

Aligning the Parasite Experiment Ontology and the Ontology for Biomedical Investigations Using AgreementMaker

Valerie Cross¹, Cosmin Stroe², Xueheng Hu¹, Pramit Silwal¹, Maryam Panahiazar³, Isabel F. Cruz², Priti Parikh³, and Amit Sheth³

¹ Computer Science and Software Engineering Department
Miami University, Oxford, OH 45066
{crossv, hux3, silwalp}@muohio.edu

² ADVIS Laboratory, Department of Computer Science
University of Illinois at Chicago, Chicago, IL 60607[†]
{cstroel, ifc}@cs.uic.edu

³ Kno.e.sis Center, Department of Computer Science and Engineering
Wright State University
3640 Colonel Glenn Hwy
Dayton, OH 45435

Abstract. Tremendous amounts of data exist in life sciences along with many bio-ontologies. Though these databases contain important information about gene, proteins, functions, etc., this information is not well utilized due to the heterogeneous formats of these databases. Therefore, ontology alignment (OA) is now very critical for life science domain. Our work utilizes AgreementMaker for OA and describes results, difficulties faced in the process, and lessons learned. We aligned two real-world ontologies, the Parasite Experiment Ontology (PEO) and the Ontology for Biomedical Investigations (OBI). The former is more application-oriented and the latter is a reference ontology for any biomedical or clinical investigations. Our study led to several enhancements to AgreementMaker: annotation profiling, mapping provenance information, and tailored lexicon building. These enhancements, which are applicable to any OA system, greatly improved the alignment of these real world ontologies, producing 90% precision with 60% recall from the BSM^{lex+}, the Base Similarity Matcher, and 57% precision with 67% recall from the PSM^{lex+}, the Parametric String Matcher, both using lexicon lookup for synonyms. The mappings obtained through this study are posted on BioPortal site for public use.

[†] Partially supported by NSF Awards IIS-0513553 and IIS-0812258 and by the Intelligence Advanced Research Projects Activity (IARPA) via Air Force Research Laboratory (AFRL) contract number FA8650-10-C-7061. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, AFRL, or the U.S. Government.

Keywords: ontology alignment, biomedical ontologies, ontology profiling, mapping provenance, lexicons.

1 Introduction

Ontology alignment (OA) is a well-recognized need for bioinformatics [10] and biomedical researchers. Currently around 260 bio-ontologies exist on the NCBO BioPortal⁴ and a number of databases exist that contain information about genes and their sequence/functions, proteins and pathway information. This knowledge, all related but modeled with heterogeneous ontologies, if better connected would greatly benefit researchers. OA addresses this challenge by identifying semantically identical or related entities in different ontologies. The resulting alignments can then be used for exchanging data and information.

Over the past decade sophisticated algorithms to improve OA have been developed. The Ontology Alignment Evaluation Initiative⁵ (OAEI) [8] is a coordinated international effort providing standard methods for assessing the performance of OA systems. It has facilitated the advancement of OA techniques with its standard set of cases and evaluation methods that developers can use to learn from and improve their OA systems. The 2010 OAEI challenge consists of 4 different tracks including one subtrack to align two anatomy ontologies, the Adult Mouse Anatomy (MA) and the NCI Thesaurus Human Anatomy (HA). This subtrack is most relevant to our research.

This paper describes the process of, lessons learned from, and results of aligning two real-world biomedical ontologies, the Parasite Experiment Ontology⁶ (PEO) and the Ontology for Biomedical Investigations⁷ (OBI). These two present an interesting but different scenario from the OAEI anatomy subtrack. The PEO is being collaboratively developed as part of an NIH funded project to develop and deploy an ontology-driven semantic problem-solving environment for parasite research [7]. The PEO models experiment details and provenance information of experimental data. It is an application-oriented and more specific domain ontology than OBI, which describes biomedical investigations. In contrast, the OAEI Anatomy Track ontologies have many common entities since both describe the same domain (i.e., anatomy) for different species (mouse vs. human). Mapping more specific ontologies, like PEO, to a more general ontology, like OBI, is important to helping establish a common point of reference, which can serve to foster cooperation and interoperability among researchers. Due to different scope of PEO and OBI, it may not be possible to have PEO absorbed into OBI therefore being able to create mappings between the two ontologies becomes essential. Given the explosion of metadata available on the web, this study to align two related ontologies in the biomedical field has the potential to

⁴ <http://bioportal.bioontology.org>

⁵ <http://oaei.ontologymatching.org>

⁶ <http://bioportal.bioontology.org/ontologies/42093>

⁷ <http://www.obofoundry.org/cgi-bin/detail.cgi?id=obi>

impact not just the biomedical field but also any research field using semantic web technologies.

