To repair or not to repair: reconciling correctness and coherence in ontology reference alignments

Catia Pesquita[1], Daniel Faria[1], Emanuel Santos[1], and Francisco M. Couto[1]

[1]Dept. de Informática, Faculdade de Ciências, Universidade de Lisboa, Portugal
`cpesquita@di.fc.ul.pt`

Abstract. A recent development in the field of ontology matching is the alignment repair process, whereby mappings that lead to unsatisfiable classes are removed to ensure that the final alignment is coherent. This process was showcased in the Large Biomedical Ontologies track of OAEI 2012, where two repair systems (ALCOMO and LogMap) were used to create separate coherent reference alignments from the original alignment based on the UMLS metathesaurus. In 2013, the OAEI introduced new reference alignments for this track, created by using the two repair systems in conjunction and manual curation when necessary. In this paper, we present the results of a manual analysis of the OAEI 2013 Large Biomedical Ontologies reference alignments, focused on evaluating the equivalence mappings removed by the repair process as well as those that were replaced by subsumption mappings. We found that up to two-thirds of the removed mappings were correct and that over 90% of the analyzed subsumption mappings were incorrect, since in most cases the correct type of relation was the original equivalence. We discuss the impact that disregarding correctness to ensure coherence can have on practical ontology matching applications, as well as on the evaluation of ontology matching systems.

Keywords: Ontology Matching, Alignment Repair, Reference Alignment, Biomedical Ontologies

1 Introduction

With ontologies growing in size and complexity, the interest in efficient and effective matching methods capable of handling large and heterogeneous ontologies is also on the rise. This is evidenced by the recent introduction of the Large Biomedical Ontologies track in the Ontology Alignment Evaluation Initiative (OAEI) [1], currently the major benchmark for ontology alignment evaluation [2].

The OAEI large biomedical track consists of finding alignments between the Foundational Model of Anatomy (FMA)[3], SNOMED CT [4], and the National Cancer Institute Thesaurus (NCI) [5]. These ontologies are semantically rich and contain tens of thousands of classes.

However, evaluating the matching of very large ontologies is in itself a recognized challenge [6], since the most common type of ontology matching evaluation relies on the comparison of an alignment produced by an ontology matching system against a reference alignment. For smaller ontologies, reference alignments are manually built, and can then be subject to debugging and quality checking steps [7, 8]. However for very large ontologies this is unfeasible since the number of mappings that need to be manually evaluated grows quadratically with the number of classes in an ontology. Even if some heuristics are used to reduce the search space, the human effort is still too demanding, especially when we are facing ontologies with tens or even hundreds of thousands of classes [9].

Consequently, efforts have been made to create reference alignments in an automated or semi-automated fashion [9–11]. One possible strategy to achieve this is based on existing resources from which the reference alignment can be derived. For the three tasks in the large biomedical track in OAEI, the reference alignments were created by processing UMLS metathesaurus entries. UMLS combines expert assessment with automated methods to connect classes from distinct biomedical ontologies and thesaurii according to their meaning.

However, the produced reference alignments lead to a considerable number of unsatisfiable classes when they are integrated with the input ontologies, and while the integration of FMA with NCI generates only 655 unsatisfiable classes, the integration of SNOMED CT and NCI leads to more than 20,000 unsatisfiable classes [12]. To address this issue, in OAEI 2012, in addition to the original reference alignment, two additional references were created by employing two different techniques to repair the logical inconsistencies of the original alignment, ALCOMO [13] and the repair facility of the ontology matching system LogMap [14, 10] (LogMap-Repair).

Ensuring that the alignment between two ontologies is coherent, i.e., that no class or property is unsatisfiable, has recently become a major focus for ontology matching. This is especially relevant when matching very large ontologies, which typically produce more unsatisfiable classes. To ensure the coherence of the alignment, a system needs to first detect the incoherencies and then repair them, by removing or altering them, in order to improve the coherent alignment with minimum intervention. However, different repair methods can produce different alignments. For instance, Figure 1 depicts three conflicting mappings in the original UMLS reference alignment for FMA-NCI. Each system removed two mappings to solve the inconsistencies caused by the disjoint clauses in NCI, but while ALCOMO removed mappings 2 and 3, LogMap removed 1 and 3. In this case, mapping 2 is correct. However the systems have no way of inferring this from the ontologies and alignment, since there are no mappings between the superclasses. For instance, if Anatomy_Kind was mapped to Anatomical_Entity, then this information could be used to disambiguate between Gingiva and Gum.

