Debugging Taxonomies and their Alignments: the ToxOntology - MeSH Use Case

Valentina Ivanova[1], Jonas Laurila Bergman[2], Ulf Hammerling[3], Patrick Lambrix[1]

(1) Department of Computer and Information Science,
and Swedish e-Science Researche Centre, Linköping University, SE-581 83 Linköping, Sweden
(2) Division of Information Technology, National Food Agency, SE-75126 Uppsala, Sweden
(3) Department of Risk Benefit Assessment, National Food Agency, SE-75126 Uppsala, Sweden

Abstract. As part of an initiative to facilitate adequate identification and display of substance-associated health effects a toxicological ontology - ToxOntology - was created. Further, an alignent with MeSH was accomplished to obtain an indirect index to the scientific literature.

To arrive at satisfactory results in the semantically-enabled applications, high-quality ontologies and alignments are both necessary. A key step towards high quality in this area is debugging the ontologies and their alignments. In this paper we present an experience report on the debugging of ToxOntology and MeSH as well as an alignment.

1 Introduction

Toxicology information, publicly available via Internet, has grown immensely over the last decade and represents a major fundament to risk assessment in a range of regulatory applications, including that of food toxicology. This corpus is commonly referred to as the Internet-based toxicology landscape [21, 10, 17]. The accordingly deposited information is, however, heterogeneous i.e. appears in various forms and formats and is distributed across a rich variety of databases. Several harmonization initiatives have, however, been launched to help extracting such information from disparate sources, typified by the construction of Internet portals (e.g. Toxnet and eChemPortal) and data format standardization [20, 26]. Moreover, the demarcation between data holding classical toxicology actions of substances and that of their general biological activity has become less sharp in recent years. Notably, the ToxCast and Tox21 initiatives have provided gargantuan amounts of data - freely available through the PubChem repository - encompassing results from a wide range of *in vitro* biological assay outputs on noxious chemicals, and the Computational Toxicogenomics Database merges molecular data on chemical health effects at various levels of resolution [18, 1, 22]. Actually, even interaction-type data has recently witnessed exploitation in computational toxicology [8, 2]. Moreover, the OpenTox project, funded by the 7th EU Framework Programme for research, aims at facilitating informatics work in toxicology, through providing an interoperable and standardized framework to support predictive toxicology [4]. Nonetheless, exhaustive toxicology data search and crosswise comparison can still be a cumbersome undertaking.

As part of a slightly broader initiative to facilitate the identification of adequate substance-associated health effects a toxicological ontology - ToxOntology - was created within an informatics system development at the Swedish National Food Agency (NFA). It is inspired by and incorporates several toxicology endpoints of the REACH chemical legislature framework, on which a considerably larger endpoints ontology has been built, as developed within the OpenTox community [16, 25]. While OpenTox vocabularies are mainly designed for advancing predictive toxicology - especially QSAR modeling - the purpose of ToxOntology, however, is to support the identification and presentation of health effects associated with (chemical) substances, as appearing in databases and the scientific literature. Terms and architecture of ToxOntology were created manually by expert toxicologists using various relevant regulatory documents as well as scientific papers in the field. ToxOntology is used in an in-house tagging service to mark textual records where existing classification systems lack coverage, and in an ontology-based text mining application. It is supported by a navigation tool for accessing databases and literature.

Further, the scientific literature is a major source of toxicology information not yet being curated and rendered available in databases. A key source of such documentation is MEDLINE, using Medical Subject Headings (MeSH, [5]) as a classification system. Although the previously mentioned tagging service could be used here for indexing articles relating to a substance of interest, a more precise connection to an already curated index was desired, implicating a need of an alignment [3] between ToxOntology and MeSH.

To obtain high-quality results in semantically-enabled applications (such as the ontology-based text mining and search applications), high-quality ontologies and alignments are both necessary. A key step towards higher quality is to debug the ontologies and their alignments. In this paper we present an experience report on the debugging of ToxOntology and MeSH as well as an alignment. In section 2 we briefly describe ToxOntology and MeSH, as well as the ontology alignment and the ontology debugging systems that were used. Section 3 describes the actual debugging experience, including the creation of an initial alignment of ToxOntology and MeSH, the detection of possible defects using RepOSE [11], two independent repairing sessions - manual and using RepOSE, as well as an experiment using a non-validated initial alignment. The paper concludes in section 4.