Aligning these ontologies requires selecting a suitable OA tool. Examining the 2010 OAEI results showed that AgreementMaker [2–4] ranked first of the nine OA systems used on the anatomy test case and performed successfully in other tracks. Since AgreementMaker had performed very well on the anatomy subtrack, had good developer support, and was readily available for our use and/or modification, it was selected. Aligning PEO and OBI exposed the need for more flexible and configurable OA algorithms. To address this need, the features of annotation profiling, mapping provenance information, and tailored lexicon building were developed in the ADVIS Laboratory, added to AgreementMaker and experimentally validated through the PEO and OBI alignment process.

In what follows, the two ontologies, an overview of AgreementMaker, the alignment process, and the results of this study are described.

2 Ontologies

The OBI, a part of the Open Biological and Biomedical Ontologies (OBO) Foundry,⁸ describes biological and clinical investigations (e.g., designs, protocols, instrumentation). It supports the integration of experimental data across various domains such as transcriptomics, proteomics, and metabolomics through its broader scope and controlled vocabulary. The OBI incorporates concepts from the Information Artifact Ontology⁹ (IAO) and also uses annotation properties defined in the IAO.

The PEO is currently not a part of the OBO Foundry, but is found in the NCBO BioPortal. It models provenance information of experiment protocols used in parasite research and other experiment details to support annotation and querying of parasite experiment data and other databases. It references the Parasite Lifecycle Ontology¹⁰ (PLO). Both PEO and OBI are represented in OWL but differ greatly in size and structure (110 vs. 3060 classes) while the OAEI MA and HA are more similar (2744 vs. 3304 classes).

3 AgreementMaker

The AgreementMaker OA system has many useful features including a well designed user interface, a diverse selection of matching algorithms (matchers), and mapping quality metrics to filter and combine the results of its matchers into a final best alignment. AgreementMaker provides an extensible architecture permitting new matching or weighting algorithms to be easily integrated and adjusted based on their performance. The user can easily evaluate, compare, and combine different strategies and matching results using its interface.

⁸ <http://www.obofoundry.org>

⁹ <http://code.google.com/p/information-artifact-ontology/>

¹⁰ <http://bioportal.bioontology.org/ontologies/39544>

The matchers fall into two main categories: concept-based, which employ multiple string similarity measures, and structural, which search for shared patterns in the hierarchical structure of the ontologies. Our work used the Base Similarity Matcher (BSM), the Advanced Similarity Matcher (ASM), the Parametric String-based Matcher (PSM), and the Vector-based Multi-Word Matcher (VMM). The BSM calculates the similarity between two concepts by comparing all the strings associated with those two concepts including the concept name, label, and comments. PSM is also a string-based matcher but more complicated since it uses a substring measure and an edit distance measure. VMM compiles a virtual document for every concept of an ontology by concatenating the strings of related concepts and annotations, transforms the resulting strings into TF-IDF vectors, and computes the similarity between those vectors using the cosine similarity measure [2]. AgreementMaker version 0.22 extended these string-based matchers by integrating two lexicons:

- (1) the Ontology Lexicon, built from synonym and definition annotations existing in the ontologies themselves, and
- (2) the WordNet Lexicon, created by starting with the ontology lexicon and adding any non-duplicated synonyms/definitions found in WordNet.

The Ontology Lexicon for each ontology (source and target) is built starting from a list of all the concepts (classes and properties) defined in an ontology. We then iterate through the list and inspect the definition of the concepts for synonym and definition annotations (*hasSynonym* and *hasDefinition* respectively). If these are found, they are added to the Ontology Lexicon entry for that concept.

The WordNet Lexicon is built starting with a previously built Ontology Lexicon. We then perform a WordNet lookup for every entry in the Ontology Lexicon and add any non-duplicated synonyms and definitions to that concept's entry in the WordNet Lexicon. It must be noted that the Ontology Lexicon and the WordNet Lexicon are kept in separate data structures. While the Ontology Lexicon contains information directly defined in the ontology, the entries in the WordNet Lexicon can contain ambiguous information. This ambiguity must be taken into account by the matching algorithms.

