Matching BFO, DOLCE, GFO and SUMO: an evaluation of OAEI 2018 matching systems

Daniela Schmidt[1], **Cassia Trojahn**[2], **Renata Viera**[1]

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

[2]Institut de Recherche en Informatique de Toulouse (France)
cassia.trojahn@irit.fr

Abstract. *Foundational ontologies play an important role in the construction and integration of domain ontologies, providing a well-founded reference model that can be shared across domains. Different foundational ontologies have been developed, under different philosophical perspectives. Interoperability across domain ontologies relying on different foundational ontologies depends hence on the ability of interchanging foundational ontologies. The first step toward this task is to find correspondences between them. This paper extends previous work in the analysis of automatic matching system in the task of matching foundational ontologies. We discuss the weaknesses of existing proposals and highlight the challenges to be addressed in the the field.*

1. Introduction

Foundational ontologies play an important role in the construction and integration of domain ontologies, providing a well-founded reference model that can be shared across domains. They describe general concepts (e.g., physical object, event) and relations (e.g., parthood, participation), which are independent of a particular domain. These ontologies, also named upper or top-level, are usually equipped with a rich axiomatic layer. While the clarity in semantics and a rich formalization of foundational ontologies are important requirements for ontology development [Mika et al. 2004, Keet 2011] improving ontology quality, they may also act as semantic bridges supporting interoperability between ontologies [Mascardi et al. 2010, Keet 2011, Nardi et al. 2013].

While the purpose of a foundational ontology is to solve interoperability issues among ontologies, the development of different foundational ontologies re-introduces the ontology interoperability problem, as stated in [Khan and Keet 2013a]. Early works addressed this problem [Grenon 2003, Seyed 2009, Temal et al. 2010] on different perspectives. While fundamental issues and primitive relations between BFO and DOLCE have been studied in [Grenon 2003] and [Seyed 2009], respectively, [Temal et al. 2010] established an alignment between these ontologies in order to conciliate their respective realistic and cognitive points of view. In [Muñoz and Grüninger 2016], the core characterization of mereotopology of SUMO and DOLCE has been studied, relating their axiomatizations via ontology alignments, while in [Oberle et al. 2007] alignments between DOLCE and SUMO have been established for supporting domain ontology integration. In [Khan and Keet 2013a, Khan and Keet 2013b], alignments between BFO, DOLCE and GFO were built both with automatic matching tools and manually, with substantially fewer alignments found by the matching tools.

In this paper we analyze the behaviour of automatic matching systems in the task of matching foundational ontologies. This work extends the work from [Khan and Keet 2013a, Khan and Keet 2013b] in two ways: it considers more recent matching systems, those participating in the Ontology Alignment Evaluation Initiative (OAEI) 2018, and it considers a new pair of aligned foundational ontologies SUMO and DOLCE [Oberle et al. 2007], which consists of subsumption relations. The alignments in [Khan and Keet 2013a] and [Oberle et al. 2007] served as a reference alignment in order to evaluate the matchers.

The aim here is not to evaluate the matchers themselves (as to point out the best matcher) but rather to analyse how they behave in the task, and specially to compare the new results to those obtained in [Khan and Keet 2013a]. As these previous [Khan and Keet 2013a, Khan and Keet 2013b] works have pointed out the weaknesses of matchers to deal with the task, we aim at analysing (quantitatively) whether a progress towards that question has been made so far. We discuss their results, point out their weaknesses, and highlight the challenges to be addressed in the the field.

The rest of this paper is organized as follows. Section 2 introduces the main foundational ontologies and Section 3 discusses the related work. We present the experiments in Section 4 and discuss the results in Section 5. Finally, Section 6 ends the paper pointing out directions for future work in the field.

