

NLM Anatomical Ontology Alignment System Results of the 2006 Ontology Alignment Contest

Songmao Zhang¹, Ph.D., Olivier Bodenreider²

¹ Institute of Mathematics, Academy of Mathematics and Systems Science
Chinese Academy of Sciences, Beijing, P. R. China

² U.S. National Library of Medicine, NIH, Bethesda, Maryland, USA

¹ smzhang@math.ac.cn, ² olivier@nlm.nih.gov

Abstract. An ontology is a formal representation of a domain modeling the entities in the domain and their relations. When a domain is represented by multiple ontologies, there is need for creating mappings among these ontologies in order to facilitate the integration of data annotated with and reasoning across these ontologies. The objective of this paper is to present our experience in aligning two large anatomical ontologies and to reflect on some of the issues and challenges encountered along the way. The anatomical ontologies under investigation are the Foundational Model of Anatomy (FMA) and GALEN. Our approach to aligning concepts is automatic, rule-based, and operates at the schema level, generating mostly point-to-point mappings. It uses a combination of domain-specific lexical techniques and structural and semantic techniques (to validate the mappings suggested lexically). It also takes advantage of domain-specific knowledge (lexical knowledge from external resources such as the Unified Medical Language System, as well as knowledge augmentation and inference techniques). Overall, the lexical alignment followed by structural validation identified 3,029 pairs of equivalent concepts in the FMA and GALEN, accounting for about 4% of all FMA concepts and 32% of all GALEN concepts.

1 Presentation of the system

Over the past four years, as part of the Medical Ontology Research project at the U.S. National Library of Medicine, we have developed domain knowledge-based techniques for aligning large anatomical ontologies, with the objective of exploring approaches to aligning representations of anatomy differing in formalism, structure, and domain coverage. We started by aligning concepts point-to-point in two large ontologies of human anatomy, using lexical and structural techniques [1]. We later tested these techniques on other pairs of anatomical ontologies, both within and across species [2, 3]. We also investigated the complex alignment of groups of concepts [4] and that of relationships [5]. Finally, we investigated the possibility of deriving the indirect alignment of two ontologies through their direct alignment to a reference ontology [6]. The objective of this paper is to present our experience in aligning two large anatomical ontologies point-to-point and to reflect on some of the issues and challenges encountered along the way. In particular, we want to show the importance

of domain-specific knowledge in our alignment strategies. Many features of our system are specific to the domain of anatomy, which is why ‘anatomy’ is the only OAEI data set to which we applied our methods.

1.1 State, purpose, general statement

The approach to aligning anatomical ontologies presented here is automatic, rule-based, and operates at the schema level, generating point-to-point mappings. It uses a combination of domain-specific lexical techniques (to map entities at the element, not instance level) and structural and semantic techniques (to validate the mappings suggested lexically). It also takes advantage of domain-specific knowledge (lexical knowledge from external resources such as the Unified Medical Language System[®] (UMLS[®]) [7], as well as knowledge augmentation and inference techniques).

The many ontology alignment systems available include PROMPT [8], CUPID [9], FCA-Merge [10], HCONE-Merge [11], and GLUE [12]. With AnchorPrompt [13], we share the notion of “anchor” (i.e., a pair of related terms across ontologies, established by lexical similarity in our case) and the use of shared paths between anchors across ontologies to validate the similarity among related terms. Therefore, AnchorPrompt is undoubtedly the system to which our approach is the most closely related. The major differences between AnchorPrompt and our approach can be summarized as follows. AnchorPrompt creates a sophisticated similarity score based on path length and other features. In contrast, we use a simpler validation scheme based on paths restricted to combinations of taxonomic and partitive relations, suitable for the anatomical domain. Unlike AnchorPrompt, our approach does not rely on path length and is therefore less sensitive to differences in granularity between ontologies. AnchorPrompts requires ontologies to be accessible in the frame-based Open Knowledge Base Connectivity (OKBC) protocol, while our approach is not specific to any particular formalism.

