September 2017

A note on the compatibility of part-whole relations with foundational ontologies

C. Maria KEET [1]
Department of Computer Science, University of Cape Town, South Africa

Abstract. Parthood in mereology is one relation, and typically is included in foundational ontologies. Some of these foundational ontologies and many domain ontologies use a plethora of parthood and part-whole relations, such as 'sub process' and 'portion'. This poses requirements on the foundational ontologies and, perhaps, Ontology, on what to do with these two different approaches to part-whole relations. We present an analysis of DOLCE, BFO, GFO, SUMO, GIST, and YAMATO on their inclusion and use of part-whole relations. It demonstrates there is no perfect fit with either for various reasons. We then aim to bridge this gap with an orchestration of ontologies of part-whole relations that are aligned to several foundational ontologies and such that they can be imported into other ontologies.

Keywords. Mereology, part-whole relations, ontology engineering

1. Introduction

Foundational ontologies (FOs) differ in their inclusion of part-whole relations. For instance, DOLCE [12] and GFO [4] have only mereological parthood and proper parthood, and one or two meronymic relations, such as participation and membership. However, even in ontology literature, more specific parthood relations are recognised [8], notably portions [2,6], relating processes [12,20], and, more generally, the 'multitude' approach to part-whole relations as first proposed comprehensively in [24]. SUMO [17] has many part-whole relations, including having them tailored to what might be considered domain entities, such as a sub plan for stating that one plan is a proper part of another plan, and parts with a temporal flavour, like initially contains part. Domain ontologies also use a plethora of variations, with the medical terminology GALEN a case in point with 23 part-whole relations [19]. Note that this does not yet even concern inclusion of various formalisations of mereological theories in FOs, which is assessed in [3], but at most the bare minimum of ground mereology together with all those variants, such as sub-processes, portions, and pieces and so on, which is what we focus on here.

From a minimalist perspective, the choice to include in one's FO the one single mereological parthood relation may sound appealing (and may be the only

[1]Corresponding Author: C. Maria Keet, Department of Computer Science, University of Cape Town, Cape Town, South Africa; E-mail: mkeet@cs.uct.ac.za.

option according to an ontologist), but it is not workable in domain ontologies, for it permits many modelling errors and can cause many undesirable deductions as a result, which should be prevented. For instance, relating an object to a process with a parthood relation would violate mereology—an object *participates in* a process instead—but can be declared in an ontology if the domain and range are declared to be just any entity (or `owl:Thing` in an OWL ontology). Then the reasoner will not detect it as a mistake and instead return *undesirable* deductions; e.g., if an atom is part of a protein (both objects) that is part of a catalytic reaction (a process), then it would undesirably infer that the atom is part of the catalytic reaction. At the other end of the spectrum is the proliferation of part-whole relations such that it becomes hard to obtain any inferences at all. If there is a 'part of X' way of naming relations (e.g., hasHand, hasArm), then one cannot obtain deductions from a partonomy thanks to part of's transitivity (e.g., that the hand is part of the human).

The questions these issues raise thus concern 1) how to bridge the Ontological 'cleanliness' of mereological theories with the practicalities of domain ontology development were multiple part-whole relations are used, and 2) where, or when, to stop proliferating naming part-whole relations. In this paper we aim to contribute to resolving such a balancing act toward a 'sweet spot' for part-whole relations. Because those more specialised part-whole and parthood relations distinguish themselves by their domain and range, we conduct a feature comparison among six FOs. First, we zoom in on their inclusion of mereological parthood and other part-whole relations, and, second, assess their coverage of top-level categories used to specify the meaning of part-whole relations. Based on the outcome of the comparison, we take an engineering approach by developing ontology modules that can be imported into FOs and thus be used for practical ontology development. These modules contain the taxonomy of part-whole relations proposed in [8], stuff relations [6], and mereotopological ones [9] and are linked to DOLCE and SUMO (manually) and to BFO and GFO (automatically).

The remainder of the paper is structured as follows. Section 2 summarises part-whole relations. Section 3 discusses the comparison and Section 4 presents the orchestration of the OWL files. We conclude in Section 5.

