Representing Mereological Relations in Weighted Description Logics

Gabriele Sacco[1,2], Loris Bozzato[1] and Oliver Kutz[2]

[1]Fondazione Bruno Kessler, Via Sommarive 18, 38123 Trento, Italy

[2]Free University of Bozen-Bolzano, Piazza Domenicani 3, 39100, Bolzano, Italy

Abstract

Perceptron operators have been recently introduced in Description Logics: they define a concept by listing features of the concept together with associated weights, as well as a threshold that needs to be reached to classify an individual as a member of the defined concept. These formal operators have been shown to be useful in simplifying the representation of complex classifications and for defining combinations of concepts. In this paper, we discuss how perceptron operators can provide a tool to capture some of the complex relationships that arise in the conceptual use and definition of a variety of mereological relations. We here study the use case of the concept of "bike" as a mereologically/structurally complex object, analyse the most relevant difficulties that emerge from the modelling approach using percepton operators, and propose several starting points to overcome these difficulties.

Keywords

Mereology, Weighted Logics, Perceptron Operators, Mereological Similarity

1. Introduction

In the field of knowledge representation, and in particular in Description Logics (DLs), *concepts* are the central element of the languages. Thus, the problem of classification, that is, of recognising whether an individual is an instance of a certain concept, is an interesting and relevant task in these languages. Recently, *perceptron operators* (or *tooth operators*) [1, 2, 3] have been proposed in the DL literature as a tool for introducing weights in the definition of concepts: a perceptron operator gives a weight to the features that an individual instantiates and, if a threshold is reached, then the individual is classified under the defined concept. The tooth operator introduced in [1] is an expression of the form:

$$\mathbb{W}^t(C_1 : w_1, \ldots, C_p : w_p)$$

where $\mathbb{W}$ is the symbol of the operator, t is a numerical threshold and $C_1, \ldots, C_p$ are sub-concepts with the associated weights $w_1, \ldots, w_p$. The expression defines a concept which is interpreted by the weighted sum of the sub-concepts: intuitively, an individual is classified as belonging to the defined concept if and only if the sum of the weights w_i of the listed concepts C_i it satisfies is greater or equal than the threshold t. For example, one can define the concept

The Eighth Joint Ontology Workshops (JOWO'22), August 15-19, 2022, Jönköping University, Sweden

gsacco@fbk.eu (G. Sacco); bozzato@fbk.eu (L. Bozzato); Oliver.Kutz@unibz.it (O. Kutz)

0000-0001-5613-5068 (G. Sacco); 0000-0003-1757-9859 (L. Bozzato); 0000-0003-1517-7354 (O. Kutz)

of Elephant as $\mathbb{W}^{1.4}(\texttt{HasTrunk} : 1.0, \texttt{IsBig} : 0.4, \texttt{IsGrey} : 0.4)$: if we have an individual dumbo which HasTrunk and IsGrey, then the sum of weights meets the threshold and it will be classified as Elephant.

With respect to the definition of concepts, in recent years *mereology* and *mereotopology*, that is, the study of the part-whole relation and additional (spatial) relations, has gained great attention and has became quite developed [4, 5, 6]. In this regard, we propose to use mereology in order to make classifications using perceptron operators: the idea is to take as input the mereological structure of an object and, on the basis of its similarity to a prototypical structure, decide if it is an instance of a particular concept. However, given a number of complex aspects related to mereology, realising a combination of these approaches is not trivial. The goal of this paper is to discuss this general idea, to point out some of the problems that emerge from this approach at the theoretical level, and to begin studying possible solutions.

In particular, in the next section we will propose and discuss a simple example of an informal application of the tooth operator to a mereologically complex object, namely a bicycle. The goal is to show how the operator works and to see more clearly what the limits in modelling are when applying the standard operator to a structured object. Then, we will discuss the problems emerging from a more technical point of view, referring explicitly to *classical mereology*, which is the standard formalisation of the parthood relation. This then allows us to sketch a DL formalisation of the proposed setting that can work as a starting point for a more detailed technical discussion.

2. A simple classification game

In order to study the problems of modelling mereological relations using perceptron operators [1], we introduce a simple scenario for our task in the form of an elementary game. We will start with a simplified scenario and gradually move towards the limits of this operator when it is used for classification using mereological relations.

The game consists in deciding if an object is or is not a bicycle according to the parts we know it has. The rules of the game are simple: (i) we have to decide what kinds of parts a bicycle has, (ii) we give to these kinds of parts a weight, indicating how important each of them is in order to classify the object composed of them as a bicycle, and (iii) we need to set a threshold, this will indicate when we can consider the object a bicycle and when we cannot.

