

Representing Mereotopological Relations in OWL Ontologies with ONTOPARTS

C. Maria Keet¹, Francis C. Fernández-Reyes², and Annette Morales-González³

¹ School of Computer Science, University of KwaZulu-Natal, South Africa

² Instituto Superior Politécnico “José Antonio Echeverría” (CUJAE), Cuba

³ Advanced Technologies Application Center, CENATAV, Havana City, Cuba
keet@ukzn.ac.za, ffernandez@ceis.cujae.edu.cu, amorales@cenatav.co.cu

Abstract. Representing and reasoning over mereotopological relations (parthood and location) in an ontology is a well-known challenge, because there are many relations to choose from and OWL has limited expressiveness in this regard. To address these issues, we structure mereotopological relations based on the KGMT mereotopological theory. A correctly chosen relation counterbalances some weaknesses in OWL’s representation and reasoning services. To achieve effortless selection of the appropriate relation, we hide the complexities of the underlying theory through automation of modelling guidelines in the new tool ONTOPARTS. It uses, mainly, the categories from DOLCE [12], which avoids lengthy question sessions, and it includes examples and verbalizations. ONTOPARTS was experimentally evaluated, which demonstrated that selecting and representing the desired relation was done efficiently and more accurately with ONTOPARTS.

1 Introduction

Part-whole relations are essential for knowledge representation, in particular in terminology and ontology development in subject domains such as biology, medicine, GIS, and manufacturing. Usage of part-whole relations are exacerbated when part-whole relations are merged with topological or mereotopological relations, such as *tangential proper part* where the part touches the boundary of the whole it is part of; e.g., the FMA has 8 basic locative part-whole relations [14] and GALEN has 26 part-whole and locative part-whole relations¹. It is also useful for annotating and querying multimedia documents and cartographic maps; e.g., annotating a photo of a beach where the area of the photo that depicts the sand *touching* the area that depicts the seawater so that, together with the knowledge that, among other locations, Varadero is a *tangential proper part of* Cuba, the semantically enhanced system can infer possible locations where the photo has been taken, or vv., propose that the photo may depict a beach scene.

Efforts have gone into figuring out which part-whole relations there are [24,11], developing a logic language with which one can represent the semantics of the

¹ <http://www.opengalen.org/tutorials/crm/tutorial9.html>
up to /tutorial16.html

relations [2,8], and how to use the two together [20,25,26]. The representation of mereotopology in Description Logics (DL) has not been investigated, but related efforts in representing the Region Connection Calculus (RCC) in DLs have [17,9,21,23,6]. Currently, the advances in mereotopology are not directly transferrable to a Semantic Web setting due to the differences in languages and theories and they miss software support to make it usable for the ontology developer. Yet, ontologists require a way to effectively handle these part-whole relations during ontology development without necessarily having to become an expert in theories about part-whole relations, mereotopology, and expressive ontology languages. Moreover, structured and guided usage can prevent undesirable deductions and increase the amount of desirable deductions even without the need to add additional expressiveness to the language. For instance, instance classification: let NTPLI be a ‘non-tangential proper located in’ relation and $\text{EnclosedCountry} \equiv \text{Country} \sqcap \exists \text{NTPLI}.\text{Country}$, and instances $\text{NTPLI}(\text{Lesotho}, \text{South Africa}), \text{Country}(\text{Lesotho}), \text{Country}(\text{South Africa})$, then it will correctly deduce $\text{EnclosedCountry}(\text{Lesotho})$. With merely ‘part-of’, one would not have been able to obtain this result.

Thus, there are three problems: (i) the lack of oversight on plethora of part-whole relations, that include real parthood (mereology) and parts with their locations (mereotopology), (ii) the challenge to figure out which one to use when, and (iii) underspecified representation and reasoning consequences. To solve these problems we propose the **ONTOPARTS** tool to guide the modeller. To ensure a solid foundation, transparency, a wide coverage of the types of part-whole relations, and effectiveness during ontology development, we extend the taxonomy of part-whole relations of [11] with the novel addition of mereotopological relations, driven by the **KGEMT** mereotopological theory [22], resulting in a taxonomy of 23 part-whole relations. We describe the design rationale and trade-offs with respect to what has to be simplified from **KGEMT** to realise as much as possible in **OWL** so that **ONTOPARTS** can load **OWL/OWL2**-formalised ontologies, and, if desired, modify the **OWL** file with the chosen relation. To enable quick selection of the appropriate relation, we use a simplified **OWL**-ized **DOLCE** ontology for the domain and range restrictions imposed on the part-whole relations and therewith let the user take ‘shortcuts’, which reduces the selection procedure to 0-4 options based on just 2-3 inputs. The usability of **ONTOPARTS** and effectiveness of the approach was evaluated and shown to improve efficiency and accuracy in modelling.

In the remainder of the paper, we first describe the theoretical foundation of the mereotopological relations and trade-offs for **OWL** (Section 2). We describe the design, implementation, and evaluation of **ONTOPARTS** in Section 3, discuss the proposed solution in Section 4, and conclude in Section 5.

