

Analysing Top-level and Domain Ontology Alignments from Matching Systems

Daniela Schmidt*, Cassia Trojahn[†], Renata Vieira*

*Pontifical Catholic University of Rio Grande do Sul (Brazil)
daniela.schmidt@acad.pucrs.br, renata.vieira@pucrs.br

[†]Université de Toulouse 2 & IRIT (France)
cassia.trojahn@irit.fr

Abstract. Top-level ontologies play an important role in the construction and integration of domain ontologies, providing a well-founded reference model that can be shared across knowledge domains. While most efforts in ontology matching have been particularly dedicated to domain ontologies, the problem of matching domain and top-level ontologies has been addressed to a lesser extent. This is a challenging task, specially due to the different levels of abstraction of these ontologies. In this paper, we present a comprehensive analysis of the alignments between one domain ontology from the OAEI Conference track and three well known top-level ontologies (DOLCE, GFO and SUMO), as generated by a set of matching tools. A discussion of the problem is presented on the basis of the alignments generated by the tools, compared to the analysis of three evaluators. This study provides insights for improving matching tools to better deal with this particular task.

1 Introduction

Guarino [5] classifies ontologies according to their “level of generality”: (i) *top-level ontologies* describe very general concepts (e.g., space, time, object, etc.), which are independent of a particular problem or domain. These ontologies, also named upper or foundational ontologies [16], are usually equipped with a rich axiomatic layer; (ii) *domain ontologies* and *task ontologies* that describe, respectively, the entities and other information related to a generic domain (e.g., biology or aeronautic), or a generic task or activity (e.g., diagnosis) by specializing the concepts represented in top-level ontologies; and finally (iii) *application ontologies*, which describe the roles played by domain entities when performing an activity (which are, respectively, described by domain and activity ontologies). While the rich semantics and formalization of top-level ontologies are important requirements for ontology design [11], they act as well as semantic bridges supporting very broad semantic interoperability between ontologies [9,10]. In that sense, they play as well a key role in ontology matching.

However, most efforts in ontology matching have been particularly dedicated to domain ontologies and the problem of matching domain and top-level ontologies has been addressed to a lesser extent. This problem poses different challenges in the field, in particular due to the different levels of abstraction of these ontologies. This is a complex task, even manually, that requires to deeply identify the semantic context of concepts.

It involves going beyond the frontiers of the knowledge encoded in the ontologies and, in particular, the identification of subsumption relations. The latter is largely neglected by most matchers. In fact, when having different levels of abstraction it might be the case that the matching process is rather capable of identifying subsumption correspondences than equivalence, since the top ontology has concepts at a higher level. Approaches dealing with this task are mostly based on manual matching [1,12].

This paper tackles the problem of matching domain and top-level ontologies in a different way. We aim at evaluating how a set of available matching tools, applying different matching strategies, performs in this task. Even though they were not exactly developed for that purpose, their output might help us to investigate the problem. We chose three well-known top-level ontologies (DOLCE, GFO, and SUMO) and one domain ontology from the OAEI Conference data set. Nine matching tools have been used in our experiments. Qualitative and quantitative analyses are based on the point of view of three evaluators at each generated alignment. The aim is to provide an analysis of the alignments provided by the tools for the task of aligning ontologies with different levels of abstraction as well as to discuss our insights on the topic and to provide directions for future improvements.

The rest of the paper is organised as follows. §2 introduces top-level ontologies and discusses related work. §3 presents the material and methods used in the experiments, the results and discussion. Finally, §4 concludes the paper and presents future work.

2 Background

2.1 Top-level ontologies

A top-level ontology is a high-level and domain independent ontology. The concepts expressed are intended to be basic and universal to ensure generality and expressiveness for a wide range of domains. It is often characterized as representing common sense concepts and is limited to concepts which are meta, generic, abstract and philosophical. There are two approaches for the use of top-level ontologies [16], *top-down* and *bottom-up*. The top-down approach uses the ontology as a foundation for deriving concepts in the domain ontology. In this way, we take the advantage of the knowledge and experience already expressed in the top-level ontology. In a bottom-up approach, one usually matches a new or existing domain ontology to the top-level ontology. This approach represents more challenges since inconsistencies may exist between domain and top-level ontologies [16]. This paper focuses on the latter approach.