1.2 Specific techniques used

We identify one-to-one concept mappings between the Foundational Model of Anatomy (FMA) and GALEN using lexical resemblance between concept names and then validate the mappings through shared hierarchical paths among concepts across ontologies.

Lexical alignment

The lexical alignment identifies shared concepts across ontologies based on lexical similarity between concept names. Both preferred concept names and synonyms, if any, are used in the lexical alignment process. Lexical similarity is assessed through exact match and after normalization. The normalization program distributed with the UMLS provides a linguistically-motivated model for lexical resemblance adapted to

the specificity of biomedical terms, abstracting away from minor differences in terms including case, hyphen, inflection and word order variations [14].

Concepts exhibiting similarity at the lexical level across ontologies are called anchors, as they are going to be used as reference concepts in the structural validation and for comparing associative relationship. Additional anchors are identified through synonymy in an external resource: the Unified Medical Language System (UMLS). More specifically, two concepts across ontologies are considered anchors if their names are synonymous in the UMLS Metathesaurus (i.e., if they name the same concept) and if the corresponding concept is in the anatomy domain (i.e., has a semantic type related to *Anatomy*).

Examples of anchors, shown in **Figure 1**, include the concepts *Cardiac valve* in FMA and *Valve in heart* in GALEN, identified as anchor concepts because *Cardiac valve* has *Valve of heart* as a synonym in FMA and *Valve in heart* matches *Valve of heart* after normalization. Additionally, *Fibrous ring of mitral valve* (with synonym *Mitral anulus*) in the FMA and *Mitral ring* in GALEN form an anchor because *Mitral anulus* and *Mitral ring* are synonyms, i.e., they are both names for the concept *Structure of anulus fibrosus of mitral orifice* in the UMLS.

Structural validation

In the structural validation of the lexical alignment, the first step is to acquire the semantic relations explicitly represented in the ontologies. Inter-concept relationships are generally represented by semantic relations $\langle \mathbf{c}_1, r, \mathbf{c}_2 \rangle$, where the relationship r links concepts \mathbf{c}_1 and \mathbf{c}_2 . Because they form the backbone of anatomical ontologies and are therefore more likely to be represented consistently across ontologies, hierarchical relationships only are considered at this step. These relationships are *IS-A* and *PART-OF*, along with their inverses *INVERSE-IS-A* and *HAS-PART*, respectively. Having extracted the relations explicitly represented in the ontologies, we then normalize the representation of the relations in each ontology in order to facilitate structural comparisons across ontologies. We first complement the hierarchical relations represented explicitly with their inverses as necessary. Implicit semantic relations are then extracted from concept names (augmentation) and various combinations of hierarchical relations (inference). Augmentation and inference are the two main techniques used to acquire implicit knowledge from the FMA and GALEN. For a detailed analysis of the contribution of each technique, the interested reader is referred to [15].

Complementation. As partial ordering relationships, hierarchical relationships are anti-symmetric. *IS-A* and *PART-OF* have inverse relationships, *INVERSE-IS-A* and *HAS-PART*. Except for *IS-A*, not every relation is represented bidirectionally. For example, $\langle \textit{Right breast proper}, \textit{HAS-PART}, \textit{Right mammary gland} \rangle$ is explicitly represented in the FMA but its inverse relation is missing. In canonical anatomy, the inverse relations are essentially always valid, although this is not necessarily the case in the real world [16]. For the sole purpose of aligning ontologies, in order to facilitate the comparison of paths between anchors across ontologies, we complement the FMA and GALEN with the inverse relations that are not explicitly represented. For

example, we generated the relation *<Right mammary gland, PART-OF, Right breast proper>*.