2 Mereotopology and OWL

Part-whole relations and the differentiation between different types of part-whole relations has its origins in cognitive science and conceptual data modelling [24,18]

and has been investigated also for Description Logics (e.g., [1,2]). There is a first important distinction between parthood versus a meronymic relation ('part' in natural language only), and, second, there is an additional aspect on parthood and location [11]. The second dividing characteristic is the domain and range of the relations (which are taken from the DOLCE foundational ontology [12] in [11]). Particularly relevant here are the containment and location axioms (Eqs. 1 and 2) [11], where $ED = \text{EnDurant}$ (enduring entity, e.g., a wall), $R = \text{Region}$ (e.g., the space that the wall occupies, the Alpine region), and has_2D and has_3D for surfaces and space are shorthand relations standing for DOLCE's qualities and qualia; note that the domain and range is *Region* that has an object occupying it, hence, this does not imply that those objects are related also by structural parthood.

$$\forall x, y(\text{contained_in}(x, y) \equiv \text{part_of}(x, y) \wedge R(x) \wedge R(y) \wedge \exists z, w(\text{has_3D}(z, x) \wedge \text{has_3D}(w, y) \wedge ED(z) \wedge ED(w))) \quad (1)$$

$$\forall x, y(\text{located_in}(x, y) \equiv \text{part_of}(x, y) \wedge R(x) \wedge R(y) \wedge \exists z, w(\text{has_2D}(z, x) \wedge \text{has_2D}(w, y) \wedge ED(z) \wedge ED(w))) \quad (2)$$

Although logically strictly not necessary, there are strong arguments to distinguish between the surface and space notions [24,11,22,18], which is especially relevant for geographic entities. For instance, a euro is *contained_in* a wallet, Paris is *located_in* France, and in the TBox of a DL-formalised ontology, e.g., $\text{City} \sqsubseteq \exists \text{locatedIn.Country}$.

2.1 Extension with Mereotopological Relations

Representation and reasoning with spatiality is very important in many applications areas such as geographical information systems, medical systems, and computer vision. Spatial relations proposed for such systems can be classified into three categories: topological, direction, and distance. Topological relations are considered to be the most important ones and several topological formalisms have been proposed in the literature. The 9-Intersection Method (9IM) is based on point-set topology, where the topological relations between two regions are characterized by the 9 intersections of interiors, boundaries and exteriors of the two regions [4]. The Region Connection Calculus (RCC) is a first order theory based on one primitive relation: the reflexive and symmetric *connection* predicate [19]. Using this primitive relation, one can define various collections of topological relations. The fundamental approach of RCC—and difference with point-set topology—is that it extends spatial entities to *regions of space* that are taken as primary rather than the points as the fundament in point-set topology. However, neither one considers the combination of the space region with the object that occupies it: this interaction is addressed by *mereotopology*, which focuses on spatial *entities*, not just regions. The challenge lies in how to realize the combination, given that no such ontology-informed categorisation for mereotopological relations exists. Concerning primitive relations [3,22], one can

Table 1. Axiomatization of KGEMT (summarised from [22]). P: partof; PP: proper part of; O: overlap, C: connection; E: enclosure; EQ: indiscernibility; IPP: interior proper part; TPP: tangential proper part; SC: self-connected; c: closure; i: interior; e: exterior; +: sum; \sim : complement.

Core axioms and definitions			
$P(x, x)$	(t1)	$P(x, y) \wedge P(y, z) \rightarrow P(x, z)$	(t2)
$P(x, y) \wedge P(y, x) \rightarrow x = y$	(t3)	$\neg P(y, x) \rightarrow \exists z(P(z, y) \wedge \neg O(z, x))$	(t4)
$\exists w\phi(w) \rightarrow \exists z\forall w(O(w, z) \leftrightarrow \exists v(\phi(v) \wedge O(w, v)))$			(t5)
$C(x, x)$	(t6)	$C(x, y) \rightarrow C(y, x)$	(t7)
$P(x, y) \rightarrow E(x, y)$	(t8)	$E(x, y) =_{df} \forall z(C(z, x) \rightarrow C(z, y))$	(t9)
$E(x, y) \rightarrow P(x, y)$	(t10)	$SC(x) \leftrightarrow \forall y, z(x = y + z \rightarrow C(y, z))$	(t11)
$\exists z(SC(z) \wedge O(z, x) \wedge O(z, y) \wedge \forall w(P(w, z) \rightarrow (O(w, x) \vee O(w, y)))) \rightarrow C(x, y)$			(t12)
$z = \sum x\phi x \rightarrow \forall y(C(y, z) \rightarrow \exists x(\phi x \wedge C(y, x)))$			(t13)
$P(x, cx)$	(t14)	$c(cx) = cx$	(t15)
$c(x + y) = cx + cy$	(t16)	$cx =_{df} \sim(e x)$	(t17)
$ex =_{df} i(\sim x)$	(t18)	$ix =_{df} \sum z\forall y(C(z, y) \rightarrow O(x, y))$	(t19)
Additional axioms, definitions, and theorems			
$PP(x, y) =_{df} P(x, y) \wedge \neg P(y, x)$	(t20)	$O(x, y) =_{df} \exists z(P(z, x) \wedge P(z, y))$	(t21)
$EQ(x, y) =_{df} P(x, y) \wedge P(y, x)$	(t22)	$TPP(x, y) =_{df} PP(x, y) \wedge \neg IPP(x, y)$	(t23)
$IPP(x, y) =_{df} PP(x, y) \wedge \forall z(C(z, x) \rightarrow O(z, y))$			(t24)
$\neg PP(x, x)$	(t25)	$PP(x, y) \wedge PP(y, z) \rightarrow PP(x, z)$	(t26)
$PP(x, y) \rightarrow \neg PP(y, x)$	(t27)		

- define parthood, P , in terms of connection, C , (i.e., $P(x, y) =_{def} \forall z(C(z, x) \rightarrow C(z, y))$) so that topology is principal and mereology a subtheory,
- consider both P and C as primitive,
- introduce a ternary relation $CP(x, y, z)$, so that $P(x, y) =_{def} \exists z CP(x, z, y)$ and $C(x, y) =_{def} \exists z CP(x, y, z)$, or
- introduce topology as a sub-domain of mereology by introducing a sorted predicate to denote region (R) and define C in terms of overlapping regions [5] ($C(x, y) =_{def} O(x, y) \wedge R(x) \wedge R(y)$).