Several top-level ontologies have been proposed in the literature. The reader can refer to [9] for a review of them. Here, we briefly introduce some well-known and largely used top-level ontologies which are used further in our evaluation:

- DOLCE [4]: Descriptive Ontology for Linguistic and Cognitive Engineering has been proposed by Nicola Guarino and his team at LOA (Laboratory for Applied Ontology). DOLCE is the first module of the WonderWeb Foundational Ontologies Library. The focus of the DOLCE is to grasp the underlying categories of human cognitive tasks and the socio-cultural environment. It is an ontology of particulars and includes concepts such as abstract quality, abstract region, physical object, process, and so on.

- GFO [6]: General Formal Ontology is a top-level ontology for conceptual modeling that has been proposed by the Onto-Med Research Group. It includes elaborations of categories such as objects, processes, time and space, properties, relations, roles, functions, facts, and situations. The work is in progress on the integration with the notion of levels of reality in order to more appropriately capture entities in the material, mental, and social areas.
- SUMO [13]: Suggested Upper Merged Ontology is an upper level ontology that has been proposed as a starter document for The Standard Upper Ontology Working Group, an IEEE working group of collaborators from the fields of engineering, philosophy, and information science. The SUMO provides definitions for general-purpose terms and acts as a foundation for more specific domain ontologies. It is being used for research and applications in search, linguistics and reasoning.

2.2 Related work

In the literature, we see the growing importance of aligning domain and top-level ontology. Recently, in [14], correspondences between DBPedia ontology and DOLCE-Zero [3], a module of DOLCE, are used to identify inconsistent statements in DBPedia. The authors focus on finding systematic errors or anti-patterns in DBPedia. For this task, they exploit previously established alignments between the DBpedia ontology and DOLCE-Zero. They argued that by aligning these ontologies and by combining reasoning and clustering of the reasoning results, errors affecting statements can be identified at a minimal human workload.

In several proposals, alignments between top and domain ontologies are manually generated. In [12], the authors align a domain ontology describing web services (OWL-S) with DOLCE, in order to overcome conceptual ambiguity, poor axiomatization, loose design and narrow scope of the domain ontology. They developed a core ontology of services to serve as middle level between the foundational and domain ontologies and used a module for DOLCE called Descriptions and Situations (D&S) previously developed. The alignment process has been manually done and combined both bottom-up and top-down approaches. First, they used DOLCE as foundational ontology and extended it with the D&S module. This basis has been then used for developing the core ontology of services. Next, they manually aligned OWL-S to the core ontology.

In [1], two domain ontologies of GeoScience (GeoSciML and SWEET) were manually aligned with DOLCE Lite. The authors discussed about the matter of aligning foundational ontologies with these domain ones as a basis for integrating knowledge in this specific domain. The aim is to produce a unified ontology in which both GeoSciML and SWEET are aligned to DOLCE. The alignment process was done in two steps. First, each domain ontology was individually aligned with DOLCE. Then, both ontologies were manually aligned to each other.

In [17], a manually generated alignment between a upper and a biomedical ontology is used for filtering out correspondences at domain level that relate two different kinds of ontology entities. The matching approach is based on a set of similarity measures and the use of top-level ontologies as a parameter for better understanding the conceptual nature of terms within the similarity calculation step. That allows for reducing the possibility of associations between terms derived from different categories. A set of initial

experiments showed an improvement on the alignment quality when using this kind of approach. Evaluation of the generated correspondences has been manually done.

A closer approach to ours has been presented in [7,8], where a repository of ontologies called ROMULUS aims at improving semantic interoperability between foundational ontologies. In order to provide the alignments available in ROMULUS, the authors aligned three foundational ontologies (DOLCE, BFO and GFO) with each other in a semi-automatic way. The alignment process used seven available matching tools (H-Match, PROMPT, LogMap, YAM++, HotMatch, Hertuda, Optima). The resulting manual alignment consists of 35 manual correspondences between DOLCE and GFO, 17 between DOLCE and BFO, and 23 between BFO and GFO. It has been used as a gold standard for comparison with the output of the tools. However, here we focus on the alignment of top and domain ontologies.