Augmentation attempts to represent with relations knowledge that is otherwise embedded in the concept names. Augmentation is based on linguistic phenomena, such as the reification of partitive relations. In this case, a relation *<P, PART-OF, W>* is created between concepts **P** (the part) and **W** (the whole) from a relation *<P, IS-A, Part of W>*, where the concept **Part of W** reifies, i.e., embeds in its name, the *PART-OF* relationships to **W**. For example, *<Neck of femur, PART-OF, Joint>* was added from the relation *<Neck of femur, IS-A, Component of joint>*, where the concept *Component of joint* reifies a specialized *PART-OF* relationship. Examples of augmentation based on other linguistic phenomena include *<Sweat gland, IS-A, Gland>* (from the concept name *Sweat gland*) and *<Extensor muscle of leg, PART-OF, Leg>* (from the concept name *Extensor muscle of leg*). The semantics of nominal modification generally corresponds to subsumption (e.g., the head noun *gland* modified by *sweat* is a hypernym of *gland*). In contrast, the semantics of prepositional clauses introduced by *of* is not necessarily a partitive relation (e.g., *glass of wine* is not part of *wine*). Here, domain knowledge was required to assess what relations can be automatically extracted with high accuracy in the particular context of anatomical terms. We determined that partitive relations could be accurately created from prepositional clauses introduced by *of* in anatomical terms containing no other prepositions.

Inference generates additional semantic relations by applying inference rules to the existing relations in order to facilitate the comparison of paths between anchors across ontologies. These inference rules, specific to this alignment, represent limited reasoning along the *PART-OF* hierarchy, generating a partitive relation between a specialized part and the whole or between a part and a more generic whole. For example, *<First tarsometatarsal joint, PART-OF, Foot>* was inferred from the relations *<First tarsometatarsal joint, IS-A, Joint of foot>* and *<Joint of foot, PART-OF, Foot>*. Analogously, *<Interphalangeal joint of thumb, PART-OF, Finger>* was inferred from the relations *<Interphalangeal joint of thumb, PART-OF, Thumb>* and *<Thumb, IS-A, Finger>*.

With these explicit and implicit semantic relations, the structural validation identifies structural similarity and conflicts among anchors across ontologies. Structural similarity, used as **positive structural evidence**, is defined by the presence of common hierarchical paths among anchors across ontologies, e.g., *<c₁, PART-OF, c₂>* in one ontology and *<c₁', PART-OF, c₂'>* in another where *{c₁, c₁'}* and *{c₂, c₂'}* are anchors across ontologies¹. The anchor concepts *Cardiac valve* in FMA and *Valve in heart* in GALEN, presented earlier, received positive structural evidence because they share hierarchical paths to some of the other anchors across ontologies. For example, as illustrated in **Figure 1**, *Cardiac valve* is related to *Heart* (*PART-OF*), to *Mitral valve* (*INVERSE-IS-A*) and to *Mitral ring* (*HAS-PART*).

Conflicts, on the other hand, are used as **negative structural evidence**. The first type of conflict is defined by the existence of hierarchical paths between the same anchors across ontologies going in opposite directions, e.g., *<c₁, PART-OF, c₂>* in one ontology and *<c₁', HAS-PART, c₂'>* in another. The second type of conflict is based on the disjointness of top-level categories across ontologies (i.e., semantic constraints).

¹ The transitive closure of hierarchical relation greatly facilitates paths comparison across ontologies, because complex paths between anchors are represented by a single relation.

For example, *Point* in FMA is a kind of *Dimensional entity*, while *Pointing* in GALEN is a *Voluntary movement of upper extremity*, which is a *Process*. *Dimensional entity* and *Process* being disjoint top-level categories, the two concepts *Point* in the FMA and *Pointing* in GALEN are considered semantically distinct, which prevents them from being aligned although they are lexically equivalent (after normalization).