2. Part-whole relations beyond just part-of

The notion of a 'multitude' of part-whole relations started in earnest with Winston, Chaffin, and Herrmann's paper [24], which deviates from the notion of a single mereological parthood relation. Over time, it has been elucidated that there are distinctions in that some relations are parthood indeed in the sense of mereology (as in, e.g., [22]), and in other cases they are just meronymic part-whole relations (in natural language language, but not necessarily ontologically), and that *for the purpose of modelling* in ontologies and conceptual models—as opposed to analytic philosophy—it serves to use different names when the part-whole relation relates objects of different categories. Specification of a relation's domain and range results in a more precise representation of its intended meaning, as also noted in, among others, [8,18,23]. This 'multitude' approach resulted in a common list of part-whole relations. The scope of the paper is not to argue for which ones are there, but accepts that there are several and that that has to be interoperable with FOs. We take the one of [8] extended with its stuff relations and portions [6] and mereotopological ones [9], which are based on extensive literature reviews. The basic set of part-whole relations has been relatively stable, and is shown informally in Fig. 1 and annotated with the domain and range types. Needing to declare the domain and range poses the question which categories from which FO should be used. For instance: (a) Are involved-in's 'processes' in Fig. 1 processes-as-in-BFO or processes-as-in-DOLCE (or whichever other FO)? (b) Is there a FO that has all the necessary domain and range categories? If a FO does not have a particular category, like portion-of's 'stuff' (amounts of matter, indicated typically with mass nouns), then what can, or should, be done? We will answer these questions in the next sections.

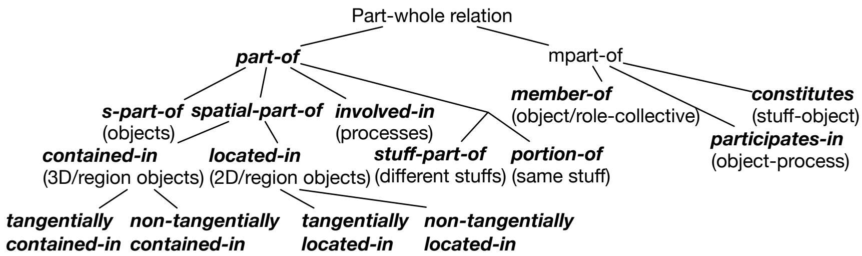

Figure 1. Summarised sketch of the part-whole taxonomy with mereotopological extension, with mpart-of merely a stub for grouping the meronymic part-whole relations together.

3. FO content comparison for part-whole relations

We conduct a brief analysis and comparison of selected FOs on their suitability for being used in, or have incorporated, part-whole relations. For this content comparison, we selected six FOs, which are:

DOLCE: The Descriptive Ontology for Linguistic and Cognitive Engineering [12] is selected because the original part-whole relations taxonomy was already aligned to it [8];

BFO: The Basic Formal Ontology [1] enjoys great uptake in the biological and biomedical domains as part of the OBO Foundry efforts [21];

GFO: The General Formal Ontology [4] is already aligned to DOLCE and BFO [11] and therewith making the comparison exercises easier;

SUMO: The Standard Upper Merged Ontology [17] was an early FO and it has relatively many relations.

GIST: The GIST minimalist upper ontology [13] was updated in 2017, and is therefore included;

YAMATO: The Yet Another More Advanced Top-level Ontology, which emphasises qualities and processes and events [14].

The UFO was not included, because we could not find an OWL file of it. Concretely, the file `DOLCE-Lite.owl` in DLP397.zip was used for DOLCE, as that is what most ontology developers will use if they use it (cf. the axiomatisation in [12]), `bfo.owl` (v2), the full GFO (cf. basic) in `gfo.owl`, the `SUMO.owl` file from 2010, `gistCore7.5.owl`, and `YAMATO20120714.owl`.