There are several problems with the first option, such as extensionality and identity [22] and that, implicitly, parthood is constrained to relating regions, because, by assumption, C is, even though normally an arbitrary part-whole relation should not be constrained to relate regions only. The third option requires a language with ternary relations, which OWL languages do not have, and it is not considered or used widely. The fourth option is, in spirit, close to Eqs. (1-2). However, adding complete definitions of both overlap and the connection is impossible in OWL; therefore, the second option with two primitives is the most appealing with the current Semantic Web languages and technologies. This fits nicely with the expressive Kuratowski General Extensional Mereotopology (KGEMT) [22], which provides both an ontological and logical underpinning of the mereotopology taxonomy. KGEMT consists of simpler mereological and topological theories. Formally, the axioms in KGEMT are as shown in Table 1, which is built up from T that is made up of (t6, t7), MT of (T + t8), M of (t1,

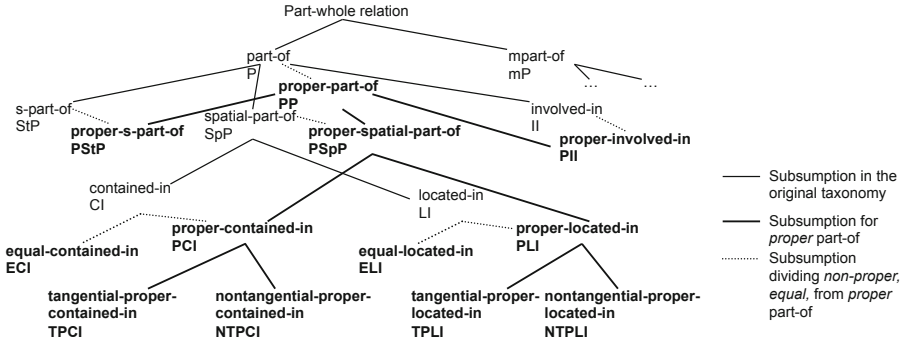


Fig. 1. Graphical depiction of the extension of the mereological branch of the basic taxonomy of part-whole relations with proper parthood and mereotopological relations (meronymic relations subsumed by *mpart-of* not shown)

t_2, t_3), GEM of ($M + t_4, t_5$), GEMT of ($MT + GEM + t_{10}, t_{12}, t_{13}$), and KGEMT of ($GEMT + t_{14}, t_{15}, t_{16}$). In addition, (t_{14} - t_{16}) require (t_{17} - t_{19}) and there are additional axioms and definitions (like t_{20} - t_{27}) that can be built up from the core ones.

Using Eqs. (1-2) and (t_{1} - t_{27}), we now extend the part-whole taxonomy of [11] with the mereotopological relations as defined in Eqs. (3-10), which is shown graphically in Fig. 1. The tangential and nontangential proper parthood relations are based on axioms 65 and 66 in [22], which are (t_{24}) and (t_{23}), respectively, in Table 1, and the same DOLCE categories are used as in Eqs. (1-2); see Fig. 1 and Table 1 for abbreviations.

$$\forall x, y (ECI(x, y) \equiv CI(x, y) \wedge P(y, x)) \quad (3)$$

$$\forall x, y (PCI(x, y) \equiv PPO(x, y) \wedge R(x) \wedge R(y) \wedge \exists z, w (has_3D(z, x) \wedge has_3D(w, y) \wedge ED(z) \wedge ED(w))) \quad (4)$$

$$\forall x, y (NTPCI(x, y) \equiv PCI(x, y) \wedge \forall z (C(z, x) \rightarrow O(z, y))) \quad (5)$$

$$\forall x, y (TPCI(x, y) \equiv PCI(x, y) \wedge \neg NTPCI(x, y)) \quad (6)$$

$$\forall x, y (ELI(x, y) \equiv LI(x, y) \wedge P(y, x)) \quad (7)$$

$$\forall x, y (PLI(x, y) \equiv PPO(x, y) \wedge R(x) \wedge R(y) \wedge \exists z, w (has_2D(z, x) \wedge has_2D(w, y) \wedge ED(z) \wedge ED(w))) \quad (8)$$

$$\forall x, y (NTPLI(x, y) \equiv PLI(x, y) \wedge \forall z (C(z, x) \rightarrow O(z, y))) \quad (9)$$

$$\forall x, y (TPLI(x, y) \equiv PLI(x, y) \wedge \neg NTPLI(x, y)) \quad (10)$$

Note that one also can use another foundational ontology, such as SUMO or GFO, provided it at least contains categories equivalent to the ones from DOLCE we use here. Concerning the interaction between this proposal for mereotopology and DOLCE: DOLCE includes the GEM mereological theory, of which KGEMT is an extension, and does not contain mereotopology; hence, our taxonomic categorisation of the mereotopological relations and additional KGEMT axioms do not lead to an inconsistency. Interaction with DOLCE's—or any other

foundational ontology’s—temporal parthood is unclear, as is a temporal extension to KGEMT; either way, they are orthogonal issues and it will become of practical interest only once there is a temporal OWL.

As with parthood and topology separately, the mereotopological relations are independent of the subject domain and thus may be applied across ontologies in domains such as medicine, manufacturing, and biology.