Figure 1. Structural validation following lexical alignment

1.3 Adaptations made for the evaluation

Extracting FMA and GALEN from OWL Full. In previous alignment experiments involving the FMA and GALEN, we purposely stayed away from any particular formalism in order to avoid distorting the source ontologies during the conversion process. Instead, we simply extracted <concept, relationship, concept> triples from the two ontologies in their original formalism (frames in Protégé for the FMA and the description logic language GRAIL for GALEN) using the Application Programming Interface provided with the ontology. In the OAEI experiment, the FMA and GALEN were converted into OWL Full by the organizers. Since our alignment approach is not designed to specifically take advantage of the OWL formalism, we again extracted concept names and <concept, relationship, concept> triples from the class definitions in OWL.

More specifically, for the FMA, we used *rdf:ID* to identify concepts, *rdf:label* and *Preferred_name* to acquire concept names, and *rdfs:subClassOf* to acquire taxonomic relations. The various partitive relations represented in the FMA (e.g., *part_of*, *constitutional_part_of*, *regional_part_of*) were acquired using the corresponding properties. An inverse property is specified for 29 properties in the OWL file. The other properties were not used in the alignment.

Because of the deep nesting exhibited by the GALEN file in OWL, we performed the extraction not from the OWL file itself, but from the *.pont* and *.pins* files resulting from its conversion into the CLISP format by the OWL plugin in Protégé. The *.pins* file contains all the information we needed and is the only one we ended up using. (This file is available for download at: <http://mor.nlm.nih.gov/pubs/supp/2006-oeai-sz/OpenGALEN.pins>). We used *rdfs:ID* to both identify concepts and acquire concept names (after tokenization), and *rdfs:subClassOf* to acquire taxonomic relations. The various partitive relations represented in GALEN (e.g., *isPartOf*, *isComponentOf*, *IsSurfaceDivisionOf*) were acquired using the corresponding properties. No inverse properties are specified in the OWL file. Inverse properties were added manually in 13 cases for alignment purposes. The other properties were not used in the alignment.

Grouping fine-grained partitive relationships. For alignment purpose in this study, we consider as only one *PART-OF* relationship (with *HAS-PART* as its inverse) the various kinds of partitive relationships present in the FMA (e.g., *part_of*, *constitutional_part_of*, *regional_part_of*) and GALEN (e.g., *isPartOf*, *isComponentOf*, *IsSurfaceDivisionOf*).

Adding disjointness axioms between top-level classes. In order to identify semantic mismatches, we added pairwise disjointness axioms between their top-level classes across the two ontologies. For example, as shown earlier, because we define *Dimensional entity* in the FMA and *Process* in GALEN as being disjoint top-level categories, *Point* in FMA – a kind of *Dimensional entity* – and *Pointing* in GALEN – a kind of *Process* – are prevented from being aligned although they are lexically equivalent after normalization.

1.5 Link to the set of provided alignments (in align format)

The result of our alignment for the ‘anatomy’ data set is available at: <http://mor.nlm.nih.gov/pubs/supp/2006-oeai-sz/Zhang&Bodenreider.rdf> in the format specified by the OAEI organizers at: <http://oeai.ontologymatching.org/2006>.

2 Results

The only data set for which we report results is ‘anatomy’. Overall, the lexical alignment followed by structural validation identified 3,029 pairs of equivalent concepts in the FMA and GALEN, accounting for about 4% of all FMA concepts and 32% of all GALEN concepts.

Acquiring concept names and relations for the FMA and GALEN. The main characteristics of the data sets under investigation are listed in **Table 1**, including the number of classes, concept names, and types of partitive relationships. The number of *IS-A* and partitive relations extracted from the OWL file and generated by complementation, augmentation and inference is shown in **Table 2**. Not surprisingly, in both ontologies, a majority of relations come from inference, which performs similarly to a transitive closure of the hierarchical relations. Also listed in **Table 2** is the small number of relations removed from the ontologies because they create cycles.