In a formal way, this corresponds to the use of the perceptron operator as:

$$\mathbb{W}^t(C_1 : w_1, \dots, C_p : w_p)$$

where, as above, $\mathbb{W}$ is the symbol of the operator, t stands for the threshold and $C_1 : w_1$ are the concepts, in this case the kinds of parts, with the associated weights.

For example, let us consider a bicycle like the one sketched in Figure 1. It can reasonably be considered a bicycle since the most essential kinds of parts of a bicycle are present. Intuitively, we can define something an *essential* kind of part if by removing it, we are not willing to consider the whole as a bicycle anymore. In this case, we have only three essential kinds of parts, namely the frame, the handlebar grip, and the wheels.

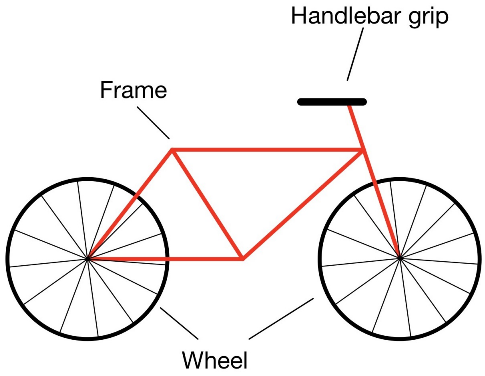

Figure 1: A sketch of a bicycle with its most essential parts.

It is important to notice that we are interested in the *kinds* of parts we have, i.e. the concepts we see instantiated in the actual parts being present.

Now we can attribute the weights. As we have said, these are all essential kinds of parts for a bicycle, so it is reasonable to give to all of them the same weight. We can decide arbitrarily for a weight of 5.

Lastly, we have to set the threshold. Again, considering that we are dealing with essential kinds of parts we want to make all of them indispensable for classifying the object in Figure 1 as a bicycle: therefore, we can set the threshold to 15, i.e., to the sum of the weights of the three kinds of parts.

Intuitively, this means that we can add other parts to it and still consider the whole a bicycle, because the threshold will be reached by having even only these kinds of parts. However, notice that atypical bike feature might receive negative weights, meaning that exhibiting essential feature is not always enough to fall under a concept. Further, this means that if we change the definition of a concept by adding some new kinds of parts or features, with certain weights, we also need to recompute the weight of the features we considered necessary, so that they are still essential from the numerical threshold perspective.

With the above tooth notation, we can represent the concept of `Bicycle` in this example as:

$$\mathbb{W}^{15}\big(\exists \texttt{hasPart.Frame} : 5, \exists \texttt{hasPart.Handlebar\_grip} : 5, \exists \texttt{hasPart.Wheel} : 5\big) \quad (1)$$

2.1. Counting instances

This classification can be further refined: for instance, in the previous game session, we deliberately overlooked the fact that a bicycle should have two wheels, that is, two parts of the same kind. In fact, with the previous representation, an object with a frame, a handlebar grip and only one wheel can still be classified as a bicycle, because the three kinds are still all there.

We could solve this by quantifying over the instances: that is, we say that the weight is computed if we have *at least two parts* that are wheels. We can proceed analogously with the inverse issue, the one about having at most two parts that are wheels. In doing this, however, we should divide the weight between the two new ‘kinds’, that is `having_at_least_two_wheels`

and having_at_most_two_wheels, attributing to them a weight of 2.5.
We can update also our example formalisation in formula (1):

$$Bicycle = \mathbb{W}^{15}\big(\exists \texttt{hasPart.Frame} : 5, \exists \texttt{hasPart.Handlebar\_grip} : 5, \geqslant 2\, \texttt{hasPart.Wheels} : 2.5, \leqslant 2\, \texttt{hasPart.Wheels} : 2.5\big) \tag{2}$$

2.2. Parts of parts

For a more precise representation, considering that we are dealing with parts, we would like to consider the transitivity of the parthood relation: intuitively, we want to provide assertions also over parts of parts. For example, we can characterize "bicycle wheels" as having spokes: thus, we can assign a weight for the spokes of the wheels, say 4. We are not interested in the number of spokes, therefore we can consider them as a whole for the moment. Following the procedure we have described above, the refined definition of our concept becomes:

$$\texttt{Bicycle} = \mathbb{W}^{15}\big(\exists \texttt{hasPart.Frame} : 5, \exists \texttt{hasPart.Handlebar\_grip} : 5, \geqslant 2\, \texttt{hasPart.Wheels} : 2.5, \leqslant 2\, \texttt{hasPart.Wheels} : 2.5, \exists \texttt{hasPart.Spokes} : 4\big) \tag{3}$$

where the threshold is exceeded by 4 for our example bike.