2. Foundational ontologies

A foundational ontology is a high-level and domain independent ontology whose concepts (e.g., physical object, event, quality, etc.) and relations (e.g., parthood, participation, etc.) are intended to be basic and universal to ensure generality and expressiveness for a wide range of domains. It is often characterized as representing commonsense concepts. Diverse foundational ontologies have been developed, influenced by different philosophies and views on the reality. Several comparisons can be found in the literature [Semy et al. 2004, Mascardi et al. 2007, Khan and Keet 2012], which mostly focus on software engineering criteria (dimensions, representation languages, modularity) [Mascardi et al. 2007] or ontological commitments and subject domain and applications [Khan and Keet 2012]. We introduce here the main insights behind the foundational ontologies most cited in the literature. We are aware however that other ones have been proposed, such as SOWA's ontology, YAMATO, GIST, and KIOTO, but were left out of the scope.

- **BFO** [Grenon et al. 2004, Arp et al. 2015][1] (*Basic Formal Ontology*) is a foundational ontology that adopts a *realistic approach* in terms of the existence in time of entities populating the world. It represents the reality into two disjoint categories of *continuant* (objects, attributes, and locations) and *occurrent* (processes and temporal regions. With different versions, BFO 2.0 represents major updates to BFO not strictly backwards compatible with BFO 1.1 and a manual alignment was required to express their incompatibilities.
- **DOLCE** [Gangemi et al. 2002][2] (*Descriptive Ontology for Linguistic and Cognitive Engineering*) is an ontology of *particulars* which adopts a *descriptive ap-*

[1] https://github.com/bfo-ontology/BFO/wiki
[2] http://www.loa.istc.cnr.it/old/DOLCE.html

proach with a clear cognitive bias, as it aims at capturing the ontological categories underlying natural language and human commonsense. DOLCE is based on a fundamental distinction between *endurant* and *perdurant* entities. *Endurants* represent objects or substances while *perdurants* corresponds to events or processes. The main relation between *endurants* and *perdurants* is that of participation. DOLCE and its different variants have been used in diverse proposals, as many efforts have been dedicated to the development of this ontology. DOLCE has been exposed with reduced axiomatization and extensions with generic or domain plugins, such as for DOLCE-Lite [Gangemi et al. 2003], DOLCE-Lite-Plus[3] or still DOLCE+DnS Ultralite[4].

- **OpenCyc** [Guha and Lenat 1993] is a foundational ontology involving thousands of "microtheories". It is meant for the representation of facts, rules, and heuristics to reason about the objects and events of everyday life in the Cyc knowledge base. It is the open source version of Cyc, a commercial product by Cycorp.
- **GFO** [Herre et al. 2007][5] (*General Formal Ontology*) is a foundational ontology that considers basic distinctions between individuals. Concrete individuals exist in time or space whereas abstract individuals do not. While an *endurant* is an individual that exists in time, but cannot be described as having temporal parts or phases; *processes*, on the other hand, are extended in time.
- **PROTON** [Terziev et al. 2005][6] (*PROTo ONtology*) serves as a lightweight foundational ontology organized in three levels, including four modules describing. The *top ontology* module, for instance, distinguishes entity types, such as *object* as existing entities (agents, locations, vehicles); *happening* as events and situations; and *abstract* as abstractions that are neither objects, nor happenings.
- **SUMO** [Niles and Pease 2001][7] (*Suggested Upper Merged Ontology*) is as ontology of particulars and universals which has two top-level concepts. *Physical* represents an entity that has a location in space-time. An *abstract* can be said to exist in the same sense as mathematical objects such as sets and relations, but they cannot exist at a particular place and time without some physical encoding.
- **UFO** [Guizzardi 2005, Guizzardi and Wagner 2010][8] (*Unified Foundational Ontology*) started as an unification of the GFO and the foundational ontology of universals underlying OntoClean[9]. UFO is divided in three parts representing different aspects of reality: an ontology of endurants (objects), an ontology of perdurants (events and processes), and an ontology of social entities, with notions such as beliefs, desires, intentions, etc.