Table 1. Main characteristics of the FMA and GALEN

#	FMA	GALEN
Concepts	72,560	9,566
Synonyms	44,597	0
Anonymous concepts	0	1,035
Part-of relationships	7	8
Has-part relationships	7	8
Inter-concept associative relationships	67	13
Datatype properties defined in the owl file	19	2

Table 2. Number of relations in the FMA and GALEN

Types of relations	FMA	GALEN
Explicitly represented is-a	72,560	18,091
Explicitly represented partitive relations	101,161	12,830
Explicitly represented associative relations	48,804	8,341
Complemented inverse-isa	72,560	18,091
Complemented partitive relations	3,561	4,364
Complemented associative relations	11,697	762
Removed because of cycles	- 40	- 2
Augmented	169,378	29,780
Inferred	5,169,034	243,436
Total	5,648,715	335,693

Table 3. Results of structural validation for the FMA-GALEN alignment

Structural evidence		3,132 anchors		
Positive evidence	Shared paths to other anchors(same type)	2056	3029	96.7%
	Shared paths to other anchors (“compatible”)	973		
Negative evidence	Conflicting paths to other anchors	9	22	0.7%
	Semantic disjointness	13		
No evidence	No paths to other anchors	61	81	2.6%
	No shared paths to other anchors	20		

Lexical alignment. 3,132 matching anchor concepts were identified lexically, accounting for about 4.3% of the FMA concepts and 32.7% of GALEN concepts. Of these, 378 anchors were identified through UMLS synonymy.

Structural validation. For the 3,132 anchors, the presence of paths to other anchors is searched in both ontologies, as well as the existence of conflicts is assessed. This information is used for the structural validation of the alignment, the results of which are summarized in **Table 3**.

Anchors with positive structural evidence. 96.7% of the anchors receive positive evidence, most of them sharing hierarchical paths of the same type (e.g., *Cardiac valve*

in FMA and *Valve in heart* in GALEN, presented earlier). An example of shared “compatible” hierarchical relations is the anchor *Pelvic fascia*. In both ontologies, this concept is linked to *Visceral pelvic fascia*, but, although going in the same direction, the relationship is *INVERSE-ISA* in GALEN and *HAS-PART* in FMA. For alignment purposes, sharing compatible hierarchical relations is deemed a sufficient condition. Anchors with positive structural evidence are presented in our result file with “=” as value for “relation” and 1.0 as value for “measure”.

Anchors with negative structural evidence. About 0.7% of the anchors represent conflicts between the two ontologies. For example, the relationship between the anchors *Apex of bladder* and *Urinary bladder* is *PART-OF* in GALEN but *HAS-PART* in FMA. Another type of conflict is represented by the semantic incompatibility between *Point* (the dimensional entity) in the FMA and *Pointing* (the process) in GALEN presented earlier. Anchors with negative structural evidence are presented in our result file with “!=” as value for “relation” and 1.0 as value for “measure”.

Anchors with no structural evidence. 2.6% of the anchors do not receive any structural evidence. For example, although linked to *Body Cavity (is-a)* and *Peritoneal Cavity (inverse-isa)* in GALEN, *Serous Cavity* has no connections to other anchors in FMA. The absence of *any* paths to other anchors represents about two thirds of the cases. The remaining cases correspond to the absence of *shared* paths to other anchors across ontologies. For example, although *Eyebrow* is linked to two anchors in FMA (e.g., *Set of hairs*), and two in GALEN (e.g., *Face*), none of these paths are shared across ontologies. Anchors with no structural evidence are presented in our result file with “=” as value for “relation” and 0.5 as value for “measure”.

3 General comments

3.1 Comments on the results

Overall, the results we obtained on this ‘anatomy’ data set are essentially similar to the results obtained in previous experiments from slightly different versions of these two large anatomical ontologies [1]. The main difference is that, in the past, we mapped all of GALEN – not just its anatomy subset – to the FMA, leading to increased ambiguity between anatomical and non anatomical concept names. In contrast to what most teams reported at OAEI 2005, for us, this ‘anatomy’ data set was actually slightly *less* complex than what we are used to.