2.2 Limitations for Mereotopology in OWL

There are several aspects to consider with respect to ‘theoretical ideal’ from an ontological perspective, as described in the previous section, versus ‘realistic implementation’ in a Semantic Web setting, which are: (i) what to represent from mereotopology, (ii) how to represent the ontological aspects in a DL language, (iii) what does one lose with the OWL 2 languages, and (iv) how does this affect reasoning. Due to space limitations, we discuss only a selection of the issues.

Representation Issues. Only a subset of the KGEMT axioms can be represented in OWL, which is summarised in Table 2. The differences for the OWL species have to do with the characteristics of the object properties, and in particular transitivity, (ir)reflexivity and (a)symmetry (see also the OWL and OWL 2 specifications [13,16,15]), which are exactly the features that affect negatively the basic reasoning scenarios. This has the consequence that some theorems in KGEMT can only be asserted as axioms, due to its partial representation. However, even its limited version is still more appealing than only RCC8 for OWL because of the notion of spatial entity and relation and the inferences with the property characteristics that is important for several subject domains [11,22,18,5] compared to region only. In addition, in some implementations, the RCC8 relations are modified into a set-based approach where regions are non-empty closed sets [17] or where concepts are generated for each RCC8 constructor and the named concepts instantiated with individuals [9,21]. While conflating sets, regions, and part-whole relations might be appealing from a logic perspective as a sufficient approximation, it is not the same neither in detail nor from an ontological or modelling viewpoint such that other trade-offs in the representation are preferred (e.g., [22,6,23]), and in particular to keep the relations as such compared to a concept-based encoding that obscures the semantics and thereby complicates its usage [11,6,23]. Put differently, KGEMT is the best (or least worst) from a modeller point of view.

Reasoning Issues. We illustrate some of the inferences we miss out on with an extended African Wildlife Ontology, which includes the knowledge that South Africa (SA) is connected with Botswana (B), Mapungubwe National Park (MNP) (home to giraffes, elephants, etc.) is tangentially proper located in (TPLI) South Africa and bordering Botswana, and in Botswana, the park is called Northern Tuli Game Reserve (NTGR). So, to this end, there are axioms in our ontology in the TBox and the ABox, among others: $\text{NationalPark} \sqsubseteq \exists \text{LI.Country}$, $\text{ContinentalCountry} \sqsubseteq \text{Country}$, $\text{ContinentalCountry} \sqsubseteq \exists \text{C.Continental}$

Table 2. Subsets of KGEMT that can be represented in the OWL species; t9, t20, t21, t22, t23, t24 can be simplified and added as primitives to each one

OWL species	Subsets of KGEMT axioms
OWL 2 DL	(t1, t2, t6, t7, t8, t10, t26) or (t1, t2, t6, t7, t8, t10, t27) or (t1, t2, t6, t7, t8, t10, t25)
OWL DL	t2, t7, t8, t10, t26
OWL Lite	t2, t7, t8, t10, t26
OWL 2 RL	t2, t7, t8, t10, t26
OWL 2 EL	t1, t2, t6, t8, t10, t26
OWL 2 QL	t1, t6, t7, t8, t10

Country, ContinentalCountry(SA), ContinentalCountry(B), NationalPark(MNP), TPLI(MNP, SA), C(MNP, NTGR), TPLI(NTGR, B). Then a query “Which South African national park is connected with Northern Tuli Game Reserve?” has as answer MNP in a language where symmetry can be asserted, and is empty in the absence of the symmetry property assertion (i.e., in OWL 2 EL). In addition, with TPLI being asymmetric, an erroneous addition of, say, TPLI(SA, MNP) to the ontology yields an inconsistent ontology, whereas without the option to represent asymmetry, the error (with respect to the subject domain semantics) would go undetected.

Parthood and connection are reflexive, which cannot be represented in OWL 2 RL. This does not really cause any major problems, other than that in certain cases its inclusion communicates to the modeller she has to be more precise. For instance, with Leaf being a StP (structural part of) Twig and StP reflexive, the automated reasoner will deduce that Leaf is a StP of Leaf, and with aforementioned axioms, that MNP is C (connected with) and LI (located in) MNP: the former is a rather odd deduction, and, in fact, should be structural *proper* part—which is irreflexive—of Twig, and the latter two are uninteresting, and, in fact, TPLI. Also, if a user queries “Which park is connected with Northern Tuli Game Reserve?”, then one would want to retrieve all parks *other than* NTGR.

Which of the characteristics of the object properties of mereotopological relations is, or are, more important depends on the desired inference scenarios; thus far, transitivity, symmetry, asymmetry, and irreflexivity appear to be the more interesting ones, which—within a Semantic Web setting—means giving precedence to OWL 2 DL and OWL 2 RL.

3 Design and Implementation of ONTOPARTS

Given the theoretical assessment on representing part-whole and mereotopological relations and feasibility to simplify it for a Semantic Web setting, we can now proceed toward the design and implementation of the ONTOPARTS tool. The tool, additional files, and demo videos can be consulted in the online supplementary material at <http://www.meteck.org/files/ontopartssup/supindex.html>.

3.1 Requirements, Design, and Core Functionality

The main requirement of the software is to hide the logic involved in the formal definition of the 23 part-whole relations, and translate it to a set of intuitive steps and descriptions that will guide the user to make the selection and decision effortlessly. The selection procedure for the 23 possible relations should be made as short as possible and present only a small relevant subset of suggestions from which the user can select the one that best fit the situation. A set of (top-level) categories should be proposed to quickly discriminate among relations since the user may be more familiar with the categories' notions for domain and range than with the relations' definitions, therewith standardizing the criteria for selecting the relations. Simple examples must be given for each relation and category. Last, the user must have the possibility also to save the selected relation to the ontology file from where the classes of interest were taken.