2 Background

ToxOntology. ToxOntology is an OWL2 ontology, encompassing 263 concepts and 266 asserted is-a relations. The ontology has ten main axes (top concepts) including Toxic effect, Route of exposure and Time of exposure. All concepts have human readable labels and synonyms attached. ToxOntology appeared after a merge of classification systems covering concepts within toxicology used by ACToR [9] and an implementation of the OpenTox API [6]. The merge was further refined and expanded manually by toxicology experts at the NFA, end-users of ToxOntology. The overall design principle can be summarized as follows: broad enough to cover almost any aspect of interest in the field and at the same time small enough to become an interactive tool in users' daily search of toxicology information.

MeSH. MeSH is a thesaurus of the National Library of Medicine (NLM). It consists of sets of terms naming descriptors in a 12-level hierarchical structure. The 2011 version of MeSH contains 26,142 descriptors. MeSH is used by NLM largely for indexing PubMed [19]. As MeSH contains many descriptors not related to the domain of toxicology, we used parts from the Diseases [C], Analytical, Diagnostic and Therapeutic Techniques and Equipment [E] and Phenomena and Processes [G] branches of MeSH. The resulting ontology contained 9,878 concepts and 15,786 asserted is-a relations. A Java program was written to parse (using the SAX parser) the XML file, filter the selected elements and create the OWL file (using Jena2.1). We note that the MeSH hierarchy is not based on subsumption relations only, and thus interpreting all structural relations as is-a relations, may lead to unintended results.

Ontology alignment system - SAMBO/KitAMO. Our ontology alignment system SAMBO (e.g. [14, 24, 12]) is based on the framework defined in [12] and implements different strategies for preprocessing, matching, combining and filtering. We briefly discuss the strategies that were used in this use case. We did not use preprocessing strategies to reduce the search space. Matchers calculate similarity values between terms. As matchers we used *TermBasic* (linguistic approach), *TermWN* (approach using WordNet [27]), *UMLSM* (approach using domain knowledge - UMLS [23]), and *NaiveBayes* (instance-based approach using scientific literature). The results of the matchers can be combined in different ways. In this use case we used the maximum-based combination strategy, which returns as final similarity value between terms, the maximum value of the similarity values computed by the individual matchers. Further, we used the single threshold filtering strategy, that retains pairs of terms with a similarity value equal to or higher than a given threshold value as mappings suggestions. The mapping suggestions should then be validated by a domain expert. KitAMO [15] is a tool for evaluating and analyzing ontology alignment strategies and their combinations. The tool covers the non-interactive part of the general framework for aligning ontologies. We have used the KitAMO tool with the SAMBO strategies mentioned above, thereby allowing us to store and analyze results from different runs of the algorithms.

Ontology debugging system - RepOSE. RepOSE (version as described in [11]) is a logic-based tool for debugging is-a structure within and mappings between taxonomies. It covers the detection and repairing of defects. It handles defects regarding missing as well as wrong is-a structure, and defects regarding missing and wrong equivalence and is-a mappings. It is based on the framework for debugging ontologies shown in Figure 1. The debugging workflow consists of 6 phases, where the first two phases are for the detection and validation of possible defects, and the last four are for the repairing. The input is a network of ontologies. The output is the set of repaired ontologies and alignments.

In the current version of RepOSE, the detection of defects uses information inherent in the network consisting of the taxonomies and the alignments. In **Phase 1** the system

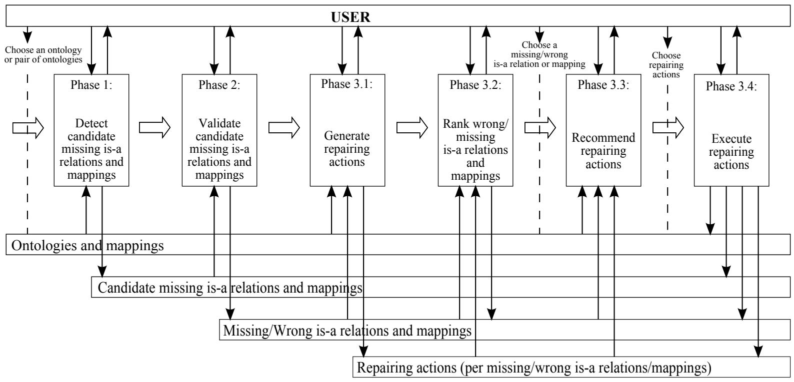

Fig. 1. Debugging workflow [11].