3.2 Discussions on the way to improve the proposed system

The strengths and weaknesses of our system have been analyzed in previous papers [17]. The major difference with other systems is that we take advantage of domain knowledge throughout the mapping process. For example, we use specific tools and resources, including normalization techniques developed for biomedical terms and synonyms from the Unified Medical Language Systems. We also developed

techniques specific to the anatomical ontologies under investigation in order to represent explicitly relations implicitly present in these ontologies. These additional synonyms and relations increase the chances of identifying matches both at the lexical and structural level. Conversely, because most of these techniques are specific to anatomical ontologies, our system is unlikely to perform well on other types of ontologies. Similar techniques would have to be developed for other domains in order to obtain similar levels of performance. Finally, specifying disjointness between top-level categories was more useful in previous experiments when all of GALEN was aligned with the FMA. For example, because we define *Anatomical structure* in the FMA and *Inert solid structure* in GALEN as being disjoint top-level categories, *Nail* in FMA – a kind of *Anatomical structure* – and *Nail* in GALEN – a kind of *Surgical fixation device* – were prevented from being aligned although they are lexically identical.

3.3 Comments on the OAEI 2006 procedure

In our opinion, one of the most controversial aspects of the OAEI procedure is the decision of the organizers to convert the anatomical data sets under investigation – the FMA and GALEN – from their original formalism to OWL Full. On the one hand, the intuition is that having to deal with only one formalism will make it easier for participants to align these two massive ontologies. On the other, the result of the conversion is dependent on the original modeling. The FMA's native environment is frame-based and makes heavy use of metaclasses [18], while GALEN's original formalism is based on the description logic language GRAIL – the GALEN Representation and Integration Language [19]. The mechanical conversion of these two ontologies to OWL Full is therefore likely to result in differing representations, e.g., because of the difference in the use of metaclasses and instances.

In practice, the OAEI participants cannot reliably take advantage of such features of the underlying representation for identifying equivalent entities across ontologies. This issue was illustrated by the example given by Kalfoglou & Hu in their attempt to align the FMA and GALEN in the 2005 OAEI campaign [20]. They specifically question the mapping established by [17] between *Pancreas* in FMA and *Pancreas* in GALEN on the basis that it “is defined as a class in FMA ... whereas in GALEN (OpenGALEN) [it is defined] as an instance of [the] class “Body Cavity Anatomy””. Looking beyond formalism, it is quite clear to domain experts that the entities referred to by *Pancreas* in the FMA and GALEN are equivalent. For example, in both ontologies, they are defined as a kind of organ and have the same parts, including exocrine and endocrine pancreas, and head, neck, body and tail of pancreas. The containment relations are slightly different in the FMA (retroperitoneal space) and GALEN (abdominal cavity), but compatible for alignment purposes.

As shown recently, converting the entire FMA to OWL is a difficult exercise because it requires some understanding of both the modeling strategy and the domain, which is beyond what can be expected from the mechanical translation provided by the “export to OWL” function in Protégé [21]. However, the mechanical conversion was sufficient for us to reliably extract the information we needed (namely, concept names and taxonomic and partitive relations). Our point here is converting these two

ontologies to OWL Full is an overkill and gives the false impression that the formalism can be relied upon.

3.4 Comments on the OAEI 2006 test cases

In 2005, only two teams participated in the anatomy challenge [20, 22]. Their reports essentially outline the difficulties encountered along the way, including the large size of the anatomical ontologies and the transformation of both ontologies from their native format into OWL Full. Hopefully, more participants will have aligned these two large ontologies in 2006. While anatomical ontologies represent good data sets for evaluating alignment strategies, important issues – beyond size and formalism – remain the absence of a gold standard alignment and the need for domain knowledge.