Given these basic functional requirements, some design decisions were made for ONTOPARTS. From a generic perspective, a separate tool is an advantage, because then it can be used without binding the ontologist to a single ontology editor. Another consideration is usability testing. We chose to use a rapid way of prototyping to develop the software to quickly determine whether it is really helpful. Therefore, we implemented a stand-alone application that works with OWL files. We also chose to use the DOLCE top-level ontology categories for the standardization of the relationships' decision criteria.

To structure the selection procedure in a consistent way in the implementation, we use activity diagrams to describe the steps that a user has to carry out to interact with ONTOPARTS and to select the appropriate relation. An activity diagram for the selection process of the mereotopological relations is available in the online supplementary material. The selection of the appropriate relation incorporates some previous ideas of a decision diagram and topological principles as an extension of mereological theories [10,25,26], and questions and decision points have been added to reflect the extended taxonomy. For the mereotopological relations considered, in principle, the decision for the appropriate one can be made in two separate ways: either find mereotopological relations and then asking to distinguish between *located_in* and *contained_in*, or vv. In the ONTOPARTS interface, we have chosen to reduce the sequence of questions to a single question (check box) that appears only when the domain and range are regions, which asks whether the classes are geographical entities.

Concerning the most expressive OWL 2 DL species and KGEMT, then antisymmetry (t3), the second order axioms (t5, t13), and the closure operators (t14-t19) are omitted, and definitions of relations are simplified to only domain and range axioms and their position in the hierarchy (recollect Table 2). In addition, OWL's *IrreflexiveObjectProperty* and *AsymmetricObjectProperty* can be used only with simple object properties, which runs into problems with a taxonomy of object properties and transitivity (that are deemed more important), therefore also (t25, t27) will not appear in the OWL file. The combination of the slimmed KGEMT and extended taxonomy of part-whole object properties together with a DOLCE ultra-ultra-light is included in the online supplementary

material as `MereoTopoD.owl`. This OWL file contains the relations with human-readable names, as included in the taxonomy of [11] and the mereotopological extension depicted in Fig. 1, as it is generally assumed to be better workable to write out abbreviations in the domain ontology as it increases readability and understandability of the ontology from a human expert point of view.

Observe that at the class-level, we have the so-called “all-some” construction for property axioms, and if the modeller wants to modify it with a min-, max-, or exact cardinality (e.g., ‘each spinal column is a proper part of exactly one human’), then it goes beyond OWL 2 DL because the `properPartOf` object property is not simple. Further, transitivity is a feature of OWL-DL, OWL Lite, OWL 2 DL, DL, EL and RL, but not QL. Because one cannot know upfront the setting of the ontology, we keep the hierarchy of relations but do not add the relational properties when writing into the `.owl` file, but the user can add them afterward.

3.2 Description of the Selection Procedure

ONTOPARTS guides the user in the process of making the decision in a very intuitive way. To select the correct part-whole relation between two classes, one first has to load the ontology into the tool and select the class that represent the part and the class that represent the whole, upon which the selection procedure commences. To not overload the user at once with a choice of 23 possible relations or having to answer 22 questions first, ONTOPARTS narrows down the set of available relations when the user specifies the category of the domain and of the range involved in the relation, which are categories from DOLCE. If the user selects regions, then she enters the branch of mereotopological relations where the interface contains one decision point to distinguish between the 2D and 3D case. Upon finishing the above steps, the user clicks the button to let the tool suggest appropriate relation(s). All possible relations are computed using the DOLCE categories and the optional check box for geographic entities. Each proposed relation is verbalized in a pseudo-natural language sentence with the selected part-class and whole-class and an example as an additional guide to make the relation more understandable. Once a user has analysed each proposed relation, she can select one and proceed to save it in the OWL file by clicking on the save button. The selection procedure is illustrated in Example 1.

Example 1. Suppose a modeller or domain expert is developing a tourism ontology and has to figure out the relation between the classes `Campground` and `RuralArea`. She loads the OWL file in ONTOPARTS, and proceeds to select two entities to be related, as shown in Fig. 2. Then, she selects the categories for each entity from the ones provided by the software, depicted in Fig. 3. Examples for each DOLCE category are shown by hovering the pointer over the terms in the taxonomy; e.g., hovering over `Process` provides “Ex: Running, Writing”. If the selected categories are regions (or any of its subtypes), then the software provides the option to specify whether the regions correspond to geographical entities. After all the required information is selected, she clicks on the button “Suggest relationships”. The amount of relations suggested depends on the chosen categories; in the example there are still four options (Fig. 4), but if we

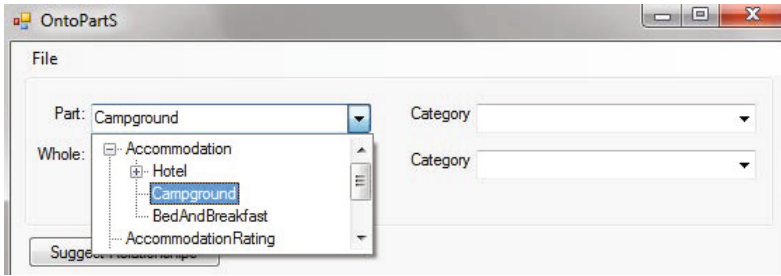


Fig. 2. Selecting the part and whole classes in ONTOPARTS from the loaded OWL file