3. Related work

Early works have addressed the problem of comparing or aligning foundational ontologies [Grenon 2003, Seyed 2009, Temal et al. 2010] on different perspectives. In [Grenon 2003], fundamental issues (as significant discrepancies related to universals and

[3] http://www.loa.istc.cnr.it/old/ontologies/DLP_397.owl
[4] http://www.ontologydesignpatterns.org/ont/dul/DUL.owl
[5] http://www.onto-med.de/ontologies/gfo/
[6] http://ontotext.com/proton
[7] https://github.com/ontologyportal/sumo
[8] http://dev.nemo.inf.ufes.br/seon/UFO.html
[9] http://www.ontoclean.org

particulars, qualities, constitution and spatio-temporality, etc.) and how similar notions apply differently in BFO and DOLCE have been studied. This manual analysis was based on the preliminaries versions of BFO (BFO 1.0). A comparison between these ontologies has also carried out in [Seyed 2009], where the primitive relations (dependence, quality, and constitution) between these BFO 1.0 and DOLCE have been discussed. While these works studied the discrepancies and similarities between the ontologies under their philosophical perspectives and conceptual points of views, giving some insights to corresponding their concepts, in [Temal et al. 2010], a manual alignment in terms of equivalences and subsumption correspondences between BFO (BFO 1.0) and DOLCE concepts has been established. This alignment has been used to integrate an ontology of telecardiology based on DOLCE to ontologies based on BFO such as FMA (*Foundational Model of Anatomy*).

Other studies have addressed other foundational ontologies. In the work from [Oberle et al. 2007], alignments between DOLCE and SUMO have been generated, where two core ontologies, the SmartDOLCE and SmartSUMO ontologies, have been developed on the basis of DOLCE and SUMO, respectively. The alignment of the SUMO taxonomy to DOLCE involved prunning the upper-level of the SUMO taxonomy and the non-trivial task of aligning the remaining concepts to appropriate DOLCE categories. In [Muñoz and Grüninger 2016], the core characterization of mereotopology of SUMO and DOLCE has been studied, relating their axiomatizations via ontology alignments. This included corrections and additions of axioms to the analyzed theories which eliminate unintended models and characterize missing ones. The resulting alignments have been expressed in FOL.

The closest work to ours is from [Khan and Keet 2013a], where alignments between BFO (BFO 1.1), DOLCE (DOLCE-Lite) and GFO have been established with automatic matching tools (H-Match, PROMPT, LogMap, YAM++, HotMatch, Hertuda and Optima) and also manually. During the process, it was found that differences in foundational ontologies, such as their hierarchical structure, conflicting axioms due to complement and disjointness, and incompatible domain and range restriction, cause logical inconsistencies in foundational ontology alignments, thereby greatly reducing the number of correspondences. While the accuracy and percentage of alignments that were found vary greatly among the tools, exploiting the aligned entities whilst keeping a consistent ontology reduces the feasible set of alignments. The resulting alignments have been made available at the ROMULUS platform [Khan and Keet 2013b] [10]. From this experiment and a set of manually curated alignments, [Khan and Keet 2014] developed the SUGOI tool (*Software Used to Gain Ontology Interchangeability*) which allows a user to interchange automatically a domain ontology, by choosing the foundational one.

Aligning foundational ontologies reveals also the problem of matching their different versions. In [Seppälä et al. 2014], a method for tracking, explaining and measuring changes between successive versions of BFO 1.0, BFO 1.1, and BFO 2.0 was applied. The aim was to provide a more comprehensive analysis of the changes with respect to the *BFOConvert* tool[11] which provides an alignment between previous BFO versions, as this resource is limited to allow for a full understanding of the impact of the changes.

[10] `http://www.thezfiles.co.za/ROMULUS/`

[11] `http://ontobull.hegroup.org/bfoconvert` (last viewed on April 1st, 2019)

4. Experiments

The aim of our experiment is to analyse how current matchers behave in the task of matching foundational ontologies. In our experiments, we used the set of matchers participating in the OAEI 2018 campaign. Previous evaluation [Khan and Keet 2013c] has considered a different set of matching tools (as described below) and manually evaluated the generated alignments. Here, we analyse whether matchers have evolved in that task. In the following, we describe the ontologies and the alignments used as 'reference' for automatically evaluating the alignments generated by the matchers.