The application of these techniques reduced the number of unsatisfiable classes to a few [1]. However, this automated process for repair is rather agressive, removing a significant number of mappings (up to 10%). In an effort to counteract this, in OAEI 2013, the three reference alignments were refined by using the two repair systems in conjunction and manual curation when necessary to ensure all

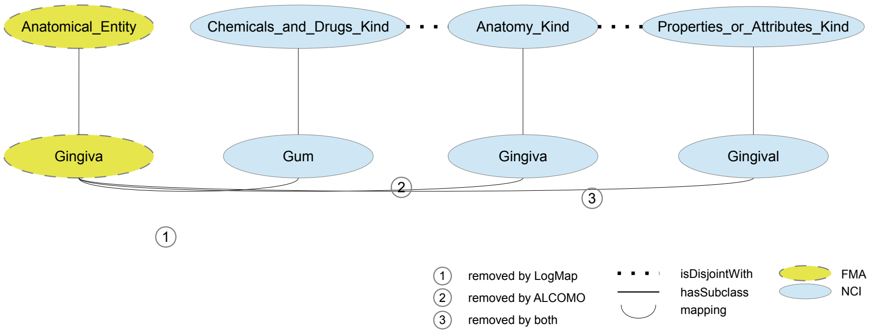

Fig. 1: An example of different repairs made by LogMap and ALCOMO

inconsistencies were solved. This resulted in more complete and fully coherent reference alignments (see Table 1).

Table 1: Reference alignment sizes

Task	Original	LogMap 2012	ALCOMO 2012	Repaired 2013
FMA-NCI	3,024	2,898	2,819	2,931 (13 <, 28 >)
FMA-SNOMED	9,008	8,111	8,132	8,941 (670 <)
SNOMED-NCI	18,844	18,324	N.A.	18,476 (7 <, 540 >)

> and < indicate subsumption mappings

One of the strategies employed to achieve coherence and decrease the number of removed mappings is provided by LogMap. LogMap splits the equivalence mappings into two subsumption mappings and keeps the one that does not violate any logical constraints. This however, may result in mappings that do not reflect the real relationship between classes. Taking again as an example Figure 1, in the repaired alignment in OAEI 2013 all three mappings were replaced by subsumptions: FMA:Gingiva > NCI:Gingival, FMA:Gingiva > NCI:Gingiva and FMA:Gingiva > NCI:Gum. With this solution, the alignment becomes coherent since the relation is directional and the inconsistency is only caused by the disjoint clauses in NCI. However, none of the mappings are correct.
These examples showcase that: 1) different repair techniques produce different repaired alignments; and 2) that solving inconsistencies with subsumption mappings can result in an erroneous alignment. In this paper, we discuss the results of a manual analysis of the OAEI Large Biomedical track reference alignments. We focused our analysis on the differences between the original UMLS and the repaired alignments, in particular on the removed mappings and the ones al-

tered to subsumptions. We also investigated the influence of using the same repair technique to repair both the matching result and to repair the reference alignment.

The paper is organized as follows: Section 2 describes how we conducted our evaluation, Section 3 presents and discusses the evaluation; and finally Section 4 proposes future alternatives for the discussed issues.

2 Methods

To compare the repaired alignments of OAEI 2013 against the original UMLS, we manually evaluated all 41 subsumption mappings in FMA-NCI and 100 randomly chosen subsumption mappings of both FMA-SNOMED and SNOMED-NCI. The evaluation was conducted by two researchers with a biomedical background. We classified each mapping as: correct, incorrect or debatable. We consider mappings correct, not based on their compliance with ontological constraints, but based on their depiction of a real existing relation. For instance, we consider the FMA-NCI mappings between Visceral Pleura, Lung and Thoracic Cavity to be correct even if their integration with the ontologies leads to unsatisfiable classes.

Furthermore, we discerned between mappings where the right relationship would have been equivalence, from those that would have been incorrect with either a subsumption or an equivalence relation. We chose to include a debatable category for those mappings that raised disagreement between the experts, or that they deemed subject to interpretation. For instance, the mappings FMA:Hormone to NCI:Therapeutic_Hormone or SNOMED:Child to NCI:Children.