The matchers using the lexicons in their algorithms are annotated with a *lex* superscript, as in BSM^{lex} , PSM^{lex} , and VMM^{lex} [4]. The Linear Weighted Combination (LWC) matcher [3] produces a single combined alignment by using mapping quality measures [3, 5] to choose the best mappings from each matcher.

4 Aligning PEO with OBI

The alignment process between the PEO and OBI is performed on the import closure of the resources, i.e., taking into account all the files each imports. The first alignment used AgreementMaker (version 0.22) with the 2010 OAEI anatomy configuration and produced only two mappings due to an inconsistency in entity descriptions of the PEO and OBI. PEO's URIs use a textual fragment identifier (e.g., <http://knoesis.wright.edu/ParasiteExperiment>).

owl#transfection), while OBI's entities use numerical identifiers (e.g., http://purl.obolibrary.org/obo/OBI_0600060). The PEO's use of the `rdfs:label` field does not follow the specification guidelines since when PEO happens to use this field (only 19.1% of the PEO classes have a label), it contains a PLO identifier. For the OBI, on the other hand, it uses the `rdfs:label` field to contain a descriptive string on almost 100% of its classes. The comment field for the PEO is used on 99% of its classes and typically provides a definition. The OBI only uses the comment field on about 4% of its classes. Although some common annotations exist between them, either PEO or OBI has low coverage. For example, OBI has high coverage for label annotations while PEO has high coverage for comment annotations. This heterogeneity and matchers aligning the same annotations to each other (i.e., class ID with class ID, label with label, etc.) resulted in almost no alignment.

The OAEI ontologies, in comparison, both use label and *hasRelatedSynonym* annotations and have descriptive local names. Experimental results show lexicon synonyms and definitions greatly benefited the MA and HA alignment [4]. Their usage for the PEO to OBI alignment could also be beneficial.

As PEO and OBI do not rely on the same metadata set to describe their entities, *annotation profiling* was implemented to allow users to select and combine different annotations of the source or target ontology for the alignment. Several synonym-like annotations exist but are not found across the ontologies. The issue becomes how to match synonym annotations between the two ontologies and how to handle the mismatch in the usage of identifier, label, and definition fields. Figure 1 illustrates this feature with PEO as the source and OBI as the target. All existing annotations in the ontologies are shown. The user selects which to use for aligning. For our experiments, the BSM and ASM matchers take advantage of annotation profiling.

Evaluating mapping results requires laboriously searching the ontologies to find descriptive names, labels, definitions, etc. This task motivated implementing the mapping provenance feature. Mapping provenance information can be automatically generated on a mapping-by-mapping basis for matchers supporting this feature. The provenance information can be interactively viewed, saved to the alignment result, and later imported.

The provenance information, an example shown in Figure 2, provides descriptive information for both entities that may come from the names, labels, definitions, and their other associated annotations. Given provenance information, a user does not have to look up what entity OBI_0302722 actually represents but can easily see it is "organ." Provenance capabilities are provided for ASM, BSM, PSM and VMM.

Finally, since AgreementMaker's previous lexicon builder used a fixed name for the synonym and definition annotations (*hasSynonym* and *hasDefinition*), it was modified to exploit the synonym annotations in PEO and OBI by having the user choose the annotation names used to create the lexicons. For synonyms, OBI does not use *hasSynonym* but uses IAO annotation properties IAO_0000111 ("editor preferred term") and IAO_0000118 ("alternative term"); they serve the