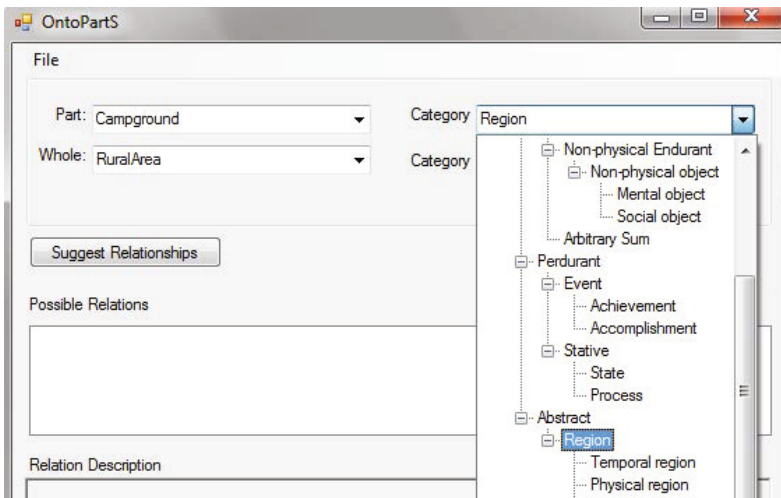


Fig. 3. Selecting the categories for each entity

would have had two classes that are both, say, processes, then there is only one option (*involved_in*, as ONTOPARTS includes the taxonomy of [11]). The software provides the suggestions, a verbalization of the possible relationship(s), e.g., “Campground coincides exactly with RuralArea”, and typical examples, as can be seen in Fig. 4. Once she selects the desired relation from the ones proposed by the software, she can choose to add this relationship to the OWL file by simply clicking the button labelled “Save relationship to file” (Fig. 4 bottom) and continue either with other classes and selection of a part-whole relation or with developing the ontology in the ontology development environment of choice. \diamond

3.3 Preliminary Experimental Assessment of ONTOPARTS

The main objectives of the experiments are to assess usability of the tool and to validate the hypothesis that the use of automated guidelines assists with representation of part-whole relations between classes during the ontology design

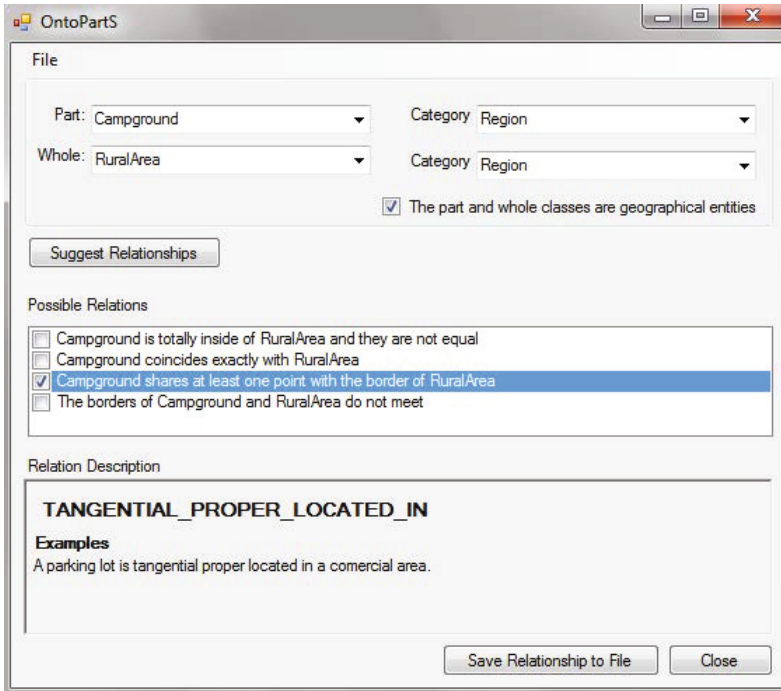


Fig. 4. Relationships suggested by ONTOPARTS

phase such that it can be done more efficiently and with less errors. To this end, a qualitative and two preliminary quantitative evaluations have been carried out.

Materials & Methods. 17 third year students in Informatics Engineering at the Instituto Superior Politécnico “José Antonio Echeverría” participated in the first test, who are well-versed in logics and knowledge representation with frames, nets and rules systems, but not with ontologies. The methodology for this experiment was as follows. The students receive an overview of foundational ontologies and part-whole relations (30 minutes). Then, from the provided computer ontology (designed by us) and ONTOPARTS, they must select the most appropriate DOLCE category for each subject domain class and the relation that holds between them. The evaluation is performed by assessing the OWL files, relations detected, selected DOLCE categories, the errors made, and why.

A second experiment was carried out with 6 honours students in Computer Science at the University of KwaZulu-Natal, with the main difference the group’s longer introduction into foundational ontologies (2h) and part-whole relations (2h), the division into two groups, one using ONTOPARTS, the other not, and the limitation of 40 minutes. The third group for qualitative evaluation consisted of 4 colleagues from the Free University of Bozen-Bolzano, two of whom experts in RCC8, part-whole relations, and logic.

The materials used for the experiments were the DOLCE taxonomy, the taxonomy with the 23 part-whole relations, and the beta version of the ONTOPARTS tool. The domain ontology about computers was developed using Protege 4.0, which was divided into two versions, one with and one without part-whole relations.