computes for every taxonomy the is-a relations that can be derived from the network but not from the taxonomy alone. These are called *candidate missing is-a relations* (CMIs). Similarly, it computes for every pair of taxonomies and their alignment the mappings that can be logically derived from the network but not from the taxonomies and their alignment alone. These are called *candidate missing mappings* (CMMs). As these CMIs and CMMs may be derived using erroneous information in the network, a domain expert is needed to validate and classify them into missing is-a relation, wrong is-a relation, missing mapping or wrong mapping (**Phase 2**). The CMIs and CMMs are shown to the domain expert using arrows together with their justification[1]. Related items are shown together. The user can validate by clicking the arrows and toggle the label to 'W' or 'M' (e.g. Figure 2). There is also a recommendation algorithm that uses external knowledge. We note that each of the validated CMIs and CMMs gives rise to a debugging opportunity. Missing is-a relations and mappings should be repaired by adding information to taxonomies or alignments. Wrong is-a relations and mappings are repaired by removing information from taxonomies or alignments.

Ontologies and alignments are repaired one by one. For the selected taxonomy or for the selected alignment and its pair of taxonomies, a user can choose to repair the missing or the wrong is-a relations/mappings **(Phase 3.1-3.4)**. Although the algorithms for repairing are different for missing and wrong is-a relations/mappings, the repairing goes through the phases of generation of repairing actions, the ranking of is-a relations/mappings, the recommendation of repairing actions and finally, the execution of repairing actions. In **Phase 3.1** repairing actions are generated. For wrong is-a relations and mappings, the repairing actions are is-a relations or mappings to remove. For each wrong is-a relation and mapping the justifications in the network are computed. The defect can be repaired by removing at least one is-a relation or mapping in each jus-

---

[1] A justification for an is-a relation or mapping can be seen as an explanation for why this is-a relation or mapping is derivable from the network. It is a minimal set of is-a relations and mappings that allows for the derivation of the given is-a relation or mapping. For a formal definition, see e.g. [11, 7].

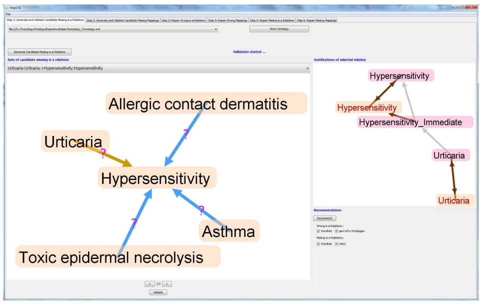

file:///C:/ToxicologyOntology/ExperiencePaper/toxicology_Toxicology.owl

Urticaria:Urticaria->Hypersensitivity:Hypersensitivity

Wrong is-a Relations :
☑ WordNet ☑ part-of in Ontologies
Missing is-a Relations :
☑ WordNet ☑ UMLS

Fig. 2. Generating and validating CMIs.

tification. RepOSE shows for each wrong is-a relation or mapping the justifications as directed graphs (Figure 3). The domain expert can repair by choosing edges in the graph and commit to removing them. For each missing is-a relation or mapping, a Source set and a Target set are computed.[2] It is guaranteed that when an is-a relation/mapping is added between any element in the Source set and any element in the Target set, the defect is repaired. The algorithm also guarantees solutions adhering to a number of heuristics [13]. The Source and Target sets are displayed in two panels to the domain expert (together with the justification of the missing is-a relation or mapping) allowing the user to conveniently repair defects by selecting elements in the panels (Figure 4). In general, there will be many is-a relations/mappings needing repairment and some of them may be easier to embark on such as those with few repairing actions. We therefore rank them with respect to the number of possible repairing actions **(Phase 3.2)**. After this, the user can select an is-a relation/mapping to repair and choose among possible repairing actions. To facilitate this process, we developed methods to guide the user by means of advised repairing actions **(Phase 3.3)**. Once the user decides on repairing actions, the chosen repairing actions are then applied to the relevant taxonomies and alignments and the consequences are computed **(Phase 3.4)**. We also note that the user can switch between different ontologies and phases at any time during the process.

3 Debugging ToxOntology, MeSH and their alignment

3.1 Aligning ToxOntology and MeSH

As an alignment of ToxOntology and MeSH was deemed necessary, and as RepOSE uses an alignment in the detection phase of defects, the first step of our process was to create an initial alignment between ToxOntology and MeSH. Moreover, due to a preference for an as complete as possible, high-quality alignment, preprocessing to reduce the search space was excluded from the procedure; we used different types of matchers; and as combination strategy we used the maximum-based strategy. We generated the similarity values for all pairs of terms. Further, we used single threshold filtering

---

[2] Essentially, for missing is-a relation a → b, Source(a,b) = super-concepts(a) $\setminus$ super-concepts(b) and Target(a,b) = sub-concepts(b) $\setminus$ sub-concepts(a).