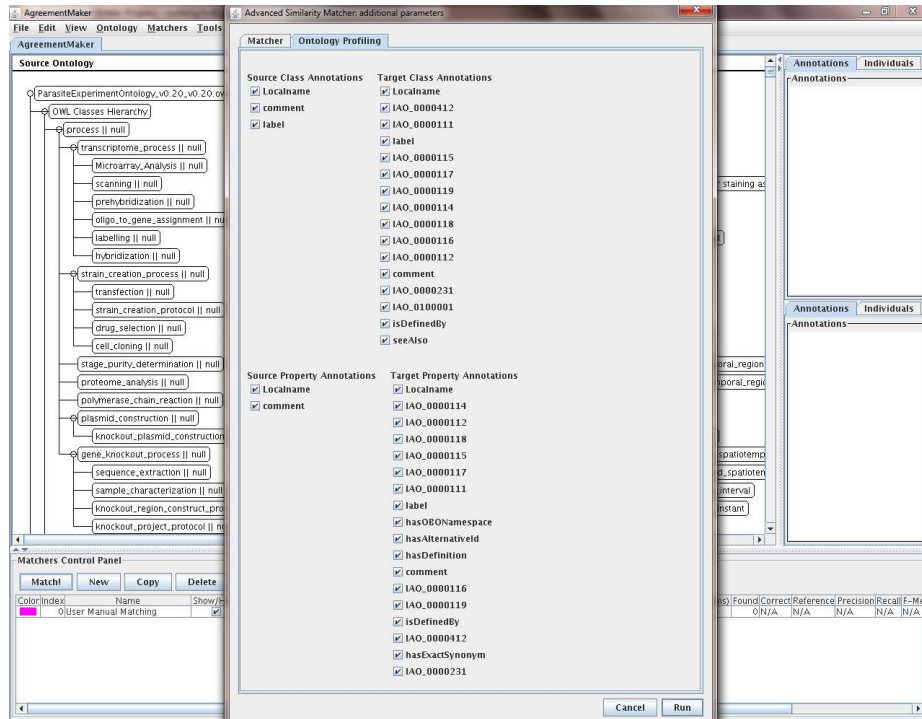


Fig. 1: Profiling interface added to AgreementMaker.

same function as synonyms for the OBI. The PEO does not use synonyms but uses the comment annotation for a definition in most cases.

5 Experimental Results, Evaluation, and Discussion

Our domain expert evaluated the possible mappings between the entities of the PEO and OBI ontologies and produced a set of mappings. The expert provided a confidence score in the range (0.0 1.0] for each mapping. A mapping with a confidence score of 0.8 or higher is considered a correct mapping and included in the reference alignment. A total of 30 PEO to OBI reference mappings were produced. The number of mappings is low since the PEO is a more specific ontology than OBI, and only their overlapping concepts can be mapped. This reference alignment and the data for these experiments can be found on the Kno.e.sis website in the section “Alignment of PEO and OBI.”¹¹ Each matcher was evaluated against the reference alignment to compute precision, recall, and F-measure.

¹¹ http://wiki.knoesis.org/index.php/Parasite_Experiment_ontology

```

<Cell>
  <entity1 rdf:resource="http://paige.ctegd.uga.edu/ParasiteLifecycle.owl#organism"/>
  <entity2 rdf:resource="http://purl.obolibrary.org/obo/OBI_0302722"/>
  <measure rdf:datatype="http://www.w3.org/2001/XMLSchema#float">1.0</measure>
  <relation>=</relation>
  <provenance>sim("organism", "organ") = 1.0</provenance>
</Cell>

```

Fig. 2: An example of mapping provenance information.

Given a reference alignment R and a computed alignment A , the precision of alignment A is calculated as

$$\text{Precision}(A, R) = \frac{|A \cap R|}{|A|}$$

and the recall is calculated as

$$\text{Recall}(A, R) = \frac{|A \cap R|}{|R|}$$

where $|A|$ represents the number of mappings in alignment A . The F-Measure is the harmonic mean of precision and recall and is calculated as

$$\text{F-Measure}(A, R) = \frac{2 \cdot \text{Precision}(A, R) \cdot \text{Recall}(A, R)}{\text{Precision}(A, R) + \text{Recall}(A, R)}.$$

Once the annotation profiling feature was implemented for the ASM, an alignment using all the annotations declared in each ontology was produced. The reasoning was that considering all the available annotation information in the matching process would lead to the best possible result, labeled ASM_{ALL} in Table 1. It shows an overall inconclusive alignment, due to low precision and medium recall. This experiment shows that matching ontology entities without discriminating between their annotations is not a viable approach—unless their annotations are semantically compatible, as seen for the OAEI ontologies.