3.5 Comments on the OAEI 2006 measures

While the format required for reporting results is adequate for point-to-point matches across ontologies, it is not clear how other kinds of mappings could be reported. For example, although not reported here, our system can identify *group-to-group matches* [4] and so do other systems such as AnchorPrompt [13]. Along the same lines, one important feature of our system is to identify *non-matches*, i.e., concepts exhibiting resemblance at the lexical level, but semantic differences. Such anchors with negative structural evidence are presented in our result file with “!=” as value for “relation” and 1.0 as value for “measure”. However, an alternative representation for such non-matches would be to use “=” as value for “relation” and a negative value (e.g., -1.0) for “measure”. In this case, the meaning of the value parameter would be somewhat equivalent to that of a correlation coefficient.

In the absence of a reference mapping, it is difficult to evaluate alignment systems on this data set. As a biomedical terminology integration resource, the Unified Medical Language System (UMLS) could provide some elements for evaluating the mappings. However, unfortunately, neither the FMA nor GALEN is part of the UMLS. Moreover, the framework for representing anatomical entities in SNOMED CT – one of the source vocabularies in the UMLS – leads to multiple mappings for many anatomical terms (ambiguity) [3]. In previous work, we analyzed the mapping results produced automatically by several systems [17]. Most mappings were identified by the two systems under investigation, which was reassuring. Analyzing the differences, i.e., mappings identified by one system but not by the other, provided interesting insights about the strengths and weaknesses of each system. Cross-validation among the various systems in competition in the OAEI campaign, would not provide a definitive evaluation metric, but could help system developers analyze the mappings specific to their systems and those they missed. The idea of evaluating the mappings in the context of an application suggested by [20] is very interesting, but difficult to realize. Since ontology alignment is key to the Semantic Web, a collaboration with the W3C Semantic Web Health Care and Life Sciences Interest Group could provide the context for such an evaluation.

4 Conclusion

Aligning large anatomical ontologies has generated a lot of interest in the biomedical and computer science communities and successful methods will contribute to create the biomedical component of the Semantic Web. Despite promising advances, aligning anatomical ontologies remains challenging, because of the sheer size of these resources and the need for domain knowledge. Differences in modeling and representation formalism are also an issue, as well as the lack of a reference mapping for evaluation purposes. Competitive evaluation campaigns such as OAEI represent for ontology alignment the same kind of driving force BioCreAtIvE has been for information extraction in biology [23].

Acknowledgments. This research was supported in part by the Intramural Research Program of the National Institutes of Health (NIH), National Library of Medicine (NLM), and by the Natural Science Foundation of China (No.60496324), the National Key Research and Development Program of China (Grant No. 2002CB312004), the Knowledge Innovation Program of the Chinese Academy of Sciences, MADIS of the Chinese Academy of Sciences, and Key Laboratory of Multimedia and Intelligent Software at Beijing University of Technology.

References

1. Zhang, S., Bodenreider, O.: Aligning representations of anatomy using lexical and structural methods. Proc AMIA Symp (2003) 753-757
2. Bodenreider, O., Hayamizu, T.F., Ringwald, M., de Coronado, S., Zhang, S.: Of mice and men: Aligning mouse and human anatomies. Proc AMIA Symp (2005) 61-65
3. Bodenreider, O., Zhang, S.: Comparing the representation of anatomy in the FMA and SNOMED CT. Proc AMIA Symp (2006) (in press)
4. Zhang, S., Bodenreider, O.: Aligning anatomical ontologies: The role of complex structural rules. Proceedings of the 18th International Conference on Systems Research, Informatics and Cybernetics (InterSymp 2006) (2006) (in press)
5. Zhang, S., Bodenreider, O.: Comparing associative relationships among equivalent concepts across ontologies. Medinfo (2004) 459-463
6. Zhang, S., Bodenreider, O.: Alignment of multiple ontologies of anatomy: Deriving indirect mappings from direct mappings to a reference. Proc AMIA Symp (2005) 864-868
7. Bodenreider, O.: The Unified Medical Language System (UMLS): integrating biomedical terminology. Nucleic Acids Res **32 Database issue** (2004) D267-270
8. Noy, N.F., Musen, M.A.: PROMPT: algorithm and tool for automated ontology merging and alignment. Proceedings of AAAI (2000) 450-455