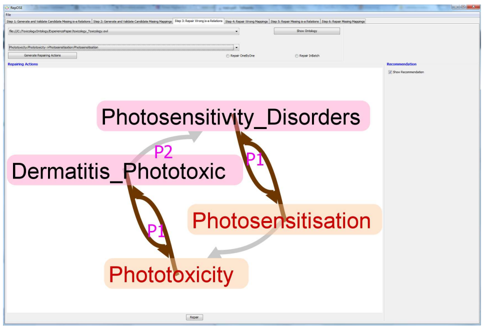

Fig. 3. Repairing wrong is-a relations.

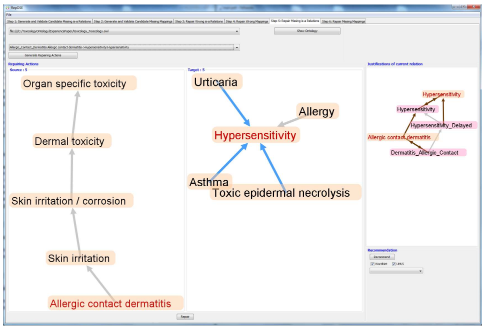

Fig. 4. Repairing missing is-a relations.

similarity value	suggestions	equivalence	ToxOntology is-a MeSH	MeSH is-a ToxOntology	related	wrong
≥ 0.8	41	29	2	2	1	7
$\geq 0.5, < 0.8$	419	9	18	31	42	319
$\geq 0.4, < 0.5$	906	2	21	14	83	786
$\geq 0.35, < 0.4$	146	1	2	2	117	24

Fig. 5. Validation of mapping suggestions - initial alignment.

with threshold 0.35 for the filtering strategy. These choices would lead to a high recall, although there would be many mapping suggestions to validate.

During the validation phase the domain expert classified the mapping suggestions into: equivalence mapping, is-a mapping (ToxOntology term is-a MeSH term and MeSH term is-a ToxOntology term), related terms mapping and wrong mapping. The mapping suggestions were shown to the domain expert in different steps based on the similarity values. The results are summarized in Figure 5. The validated alignment consists of 41 equivalence mappings, 43 is-a mappings between a ToxOntology term and a MeSH term, 49 is-a mappings between a MeSH term and a ToxOntology term and 243 related terms mappings. Further, there is information about 1,136 wrong mappings.

3.2 Debugging using validated alignment

It was not considered feasible to identify defects manually. Therefore, we used the detection mechanisms of RepOSE. RepOSE computed CMIs, which were then validated by domain experts. As there initially were only 29 CMIs, we decided to repair the ontologies and their alignment independently in two ways. First, the CMIs and their justifications were given to the domain experts who manually repaired the ontologies and their alignment. Second, the repairing mechanisms of RepOSE were used. A summary of the changes in the alignment and in ToxOntology due to the debugging sessions are summarized in Figure 6 columns 'original alignment' and 'final alignment'[3], and Fig-

---

[3] The final alignment contains changes from the two debugging sessions and is the one that is now used.

ToxOntology	MeSH	original alignment	final alignment	final alignment manual	final alignment RepOSE
metabolism	metabolism	≡	→	→	removed ←
photosensitisation	photosensitivity disorders	≡	R	R	removed ←, →
phototoxicity	dermatitis phototoxic	≡	R	R	removed ←, →
inhalation	administration inhalation	≡	W	W	removed ←, →
urticaria	urticaria pigmentosa	←	W	W	removed ←
autoimmunity	diabetes mellitus type 1	←	R	R	removed ←
autoimmunity	hepatitis autoimmune	←	R	R	removed ←
autoimmunity	thyroiditis autoimmune	←	R	R	removed ←
gastrointestinal metabolism	carbohydrate metabolism	←	W	W	removed ←
gastrointestinal metabolism	lipid metabolism	←	W	W	removed ←
cirrhosis	fibrosis	≡	R	R	removed ←, →
cirrhosis	liver cirrhosis	←	≡	≡	-
metabolism	biotransformation	←	≡	≡	-
metabolism	carbohydrate metabolism	←	W	W	-
metabolism	lipid metabolism	←	W	W	-
hepatic porphyria	porphyrias	≡	→	W	removed ←
hepatic porphyria	drug induced liver injury	→	R	-	removed →

Fig. 6. Changes in the alignment (equivalence mapping (≡), ToxOntology term is-a MeSH term (→), MeSH term is-a ToxOntology term (←), related terms (R), wrong mapping (W)).

ure 7 column ’final’, respectively. There are also 5 missing is-a relations for MeSH. In the remainder of this subsection we describe the detection and repairing in more details and compare the manual repairing with the repairing using RepOSE.