After experimenting with ASM_{ALL} , we decided to use only the most useful and compatible annotations, reasoning this approach should give better results. Since the ASM computes alignments using a string matching similarity more suitable for short strings and compound words, the next experiment labeled ASM_{SYN} used only the synonym annotation properties declared in the ontologies. For example, the local names of the OBI ontology were not used since they are mostly ID numbers. Instead for the OBI, `IA0_0000111`, `IA0_0000118`, and `label` were used for class annotations and `hasExactSynonym` was added to these for the property annotations. The resulting alignment contained fewer mappings but only reduced the number of correct mappings by one mapping, leading to a 13% increase in precision while losing only 3% recall. Our reasoning was correct; however, using only a string matching algorithm was not enough to match the ontologies.

Matcher	Precision	Recall	FMeasure
ASM _{ALL}	0.25	0.53	0.34
ASM _{SYN}	0.38	0.50	0.43
BSM ^{lex+}	0.90	0.60	0.72
PSM ^{lex+}	0.57	0.67	0.62
VMM ^{lex+}	0.50	0.03	0.06
LWC	0.49	0.70	0.58
Combined	0.26	0.80	0.39

Table 1: Precision, Recall and FMeasure results for each alignment experiment.

Next, the lexicons based on the modification to lexicon building process as previously described were incorporated. All synonym and definition annotations in the ontologies were selected for use in the building of the lexicons. The BSM^{lex+}, PSM^{lex+}, and the VMM^{lex+} use the user customized lexicons in the matching process.

As suspected, the lexicons greatly improved the alignment quality. BSM^{lex+} achieves high precision with good recall; a similar performance was observed when matching the OAEI ontologies. PSM^{lex+} further improves recall by applying more sophisticated string matching algorithms. However more incorrect mappings are produced. VMM^{lex+}, which uses definition annotations, found two mappings but only one was correct. All other matchers found the same correct mapping.

With these promising individual matcher results, the next experiment combined these individual alignments into one final alignment. LWC combined the ASM_{SYN}, PSM^{lex+}, VMM^{lex+}, and BSM^{lex+} into one alignment result using the “local confidence” quality metric [3] and a mapping selection threshold of 0.5.

The LWC alignment in Table 1 has the best recall but cannot avoid including incorrect mappings, leading to lower precision. The LWC combines alignments by applying a quality weight to each mapping in the input alignment; if correct mappings are only slightly better than other incorrect mappings, the combined alignment can be less precise. Although the input matchers produce good alignments, improvement is needed to better discern correct from incorrect mappings. This result, not observed for the OAEI ontologies, is most likely due to the high level of heterogeneity between PEO and OBI annotations. Aligning the PEO and OBI, as real-world test cases, showed our lexicon based matching algorithms greatly improve alignment results, but more domain specific lexicons are needed to aid in the disambiguation of very similar entities.

Finally, to gauge LWC performance in producing a better alignment than just a simple combination of the input alignments, the input mappings to LWC were examined for overlap. Figure 5 shows that LWC usually can choose most of the correct mappings from each matcher. To further evaluate LWC, we manually produced a “Combined” alignment consisting of all distinct mappings produced by the six experiments, labelled “Combined” in Table 1. The Combined align-

	ASM _{ALL}	ASM _{SYN}	BSM ^{lex+}	PSM ^{lex+}	LWC
ASM _{ALL}	16	14	13	13	15
ASM _{SYN}	-	15	13	14	14
BSM ^{lex+}	-	-	18	16	17
PSM ^{lex+}	-	-	-	20	19
LWC	-	-	-	-	21

Fig. 3: *

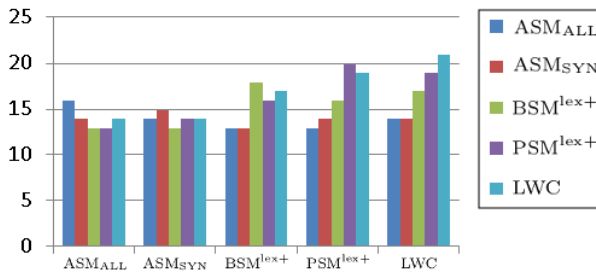


Fig. 4: *

(a)

(b)

Fig. 5: The number of overlapping correct mappings for each pair of matchers shown in (a) tabular and (b) graphical formats, showing that LWC combination of multiple matching algorithms produces a better alignment than a standalone matching algorithm.

ment was then compared with the reference alignment. The Combined alignment is only 10% better than the one produced by LWC in terms of recall but precision is 23% lower. This result shows that LWC can indeed discriminate between correct and incorrect mappings because of its use of an intrinsic quality measure [3], albeit not perfectly. Research may be needed to develop a more robust quality measure.

