignments in a preprocessing step, to define new matchers and in new filtering steps.⁴

In the preprocessing approaches we investigate whether we can use a partial reference alignment to partition the ontologies into mappable parts and test whether, in addition to the fact that we do not have to compute similarity values between all terms from the first ontology and all terms from the second ontology, this also leads to a better quality of the mapping suggestions. In the first approach we partition the ontologies into mappable parts using the partitioning step of the double threshold filtering described in section 2.2 and [2]. A part of the partial reference alignment satisfying the consistent group property is used as a consistent group. Further, according to our experience in aligning ontologies we know that the structure of the source ontologies is not always perfect. For instance, given the two ontologies and the partial reference alignment in the anatomy case of OAEI 2008, it can be deduced that many is-a relations are missing in at least one of the source ontologies. Based on this observation we experiment with a second approach where we add to the source ontologies the missing is-a relationships that can be deduced from the source ontologies and the partial reference alignment. After this 'fixing' of the source ontologies the partial reference alignment will satisfy the consistent group property. As the intuition of the preprocessing step is to partition the ontologies into mappable parts, we can only generate mapping suggestions that are reasonable from a structural point of view. This suggests that, when using a preprocessing step, the precision may become higher as suggestions that do not conform to the structure of the source ontologies cannot be made. As we add the partial reference alignment to the result, the recall may be increased as some of the partial reference alignment mappings may not be found by the base systems. However, the similarity values between the terms do not change and it is therefore not likely that new mappings are found. For thresholds 0.6 and above our experiments corroborate this intuition. Another observation is that, contrary to the intuition, fixing the source ontologies may lead to a decrease in recall. The reason for this is the quality of the underlying ontologies where 'is-a' is not always properly used.

One way to create a matcher based on a partial reference alignment, is to use underlying properties of

⁴Thanks to Christian Meilicke of the organization committee of OAEI Anatomy for running our newly developed algorithms on the anatomy data set.

the mappings in the partial reference alignment. We have previously observed that sometimes for two given source ontologies, common patterns can be found between the correct mappings. For instance, in the partial reference alignment of the OAEI 2008 anatomy we find the mappings <lumbar vertebra 5, 15 vertebra> and <thoracic vertebra 11, t11 vertebra> which share a similar linguistic pattern. Based on this observation we developed a matcher that augments previously generated similarity values for term pairs when these term pairs display a similar (linguistic) pattern as mappings in the partial reference alignment. Several new correct mappings were found.

Finally, we also experimented with a filter strategy that removes suggestions that do not have similar linguistic patterns than the mappings in the partial reference alignment. We expect therefore that some correct suggestions obtained through UMLS will be removed and therefore the recall may go down. This is indeed the case in our experiments. The precision when using this filter approach is, however, always higher or equal to the precision for SAMBO. This is because the suggestions that had a linguistically similar pattern as mappings in the partial reference alignment were usually correct.

5 Conclusion

We have briefly described our ontology alignment systems SAMBO and SAMBOdtf and their results for the anatomy alignment tasks of OAEI. We have used a combination of UMLSKSearch and TermWN and obtained the best results in OAEI anatomy 2008 with respect to quality of the suggestions. However, as the recall+ of the best system is still around 0.6, work still needs to be done to find non-trivial mappings.

Another problem that we investigate is whether systems that do well in the anatomy case will also perform well for other cases. More large-scale evaluation is needed in the area.

Further, the OAEI cases only provide a benchmark for part I of the framework described in section 2.1. Not so much work has been done on user involvement, user interfaces and ontology and ontology alignment visualization [9, 5].

Also, given the fact that different algorithms seem to do differently well for different kinds of ontologies and evaluation measures, a major problem is deciding which algorithms should be used for a given alignment

task. This is a problem that users face, and that we have also faced in the evaluation. Recommendation strategies [19, 14, 3] may alleviate this problem.

Other challenges for the ontology alignment field are given in [17].