Analysing the impact of using top ontologies as semantic bridges (as in [17]) has been done in [10]. A set of algorithms exploiting such semantic bridges are applied and the authors studied under which circumstances upper ontologies improves traditional matching approaches that do not exploit them. They developed different algorithms : one that does not look at the ontology structure; one that looks at the identity and structural information of concepts to decide when the concepts are related; and another one aggregating the structural algorithm with another that does not use upper ontologies. The experiments involved 17 ontologies and 3 top-level ontologies (SUMO, Cyc and DOLCE) used as bridges for matching domain ontologies. 10 tests cases were designed and for each, a reference alignment was manually created including only concepts.

These works use top-level ontologies as a resource for producing better domain ontologies and alignments. Some have used alignments with top-level ontologies that were manually made and others apply automatic approaches for matching ontologies of same level or for analysing the impact of using top ontologies as semantic bridges in the matching process. In fact, the best part of efforts in ontology matching research are targeted to align same domain ontologies while matching domain ontologies with top-level ontologies poses different challenges. This paper tackles the problem in a different way and analyse the behaviour of available matching tools when aligning domain with top-level ontologies. The analysis is more qualitative than quantitative, therefore it is based on a reduced data set, one domain ontology against three of the most well known top-level ontologies available. The experiments are described in the next section.

3 Experiments

3.1 Data set and matchers

OAEI Conference data set. The OAEI Conference data set¹ contains 16 ontologies covering the domain of conference organization. A subset of 21 reference alignments involving 7 ontologies (Ekaw, Conference, Sigkdd, Iasted, ConfOf, Cmt and Edas) has been published. We have chosen this data set because it provides expressive ontologies and is one of the most popular data set in the ontology matching evaluation community [2]. In the experiments presented below, we have used one ontology (the *Conference*

¹ <http://oaei.ontologymatching.org>

ontology²). This ontology has 60 concepts, 46 object properties and 18 data properties. Here, we focus on the alignment of concepts.

Top-level ontologies. The top-level ontologies DOLCE Lite, GFO Basic, and SUMO-OWL were aligned with the *Conference* ontology:

- DOLCE Lite³: the lite version is freely available and it is composed by 37 Concepts and 70 Object properties.
- GFO Basic⁴: the basic version is freely available and it is composed by 45 Concepts and 41 Object properties.
- SUMO⁵: the OWL version is freely available and composed by about 4.500 Concepts and 778 Object properties.

Ontology matching tools. A set of tools, publicly available, from previous OAEI campaigns (not limited to Conference track top participants), and implementing different matching strategies was selected. Even though they are not exhaustive and were not exactly developed for that purpose, their output might help us to investigate the problem of aligning domain and top-level ontologies. Aroma⁶ is a hybrid tool based on association rules; Falcon-AO⁷ applies linguistic and structural approaches, as Lily⁸, which includes debugging strategies; LogMap⁹ applies logical reasoning and repair strategies and its variant LogMap-Lite is essentially based on string similarities; MaasMatch adopts a similarity cube and a disambiguation phase as described in [15]; WeSeE-Match¹⁰ uses web search results for improving similarity measures; WikiMatch¹¹ uses Wikipedia as external knowledge source and YAM++¹² applies both linguistic and graph-based approaches together with machine learning. MaasMatch and YAM++ use WordNet as background knowledge. All the tools were run with their default configuration settings.

3.2 Results and discussion

Manual evaluation. For our experiments, we ran each of the above mentioned systems for the pairs composed by the *Conference* ontology against each top-level ontology. We then merge the alignments generated by the matchers, resulting in 28 correspondences (Table 1), and submitted the resulting merge to the analysis of three evaluators. The evaluators are researchers that have common-sense knowledge about conferences (the domain ontology), with a strong background in Computer Science and well-familiarised