The respective OWL files were examined on two aspects: 1) content pertaining to the relations themselves, being:

i) whether they include the basic parthood (P) and proper parthood (PP);
ii) their respective use of the usual characteristics in ground mereology [22], where part-of is reflexive (Ref.), antisymmetric (Anti.), and transitive (Trans.) and proper part of is irreflexive (Irr.), asymmetric (Asym.), and transitive;
iii) inclusion in the OWL file (anywhere) of property chains (Chain), qualified cardinality (Qual.), and object property hierarchies (H), because they are 'interesting' computationally, in that not all languages have those features and no Description Logics-based OWL species can have them all together with the parthood characteristics [16];
iv) which other part-whole relations they include, if any;

and 2) their main categories (classes/concepts/kinds/types) of entities, which may be more or less applicable as domain or range for the part-whole relations. The results are discussed in the next two subsections.

3.1. Part-whole relations in FOs

First, we compare the FOs' contents regarding part-whole relations, which is encoded in Table 1. There are several observations to make. First, the DOLCE and GFO releases were before OWL 2 was standardised in 2008 and they have not been updated accordingly. This means that reflexivity, irreflexivity, asymmetry, property chains, and qualified cardinality could not have been used and thus necessarily have a "–" in the table. The other four FOs have later release dates, and thus could have, with caution, more or other 'object property characteristics' than just transitivity. For instance, irreflexivity could have been asserted on parthood, at the cost of omitting transitivity[2]. Also, antisymmetry is not available in OWL 2 [16][3], hence it necessarily has a "–" for all FOs. Both issues also explain why one cannot even represent ground mereology fully in OWL 2 DL; if a modeller wants them all, one could use, e.g., the Distributed Ontology Language [15].

The second obvious observation is that BFO v2.0 has "–" everywhere: it is a bare taxonomy. While BFO v2.0 is richly annotated, this does not count from a formal (logic-based) perspective. There are three close alternatives one could consider: the BFO-associated Relation Ontology (RO) [20], the ROcore[4], and the draft release of BFO v2.1[5]. The RO's informal domains and ranges are not taken from a specific FO [20], though this is likely to be BFO; we shall return to this aspect in the next section. The more recent ROcore does indeed contain

[2]irreflexivity can be declared on 'simple' relations, but a transitive relation is not 'simple' [16]
[3]'antisymmetry of the irreflexive kind', what the original $\mathcal{SROIQ}$ paper claims [5], amounts to asymmetry
[4]`https://github.com/oborel/obo-relations/wiki/ROCore`
[5]`https://raw.githubusercontent.com/BFO-ontology/BFO/releases/2014-05-03/owl-group/bfo.owl`

Table 1. Main characteristics of part-whole relations in the selected FOs;

FO	P	PP	Ref.	Anti.	Trans.	Asym.	Irr.	H	Chain	Qual.	Other part-whole relations
DOLCE	+	+	–	–	+	–	–	+	–	–	participant, temporary part (for endurants)
BFO 2.0	–	–	–	–	–	–	–	–	–	–	–
GFO	+	+	–	–	+	–	–	+	–	–	participant, member
SUMO	+	+	–	–	–	–	–	+	–	–	geometric part, subprocess, member, sub plan, ...
GIST	+	–	–	–	+	–	–	+	–	+	direct part, member, made up of
YAMATO	+	–	–	–	–	–	–	+	–	–	very many, including 'has_X' [as part]

Table 2. Alignments of the domain and range categories used in defining part-whole relations.

Domain/range	DOLCE	BFO	GFO	SUMO	GIST	YAMATO
none	`owl:Thing`	`owl:Thing`	`owl:Thing`	`owl:Thing`	`owl:Thing`	`owl:Thing`
Particular	Particular	Entity	Individual	Entity	–	Particular
Endurant	Endurant	Independent Continuant	Presential	Object	–	± continuant
Perdurant	Perdurant	Occurrent	Occurrent	Process	± Event	occurrent
Physical Object	Physical Object	Object	Material Object	± Self-connected object	Physical identifiable item	± object
Social Object	Social Object	–	± Social role	Group	–	–
Stuff	Amount of matter	–	Amount of Substrate	Substance	Physical substance	amount of matter
Region	Region	± union of Spatial, Temporal, and Spatio-temporal region	± Spatial and Temporal region	± Region	–	space
Collective	–	± object aggregate	–	Collection	Collection	–