Detection using RepOSE - first run. As input to RepOSE we used ToxOntology and MeSH as discussed in section 2. Further, we used the validated part of the alignment discussed in section 3.1, that contains the 41 equivalence mappings, the 43 is-a mappings between a ToxOntology term and a MeSH term and the 48 is-a mappings between a MeSH term and a ToxOntology term.[4]

RepOSE generated 12 non-redundant CMIs for ToxOntology (34 in total) of which 9 were validated by the domain experts as missing and 3 as wrong. For MeSH, RepOSE generated 17 non-redundant CMIs (among which 2 relations represented one equivalence relation - 32 CMIs in total) of which 5 were validated as missing and the rest as wrong.

Manual repair. The domain experts focused on repairment of ToxOntology and the alignment. Regarding the 9 missing is-a relations in ToxOntology, these were all added to the ontology. Further, another is-a relation, *asthma → respiratory toxicity*, was added,

---

[4] The related term mappings cannot be used in logical derivation related to the is-a structure of the ontologies and are therefore not included in the alignment used in RepOSE.

Added is-a relations	final	manual	RepOSE
absorption → physicochemical parameter	Yes	Yes	Yes
hydrolysis → metabolism	Yes	Yes	Yes
toxic epidermal necrolysis → hypersensitivity	Yes	Yes	Yes
urticaria → hypersensitivity	Yes	Yes	Yes
asthma → hypersensitivity	Yes	Yes	Yes
asthma → respiratory toxicity	Yes	Yes	No
allergic contact dermatitis → hypersensitivity	Yes	Yes	Yes
subcutaneous absorption → dermal absorption	Yes	Yes	Yes
oxidation → metabolism	Yes	Yes	Yes
oxidation → physicochemical parameter	Yes	Yes	Yes

Fig. 7. Changes in the structure of ToxOntology.

in addition to *asthma* → *hypersensitivity*, based on an analogy of this case with the already existing *urticaria* → *dermal toxicity* and added *urticaria* → *hypersensitivity*. This is summarized in Figure 7 column 'manual'. The domain experts also removed two asserted is-a relations (*asthma* → *immunotoxicity* and *subcutaneous absorption* → *absorption*) for reasons of redundancy. These is-a relations are valid and they are derivable in ToxOntology.

The wrong is-a relations for MeSH and ToxOntology were all repaired by removing mappings in the alignment (Figure 6 column 'final alignment manual'). In 5 cases a mapping was changed from equivalence or is-a into related. In one of the cases (concerning *cirrhosis* in ToxicOntology and *fibrosis* and *liver cirrhosis* in MeSH) a further study also led to the change of *cirrhosis* ← *liver cirrhosis* into *cirrhosis* ≡ *liver cirrhosis*.

The wrong is-a relations involving *metabolism* in ToxOntology, invoked a deeper study of the use of this term in ToxOntology and in MeSH. The domain experts concluded that the ToxOntology term *metabolism* is equivalent to the MeSH term *biotransformation* and a subconcept of the MeSH term *metabolism*. This observation led to a repair of the mappings related to *metabolism*.

Further, some mappings were changed from an equivalence or is-a mapping to a wrong mapping.[5] In these cases (e.g. between *urticaria* in ToxOntology and *urticaria pigmentosa* in MeSH) the terms were syntactically similar and were initially validated wrongly during the alignment phase.

Repairing using RepOSE. For the 3 wrong is-a relations for ToxOntology and the 12 wrong is-a relations for MeSH, the justifications were shown to the domain experts. The justifications for a wrong is-a relation contained at least 2 mappings and 0 or 1 is-a relations in the other ontology. In each of these cases the justification contained at least one mapping that the domain expert validated to be wrong or related and the wrong is-a relations were repaired by removing these mappings (see Figure 6 column 'final

---

[5] So the domain experts changed their original validation based on the reasoning support provided by RepOSE.

alignment RepOSE', except last row). In some cases repairing one wrong is-a relation also repaired others (e.g. removing mapping *hepatic porphyria* ← *porphyrias*, repairs two wrong is-a relations in MeSH: *porphyrias* → *porhyrias hepatic* and *porphyrias* → *drug induced liver injury*).