We obtain that the object is still classified as a bicycle: however, the weight of the spokes does not influences the contribution of the weights of the wheels. The intuition behind this idea is that if we considered the contribution of a kind of part to the classification of the composite object, then we should already consider inside that contribution also the contribution of all the kinds of sub-parts of that kind of part.

The solution to this problem is not trivial. A first step can be to consider the *'hierarchy'* of kinds of parts in the attribution of the weights, by starting from the mereological atoms (that is, parts with no parts) with an arbitrary attribution, and then compute the other weights climbing up the hierarchy. But this does not solves the central problem of how to avoid to consider two times the contribution of the same kind of part, an issue related to problems with multiple inheritance in ontology.

2.3. Overlapping parts

The problem becomes even more complex when we consider the case of overlapping parts, that is, in cases where we have two parts that have a part in common. In this case, the information that we are counting two times the same part is more 'hidden': we need to look at the parts of parts and then to recognise when they are parts of the same part. The situation would be that of Figure 2, where the grey connections show where there is an overlap.

A concrete example could be, in a more complex bicycle, that of the shifter (the part of the bicycle that allows to control the gearing mechanisms). It could be considered both as part of the total shift system (composed of the shifter, the gearing mechanism and the cables going from the former to the latter) and of the handlebar. Now, if we are computing the weights of

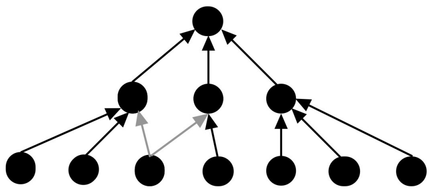

Figure 2: Schema of a structure with two overlapping parts; connections stand for the parthood relation.

the handlebar and of the shift system for the classification, we would like to take into account the fact that the shifter is a part of both: thus, in a certain sense, we are considering its weight two times.

2.4. Scattered and re-arranged parts

A more general problem is that so far we are assuming that the parts we are considering are structured in a way that allows us to consider them a bicycle. However, in the procedure we described there is no explicit reference to this aspect. Consider for example Figure 3.(a): here we have all the same parts of Figure 1, but we could be reluctant to consider it a bicycle.

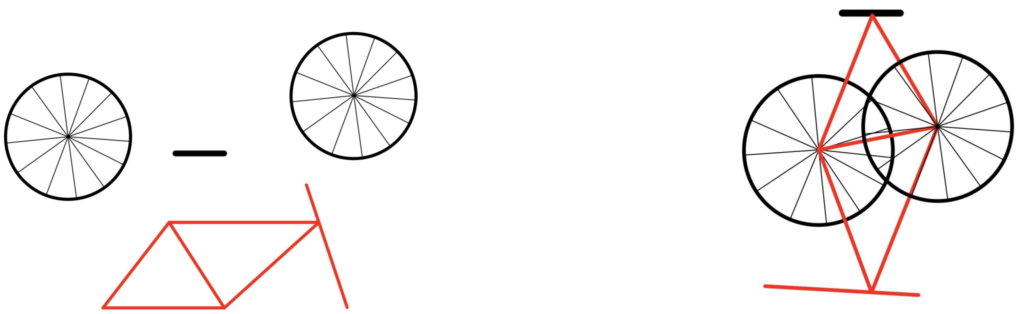

(a) Scattered bicycle

(b) Rearranged bicycle

Figure 3: Scattered and rearranged bicycle.

An easy way to resolve this problem is simply to make explicit that when we speak of the parthood relation we are assuming that the parts are connected. However, we can have cases such as the one in Figure 3.(b), where the parts are indeed connected, but where they are arranged in a way that does not allow us to classify the whole as expected, in our case, as a bicycle. Intuitively, a solution is to consider in addition to the parthood relation also a relation that allows us to represent the connection of the parts among each other. However, this solution needs a description of the object that has to be quite fine grained. In fact, consider again Figure 3.(b): if we consider the same parts as above, i.e. the frame, handlebar grip, wheels and spokes, the connections remain the same as in Figure 1. That is, they are represented in the

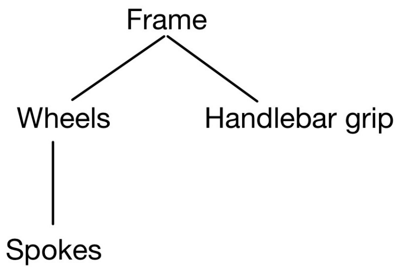

Figure 4: A schema of the connections among the parts of a bicycle as represented in Figure 1 and 3.(b).

schema of Figure 4. Therefore, even if the solution seems quite straightforward, its application even (or especially) to simple cases is not trivial.