4.1. Ontologies and reference alignments

In our experiments we have used the following foundational ontologies Table 1: BFO, DOLCE-Lite, GFO and SUMO. For BFO[12] and GFO[13], we have used the versions referenced and available at the ROMULUS repository[14]. For DOLCE-Lite, whose link was unreached at ROMULUS, the version was the one available on LOA[15] (whose version has the same base namespace than the one in the alignments from ROMULUS).

For the experiments involving SUMO, for sake of results reproducibility and compatibility, we use the same DOLCE-Lite and SUMO ontology versions used in [Oberle et al. 2007].

Table 1. Foundational ontologies. *Logical axioms count from Protege.

Ontology	Version	#Classes	#ObjProp	#Axioms*
BFO	1.1	39	0	95
DOLCE	DOLCE-Lite	82	121	356
	DOLCE-Lite [Oberle et al. 2007]	82	162	619
GFO	1.0	78	67	323
SUMO		630	236	1307

The pairs of alignments in Table 2 have been considered. As stated before, this choice is based on the available existing alignments between the ontologies. The reference alignments for the first three pairs have been obtained from ROMULUS, while the reference alignment for the pair DOLCE-Lite and SUMO has been the one from [Oberle et al. 2007]. We refer to these alignments as *reference*, as they are the alignments manually curated and available for comparison. However, an analysis of them shows that they are not exhaustive. First, the alignments in ROMULUS result from a manual evaluation, selection and enrichment of correspondences generated by automatic matchers. In that way, they may introduce a bias as we evaluate the tools on the basis of alignments that have been partially generated from automatic tools as well. The alignments from [Oberle et al. 2007] have been manually generated to support the task of ontology integration and for their purpose, they are composed of subsumptions (as shown in Table 2.Besides the fact that the ontologies are equipped with object properties, with

[12] `https://raw.githubusercontent.com/BFO-ontology/BFO/releases/1.1.1/bfo.owl`

[13] `http://www.onto-med.de/ontologies/gfo.owl`

[14] `http://www.thezfiles.co.za/ROMULUS/downloads.html` (last view 3rd April 2019)

[15] `http://www.loa.istc.cnr.it/ontologies/DOLCE-Lite.owl`

exception of the BFO 1.1 version, the alignment of this kind of ontology entity is covered to a lesser extent.

Table 2. Reference alignments

Pair	#Concepts	#ObjProp
BFO–DOLCE-Lite	7 $\equiv$	0
BFO–GFO	11 $\equiv$	0
DOLCE-Lite–GFO	9 $\equiv$	6 $\equiv$ 10 $\sqsubset$
DOLCE-Lite [Oberle et al. 2007]–SUMO	41 $\sqsubset$	0

4.2. Matchers

All tools participating in the 2018 edition of the OAEI campaigns for schema-based tracks were selected: ALIN, ALOD2Vec, AML, DOME, FCAMapX, Holontology, KEPLER, Lily, LogMap, LogMapLt, POMAP++ and XMap. These tools implement different matching strategies. The reader can refer to the OAEI papers for details on the tools[16]. All the tools were run with their default configuration settings. As stated before, the aim here is not to evaluate the matching systems themselves, for that reason, in the following we anonymized the systems.

5. Results and discussion

Table 3 shows the results reported in [Khan and Keet 2013c] for the pairs BFO–DL, BFO–GFO and DOLCE-Lite–GFO. As stated in Section 3, in that work the alignments generated by a set of matching systems have been manually evaluated in terms of precision.