² <http://oaei.ontologymatching.org/2015/conference/data/Conference.owl>

³ <http://www.loa.istc.cnr.it/old/DOLCE.html>

⁴ <http://onto.eva.mpg.de/gfo-bio/gfo-bio.owl>

⁵ <http://www.adampease.org/OP/SUMO.owl>

⁶ <https://exmo.inrialpes.fr/software/aroma/>

⁷ <http://ws.nju.edu.cn/falcon-ao/>

⁸ <http://cse.seu.edu.cn/people/pwang/lily.htm>

⁹ <https://www.cs.ox.ac.uk/isg/tools/LogMap/>

¹⁰ <http://www.ke.tu-darmstadt.de/resources/ontology-matching/wesee-match>

¹¹ <http://www.ke.tu-darmstadt.de/resources/ontology-matching/wikimatch>

¹² <http://www.lirmm.fr/yam-plus-plus/>

with ontology matching. Each of the 28 correspondences (pairs of concepts) were presented to the evaluators, separately, via an online evaluation form (Figure 1). In this form, the first concept in the pair denotes the domain concept and the second one denotes the top concept. For the top concepts, a description (as provided by the top-level ontology) is presented in the form. Checking the ontologies could be done outside the evaluation form. Figure 1 shows one example for the pair ‘Abstract’ - ‘Abstract (DOLCE)’ presented to the evaluators. The evaluators analysed each correspondence and selected one type of relation – Equivalent, Sub/Super concept, or None – according to the relation they judged as correct.

Concepts relation analysis - Part 1

* Required

Abstract -- Abstract *

Definition of abstract (in Dolce): The main characteristic of abstract entities is that they do not have spatial nor temporal qualities, and they are not qualities themselves. The only class of abstract entities we consider in the present version of the upper ontology is that of quality regions (or simply regions). Quality spaces are special kinds of quality regions, being mereological sums of all the regions related to a certain quality type. The other examples of abstract entities (sets and facts) are only indicative.

Equivalent
 Sub/Super concept
 None

83% complete

Fig. 1. Example of correspondence as shown in the online evaluation form.

A summary of the correspondences generated by the matchers together with the results of the manual annotation is presented in Table 1. In this table, the first column presents the concepts of the domain ontology for which one correspondence was found by at least one matcher. The second column shows the top-level concept that was aligned with the corresponding domain concept. The concept hierarchy is included for all concepts. The third column identifies the top-level ontology involved in the alignment. The fourth, fifth, and sixth columns are used to show the evaluators judgment about the pair of concepts. The numbers indicate how many evaluators voted for each type of correspondence. Finally, the last column summarizes how many tools aligned the corresponding pairs of concepts.

Regarding the evaluators judgement, there was total agreement among them in 20 (out of 28 correspondences). However, for 14 of them, no relation has been identified so that half of the automatically aligned concepts were considered neither equivalent nor subsumed. In 3 cases there was total agreement regarding “Subsumption”, and in 3 cases total agreement for “Equivalence”. From the 8 pairs resulting in a disagreement, only 2 of those corresponded to a full disagreement. These 2 cases of total disagreement were discussed among the evaluators, and in one case a total agreement for subsumption was reconsidered. For the other case, a partial agreement for ‘None’ (no relation) was achieved. The results in Table 1 correspond to the final agreement.

We note that, regarding the 28 correspondences, only 18 concepts of a total of 60 from the domain ontology participated in a correspondence.

Tools alignment evaluation. The evaluation of the alignments generated by the tools is based on their precision with respect to the manual analysis. We consider 4 sets of alignments:

- P_1 considers the cases of total agreement, where a correspondence is considered as correct if it has been marked either as equivalent or subsumed by the evaluators (21 correspondences regardless the type of relation – equivalence, subsumption or none – where 7 of them correspond to either equivalence or subsumption);
- P_2 considers the cases involving both total and partial agreements (28 correspondences regardless the type of relation with 14 corresponding to either equivalence or subsumption);
- P_3 considers only total agreement for equivalences (matchers have generated only equivalences) (21 correspondences with 3 equivalences);
- P_4 considers both total and partial agreements only for equivalences (28 correspondences with 4 equivalences).