Results. The students in the first experiment asserted 380 part-whole relations among the classes in the domain ontology (37 mereotopological), of which 210 were correct, i.e., on average, about 12.4 correct assertions/participant (variation between 5 and 37); for the second experiment, the numbers are, respectively, 82 (22 mereotopological), 58, and an average of 9.7 (with variation between 0 and 27). Given the controlled duration of the second experiment, this amounts to, on average, a mere 4 minutes to choose the correct relation with ONTOPARTS.

Evaluating the mistakes made by the participants revealed that an incorrect selection of part-whole relation was due to, mainly, an incorrect association of a domain ontology class to a DOLCE category. This was due to the late discovery of the tool-tip feature in ONTOPARTS by some participants and the lack of an “undo” button (even though a user could have switched back to the ontology development environment and deleted the assertion manually). Several errors were due to the absence of inverses in the beta version of the ONTOPARTS tool, leading some participants to include *Metal constitutes some Hardware*. 83% of the errors in the second experiment were made by those who did not use ONTOPARTS, which looks promising for ONTOPARTS.

The responses in the qualitative evaluation was unanimous disbelief that selection could be made this easy and quickly, and the desire was expressed to consult the formal and ontological foundations. As such, ONTOPARTS stimulated interest for education on the topic along the line of “the tool makes it easy, then so the theory surely will be understandable”.

Overall, it cannot be concluded that modelling of part-whole relations with ONTOPARTS results in statistically significant less errors—for this we need access to more and real ontology developers so as to have a sufficiently large group whose results can be analysed statistically. Given the speed with which correct relations were selected, the automated guidelines do assist with representation of part-whole relations such that it can be done more efficiently and quickly. The experimentation also aided in improving ONTOPARTS’s functionality and usability, so that it is now a fully working prototype.

4 Discussion

Despite the representation and reasoning limitations with the DL-based OWL 2 species, there are several modelling challenges that can be addressed with the mereotopological and part-whole relation taxonomy together with ONTOPARTS and they solve the three problems identified in the introduction.

Mitigating Representation Limitations with the Taxonomy. ONTOPARTS sensitizes the modeller to part-whole relations, which thereby already prevents

the is-a vs. part-of confusion—i.e., using the is-a relation where a part-of relation should have been used—common among novice ontology developers, and no such errors were encountered during our experiments either (recollect Section 3.3). By making part-whole relations easily accessible without the need for immediate in-depth investigation into the topic, it is expected that this type of error may also be prevented without any prior substantial training on the topic.

The fine-grained distinctions between the parthood relation enables, among others, proper instance classification like mentioned in the introduction for `EnclosedCountry`, thanks to being able to select the right relation and thereby capturing the intended semantics of `EnclosedCountry`. If, on the other hand, the modeller would have known only about proper part-of but not proper located in, then she could only have asserted that Lesotho is a proper part of South Africa, which holds (at best) only *spatially* but not *administratively*. Not being able to make such distinctions easily leads to inconsistencies in the ontology or conflicts in ontology import, alignment, or integration. OntoClean [7] helps with distinguishing between geographic and social entity, and, in analogy, ONTOPARTS aids *relating* the entities with the appropriate relation.

Solving the Problems. In the introduction, we identified three problems that were in need of a solution to improve modelling of part-whole relations. The first problem regarding the plethora of part-whole relations has been solved by providing logic-based definitions of the relations and structuring them in a taxonomy. In so doing, we extended the taxonomy of [11] that already makes a clear distinction between parthood (mereology) versus other so-called part-whole relations (meronymy), so that the parts with their locations (mereotopology) could be integrated in the parthood branch of that taxonomy. Their position in the taxonomy follows directly from the definitions (Eqs. 3-10), which, in turn, are based on the KGEMT mereotopology and the common modeller’s preference to distinguish between 3D parthood (containment) and 2D parthood (location).

The challenge to figure out which part-whole relations to use when has been addressed by ONTOPARTS. Determining which part-whole relation fits best for the given classes requires only the input from the modeller about the categories and, depending on the category, if they are geographic entities. Upon this input, ONTOPARTS suggests 1-4 possible relations, which is considerably less than the full 23 together.

Implementing part-whole relations with its underlying KGEMT mereotopology in OWL requires making concessions, which concern principally the choice to represent relations as relations, and the impossibility to include definitions of relations and some of the characteristics of the part-whole object properties. The latter is practically important, because the OWL species differ in language features in this regard, and its consequences for reasoning were analysed. Transitivity, symmetry, asymmetry, and irreflexivity are the more relevant ones with respect to the deductions one gains or loses. This means giving precedence to OWL 2 DL and OWL 2 RL over the other OWL species when modelling part-whole relations.

5 Conclusions

We have introduced mereotopology from a modeller's perspective and integrated it into the taxonomy of part-whole relations that is founded on the KGEMT mereotopological theory and basic categories of the DOLCE foundational ontology so as to disambiguate the relations. To make this practically usable, we transformed the mappable KGEMT axioms into OWL, added a DOLCE ultra-ultra-light and extended the taxonomy of part-whole OWL object properties, designed an additional activity diagram for the newly added mereotopological relations, added examples, and implemented this in the ONTOPARTS application. ONTOPARTS ensures that the complexities of the underlying theories and languages are hidden from the modeller by means of the automated modelling guidelines and adds the new axioms to the OWL file with a simple one-click button. This ontology-inspired part-whole relation selection tool, ONTOPARTS, can be used with different OWL languages and different ontology development tools. ONTOPARTS was evaluated with modellers, which was found to simplify the task, was performed more accurately, and generated desire to learn more about the theoretical details.