Our manual evaluation also included the verification of all removed mappings in FMA-NCI and FMA-SNOMED, and of 100 randomly chosen mappings in SNOMED-NCI. These were also classified into the three above-mentioned categories. In addition, we also repaired the original reference alignment with our novel repair technique (AML-Repair) [15] and evaluated the removed mappings.

3 Results and Discussion

Table 2 shows the results of our manual evaluation of the mappings removed or altered from equivalence to subsumption in the repair of the OAEI 2013 Large Biomedical reference alignments. Please note that for the sake of calculating statistics we chose to ignore the debatable removals and alterations.

For FMA-NCI the removal of equivalence mappings is quite successful, with 60 out of 87 removed mappings being correctly so. However, in SNOMED-NCI only half of the mappings were correctly removed, while in FMA-SNOMED this dropped to only 19 out of 65. Regarding the alteration of the mapping relation from equivalence to subsumption, the results are even poorer if more homogeneous between tasks, with 80 to 95% of the alterations being incorrect. Taking into account both removals and alterations, the percentage of correct reparations

ranges from 13% in FMA-SNOMED to 54% in FMA-NCI. Furthermore, considering that the majority of the mappings altered to subsumption by the OAEI 2013 repair are actually equivalences, these alterations do not actually improve the practical quality of the alignment, they just allow the alignment to become coherent without removing the mappings.
To complement this analysis we also repaired the original UMLS reference alignments with our own repair technique (AML-Repair). Compared to the OAEI 2013 repair, AML-Repair makes far more incorrect removals (see Table 3). However, when both removal and alteration are taken into account, AML has a higher percentage of correct repairs in both FMA-SNOMED and SNOMED-NCI.

Table 2: Evaluation of the OAEI 2013 Repair in the Large Biomedical Ontologies track

	Equivalence removal			Alteration to subsumption			
Task	Correct	?	Incorrect	Correct	?	Incorrect	Total correct
FMA-NCI	60	6	27	8	3	30 (26)	54.4 %
FMA-SNOMED	19	1	46	2	5	93 (73)	13.1 %
SNOMED-NCI	42	16	42	4	5	91 (73)	25.7 %

?: Debatable mapping. Numbers in ( ) correspond to mappings where the correct relation is equivalence.

Table 3: Evaluation of AML-Repair in the Large Biomedical Ontologies track

		Equivalence removal			
Task	Size	Correct	?	Incorrect	Total correct
FMA-NCI	2901	48	11	54	47.1%
FMA-SNOMED	8349	19	0	81	19%
SNOMED-NCI	18065	43	6	51	45.7%

?: Debatable mapping.

These results mean that a large percentage of the removed or altered mappings were correct and that both repair techniques are in fact too aggressive.
A fundamental issue here is that different ontologies can have different models of the same subject, and as such, a set of mappings that should be considered correct can render some classes unsatisfiable when the alignment is integrated with the ontologies. For instance, consider the mappings FMA:Fibrillar_Actin = NCI:F-actin and FMA:Actin = NCI:Actin. Both mappings could be considered correct, but when they are integrated with the ontologies they cause an inconsistency. Figure 2 illustrates this issue. Since in FMA F-actin is a subclass of

Actin and in NCI it is a subclass of Actin_Fillament which is disjoint with Actin, the two mappings are in conflict. However, from the biomedical perspective it is arguable that both mappings are correct: F-Actin is the polymer microfilament form of Actin. The OAEI 2013 repair technique solves this issue by changing the relation type in the FMA:Actin=NCI:Actin mapping to subsumption. Since the only constraints violated by the mapping reside in the NCI ontology, by making the mapping one-way, this strategy restores the coherence to the alignment. However, FMA:Actin > NCI:Actin does not represent the true relationship between these classes, which is equivalence.