Table 4 presents the results for the matchers which were able to generate a non empty alignment for a least one pair of ontologies. In terms of F-measure, the best results were observed for the pair involving BFO and GFO ontologies, from which 6 out 11 correspondences refer to lexically close terms (e.g., bfo:SpatialRegion and gfo:Spatial_Region) (2 out of 7 for BFO-DL and 4 out of 26 for DOLCE-Lite–GFO). For the pair BFO–DOLCE-Lite, we observe that all matchers report the same recall, as they were able to retrieve 2 out of 7 correct correspondences (bfo:Quality and dolce:quality, and bfo:Process and dolce:process). Overall, matchers still mostly output correspondences between concepts whose associated terms are similarly written (for instance, involving exact match or substring match e.g., dolce:boundary and gfo:has_boundary or involving head modifier as for bfo:Object gfo:Material_object). While some are correct, such as the examples above, the ones involving different terms are in general incorrect (e.g, dolce:overlaps and gfo:requirement_of).

There is only one common system in both experiments, but with different versions, what explains its different results in terms of precision (in particular for the pairs BFO–GFO and DOLCE-Lite–GFO). No matcher were able to generate subsumption relations, as no matcher was able to find the correspondences for DOLCE-Lite and SUMO. For this pair, only one correspondence refers to similar terms (dolce:geographical-object and sumo:GeographicArea). From the 41 correspondences in the reference alignment, 5 of them could have been found via a head modifier method (e.g., dolce:organization and sumo:EducationalOrganization or dolce:organization and sumo:PoliticalOrganization).

Table 3. Results in terms of precision, as reported in [Khan and Keet 2013c].

System	BFO–DOLCE-Lite		BFO–GFO		DOLCE-Lite–GFO	
	#	P	#	P	#	P
H-Match	4	.25	5	.16	4	.16
PROMPT	3	.25	7	.58	8	.66
LogMap	2	1.0	11	.91	3	1.0
YAM++	4	1.0	6	.85	13	.52
HotMatch	3	1.0	7	1.0	10	.83
Hertuda	3	1.0	7	1.0	11	.84
Optima	4	.30	9	.52	7	.17
Average	3.28	.69	7.42	.60	8	.72

Table 4. Classical precision (P), recall (R) and F-measure (F) of matchers.

System	BFO–DOLCE-Lite				BFO–GFO				DOLCE-Lite–GFO				DOLCE-Lite–SUMO			
	#	P	F	R	#	P	F	R	#	P	F	R	#	P	F	R
M1	2	1.0	.44	.29	7	.86	.67	.55	9	.56	.29	.20	0	.00	.00	.00
M2	4	.50	.36	.29	6	.83	.59	.45	5	.40	.13	.08	18	.00	.00	.00
M3	3	.67	.40	.29	7	.86	.67	.55	19	.26	.23	.20	22	.00	.00	.00
M4	2	1.0	.44	.29	7	.86	.67	.55	4	.75	.21	.12	11	.00	.00	.00
M5	4	.50	.36	.29	7	.71	.56	.45	13	.38	.26	.20	22	.00	.00	.00
M6	2	1.0	.44	.29	7	.71	.56	.45	9	.75	.21	.12	15	.00	.00	.00
M7	2	1.0	.44	.29	7	.86	.67	.55	9	.56	.29	.20	15	.00	.00	.00
M8	4	.50	.36	.29	7	.86	.67	.55	5	.60	.20	.12	17	.00	.00	.00
Average	2.88	.77	.41	.29	6.88	.82	.63	.51	9.13	.53	.23	.16	15	.00	.00	.00

In order to see how close the generated alignments were to the reference, we have calculated the relaxed precision and recall [Ehrig and Euzenat 2005]. That will consider if closer correspondences than the ones given by the reference were found. For the pairs BFO-DL, BFO-GFO and DL-GFO, the results for relaxed precision and recall were the same than the ones reported using the classical precision and recall. This shows that no other correspondence closer to the reference were found. For the pair DL-SUMO, however, we observed that closer correspondences have been generated, as Table 5 shows. However, the results are still poorer for this pair with respect to the others.

Overall, there is a slightly improvement in the average results (.66 from those reported results to .70 reported here). For the pairs BFO–DOLCE-Lite and BFO–GFO we can observe a relatively significant improvement in terms of precision (from .69 up to .77 and from .60 up to.82), the results for the pair DOLCE-Lite–GFO decreases (.53 and .72, respectively).