References

- [1] C Caracciolo, J Euzenat, L Hollink, R Ichise, A Isaac, V Malaise, C Meilicke, J Pane, P Shvaiko, H Stuckenschmidt, O Svab-Zamazal, and V Svatek. Results of the ontology alignment evaluation initiative 2008. In *Proceedings of the Third International Workshop on Ontology Matching*, 2008.
- [2] B Chen, H Tan, and P Lambrix. Structure-based filtering for ontology alignment. In *Proceedings of the IEEE WETICE Workshop on Semantic Technologies in Collaborative Applications*, pages 364–369, 2006.
- [3] M Ehrig, S Staab, and Y Sure. Bootstrapping ontology alignment methods with APFEL. In *Proceedings of the International Semantic Web Conference*, pages 186–200, 2005.
- [4] J Euzenat and P Shvaiko. *Ontology Matching*. Springer, 2007.
- [5] S Falconer and M-A Storey. A cognitive support framework for ontology mapping. In *Proceedings of the 6th International Semantic Web Conference*, pages 114–127, 2007.
- [6] A Gómez-Pérez. Ontological engineering: A state of the art. *Expert Update*, 2(3):33–43, 1999.
- [7] R Jasper and M Uschold. A framework for understanding and classifying ontology applications. In *Proceedings of the 12th Workshop on Knowledge Acquisition, Modeling and Management*, 1999.
- [8] Y Kalfoglou and M Schorlemmer. Ontology mapping: the state of the art. *The Knowledge Engineering Review*, 18(1):1–31, 2003.
- [9] P Lambrix and A Edberg. Evaluation of ontology merging tools in bioinformatics. In *Proceedings of the Pacific Symposium on Biocomputing*, pages 589–600, 2003.

- [10] P Lambrix and Q Liu. Using partial reference alignments to align ontologies. In *Proceedings of the 6th European Semantic Web Conference*, 2009.
- [11] P Lambrix and H Tan. SAMBO - a system for aligning and merging biomedical ontologies. *Journal of Web Semantics, Special issue on Semantic Web for the Life Sciences*, 4(3):196–206, 2006.
- [12] P Lambrix, H Tan, and Q Liu. SAMBO and SAMBOdtf results for the ontology alignment evaluation initiative 2008. In *Proceedings of the Third International Workshop on Ontology Matching*, pages 190–198, 2008.
- [13] P Lambrix, H Tan, and W Xu. Literature-based alignment of ontologies. In *Proceedings of the Third International Workshop on Ontology Matching*, pages 219–223, 2008.
- [14] M Mochol, A Jentsch, and J Euzenat. Applying an analytic method for matching approach selection. In *Proceedings of the Workshop on Ontology Matching*, 2006.
- [15] NF Noy. Semantic integration: A survey of ontology-based approaches. *Sigmod Record*, 33(4):65–70, 2004.
- [16] P Shvaiko and J Euzenat. A survey of schema-based matching approaches. *Journal on Data Semantics*, IV:146–171, 2005.
- [17] P Shvaiko and J Euzenat. Ten challenges for ontology matching. In *Proceedings of the 7th International Conference on Ontologies, Databases, and Applications of Semantics*, 2008.
- [18] H Tan, V Jakonienė, P Lambrix, J Aberg, and N Shahmehri. Alignment of biomedical ontologies using life science literature. In *Proceedings of the International Workshop on Knowledge Discovery in Life Science Literature, LNBI 3886*, pages 1–17, 2006.
- [19] H Tan and P Lambrix. A method for recommending ontology alignment strategies. In *Proceedings of the 6th International Semantic Web Conference, LNCS 4825*, pages 494–507, 2007.
- [20] H Tan and P Lambrix. SAMBO results for the ontology alignment evaluation initiative 2007. In *Proceedings of the Second International Workshop on Ontology Matching*, pages 236–243, 2007.
- [21] T Wächter, H Tan, A Wobst, P Lambrix, and M Schroeder. A corpus-driven approach for design, evolution and alignment of ontologies. In *Proceedings of the Winter Simulation Conference*, pages 1595–1602, 2006. Invited contribution.
- [22] P Wang and B Xu. Lily: ontology alignment results for OAEI 2008. In *Proceedings of the Third International Workshop on Ontology Matching*, pages 167–175, 2008.
- [23] S Zhang and O Bodenreider. Hybrid alignment strategy for anatomical ontologies: results of the 2007 ontology alignment contest. In *Proceedings of the Second International Workshop on Ontology Matching*, pages 139–149, 2007.
- [24] X Zhang, Q Zhong, J Li, and J Tang. RiMOM results for OAEI 2008. In *Proceedings of the Third International Workshop on Ontology Matching*, pages 182–189, 2008.