the relevant BFO classes, it has a transitive part-of, property hierarchies, and also location (subsuming both the contained-in and located-in of [8]), participates in with domain continuant and range occurrent, and an irreflexive has-member for items and collections. Thus, while it does not include all proposed part-whole relation variants, there is a notable commonality. BFO v2.1 is problematic in that its relations have no compatibility with either the RO or ROcore, and very limitedly with or any other ontology, due to the fact that most names of the relation would suggest temporality, yet OWL is atemporal. The part of occurrent matches involved-in of the aforementioned part-whole taxonomy, but any other part-whole relation in v2.1 is 'temporalised' in its name, such as part of continuant at some time, and as such different in intended meaning. This is similar for DOLCE's temporary-part-of between endurants that was a time-indexed part in [12] (DOLCE's parthood between perdurants in [12], the same as involved-in in [8], is not included as such in `DOLCE-Lite`).

SUMO, GIST, and YAMATO stand in stark contrast to DOLCE, BFO, and GFO, for they do contain multiple part-whole relations. For instance, SUMO's subProcess matches with the taxonomy's involved-in, member with member-of, and material with constitutes. The case is similar for GIST, such as its made up of matching constituted-of (the inverse of constitutes). While both SUMO and GIST are FOs, they also contain a substantial amount of knowledge that would typically not be considered within the realm of FOs but of domain ontologies, such as types of organisations, manufacturer, types of groups, or a building address. YAMATO lies at the other extreme for part-whole relations. It essentially has introduced relations for each part-whole relation between distinct domain objects; e.g. has-arm, has-brain, and has-mouth are subproperties of hasPart. It is clear that these domains and ranges are domain entities rather than categories in FOs, such as 'process' or 'object'. However, the boundary between what a subject domain entity is and what a FO one is probably not crisp. In the interest of reusability of FOs with some part-whole relations, one would assume that only the domain-independent ones should be included.

Overall, based on just the analysis of the contents on part-whole relations, one could argue for aligning the part-whole relations with SUMO, for there is the most overlap, or BFO, for it would not need any alignment of the relations.

3.2. Possible domain and range categories in FOs

The second step is to examine whether the common part-whole relations could perhaps be easily slotted in—i.e., the FOs extended—or be designed as an importable 'foundational ontology module of part-whole relations' that would easily align with the FO. To assess this, we list first which domains and ranges are needed in the first column in Table 2. The first eight categories follow trivially from the literature and Fig. 1. The last one, collective, was added because it came afore in the assessment of verbalising part-whole relations in isiZulu, the largest language by first language speakers in South Africa, and its subsequent ontological analysis [7]. In short: while participates-in is understood as relating any endurant participating in a perdurant, in isiZulu (the Zulu language) or by the amaZulu (the Zulu people), it has that general notion of 'participates in' as well as one for collectives specifically (*-hlanganyele*), such as an electorate participating in an election (cf., e.g., a voter participating in an election) [7].

We availed of the FO comparison of [11] for DOLCE, BFO, and GFO, and added a tentative alignment for the others. The outcome is listed in Table 2, where a "±" indicates that we deemed it 'roughly the same' based on the available information (documentation, annotations, and use), but not fully, and a "–" that no match could be found. Among the latter, two refinements might be made, in the sense of 'some of it is there' vs. 'nothing of the like is there'. For instance, endurant, social object, and region do have a very limited presence in GIST: there are several classes that, when joined and asserted as subclass, would approximate it. This is in contrast to not having collective in DOLCE and no stuff in BFO. From a pedantic stance, one may assume that the DOLCE and BFO authors intentionally excluded it, i.e., the foundational ontological choice was made to, by their absence, assert that those kind of entities do not exist. Extending the FO with such entities would violate those principles. But, perhaps, it was an oversight or the FO is considered still under construction.

Based on the comparison in the table, SUMO seems a good candidate. Zooming in on the two "±"s, a SelfConnectedObject is "any Object that does not consist of two or more disconnected parts", which is not the same as a physical object in DOLCE that is an "... endurant[s] with unity. However, [it has] no common unity criterion...". To illustrate: a bikini is a physical object, but not a self-connected one. SUMO's Region is "A topographic location", i.e., just the '2D' part of space and is a subclass of Object, whereas DOLCE's Region is subsumed by Abstract and subsumes, among others, any physical region, hence, also the '3D' space. Logically, they could be aligned through subsumption because SUMO does not have a disjointness axiom between Physical (the parent of Object) and Abstract. However, ontologically by their parent categories, they cannot. Whether that distinction is a disagreement among the Ontology of the developers or just some artefact of modelling remains to be seen. From a user-based perspective to FOs, as is the case here, one cannot ignore such distinctions.