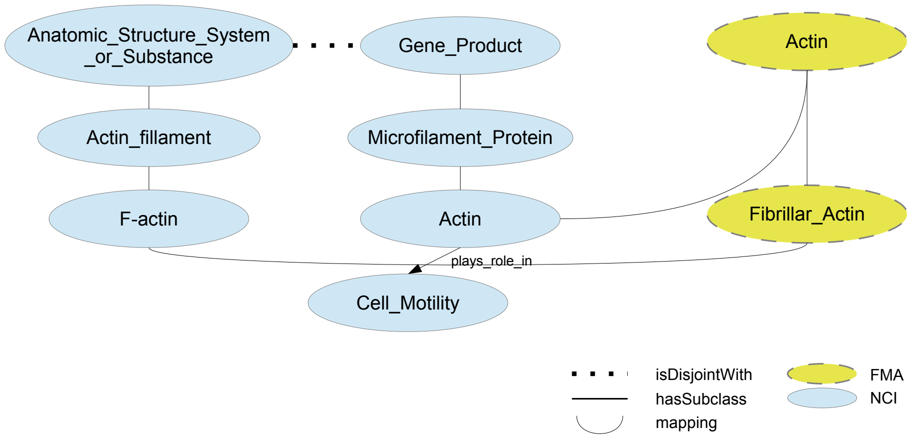

Fig. 2: Two correct mappings causing an inconsistency

So the question is: when creating a reference alignment through automated methods, what is best, an incomplete but coherent reference alignment, or a complete but incoherent one? The answer, we think, depends on the application of the alignment. If the final goal of creating an alignment is to support the integration of two ontologies, then it is necessary to ensure coherence, so that the derived ontology is logically correct and supports reasoning. However, if the goal is supporting the establishment of cross-references between the ontologies to allow navigation between them, then an alignment that does not support linking FMA:Actin to NCI:Actin or reduces the relation to a subsumption would prevent a user from reaching the information that actin plays a role in cell motility. One of the underlying problems is that the existing repair techniques are not guaranteed to remove the incorrect mappings and may erroneously remove correct mappings. The reason for this is that the premise of removing the minimum number of mappings possible (either locally or globally) can fail in cases where there are as many or more incorrect mappings than correct mappings leading to unsatisfiable classes. Indeed, this is exemplified in Figure 1, where ALCOMO erroneously removed the correct mapping. If we evaluated an alignment containing the correct mapping and not the incorrect ones against the ALCOMO-repaired reference, the alignment would be penalized twice: first for having a mapping not present in the reference, and second for not including the erroneous mapping. This means that, even if the true alignment between two ontologies is coherent, by employing an automated repair technique to create a coherent reference alignment we risk excluding correct mappings, and thus providing a more misleading evaluation than if we used the unrepaired reference alignment.

This problem is amplified by the fact that two repair techniques may remove different mappings and arrive at different coherent alignments of comparable size, as exemplified in Figure 1. Without knowing the true alignment, it is impossible to assess which repair technique produces the more correct alignment. However, if the differences between the techniques are statistically significant, in choosing one technique to repair the reference alignment we may bias the evaluation towards that technique. More concretely, if two matching systems produce a similar unrepaired algorithm but use different repair techniques, the one that uses the same repair technique used to repair the reference alignment is likely to produce better results. This is illustrated in Figure 3, which shows two different repairs with techniques 1 and 2 of the same original reference alignment (A). When technique 1 is used to repair the alignment produced by a matching system, its overlap with the reference alignment repaired by 1 (B) is considerable greater than its overlap with the reference alignment repaired by 2 (C).

Table 4: McNemar's exact test for differences between alignments

Task	ALCOMO - LogMap-Repair	OAEI 2013 Repair - AML-Repair
FMA-NCI	2.80E-4	9.01E-4
FMA-SNOMED	2.97E-09	<1.00E-15
SNOMED-NCI	<1.00E-15	2.08E-08

Values shown are two-sided exact p-values

A related work argued that the differences between repair techniques were on average negligible, by comparing the results of applying LogMap-Repair and ALCOMO to the top three systems that participated in the Large Biomedical track of OAEI 2012 [16]. Although the differences between the repair techniques were indeed generally small in percentage, they reflect differences in tens or even hundreds of mappings and can be significant in the context of the OAEI competition.