For two ontology pairs we observe that more recent systems are performing with better precision. However it has dropped in one case. However, for a new added pair involving SUMO, annotated with subsumption relations only, the results were clearly worse. Although we can see some progress, it is not consistent throughout the entire alignments base, and no progress was made towards subsumption relations.

[16] http://www.om2018.ontologymatching.org/#ap

Table 5. Relaxed precision (P), recall (R) and F-measure (F) of matchers.

System	DOLCE-Lite–SUMO		
	P	F	R
M1	.00	.00	.00
M2	.33	.18	.15
M3	.39	.27	.21
M4	.77	.34	.21
M5	.32	.25	.17
M6	.28	.14	.12
M7	.57	.31	.21
M8	.50	.42	.21
Average	.40	.23	.16

Overall, we could observe that:

- besides the fact that they are not specifically designed to the task – to the best of our knowledge there is no matcher designed to match foundational ontologies – general purpose matchers are not able to correctly deal with the level of abstraction of foundational ontologies;
- in general there was only a small quantity of aligned concepts by the matchers (column # in Table 4);
- there were many incorrect correspondences (in particular for DL-SUMO), many string matching cases which are usually safe in same domain correspondences do not has the same impact here;
- there is a lack of comprehensive evaluation data sets to evaluate this task, and the reference alignments used in the experiments reported in this paper should be further extended in order to be exploited in OAEI evaluation campaigns;
- knowledge on foundational ontologies is highly specialized, it is crucial that such evaluation considers an overview of experts in this area;
- matching strategies for dealing with this task should consider a variety of input, such as 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 an essential point.

6. Conclusions and perspectives

This paper presented an analysis of the alignments between four top-level ontologies, BFO, DOLCE, GFO and SUMO. Our goal was to analyse the behaviour of current state of the art tools, which apply diverse matching techniques, with respect to this task. We could observe that matching top-level ontologies automatically is a challenging task, in particular when involving subsumption relations.

Overall, the results found here are in line to what has been reported when evaluating the behaviour of matchers in the task of matching domain and foundational ontologies, which would also require identification of subsumption relations

[Schmidt et al. 2016a]. Current tools fail on correctly capturing the semantics behind the ontological concepts, what requires deeper contextualization of the concepts on the basis of their hierarchy and axioms.

Besides that, the task requires the identification of other relations than equivalences, such as subsumption and meronym. The latter is largely neglected by current matchers. Most of them typically still rely on string-based techniques as an initial estimate of the likelihood that two elements refer to the same real world phenomenon, hence the found correspondences represent equivalences with concepts that are equally or similarly written. However, in many cases, this correspondence is wrong [Schmidt et al. 2016b].

Hence, matching systems need to be improved to better exploit the knowledge encoded in the ontologies, to include more abstract and philosophical semantic relations and semiotic matching, to take advantage of structural features of the ontologies and axioms in order to better compare their formal definitions, and to take advantage of background knowledge from external resources targeting subsumption and other semantic relations. These have to be combined with logical reasoning techniques for guarantee the consistency of the generated alignments. The current approaches have to be hence revised to better deal with the specificities of matching foundational ontologies.

While the automatic approaches have been mostly manually evaluated, with few exceptions [Damova et al. 2010, Schmidt et al. 2018], systematically evaluations of matching systems have been so far dedicated to domain ontologies. Despite the variety of tasks in the OAEI campaigns[17], the evaluation of matching involving foundational ontologies has not been addressed. Producing comprehensive evaluation data sets on which matching solutions can be evaluated would foster the development of approaches involving foundational ontologies and support a next generation of semantic matching approaches.

Another aspect refers to the evolution or the consistency of alignments with respect to the evolution or the different variants of the ontologies. Evolving alignments to cope with the different versions of the ontologies is still an open challenge.

Last, but not least, very few foundational ontologies are equipped with lexical layers in other natural languages than English (e.g., BFO has been enriched with a lexical annotation in Portuguese). However, with the increasing amount of multilingual data on the Web and the consequent development of ontologies in different natural languages, foundational ontologies should also be equipped with richer multilingual annotations in order to facilitate the multilingual and cross-lingual ontology matching tasks.