Overall, there is no exact match for either FO regarding declaring domains and ranges for the part-whole relations, but DOLCE and SUMO come closest.

4. Operationalising part-whole relations in OWL files

The assessment in the previous section indicates that it will be easier to align the part-whole relations to DOLCE, BFO, or SUMO than the others. The orchestration of the ontology modules is depicted in Fig. 2. Their basic statistics are listed in Table 3. All files are available at `http://www.meteck.org/files/ontologies/`.

We commenced with the alignment with DOLCE, for it existed in part already. First, the file named `PW` was created and the part-whole taxonomy ([8] updated with stuff relations [6]), their inverses (`has-part` etc.), the relevant proper part versions of them (with inverses), transitivity, and annotations were added. A module of DOLCE, called, `dolcemini`, was then imported, domains and ranges declared according to Table 2, where applicable, and equivalences to some DOLCE relations were asserted, constituting the complete `PW.owl` file. We did not use the

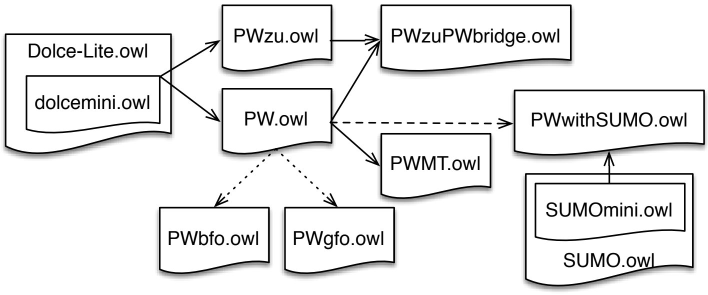

Figure 2. Orchestration of OWL files, where the solid arrow denotes 'imported into', the arrow with dashed shaft denotes a 'conceptual' import (in essence, but implemented differently), and the dotted arrow stands for automated generation with SUGOI.

full DOLCE mainly because of performance reasons. `DOLCE-Lite` was modularised by manually removing unnecessary classes and axioms, and social object and agentive and non-agentive physical objects were added, for they are in the DOLCE documentation [12] and social object is needed for the part-whole relations (recall Table 2), resulting in `dolcemini`. Subsequently, we used the SUGOI tool [10] to interchange DOLCE automatically for BFO and GFO, resulting in `PWbfo` and `PWgfo`, respectively. In turn, `PW` was imported into `PWMT` for the mereotopological extension. This was done rather than adding the basic mere topological relations to `PW`, because they may be used less often, and there were relatively many relations in `PW` already.

The alignment of the part-whole relations with SUMO was carried out manually as follows. First, a small module was created, because the 35MB OWL file[6] from 2010 is practically unworkable. We deleted all instances and most classes and object properties that were subject domain entities; the resultant module is `SUMOmini`. Then, to ensure correct IRIs and avoiding unintentional original file loading, `PW` was saved as `PWwithSUMO` and `sumomini` imported. The swap of `dolcemini` for `sumomini` was carried out following SUGOI's algorithm [10].

Finally, `PWzu` is the part-whole taxonomy informed by isiZulu culture and language [7] and it was aligned to `PW` in the `PWzuPWbridge` file. The main distinctions are some non-1:1 alignments and transitivity was not asserted for its generic parthood (*-ingxenye*), because it does not hold always (see [7] for a discussion).

The notable differences in number of relations and logical axioms for the ontologies in the orchestration are listed in the last two columns of Table 3: just declaring the part-whole relations already count for a third of the axioms in the `PW` ontology and a quarter in `PWwithSUMO`, and between 25-100% of the relations in the ontology files. The largest difference is between `PWMT` and `dolcemini`: it has a total of 103 relations, of which 59 domain-independent part-whole relations, 1 deprecated, and 43 from DOLCE. This may seem off-putting to use, but effective tooling to help selecting the right relation is possible with OntoPartS [9].