To demonstrate that the alignments produced by different repair techniques are statistically different, we performed a McNemar's exact test [17] comparing two sets of reference alignments: the OAEI 2012 reference alignments repaired by LogMap and ALCOMO, and the OAEI 2013 reference alignment with the original UMLS reference alignment repaired by AML-Repair. LogMap and ALCOMO

disagree over 177 mappings and UMLS original and repaired differ in 78 mappings. The results in Table 4 show that there is indeed a statistical difference between these sets of alignments, as the p-values obtained are clearly below the lowest significance intervals typically considered (0.01).
To empirically test the possibility that the repair technique selected to repair the reference alignment may lead to a bias in evaluation, we produced simple lexical-based alignments for the three tasks of the Large Biomedical Ontologies (by using AML on the small overlapping ontology fragments [18]). Then, we repaired these alignments using either LogMap-Repair or AML-Repair, and evaluated the repaired alignments against a set of reference alignments: original (UMLS unrepaired), LogMap-Repair (the original repaired with LogMap, as provided in OAEI 2012), and AML-Repair (the original repaired with AML-repair). The results of this evaluation are shown in Table 5. With the sole exception of the AML + LogMap-Repair in the FMA-SNOMED task, the best evaluation results in each task were obtained when the repair technique used to repair the alignment was the same that was used in the reference. Although the differences between the various reference alignments were relatively small (usually below 1%) they are not irrelevant from the perspective of the OAEI evaluation, as the differences between matching systems are often in this range. Thus, the repair technique used to repair the reference alignment can indeed lead to a biased evaluation. What is more, this encourages systems competing in OAEI to adopt existing repair techniques, rather than try to develop novel and potentially better alternatives.

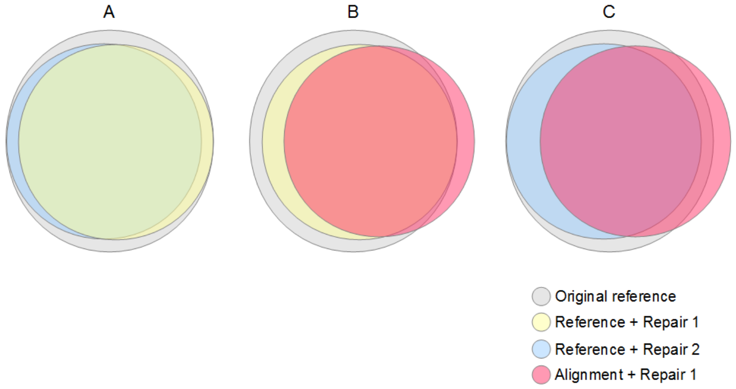

Fig. 3: Comparing a repaired alignment with two different repaired references

We posit that a reference alignment for evaluating ontology matching systems should not exclude potentially correct alignments. As we have shown in Figure 2, it is possible that the true alignment between two ontologies is not coherent. In such cases, repairing the alignment should only be considered if the

Table 5: Influence of different repair techniques on the evaluation of matching systems

Reference	Precision	Recall	F-measure	Size	Correct	Reference
AML + AML-Repair (FMA vs NCI small)						
Original	96.9%	78.8%	87.4%	2457	2382	3024
LogMap-Repair	95.2%	80.7%	87.7%	2457	2339	2898
AML-Repair	95.9%	81.2%	**88.2%**	2457	2356	2901
AML + LogMap-Repair (FMA vs NCI small)						
Original	96.8%	78.8%	87.4%	2461	2383	3024
LogMap-Repair	95.4%	81%	**87.9%**	2461	2347	2898
AML-Repair	95.2%	80.8%	87.7%	2461	2343	2901
AML + AML-Repair (FMA vs SNOMED small)						
Original	95.2%	65.4%	78.9%	6187	5889	9008
LogMap-Repair	86.1%	65.7%	75.2%	6187	5329	8111
AML-Repair	93.2%	69%	**80.2%**	6187	5764	8349
AML + LogMap-Repair (FMA vs SNOMED small)						
Original	94.9%	66.4%	**79.4%**	6298	5978	9008
LogMap-Repair	86.4%	67.1%	76.1%	6298	5439	8111
AML-Repair	89.9%	67.8%	78.1%	6298	5660	8349
AML + AML-Repair (SNOMED vs NCI small)						
Original	92.6%	60.4%	74.8%	12305	11390	18844
LogMap-Repair	91.6%	61.5%	75.1%	12305	11275	18324
AML-Repair	91.5%	62.3%	**75.5%**	12305	11255	18065
AML + LogMap-Repair (SNOMED vs NCI small)						
Original	92.6%	61.3%	75.3%	12474	11550	18844
LogMap-Repair	91.7%	62.4%	**75.7%**	12474	11439	18324
AML-Repair	90.7%	62.6%	75.4%	12474	11312	18065

Best F-score values in bold face

ontologies are to be merged into an integrated resource, as otherwise repairing it implies losing correct mappings. However, even in the cases where the true alignment between two ontologies is expected to be coherent, the use of automatic repair techniques to build a reference alignment is likely to lead to the loss of some correct mappings. Penalizing a system that finds true hard-to-find mappings because these happened to be removed during the repair of the reference alignment is certainly not desirable. The OAEI 2013 reference alignments attempt to minimize the number of mappings removed while still maintaining coherence by replacing equivalence relations with subsumption relations where necessary. But as we have shown, only a small fraction of these relationships are correct as subsumptions. In most cases, the original equivalence relation was correct, and in some other cases the mappings should not exist at all.