As future work, we plan to work on an approach that makes use of the knowledge encoded in the ontologies, using hypernym relation extraction strategies such as lexico-synthactic patterns. For instance, when applied on the definition below, the hypernym relations (Self Connected Object, planet), (Self Connected Object, star), (Self Connected Object, asteroid) can be identified through such patterns.

```
<owl:Class rdf:about="#AstronomicalBody">
  <rdfs:comment> The Class of all astronomical objects of significant
     size. It includes Self Connected Objects like planets, stars, and
     asteroids ...
```

[17]http://oaei.ontologymatching.org/2018/

```
 </rdfs:comment>
</owl:Class>
```

We also plan to involve evaluators with expertise in foundational ontologies to extend the alignments used here as reference; and to propose an OAEI task in order to promote the development of matchers able to deal with the task, within an interactive matching process.

Acknowledgments

We warmly thank D. Oberle for sending us all the generated alignments between SUMO and DOLCE-Lite.

Referências

Arp, R., Smith, B., and Spear, A. (2015). *Building Ontologies with Basic Formal Ontology*. MIT Press.

Damova, M., Kiryakov, A., Simov, K. I., and Petrov, S. (2010). Mapping the central LOD ontologies to PROTON upper-level ontology. In *Proceedings of the Workshop on Ontology Matching*.

Ehrig, M. and Euzenat, J. (2005). Relaxed precision and recall for ontology matching. In *Integrating Ontologies '05, Proceedings of the K-CAP 2005 Workshop on Integrating Ontologies, Banff, Canada, October 2, 2005*.

Gangemi, A., Guarino, N., Masolo, C., Oltramari, A., and Schneider, L. (2002). Sweetening Ontologies with DOLCE. In *Proceedings of the 13th Conference on Knowledge Engineering and Knowledge Management*, pages 166–181.

Gangemi, A., Navigli, R., and Velardi, P. (2003). *The OntoWordNet Project: Extension and Axiomatization of Conceptual Relations in WordNet*, pages 820–838.

Grenon, P. (2003). BFO in a Nutshell: A Bi-categorial Axiomatization of BFO and Comparison with DOLCE. Technical report, Leipzig.

Grenon, P., Smith, B., and Goldberg, L. (2004). Biodynamic Ontology: Applying BFO in the Biomedical Domain. In *Stud. Health Technol. Inform*, pages 20–38. IOS Press.

Guha, V. and Lenat, D. (1993). Cyc: A Midterm Report. In *Readings in Knowledge Acquisition and Learning*, pages 839–866. Morgan Kaufmann Publishers Inc.

Guizzardi, G. (2005). *Ontological foundations for structural conceptual models*. PhD thesis, University of Twente, Enschede, The Netherlands, Enschede.

Guizzardi, G. and Wagner, G. (2010). Using the unified foundational ontology (ufo) as a foundation for general conceptual modeling languages. In *Theory and Applications of Ontology: Computer Applications*, pages 175–196.

Herre, H., Heller, B., Burek, P., Hoehndorf, R., Loebe, F., and Michalek, H. (2007). General Formal Ontology (GFO): A Foundational Ontology Integrating Objects and Processes. In *Basic Principles, Research Group Ontologies in Medicine*.

Keet, C. (2011). The use of foundational ontologies in ontology development: An empirical assessment. In *ESWC*, pages 321–335.

Khan, Z. and Keet, C. (2012). ONSET: Automated Foundational Ontology Selection and Explanation. In *Proceedings of the 18th International Conference on Knowledge Engineering and Knowledge Management*, pages 237–251.

Khan, Z. and Keet, C. (2013a). Addressing issues in foundational ontology mediation. In *Proceedings of the International Conference on Knowledge Engineering and Ontology Development*, pages 5–16.

Khan, Z. and Keet, C. (2013b). The Foundational Ontology Library ROMULUS. In *Proceedings of the 3rd International Conference on Model and Data Engineering*, pages 200–211.