For the 9 missing is-a relations in ToxOntology and the 5 missing is-a relations in MeSH, possible repairing actions (using Source and Target sets) were generated. For most of these missing is-a relations the Source and Target sets were small, although for some there were too many elements in the set to provide for good visualization. For all these missing is-a relations the repairing constituted of adding the missing is-a relations themselves (Figure 7 column 'RepOSE'). In all but three cases this is what RepOSE recommended based on external knowledge from WordNet and UMLS. In 3 cases the system recommended to add other is-a relations, that were not considered correct by the domain experts (and thus wrong or based on a different view of the domain in the external domain knowledge).

After this repairing, we detected one new CMI in MeSH. This was validated as a wrong is-a relation and resulted in the removal of one more mapping (see Figure 6 column 'final alignment RepOSE' last row).

Discussion. Generally, detecting defects in ontologies without the support of a dedicated system is cumbersome and unreliable. In the case outlined in this paper RepOSE clearly provided a necessary support. Further, visualization of the justifications of possible defects was very helpful to have at hand as well as a graphical display of the possible defects within their contexts in the ontologies addressed. Moreover, RepOSE stored information about all changes made and their consequences as well as the remaining defects needing amendment.

As the set of CMIs was relatively small, it was possible for domain experts to perform a manual repair. They could focus on the pieces of ToxOntology that were related to the missing and wrong is-a relations. This allowed us to compare results of manual repair with those of repairment using RepOSE.

Regarding the changes in the alignment, for 11 term pairs the mapping was removed or changed in both approaches. For 2 term pairs the manual approach changed an is-a relation into an equivalence and for 2 other term pairs an is-a relation was changed into a wrong relation. These changes were not logically derivable and could not be found by RepOSE. For 3 of these term pairs the change came after the domain experts realized (using the justifications of the CMIs) that *metabolism* in MeSH has a different meaning than *metabolism* in ToxOntology. For 1 term pair (one but last row in Figure 6) the equivalence mapping was changed into wrong by the domain experts, while using RepOSE it was changed into an is-a relation. In the final alignment the RepOSE result was used. Further, through a second round of detection, using RepOSE an additional wrong mapping was detected and repaired, which was not found in the manual approach.

Regarding the addition of is-a relations to ToxOntology, the domain experts added one more is-a relation in the manual approach than in the approach using RepOSE. It could not be logically derived that *asthma* → *respiratory toxicity* was missing, but it was added by the domain experts in analogy to the repairing of another missing is-a relation.

In some cases, when using RepOSE, the justification for a missing is-a relation was removed after a wrong is-a relation was repaired by removing a mapping. For instance, after removing *metabolism (ToxicOntology)* ← *metabolism (MeSH)*, there was no more justification for the missing is-a relation *hydrolysis* → *metabolism*. However, an advantage of RepOSE is that once a relation is validated as missing, RepOSE requires that it will be repaired and thus, this knowledge will be added, even without a justification.

Another advantage of RepOSE is that, for repairing a wrong is-a relation, it allows to remove multiple is-a relations and mappings in the justification, even though it may be sufficient to remove one. This was used, for instance, in the repair of the wrong is-a relation *phototoxicity* → *photosensitisation* in ToxOntology where *photosensitisation* ≡ *photosensitivity disorders* and *phototoxicity* ≡ *dermatitis phototoxic* were removed. Further, the repairing of one defect can lead to other defects being repaired. For instance, the removal of these two mappings also repaired the wrong is-a relation *photosensitivity disorders* → *dermatitis phototoxic* in MeSH. In general, RepOSE facilitates the computation and understanding of the consequences of repairing actions.

Interestingly, in this use case only mappings were removed to repair wrong is-a relations. This indicates that the ontology developers modeled the is-a structure decently. This kind of repair is not, however, a consistent outcome. For instance, in the experiment outlined in [11] involving debugging two ontologies and their alignment from the Anatomy track in OAEI 2010 (Adult Mouse Anatomy Dictionary (AMA) and the NCI Thesaurus anatomy (NCI-A), 14 is-a relations were removed from AMA and 11 from NCI-A, as well as 5 mappings. Further, in this use case all missing is-a relations were repaired by adding the missing is-a relations themselves. In the experiment in [11] in 27 cases in AMA and 11 cases in NCI-A a missing is-a relation was repaired using a more informative repairing action, thereby adding new knowledge that was not derivable from the ontologies and their alignment.