On the other hand, using the original (unrepaired) reference alignments is not without issues because these do contain erroneous mappings. Going back to the example in Figure 1, a system that finds only the correct mapping would get a worst result than a system that found the two incorrect mappings if it were evaluated with the original reference alignment. The same would also be true if the system were evaluated with the OAEI 2013 reference alignment, as all three mappings are present in this alignment in the form of subsumptions (assuming the evaluation only considers the presence/absence of mappings and not their relationships).

We propose that a more impartial evaluation could benefit from the fact that existing alignment repair algorithms compute the sets of conflicting mappings as part of their process. Mappings within these sets would be tagged as uncertain, and their presence or absence in the evaluated alignments would not be taken into account when calculating performance metrics. A similar approach has been proposed for cases where only a fraction of the possible mappings have been manually evaluated [19]. Coupling this approach with a satisfiability check on the alignment would allow a more impartial evaluation w.r.t. the repair approach chosen by the matching systems. To illustrate this we have evaluated the AML, AML+AML-Repair and AML+LogMap-Repair alignments for FMA-NCI against an unbiased reference alignment where all conflicting mappings (due to disjointness clauses) have been identified and their presence or absence is not considered in the evaluation. Table 6 presents these results, showing that repaired alignments have a higher precision without losing recall.

Table 6: Evaluaton of different repair techniques against an unbiased reference

Repair Technique	Precision	Recall	F-measure	Size	Correct	Reference
No Repair	95.2%	81.8%	88.2%	1845	1756	2147
AML-Repair	95.9%	81.8%	88.6%	1831	1756	2147
LogMap-Repair	95.7%	81.8%	88.5%	1834	1756	2147

Size and Reference do not include uncertain mappings.

4 Conclusions

As ontologies become more prevalent, large and complex, so must ontology matching systems evolve and with them their evaluation strategies. A recent step in this direction has been the introduction of the large biomedical track in OAEI 2012, where the reference was automatically created by processing an external set of integrated vocabularies and then taking this unrefined alignment and repairing it to diminish its incoherence.

We have found that the repair technique employed to create the OAEI 2013 reference alignment, although less aggressive than the ones used in 2012, still removes a considerable portion of correct mappings and incorrectly alters equivalence mappings to subsumptions. Furthermore, we have shown that alignments repaired with different techniques are significantly different, which can have an impact on the evaluation of ontology matching systems. To decrease the impact of these issues on the evaluation of ontology matching systems, we have proposed an alternative for the evaluation of repaired alignments, where the presence or absence of conflicting mappings is not accounted for. We consider that an alignment between two ontologies should enforce coherence, when the advantages of doing so outweigh the disadvantages, which depends on the application of the alignment and on the ontologies themselves. For instance, if the goal of an alignment is to support integration, then coherence is paramount. However, if the alignment is only intended to support a "lighter" connection between the ontologies (e.g., cross-references), then coverage is likely more relevant than coherence, especially if we consider the error rates of repair techniques. Moreover, when ontologies do not model conflicting views of their domain, then a fruitful alignment between them should be coherent, and ensuring coherence can be a crucial step in filtering out errors. However, when ontologies have incompatible ontological models, their complete integration is impossible and enforcing coherence in their alignment will necessarily remove or alter correct mappings.

How to best integrate ontologies with conflicting views is still a debated question [20], and in some cases the goal might not even be a full-fledged integration. We agree with the opinion expressed in [21] that to solve inherent incompatibilities between ontologies, expert intervention is necessary. However, some incompatibilities are unsolvable, and consequently a full coherent integration of the ontologies is impossible. To promote the usefulness of the alignments there should be room for alignments to contain mappings that violate constraints but are ultimately relevant. A next logical step is to investigate the best approach to support the encoding of these conflicts in the alignment.