Khan, Z. and Keet, C. (2013c). Toward Semantic Interoperability with Aligned Foundational Ontologies in ROMULUS. In *Proceedings of the 7th International Conference on Knowledge Capture*, pages 23–26.

Khan, Z. and Keet, C. (2014). Feasibility of Automated Foundational Ontology Interchangeability. In *Proceedings of the 19th Conference on Knowledge Engineering and Knowledge Management*, pages 225–237.

Mascardi, V., Cordì, V., and Rosso, P. (2007). A Comparison of Upper Ontologies. In *Proceedings of the 8th AI*IA/TABOO Workshop on Agents and Industry*, pages 55–64.

Mascardi, V., Locoro, A., and Rosso, P. (2010). Automatic Ontology Matching via Upper Ontologies: A Systematic Evaluation. *IEEE Transactions on Knowledge and Data Engineering*, 22(5):609–623.

Mika, P., Oberle, D., Gangemi, A., and Sabou, M. (2004). Foundations for Service Ontologies: Aligning OWL-S to DOLCE. In *Proceedings of the 13th Conference on World Wide Web*, pages 563–572.

Muñoz, L. and Grüninger, M. (2016). Verifying and mapping the mereotopology of upper-level ontologies. In *Proceedings of the International Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management*, pages 31–42.

Nardi, J. C., de Almeida Falbo, R., and Almeida, J. P. A. (2013). Foundational ontologies for semantic integration in EAI: A systematic literature review. In *Proceedings of the 12th IFIP WG Conference on e-Business, e-Services, and e-Society, I3E*, pages 238–249.

Niles, I. and Pease, A. (2001). Towards a Standard Upper Ontology. In *Proceedings of the Conference on Formal Ontology in Information Systems*, pages 2–9.

Oberle, D., Ankolekar, A., Hitzler, P., Cimiano, P., Sintek, M., Kiesel, M., Mougouie, B., Baumann, S., Vembu, S., Romanelli, M., Buitelaar, P., Engel, R., Sonntag, D., Reithinger, N., Loos, B., Zorn, H., Micelli, V., Porzel, R., Schmidt, C., Weiten, M., Burkhardt, F., and Zhou, J. (2007). DOLCE Ergo SUMO: On Foundational and Domain Models in the SmartWeb Integrated Ontology (SWIntO). *Web Semantics*, 5(3):156–174.

Schmidt, D., Basso, R., Trojahn, C., and Vieira, R. (2018). Matching domain and top-level ontologies exploring word sense disambiguation and word embedding. In *Emerging Topics in Semantic Technologies [best papers from the workshops at ISWC 2018]*, pages 27–38.

Schmidt, D., Trojahn, C., and Vieira, R. (2016a). Analysing Top-level and Domain Ontology Alignments from Matching Systems. In *Proceedings of the Workshop on Ontology Matching*, pages 1–12.

Schmidt, D., Trojahn, C., and Vieira, R. (2016b). Analysing Top-level and Domain Ontology Alignments from Matching Systems. In *Proceedings of the 11th Workshop on Ontology Matching*, pages 1–12.

Semy, S., Pulvermacher, M., and Obrst, L. (2004). Toward the use of an upper ontology for U.S. government and U.S. military domains: An evaluation. Technical report, MTR 04B0000063, The MITRE Corporation.

Seppälä, S., Smith, B., and Ceusters, W. (2014). Applying the realism-based ontology-versioning method for tracking changes in the basic formal ontology. In *Proceedings of FOIS*, pages 227–240.

Seyed, A. (2009). BFO/DOLCE Primitive Relation Comparison. In *Nature Proceedings*.

Temal, L., Rosier, A., Dameron, O., and Burgun, A. (2010). Mapping BFO and DOLCE. In *Proceedings of the World Congress on Medical Informatics*, pages 1065–1069.

Terziev, I., Kiryakov, A., and Manov, D. (2005). Base Upper-level Ontology (BULO) Guidance. Deliverable 1.8.1, SEKT project.