ONTOPARTS and the OWLized hierarchy of relations are freely available in the online material (<http://www.meteck.org/files/ontopartssup/supindex.html>). Current and future work pertains to adding features, such as better interaction with property chaining and language choices in conjunction with the object property characteristics, and handling ABox assertions.

References

1. Artale, A., Franconi, E., Guarino, N., Pazzi, L.: Part-whole relations in object-centered systems: An overview. *DKE* 20(3), 347–383 (1996)
2. Bittner, T., Donnelly, M.: Computational ontologies of parthood, componenthood, and containment. In: *Proc. of IJCAI 2005*, pp. 382–387. AAAI Press, Cambridge (2005)
3. Cohn, A.G., Renz, J.: Qualitative spatial representation and reasoning. In: *Handbook of Knowledge Representation*, ch. 13 p. 551–596. Elsevier (2008)
4. Egenhofer, M.J., Herring, J.R.: Categorizing binary topological relations between regions, lines, and points in geographic databases. Tech. Rep. 90-12, National Center for Geographic Information and Analysis, University of California (1990)
5. Eschenbach, C., Heydrich, W.: Classical mereology and restricted domains. *International Journal of Human-Computer Studies* 43, 723–740 (1995)
6. Grütter, R., Bauer-Messmer, B.: Combining OWL with RCC for spatioterminological reasoning on environmental data. In: *Proc. of OWLED 2007* (2007)
7. Guarino, N., Welty, C.: An overview of OntoClean. In: Staab, S., Studer, R. (eds.) *Handbook on Ontologies*, pp. 151–159. Springer (2004)
8. Horrocks, I., Kutz, O., Sattler, U.: The even more irresistible *SRTOIQ*. In: *Proc. of KR 2006*, pp. 452–457 (2006)
9. Katz, Y., Cuenca Grau, B.: Representing qualitative spatial information in OWL-DL. In: *Proc. of OWLED 2005*, Galway, Ireland (2005)

10. Keet, C.M.: Part-Whole Relations in Object-Role Models. In: Meersman, R., Tari, Z., Herrero, P. (eds.) OTM 2006 Workshops. LNCS, vol. 4278, pp. 1116–1127. Springer, Heidelberg (2006)
11. Keet, C.M., Artale, A.: Representing and reasoning over a taxonomy of part-whole relations. *Applied Ontology* 3(1-2), 91–110 (2008)
12. Masolo, C., Borgo, S., Gangemi, A., Guarino, N., Oltramari, A.: Ontology library. WonderWeb Deliverable D18 (ver. 1.0) (December 31, 2003), <http://wonderweb.semanticweb.org>
13. McGuinness, D.L., van Harmelen, F.: OWL Web Ontology Language Overview. W3C Recommendation (2004), <http://www.w3.org/TR/owl-features/>
14. Mejino, J.L.V., Agoncillo, A.V., Rickard, K.L., Rosse, C.: Representing complexity in part-whole relationships within the foundational model of anatomy. In: Proc. of the AMIA Fall Symposium, pp. 450–454 (2003)
15. Motik, B., Grau, B.C., Horrocks, I., Wu, Z., Fokoue, A., Lutz, C.: OWL 2 Web Ontology Language Profiles. W3c recommendation, W3C (October 27, 2009), <http://www.w3.org/TR/owl2-profiles/>
16. Motik, B., Patel-Schneider, P.F., Parsia, B.: OWL 2 web ontology language structural specification and functional-style syntax. W3c recommendation, W3C (October 27, 2009) <http://www.w3.org/TR/owl2-syntax/>
17. Nutt, W.: On the Translation of Qualitative Spatial Reasoning Problems into Modal Logics. In: Burgard, W., Christaller, T., Cremers, A.B. (eds.) KI 1999. LNCS (LNAI), vol. 1701, pp. 113–124. Springer, Heidelberg (1999)
18. Odell, J.J.: *Advanced Object-Oriented Analysis & Design using UML*. Cambridge University Press, Cambridge (1998)
19. Randell, D.A., Cui, Z., Cohn, A.G.: A spatial logic based on regions and connection. In: Proc. of KR 1992, pp. 165–176. Morgan Kaufmann (1992)
20. Schulz, S., Hahn, U., Romacker, M.: Modeling anatomical spatial relations with description logics. In: AMIA 2000 Annual Symposium, pp. 779–783 (2000)
21. Stocker, M., Sirin, E.: Pelletspatial: A hybrid RCC-8 and RDF/OWL reasoning and query engine. In: Proc. of OWLED 2009, Chantilly, USA, October 23-24. CEUR-WS, vol. 529 (2009)
22. Varzi, A.: Spatial reasoning and ontology: parts, wholes, and locations. In: *Handbook of Spatial Logics*, pp. 945–1038. Springer, Heidelberg (2007)
23. Wessel, M.: Obstacles on the way to qualitative spatial reasoning with description logics: some undecidability results. In: Proc. of DL 2001, Stanford, CA, USA, August 1-3. CEUR WS, vol. 49 (2001)
24. Winston, M.E., Chaffin, R., Herrmann, D.: A taxonomy of partwhole relations. *Cognitive Science* 11(4), 417–444 (1987)
25. Yang, W., Luo, Y., Guo, P., Tao, H., He, B.: A Model for Classification of Topological Relationships Between Two Spatial Objects. In: Wang, L., Jin, Y. (eds.) FSKD 2005. LNCS (LNAI), Part II, vol. 3614, pp. 723–726. Springer, Heidelberg (2005)
26. Zhong, Z.-N., et al.: Representing topological relationships among heterogeneous geometry-collection features. *J. of Comp. Sci. & Techn.* 19(3), 280–289 (2004)