Acknowledgements

DF, CP, ES and FMC were funded by the Portuguese FCT through the SOMER project (PTDC/EIA-EIA/119119/2010) and the multi-annual funding program to LASIGE. CP was funded by the FLAD-NSF 2013 Programme under the project "Turning Big Data into Smart Data".

References

1. Eckert, K., Ferrara, A., Hollink, L., Meilicke, C., Nikolov, A., Ritze, D., Shvaiko, P., Grau, B.C., Zapilko, B.: Results of the Ontology Alignment Evaluation Initiative 2012. (2012) 73–115
2. Euzenat, J., Meilicke, C., Stuckenschmidt, H.: Ontology Alignment Evaluation Initiative : six years of experience. Volume 6720. (2011)
3. Rosse, C., Jr, L.V.M.: A reference ontology for biomedical informatics : the Foundational Model of Anatomy. Journal of Biomedical Informatics **36** (2003) 478–500
4. Schulz, S., Cornet, R., Spackman, K.: Consolidating SNOMED CT's ontological commitment. Applied Ontology **6** (2011) 1–11
5. Golbeck, J., Fragoso, G.: The National Cancer Institute's thesaurus and ontology. Web Semantics: Science, Services and Agents on the World Wide Web (2011)
6. Shvaiko, P., Euzenat, J.: Ontology Matching: State of the Art and Future Challenges. IEEE Transactions on Knowledge and Data Engineering **25**(1) (January 2013) 158–176
7. Lambrix, P., Ivanova, V.: A unified approach for debugging is-a structure and mappings in networked taxonomies. Journal of Biomedical Semantics **4**(1) (2013)
8. Beisswanger, E., Hahn, U.: Towards valid and reusable reference alignments - ten basic quality checks for ontology alignments and their application to three different reference data sets. Journal of Biomedical Semantics **3 Suppl 1** (2012) S4
9. Giunchiglia, F., Yatskevich, M., Avesani, P., Shvaiko, P.: A large scale dataset for the evaluation of matching systems. Knowledge Eng. Review (January) (2009)
10. Jiménez-Ruiz, E., Grau, B., Zhou, Y., Horrocks, I.: Large-scale Interactive Ontology Matching: Algorithms and Implementation. ECAI (ii) (2012) 444–449
11. Rosoiu, M., dos Santos, C., Euzenat, J.: Ontology matching benchmarks: generation and evaluation. In: 6th ISWC workshop on ontology matching (OM). (2011)
12. Jiménez-Ruiz, E., Grau, B.C., Horrocks, I.: Exploiting the UMLS Metathesaurus in the Ontology Alignment Evaluation Initiative. E-LKR Workshop (2012) 1–6
13. Meilicke, C.: Alignment incoherence in ontology matching. PhD thesis, University of Mannheim (2011)
14. Jiménez-Ruiz, E., Grau, B.: Logmap: Logic-based and scalable ontology matching. The Semantic WebISWC 2011 (2011)
15. Santos, E., Faria, D., Pesquita, C., Couto, F.: Ontology alignment repair through modularization and confidence-based heuristics. arXiv:1307.5322 (2013)
16. Jiménez-Ruiz, E., Meilicke, C., Grau, B., Horrocks, I.: Evaluating Mapping Repair Systems with Large Biomedical Ontologies. In: 26th International Workshop on Description Logics. (2013)
17. Liddell, F.D.: Simplified exact analysis of case-referent studies: matched pairs; dichotomous exposure. Journal of Epidemiology and Community Health **37**(1) (1983) 82–84
18. Faria, D., Pesquita, C., Santos, E., Palmonari, M., Cruz, I., Couto, F.M.: The AgreementMakerLight Ontology Matching System. In: ODBASE. (2013)
19. Autayeu, A., Maltese, V., Andrews, P.: Recommendations for better quality ontology matching evaluations. In: AISB Workshop on Matching and Meaning. (2010)
20. Schulz, S., Rector, A., Rodrigues, J., Chute, C., Üstün, B., Spackman, K.: Ontology-based convergence of medical terminologies: SNOMED CT and ICD-11. In: eHealth2012. (2012) 89–94
21. Jiménez-Ruiz, E., Grau, B.C., Horrocks, I., Berlanga, R.: Logic-based assessment of the compatibility of UMLS ontology sources. Journal of Biomedical Semantics **2 Suppl 1** (2011) S2