[6] `http://www.adampease.org/OP/SUMO.owl`

Table 3. Basic statistics of the OWL files and notable changes; |Ax|: 'logical axiom count' as measured in Protégé 5.x (i.e., excluding annotations); Other: other property characteristics, such as functional and inverses; δ: number of additions cf. its predecessor.

File	|Ax|	|C|	|OP|	|SubOP|	|Trans|	|Other|	$\delta_{\mathsf{ax}}^{prev.}$	$\delta_{\mathsf{rel}}^{prev.}$
Dolce-Lite	349	37	70	70	6	39		
dolcemini (dm)	243	40	43	42	4	21		
PW	364	40	92	98	23	47	δ_{121}^{dm}	δ_{49}^{dm}
PWMT	380	40	103	109	23	52	δ_{16}^{PW}	δ_{11}^{PW}
PWbfo	224	48	55	56	19	26		
PWgfo	341	53	92	78	27	68		
SUMO	175208	4558	778	5330	51	144		
SUMOmini (sm)	363	104	138	0	31	32		
PWwithSUMO	489	104	187	64	50	57	δ_{126}^{sm}	δ_{49}^{sm}
PWzu	272	40	60	57	4	27	δ_{29}^{dm}	δ_{17}^{dm}
PWzuPWBridge	405	40	109	121	23	57		

5. Conclusions and future work

The comparison of six foundational ontologies showed that neither is a perfect fit for typical part-whole relations found in domain ontologies and conceptual models. This is due to language limitations, the stance of relational parsimony vs abundance as fundamental modelling choice for the ontology, and limited coverage of categories of entities (classes/universals) that are needed for most part-whole relations to specify their meaning. Informed by the analysis, several Ontology and FO-informed ontology modules of part-whole relations were created and aligned to four foundational ontologies.

These files solve some issues in ontology development—extending a FO and reaching out to a domain ontology—and indicate a cut-off point for generic part-whole relations. The predominantly engineering approach taken in this paper, however, still leaves much to be investigated. Ontologically, it will be useful to examine whether the mereological theories should be extended in some way to take into account domain and range declarations and therewith accommodate the 'multitude' approach rather that offering just a single parthood relation. That is, perhaps no separate spatial theories with parthood relating regions and theories of portions relating stuff parts to wholes, but all that within one connected network of theories. Extending beyond the current scope of part-whole relations, a more precise specification of the borderline of FO vs top-domain and domain ontology may be of use, as would be revisiting the FOs on inclusion of certain categories.

References

[1] BFO. Basic formal ontology, (last accessed July 2017). http://ifomis.uni-saarland.de/bfo/.

[2] M. Donnelly and T. Bittner. Summation relations and portions of stuff. *Philosophical Studies*, 143:167–185, 2009.

[3] M. Fernández-López, A. Gómez-Pérez, and M. C. Suárez-Figueroa. Selecting and customizing a mereology ontology for its reuse in a pharmaceutical product ontology. In *Proc. of FOIS'08*, pages 181–194. IOS Press, 2008.

[4] H. Herre. General Formal Ontology (GFO): A foundational ontology for conceptual modelling. In *Theory and Applications of Ontology: Computer Applications*, chapter 14, pages 297–345. Springer, 2010.

[5] I. Horrocks, O. Kutz, and U. Sattler. The even more irresistible $\mathcal{SROIQ}$. *Proceedings of KR-2006*, pages 452–457, 2006.

[6] C. M. Keet. Relating some stuff to other stuff. In E. Blomqvist, P. Ciancarini, F. Poggi, and F. Vitali, editors, *Proc. of EKAW'16*, volume 10024 of *LNAI*, pages 368–383. Springer, 2016. 19-23 November 2016, Bologna, Italy.

[7] C. M. Keet. Representing and aligning similar relations: parts and wholes in isizulu vs english. In J. Gracia et al., editors, *Language, Data, and Knowledge 2017 (LDK'17)*, volume 10318 of *LNAI*, pages 58–73. Springer, 2017. 19-20 June, 2017, Galway, Ireland.