6 Conclusions and Future Research

Aligning the PEO and OBI exposes the problem of heterogeneous annotations in ontologies. This problem can be managed by increasing the flexibility of the state of the art matching algorithms. Our implementation of annotation profiling, mapping provenance information, and custom lexicons contribute greatly to providing this flexibility.

AgreementMaker’s past approach of extending matching algorithms using lexicons [4] has also been validated since the best results are produced by matchers using lexicons (e.g., BSM^{lex+} in Table 1). However, including more lexicons such as UMLS [1] needs to be investigated in order to achieve even better results. More lexicons would allow matchers to better disambiguate entities and, thus, improve the combination of the matching results.

Our current approach to managing heterogeneity relies on user selection of relevant annotations for the matching process. Annotation profiling and mapping provenance information support a domain expert in this process; however, research is needed to automatically identify semantically compatible annotations, possibly by applying established ontology evaluation metrics [6]. The heterogeneity present in real-world ontologies must be further addressed so that the

OA process can foster cooperation and interoperability between researchers and organizations.

7 Acknowledgements

We thank Dr. Jie Zheng, postdoc at University of Pennsylvania and one of the primary developers of OBI for her initial support and help with mappings.

References

1. O. Bodenreider. The Unified Medical Language System (UMLS): integrating biomedical terminology. *Nucleic Acids Research*, 32(Database-Issue):267–270, 2004.
2. I. F. Cruz, F. Palandri Antonelli, and C. Stroe. AgreementMaker: Efficient Matching for Large Real-World Schemas and Ontologies. *PVLDB*, 2(2):1586–1589, 2009.
3. I. F. Cruz, F. Palandri Antonelli, and C. Stroe. Efficient Selection of Mappings and Automatic Quality-driven Combination of Matching Methods. In Shvaiko et al. [9].
4. I. F. Cruz, C. Stroe, M. Caci, F. Caimi, M. Palmonari, F. Palandri Antonelli, and U. C. Keles. Using AgreementMaker to align ontologies for OAEI 2010. In Shvaiko et al. [8].
5. C. Joslyn, P. Paulson, and A. M. White. Measuring the Structural Preservation of Semantic Hierarchy Alignment. In Shvaiko et al. [9].
6. M. Mochol and A. Jentzsch. Towards a Rule-Based Matcher Selection. In A. Gangemi and J. Euzenat, editors, *EKAW*, volume 5268 of *Lecture Notes in Computer Science*, pages 109–119. Springer, 2008.
7. S. S. Sahoo, D. B. Weatherly, R. Mutharaju, P. Anantharam, A. P. Sheth, and R. L. Tarleton. Ontology-Driven Provenance Management in eScience: An Application in Parasite Research. In R. Meersman, T. S. Dillon, and P. Herrero, editors, *OTM Conferences (2)*, volume 5871 of *Lecture Notes in Computer Science*, pages 992–1009. Springer, 2009.
8. P. Shvaiko, J. Euzenat, F. Giunchiglia, H. Stuckenschmidt, M. Mao, and I. Cruz, editors. *Proceedings of the 5th International Workshop on Ontology Matching (OM-2010) collocated with the 9th International Semantic Web Conference (ISWC-2010) Shanghai, China, November 7, 2010*, volume 689 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2010.
9. P. Shvaiko, J. Euzenat, F. Giunchiglia, H. Stuckenschmidt, N. F. Noy, and A. Rosenthal, editors. *Proceedings of the 4th International Workshop on Ontology Matching (OM-2009) collocated with the 8th International Semantic Web Conference (ISWC-2009) Chantilly, USA, October 25, 2009*, volume 551 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2009.
10. H. Tan, V. Jakoniene, P. Lambrix, J. Aberg, and N. Shahmehri. Alignment of Biomedical Ontologies Using Life Science Literature. In E. G. Bremer, J. Hakenberg, E.-H. Han, D. P. Berrar, and W. Dubitzky, editors, *KDLL*, volume 3886 of *Lecture Notes in Computer Science*, pages 1–17. Springer, 2006.