Table 2 presents the precision of each tool (average of the results for the 3 pairs of ontologies). Here we have a total of 49 correspondences to be analysed, since more than one matcher may indicate a correspondence for the same pair. While some tools were able to generate alignments between *Conference* and the three top-level ontologies (LogMap, LogMapLite and YAM++), other systems have generated alignments for only one pair of ontologies (Conference-DOLCE for Aroma and Conference-GFO for Falcon-AO). Moreover, some systems were not able to generate any alignment (Lily, WeSeE and WikiMatch) and some only generate incorrect ones (Falcon-AO).

For those systems generating non empty alignments, MaasMatch and YAM++ were able to generate more correspondences than the other systems (with LogMap and its variant coming just behind). These 2 systems use WordNet in their matching approaches. This background knowledge resource is a source of lexical relations and can potentially be exploited for finding other relations than equivalence. This can explain the fact that these systems find more alignments. Their best results were obtained for P_2 (however, the best results for this set have been obtained by LogMap). Contrary to what would be expected, these systems (and all others, in fact) were not able to generate subsumption (even though some have been designed to). They generated only equivalences, even when they were in fact subsumptions.

Looking to the different sets, in P_1 , LogMap, LogMapLite and MaasMatch outperformed YAM++. In P_2 , LogMap achieves the best results followed by MaasMatch. When only equivalence (P_3) is considered, the numbers drop for some matchers. When relaxing to both partial and total agreements the results drop even more (P_4). Some matchers are doing equivalence consistently (LogMapLite), whereas others are also indicating correspondences which were in fact considered subsumption by the judges (LogMap, MaasMatch, YAM++), so that P_3 and P_4 decrease. Moreover, precision is low if we compare the results when the same systems are matching domain ontologies¹³.

¹³ <http://oaei.ontologymatching.org/2015/conference/eval.html>

Table 1. Union of the correspondences found by the tools.

Conference Ontology	Top-Level Ontology	Ontologies	Manual			Tools
			≡	⊆	None	
Conference_document/Conference_contribution/Written_contribution/Regular_contribution/Extended_abstract/Abstract	particular/abstract	DOLCE Lite			3	5
	Entity/abstract	SUMO			3	3
	Individual/Abstract	GFO Basic			3	5
Person/Committee_member/Chair	Entity/object/artifact/furniture/seat/chair	SUMO			3	3
Person/Conference_applicant	particular/spatio-temporal-particular/endurant/non-physical-endurant/non-physical-object	DOLCE Lite			3	1
Conference_document	particular	DOLCE Lite			3	1
Conference_part	particular/spatio-temporal-particular/endurant/physical-endurant/feature/relevant-part	DOLCE Lite	1	2		1
	particular/abstract/region	DOLCE Lite			3	1
Conference_proceedings	particular/spatio-temporal-particular/perdurant/stative/process	DOLCE Lite	1	2		1
	Individual/Concrete/Processual Structure/Process	GFO Basic		1	2	1
Conference/Conference_volume	particular	DOLCE Lite	2	1		1
Conference_document/Conference_contribution/Written_contribution/Regular_contribution/Extended_abstract	particular/abstract/region/abstract-region	DOLCE Lite			3	1
Organization	Entity/physical/object/agent/group/organization	SUMO	3			2
Organizer	Entity/physical/object/agent/group/organization	SUMO			3	1
	Entity/physical/object/agent/organism	SUMO		3		1
Conference_document/Conference_contribution/Written_contribution/Regular_contribution/Paper	Entity/physical/object/artifact/paper	SUMO			3	3
Person	Individual	GFO Basic		1	2	1
Conference_document/Conference_contribution/Poster	Entity/physical/content bearing physical/VisualContentBearing Object/PrintedSheet/Poster	SUMO	3			3
	Individual/Property	GFO Basic			3	1
Conference_document/Conference_contribution/Presentation	particular/abstract/proposition	DOLCE Lite		2	1	1
	Individual/Concrete/Processual Structure/Occurrent/Event	GFO Basic		3		1
Publisher	Entity/physical/object/agent/commercial-agent/publisher	SUMO	3			3
Person/Conference_applicant/Registered_applicant	particular/spatio-temporal-particular/endurant/physical-endurant/physical-object	DOLCE Lite		3		1
	Entity	GFO Basic		3		1
Topic	particular/abstract/region/temporal-region/time-interval	DOLCE Lite			3	1
	Category/Concept	GFO Basic	1		2	1
Conference_part/Workshop	particular/spatio-temporal-particular/quality/physical-quality/spatial-location.q	DOLCE Lite			3	1
	Entity/physical/object/region/geographic-area/LocalizablePlace/stationary artifact/workshop	SUMO			3	3
Total of correspondences found by the tools:						49

Table 2. Precision of each system considering their complete set of alignments.