[8] C. M. Keet and A. Artale. Representing and reasoning over a taxonomy of part-whole relations. *Applied Ontology*, 3(1-2):91–110, 2008.

[9] C. M. Keet, F. C. Fernández-Reyes, and A. Morales-González. Representing mereotopological relations in OWL ontologies with ONTOPARTS. In E. Simperl et al., editors, *Proc. of ESWC'12*, volume 7295 of *LNCS*, pages 240–254. Springer, 2012. 29-31 May 2012, Heraklion, Crete, Greece.

[10] Z. C. Khan and C. M. Keet. Feasibility of automated foundational ontology interchangeability. In K. Janowicz and S. Schlobach, editors, *Proc. of EKAW'14*, volume 8876 of *LNAI*, pages 225–237. Springer, 2014. 24-28 Nov, 2014, Linkoping, Sweden.

[11] Z. C. Khan and C. M. Keet. Foundational ontology mediation in ROMULUS. In A. Fred et al., editors, *Knowledge Discovery, Knowledge Engineering and Knowledge Management: IC3K 2013 Selected Papers*, volume 454 of *CCIS*, pages 132–152. Springer, 2015.

[12] C. Masolo, S. Borgo, A. Gangemi, N. Guarino, and A. Oltramari. Ontology library. WonderWeb Deliverable D18 (ver. 1.0, 31-12-2003)., 2003. http://wonderweb.semanticweb.org.

[13] D. McComb. Gist: The minimalist upper ontology (abstract). Semantic Technology Conference, 2010. 21-25 June 2010, San Francisco, USA.

[14] R. Mizoguchi. YAMATO: Yet Another More Advanced Top-level Ontology. In *Proceedings of the Sixth Australasian Ontology Workshop (AOW'10)*, Conferences in Research and Practice in Information Technology, pages 1–16. CRPIT, 2010. Sydney : ACS.

[15] T. Mossakowski, C. Lange, and O. Kutz. Three semantics for the core of the Distributed Ontology Language. In M. Grüninger, editor, *Proc. of FOIS'12*. IOS Press, 2012. 24-27 July, Graz, Austria.

[16] B. Motik, P. F. Patel-Schneider, and B. Parsia. OWL 2 web ontology language structural specification and functional-style syntax. W3c recommendation, W3C, 27 Oct. 2009. http://www.w3.org/TR/owl2-syntax/.

[17] I. Niles and A. Pease. Towards a standard upper ontology. In C. Welty and B. Smith, editors, *Proc of FOIS'01*, 2001. Ogunquit, Maine, October 17-19, 2001.

[18] M. Poveda-Villalón, M. C. Suárez-Figueroa, and A. Gómez-Pérez. Validating ontologies with OOPS! In A. ten Teije et al., editors, *Proc. of EKAW'12*, volume 7603 of *LNAI*, pages 267–281. Springer, 2012. Oct 8-12, Galway, Ireland.

[19] J. Rogers and A. Rector. GALEN's model of parts and wholes: experience and comparisons. In *Proceedings of the AMIA Symp 2000*, page 714718. AMIA, 2000.

[20] B. Smith, W. Ceusters, B. Klagges, J. Köhler, A. Kumar, J. Lomax, C. Mungall, F. Neuhaus, A. L. Rector, and C. Rosse. Relations in biomedical ontologies. *Genome Biology*, 6:R46, 2005.

[21] B. Smith et al. The OBO Foundry: Coordinated evolution of ontologies to support biomedical data integration. *Nature Biotechnology*, 25(11):1251–1255, 2007.

[22] A. C. Varzi. Mereology. In E. N. Zalta, editor, *Stanford Encyclopedia of Philosophy*. Stanford, fall 2004 edition, 2004. http://plato.stanford.edu/archives/fall2004/entries/mereology/.

[23] L. Vieu and M. Aurnague. Part-of relations, functionality and dependence. In M. Aurnague, M. Hickmann, and L. Vieu, editors, *Categorization of Spatial Entities in Language and Cognition*. John Benjamins, Amsterdam, 2005.

[24] M. Winston, R. Chaffin, and D. Herrmann. A taxonomy of partwhole relations. *Cognitive Science*, 11(4):417–444, 1987.