An identified constraint of RepOSE pertains to the fact that adding and removing is-a relations and mappings not appearing in the computations in RepOSE can be a demanding undertaking. Currently, these changes need to be conducted in the ontology files, but it would be useful to allow a user to do this via the system. For instance, it would have been useful to add *asthma* → *respiratory toxicity* via RepOSE.

3.3 Debugging using non-validated alignment

In the previous subsection the validated alignment was used as input. As a domain expert validated the mappings, they could be considered of high quality, although we showed that defects in the mappings were detected. In this subsection we perform an experiment with a non-validated alignment; we use the 41 mapping suggestions with a similarity value higher than or equal to 0.8 and use them initially as equivalence mappings.[6]

Using RepOSE (in 2 iterations) 16 non-redundant CMIs (27 in total), were computed for ToxOntology of which 6 were also computed in the debugging session described in 3.2. For MeSH 6 non-redundant CMIs (10 in total) were computed of which

---

[6] From the validation we know that these actually contain 29 equivalence mappings, 2 is-a mappings between a ToxOntology term and a MeSH term, 2 is-a mappings between a MeSH term and a ToxOntology term, 1 related term mapping and 7 wrong mappings.

2 were also computed earlier. As expected, the newly computed CMIs were all validated as wrong is-a relations and their computation was a result of wrong mappings. During the repairing 5 of the 7 wrong mappings were removed, and 2 initial mappings were changed into is-a mappings. RepOSE can thus be helpful in the validation of non-validated alignments - a domain expert will be able to detect and remove wrong mappings that lead to the logical derivation of wrong is-a relations, but wrong mappings that do not lead to logical derivation of wrong is-a relations, may not be found.

4 Conclusion

In this paper we presented an experience report on the debugging of ToxOntology, MeSH and an alignment. We showed the usefulness of RepOSE in detecting and repairing the structure of the ontologies and the alignment.

RepOSE is a logic-based debugging system[7] and detects defects based on logically derivable missing or wrong structure and mappings. In the future, we will investigate the integration of other detection approaches into RepOSE. Also, we will facilitate the adding and removing is-a relations and mappings that do not occur in the computation of the system. Finally, we will investigate the integration of RepOSE with SAMBO.

Acknowledgements. We thank the Swedish Research Council (Vetenskapsrådet), the Swedish e-Science Research Centre (SeRC) and the Swedish Civil Contingencies Agency for financial support.

References

1. DJ Dix, KA Houck, MT Martin, AM Richard, RW Setzer, and RJ Kavlock. The ToxCast program for prioritizing toxicity testing of environmental chemicals. *Toxicol. Sci.*, 95:5–12, 2007.
2. A Edberg, D Soeria-Atmadja, J Bergman Laurila, F Johansson, MG Gustafsson, and U Hammerling. Assessing relative bioactivity of chemical substances using quantitative molecular network topology analysis. *J. Chem. Inf. Model.*, 52:1238–1249, 2012.
3. J Euzenat and P Shvaiko. *Ontology Matching*. Springer, 2007.
4. B Hardy, N Douglas, C Helma, M Rautenberg, N Jeliazkova, V Jeliazkov, I Nikolova, R Benigni, O Tcheremenskaia, S Kramer, T Girschick, F Buchwald, J Wicker, A Karwath, M Gutlein, A Maunz, HS Sarimveis, G Melagraki, A Afantitis, P Sopasakis, D Gallagher, V Poroikov, D Filimonov, A Zakharov, A Lagunin, T Gloriozova, S Novikov, N Skvortsova, D Druzhilovsky, S Chawla, I Ghosh, S Ray, H Patel, and S Escher. Collaborative development of predictive toxicology applications. *J. Cheminform.*, 2:7, 2010.
5. Medical Subject Headings. http://www.nlm.nih.gov/mesh/.
6. N Jeliazkova and V Jeliazkova. AMBIT RESTful web services: an implementation of the OpenTox application programming interface. *J. Cheminform.*, 3:18, 2011.
7. E Jimenez-Ruiz, B Cuenca Grau, I Horrocks, and R Berlanga. Ontology integration using mappings: Towards getting the right logical consequences. In *Proceedings of the 6th European Semantic Web Conference*, volume 5554 of *LNCS*, pages 173–187. Springer, 2009.