System	P_1		P_2		P_3		P_4	
Aroma	0/2	0	1/3	.33	0/2	0	0/3	0
Falcon-AO	0/1	0	0/1	0	0/1	0	0/1	0
Lily	-	-	-	-	-	-	-	-
LopMap	3/9	.33	5/11	.55	3/9	.33	3/11	.27
LogMapLite	3/9	.33	3/9	.33	3/9	.33	3/9	.33
MaasMatch	3/10	.30	5/12	.42	0/10	0	1/12	.08
WeSeE-Match	-	-	-	-	-	-	-	-
WikiMatch	-	-	-	-	-	-	-	-
YAM++	3/11	.27	5/13	.38	2/11	.18	2/13	.15
Total	12/42	.29	19/49	.39	8/42	.19	9/49	.18

Table 3 shows the overall precision of aligned concepts for each pair of ontologies (based on the union of generated alignments). As expected, the best precision is achieved for P_2 (for the pairs involving GFO). However, if we consider only equivalences (P_3 and P_4), the best precision was achieved with SUMO. We also observe that more correspondences have been generated involving DOLCE concepts (12 pairs), but it corresponds to the lower precision across the different sets.

Table 3. Precision of the alignment union (considering all systems).

Pair of Ontologies	P_1		P_2		P_3		P_4	
Conference - DOLCE Lite	1/8	.13	5/12	.42	0/8	0	0/12	0
Conference - GFO Basic	2/4	.50	5/7	.71	0/4	0	1/7	.14
Conference - SUMO	4/9	.44	4/9	.44	3/9	.33	3/9	.33
Total	7/21	.33	14/28	.50	3/21	.14	4/28	.14

Another simpler way to look at the quality of the alignments generated by the tools is presented in Table 4. It summarizes the correspondences considered correct by at least 1, 2 or by all 3 evaluators. The table also indicates the number of times a relation of equivalence found by the matchers were considered equivalence or subsumption by the evaluators. It shows that 14 out 28 correspondences made by the tools were considered as equivalent or subsumed by at least 1 evaluator. Total agreement happened in 7 of these cases (after discussion on the cases of total disagreement).

Table 4. Number of correct correspondences according to the evaluators analysis.

	at least 1 judge	at least two judges	three judges
\sqsupseteq	10	6	4
\equiv	4	3	3
$\sqsupseteq + \equiv$	14	9	7

Discussion. Regarding the qualitative analysis of the alignments, we observe that the systems found various correspondences between concepts with the same term (“Abstract”, “Chair”, “Paper”, “Workshop”, “Organization”, “Poster” and “Publisher”). This is quite expected as all tools are based on some string-based matching strategy. However, many of them were considered as having no correspondence by the evaluators

(“Abstract”, “Chair”, “Paper”, “Workshop”). Among these concepts, the most common aligned one was ‘Abstract’ involving DOLCE and GFO (5 tools) and SUMO (3 tools). Some other concepts were aligned by three or two different tools, but most concepts were aligned just by one. The other correspondences provided by the tools which were considered no correspondent by all the evaluators can be found in Table 1.

There were correspondences with the same term which were considered equivalent by the evaluators :

- “Organization” in the top-level ontology is defined as: a group of people with a common purpose or function in a corporate or similar institution, the same as in the conference domain.
- “Poster” is defined as: a printed sheet intended to be posted on a horizontal surface, so as to make the information it displays visible to passers by.
- “Publisher” in the top-level ontology refers to: some service that includes the publication of texts, so as in the conference domain.