9. Madhavan, J., Bernstein, P.A., Rahm, E.: Generic schema matching using Cupid. Proceedings of 27th International Conference on Very Large Data Bases (2001) 49–58
10. Stumme, G., Maedche, A.: FCA-Merge: Bottom-up merging of ontologies. Proceedings of the 7th International Conference on Artificial Intelligence (IJCAI '01) (2001) 225-230
11. Kotis, K., Vouros, G.A., Stergiou, K.: Towards automatic merging of domain ontologies: The HCONE-merge approach. Web Semantics: Science, Services and Agents on the World Wide Web **4** (2006) 60-79
12. Doan, A., Madhavan, J., Domingos, P., Halevy, A.Y.: Ontology matching: A machine learning experience. In: Staab, S., Studer, R. (eds.): Handbook on Ontologies. Springer-Verlag (2004) 385-403
13. Noy, N.F.: Tools for mapping and merging ontologies. In: Staab, S., Studer, R. (eds.): Handbook on Ontologies. Springer-Verlag (2004) 365-384
14. McCray, A.T., Srinivasan, S., Browne, A.C.: Lexical methods for managing variation in biomedical terminologies. Proc Annu Symp Comput Appl Med Care (1994) 235-239
15. Zhang, S., Bodenreider, O.: Investigating implicit knowledge in ontologies with application to the anatomical domain. In: Altman, R.B., Dunker, A.K., Hunter, L., Jung, T.A., Klein, T.E. (eds.): Pacific Symposium on Biocomputing 2004. World Scientific (2004) 250-261
16. Smith, B., Ceusters, W., Klagges, B., Kohler, J., Kumar, A., Lomax, J., Mungall, C., Neuhaus, F., Rector, A.L., Rosse, C.: Relations in biomedical ontologies. Genome Biol **6** (2005) R46
17. Zhang, S., Mork, P., Bodenreider, O.: Lessons learned from aligning two representations of anatomy. In: Hahn, U., Schulz, S., Cornet, R. (eds.): Proceedings of the First International Workshop on Formal Biomedical Knowledge Representation (KR-MED 2004) (2004) 102-108
18. Noy, N.F., Musen, M.A., Mejino, J.J.L.V., Rosse, C.: Pushing the envelope: challenges in a frame-based representation of human anatomy. Data & Knowledge Engineering **48** (2004) 335-359
19. Rector, A.L., Bechhofer, S., Goble, C.A., Horrocks, I., Nowlan, W.A., Solomon, W.D.: The GRAIL concept modelling language for medical terminology. Artif Intell Med **9** (1997) 139-171
20. Kalfoglou, Y., Hu, B.: CROSI Mapping System (CMS) - Result of the 2005 Ontology Alignment Contest. In: Ashpole, B., Ehrig, M., Euzenat, J., Stuckenschmidt, H. (eds.): Proceedings of the K-CAP 2005 Workshop on Integrating Ontologies (2005)
21. Golbreich, C., Zhang, S., Bodenreider, O.: The foundational model of anatomy in OWL: Experience and perspectives. Web Semantics: Science, Services and Agents on the World Wide Web **4** (2006) 181-195
22. Jian, N., Hu, W., Cheng, G., Qu, Y.: FalconAO: Aligning ontologies with Falcon. In: Ashpole, B., Ehrig, M., Euzenat, J., Stuckenschmidt, H. (eds.): Proceedings of the K-CAP 2005 Workshop on Integrating Ontologies (2005)
23. Hirschman, L., Yeh, A., Blaschke, C., Valencia, A.: Overview of BioCreAtIvE: critical assessment of information extraction for biology. BMC bioinformatics [electronic resource] **6 Suppl 1** (2005) S1