---

[7] For an overview of current approaches to debugging we refer to [11].

8. RS Judson, RJ Kavlock, RW Setzer, EA Cohen Hubal, MT Martin, TB Knudsen, KA Houck, RS Thomas, BA Wetmore, and DJ Dix. Estimating toxicity-related biological pathway altering doses for high-throughput chemical risk assessment. *Chem. Res. Toxicol.*, 24:451–462, 2011.
9. RS Judson, AM Richard, DJ Dix, K Houck, F Elloumi, M Martin, T Cathey, TR Transue, R Spencer, and M Wolf. ACToR - aggregated computational toxicology resource. *Toxicol. Appl. Pharmacol.*, 233(1):7–13, 2008.
10. RS Judson, AM Richard, DJ Dix, K Houck, M Martin, R Kavlock, V Dellarco, T Henry, T Holderman, P Sayre, S Tan, T Carpenter, and E Smith. The toxicity data landscape for environmental chemicals. *Environ. Health Perspect.*, 117:685–695, 2009.
11. P Lambrix and V Ivanova. A unified approach for debugging is-a structure and mappings in networked taxonomies. *submitted*, 2012.
12. P Lambrix and Q Liu. Using partial reference alignments to align ontologies. In *Proceedings of the 6th European Semantic Web Conference*, volume 5554 of *LNCS*, pages 188–202, 2009.
13. P Lambrix, Q Liu, and H Tan. Repairing the missing is-a structure of ontologies. In *Proceedings of the 4th Asian Semantic Web Conference*, volume 5926 of *LNCS*, pages 76–90, 2009.
14. P Lambrix and H Tan. SAMBO - a system for aligning and merging biomedical ontologies. *Journal of Web Semantics*, 4(3):196–206, 2006.
15. P Lambrix and H Tan. A tool for evaluating ontology alignment strategies. *Journal on Data Semantics*, VIII:182–202, 2007.
16. W Lilienblum, W Dekant, H Foth, T Gebel, JG Hengstler, R Kahl, PJ Kramer, H Schweinfurth, and KM Wollin. Alternative methods to safety studies in experimental animals: role in the risk assessment of chemicals under the new European Chemicals Legislation (REACH). *Arch. Toxicol.*, 82:211–236, 2008.
17. F Maddah, D Soeria-Atmadja, P Malm, MG Gustafsson, and U Hammerling. Interrogating health-related public databases from a food toxicology perspective: Computational analysis of scoring data. *Food Chem Toxicol.*, 49:2830–2840, 2011.
18. CJ Mattingly, MC Rosenstein, AP Davis, GT Colby, JN Forrest Jr, and JL Boyer. The comparative toxicogenomics database: a cross-species resource for building chemical-gene interaction networks. *Toxicol. Sci.*, 92:587–595, 2006.
19. PubMed. http://www.ncbi.nlm.nih.gov/pubmed/.
20. AM Richard, LS Gold, and MC Nicklaus. Chemical structure indexing of toxicity data on the internet: moving toward a flat world. *Curr. Opin. Drug Discov. Devel.*, 9:314–325, 2006.
21. AM Richard, C Yang, and RS Judson. Toxicity data informatics: supporting a new paradigm for toxicity prediction. *Toxicol. Mech. Methods*, 18:103–118, 2008.
22. SJ Shukla, R Huang, CP Austin, and M Xia. The future of toxicity testing: a focus on in vitro methods using a quantitative high-throughput screening platform. *Drug Discov. Today*, 15:997–1007, 2010.
23. Unified Medical Language System. http://www.nlm.nih.gov/research/umls/.
24. H Tan, V Jakoniene, 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*, volume 3886 of *LNBI*, pages 1–17, 2006.
25. O Tcheremenskaia, R Benigni, I Nikolova, N Jeliazkova, SE Escher ans M Batke, T Baier, V Poroikov, A Lagunin, M Rautenberg, and B Hardy. OpenTox predictive toxicology framework: toxicological ontology and semantic media wiki-based OpenToxipedia. *J. Biomed. Semantics*, 3 Suppl 1:S7, 2012.
26. GM Woodall and RB Goldberg. Summary of the workshop on the power of aggregated toxicity data. *Toxicol. Appl. Pharmacol.*, 233:71–75, 2008.
27. WordNet. http://wordnet.princeton.edu/.