The 3 cases above were SUMO concepts. The concept “Organization” was aligned by 2 systems and the others by 3. For some concepts, all evaluators considered that there was a correspondence but selected subsumption instead of equivalence :

- Organizer and organism: The first concept refers to people who organizes conferences and the second refers to a living individual, then, the concept “Organizer” was considered as subsumed by “Organism”.
- Presentation and event: The first concept refers to the action of explaining about some topic for a group of people. The second refers to processual structures comprising a process. “Presentation” was considered as subsumed by “Event” by the judges.
- Registered_applicant and physical-object: The first concept refers to people who apply and is able to participate in the conference. The main characteristic of the second concept is that they are endurants with unity and most physical objects change some of their parts while keeping their identity, they can have therefore temporary parts. In this case, one “Registered_applicant” is a person who in some specific time interval assumes this role, but keeping their identity as person, then, it was considered as subsumed by “Physical-object”.
- Registered_applicant and entity: The first concept was interpreted in the same way as above. The second concept refers to everything that exists in the broadest sense. In this case, one “Registered_applicant” is something that exists, then, the first concept was considered as subsumed by “Entity”.

An important aspect is that finding subsumption correspondences is in fact highly desirable when matching domain and top-level ontologies. Ideally, such a matcher should try to find the closest super concept. However the matchers we tested in this experiment were not able to generate subsumption, even if some of them (Aroma, for instance) are supposed to do so. They generated only equivalences, even when they were in fact subsumptions. This is however an important distinction. Finally, our analysis does not take into account the inconsistencies introduced in the merging alignments from all tools. In fact, it is contradictory that *Conference_applicant* aligns to *Non_physical_object* and *Registered_applicant* to *Physical_object*, considered that the

latter is a subclass of *Conference_applicant* in the domain ontology. This could be exploited for further filtering out correspondences. We could as well enrich the set of manually validated correspondences by introducing simple hierarchical reasoning.

To sum up, although the number of evaluators is relatively small, it allowed us to establish a first evaluation of available tools on the task. Our study was useful to observe various questions in the task of matching ontologies of different levels of abstraction :

- there was a small quantity of aligned concepts by the tools in general (in total, 18 of 60 concepts), even considering all concepts provided by the top ontologies;
- there were many produced correspondences which were not considered as correspondences by the specialists, many string matching cases which are usually safe in same domain correspondences did not apply, according to our study;
- there is a lack of comprehensive evaluation data sets (regarding domain vs. top-level ontologies) to evaluate the systems, and to overcome that we presented an analysis of the output generated by current systems;
- knowledge on top level ontologies is highly specialized, it is important that such evaluation considers an overview of experts in this area;
- both domain and top ontologies may lack further context or documentation that is appropriate to help identifying the right correspondences;
- manual analysis or correspondences generation by specialists is a hard and expensive work, in this work we ran experiments on a small set of concepts and this problem has been reduced; bigger data sets would require more efforts;
- matching strategies for dealing with this task should take advantage of structural features of the ontologies, background knowledge from external resources targeting subsumption correspondences, and logical reasoning techniques for guarantee the consistency of the generated alignments;
- at last, but not least, current tools do not distinguish between subsumption and equivalence correspondences, which in this kind of task is a crucial point, finding the closest super-concept is quite desirable when aligning to a top-level ontology.

4 Concluding remarks and future work

This paper presented an analysis of the alignments between three top-level ontologies with one domain ontology as produced by a set of matching tools. Our goal was to analyse the behaviour of these tools, which apply diverse matching techniques, with respect to this task. We could observe that matching top-level and domain ontologies automatically is an interesting and challenging task. Top-level ontologies focus on the standardisation of more general concepts to be easily reused in a large amount of domains. On the other hand, there are a lot of domain ontologies available in different fields. Therefore, we claim that it is important to reuse the well-founded knowledge available in the top-level ontologies together with the domain ontologies to reduce the time of ontology modeling, the heterogeneity problem of the knowledge representation, and the complexity of ontology modeling. Hence the automatic matching process should be an alternative. Furthermore, top-level ontologies are semantic bridges for helping solving the heterogeneity between domain ontologies that have to be integrated.

As future work, we plan to run experiments exploiting the whole space of possible alignments (regarding a data set) and to extend the evaluation taking into account

matching tools participating in more recent OAEI campaigns. We plan as well to involve evaluators experimented in top-level ontologies and with different backgrounds (Computer Scientists, Philosophers) in the manual evaluation process. We intend also to exploit background knowledge from external resources (like BabelNet) in order to improve the results reported here, paying special attention to subsumption relations. Combining it with logical reasoning is another aim. Finally, we intend to exploit other data sets such as the ones available on the BioPortal, which contain manually validated alignments between biomedical ontologies and the top level ontologies GFO and BFO.

References

1. Brodaric, B., Probst, F.: DOLCE ROCKS: Integrating Geoscience Ontologies with DOLCE. In: *Semantic Scientific Knowledge Integration*. pp. 3–8 (2008)
2. Cheatham, M., Hitzler, P.: Conference v2.0: An Uncertain Version of the OAEI Conference Benchmark. In: *Proc. of the 13th Intern. Semantic Web Conference*. pp. 33–48 (2014)
3. Gangemi, A., Guarino, N., Masolo, C., Oltramari, A.: Sweetening WORDNET with DOLCE. *AI Magazine* 24(3), 13–24 (2003)
4. Gangemi, A., Guarino, N., Masolo, C., Oltramari, A., Schneider, L.: Sweetening Ontologies with DOLCE. In: *Proc. of the 13th Intern. Conf. on Knowledge Engineering and Knowledge Management*. pp. 166–181 (2002)
5. Guarino, N.: *Formal Ontology in Information Systems: Proceedings of the 1st Intern. Conference*. IOS Press, 1st edn. (1998)
6. Herre, H., Heller, B., Burek, P., Hoehndorf, R., Loebe, F., Michalek, H.: *General Formal Ontology (GFO): A Foundational Ontology Integrating Objects and Processes*. In: *Basic Principles, Research Group Ontologies in Medicine* (2007)
7. Khan, Z., Keet, C.M.: The Foundational Ontology Library ROMULUS. In: *Proc. of the 3rd Conf. on Model and Data Engineering*. pp. 200–211 (2013)
8. Khan, Z., Keet, C.M.: Toward semantic interoperability with aligned foundational ontologies in ROMULUS. In: *Proc. of the 7th Conf. on Knowledge Capture*. pp. 23–26 (2013)
9. Mascardi, V., Cordi, V., Rosso, P.: A Comparison of Upper Ontologies. In: *Proc. of the 8th AI*IA/TABOO Joint Workshop on Agents and Industry*. pp. 55–64 (2007)
10. Mascardi, V., Locoro, A., Rosso, P.: Automatic Ontology Matching via Upper Ontologies: A Systematic Evaluation. *IEEE Trans. on Knowl. and Data Eng.* 22(5), 609–623 (2010)
11. Mika, P., Oberle, D., Gangemi, A., Sabou, M.: Foundations for Service Ontologies: Aligning OWL-S to Dolce. In: *Proc. of the 13th Conf. on World Wide Web*. pp. 563–572 (2004)
12. Mika, P., Oberle, D., Gangemi, A., Sabou, M.: Foundations for Service Ontologies: Aligning OWL-S to DOLCE. In: *Proc. of the 13th Conf. on World Wide Web*. pp. 563–572 (2004)
13. Niles, I., Pease, A.: Towards a Standard Upper Ontology. In: *Proc. of the Intern. Conf. on Formal Ontology in Information Systems*. pp. 2–9 (2001)
14. Paulheim, H., Gangemi, A.: Serving DBpedia with DOLCE - More than Just Adding a Cherry on Top. In: *The Semantic Web*, pp. 180–196 (2015)
15. Schadd, F.C., Roos, N.: Coupling of Wordnet Entries for Ontology Mapping Using Virtual Documents. In: *Proc. of the 7th Conf. on Ontology Matching*. pp. 25–36 (2012)
16. Semy, S., Pulvermacher, M., Obrst, L.: Toward the use of an upper ontology for U.S. government and U.S. military domains: An evaluation. Tech. rep., Sub. to Workshop on Information Integration on the Web, in conjunction with VLDB-2004 (2004)
17. Silva, V., Campos, M., Silva, J., Cavalcanti, M.: An Approach for the Alignment of Biomedical Ontologies based on Foundational Ontologies. *Information and Data Management* 2(3), 557–572 (2011)