The discussion so far focused generally on the approach we would like to pursue. However, there are also more specific problems and issues that are due to the specific choices made for the formalisation of the parthood relation.

3. Classification and classical mereology

In the following, we introduce the axioms that define our understanding of the parthood relation. The classical theory that formalises the parthood relation is called *classical mereology* [4]. Generally, it is based on five types of axioms: the first three are called *ordering axioms*, namely irriflexivity, asymmetry and transitivity, the fourth one is called *decomposition* axiom and the fifth the *composition* axiom. There are different equivalent versions of classical mereology: here we are using the one proposed in [4, section 2.4.1] for proper parthood. Using this version, we can identify the more problematic axioms for our goals as the last three, which in this version are called *Transitivity*, *Weak Supplementation* and *Unrestricted Fusion′*.

3.1. Transitivity

Considering PP as the *proper parthood* relation, transitivity is formulated as:

$$\forall x \forall y \forall z ((PPxy \wedge PPyz) \rightarrow PPxz). \qquad \text{Transitivity (PP)}$$

We are using proper parthood as a primitive for convenience: proper and improper parthood are inter-definable, thus we do not lose generality. The problem with transitivity is essentially what we have discussed above in terms of parts of parts in Section 2.2. However, about the first step we mentioned above, i.e. to exploit the hierarchy of parts, we can define a new parthood relation. The idea is to use it for the computation of the weights starting from the atomic parts, therefore we should consider only the parts of the 'level' immediately before. Consequently, the relation we define is that of *immediate parthood*: it defines an immediate part as a proper part which is not a proper part of another proper part. Formally:

$$IPxy :\equiv PPxy \wedge \neg \exists z (PPxz \wedge PPzy). \qquad \text{Immediate part}$$

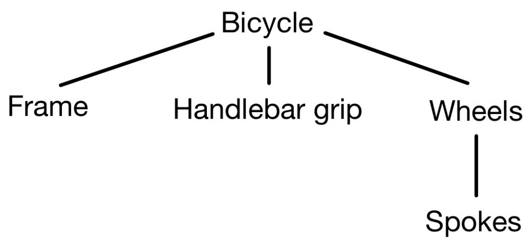

Figure 5: A schema of the parthood relations of the bicycle example.

Using our example again, the spokes are immediate parts of the wheel, but not of the bicycle, because they are proper parts of a proper part (namely, the wheel) of the bicycle. It is important to notice that immediate parts are *antitransitive*

$$\forall x \forall y \forall z ((IPxy \land IPyz) \rightarrow \neg IPxz). \qquad \text{Antitransitivity (IP)}$$

This is desirable in our case: in fact, antitransitivity "blocks the access" to other levels of the mereological hierarchy.

3.2. Decomposition

In our formalisation of classical mereology, the axiom regulating decomposition is called *weak supplementation.* Quite intuitively, this axiom states that if one removes a proper part from a whole, then there is another, distinct proper part that remains. Formally:

$$\forall x \forall y (PPyx \rightarrow \exists z (PPzx \land Dzy)), \qquad \text{Weak Supplementation}$$

where Dzy means that z and y have no parts in common

$$Dzy :\equiv \neg \exists x ((PPxz \land PPxy) \lor (x = z \land x = y)). \qquad \text{Disjointness}$$

To see what the problem is for our approach, consider the schema in Figure 5 representing the parthood relations in the bicycle example. Note that wheels have only spokes as a kind of part. This is not directly a problem given the fact that we are still moving at the conceptual level, whereas the parthood relation holds among the individuals. However, the bicycle we have described so far can have a model with only one spoke, or with an individual corresponding to spokes in the sense of a plural or mass object: in this case weak supplementation would not be respected.

3.3. Composition

The axiom that regulates composition is *unrestricted fusion*′:

$$\exists x \varphi x \rightarrow \exists z F'_{\varphi} z \qquad \text{Unrestricted Fusion}'$$

with $F'_\varphi z$ defined as

$$F'_\varphi z :\equiv \forall x(\varphi x \rightarrow Pxz) \wedge \forall y(Pyz \rightarrow \exists x(\varphi x \wedge Oyx)). \qquad \text{Fusion}'$$

The axiom simply states that there is a "fusion" of everything, where a fusion is the whole comprising all the things that are φ and that has as parts only things that have at least a part in common with the things that are φ. In fact, Oyx means that y and x have at least one part in common:

$$Ozy :\equiv \exists x((PPwy \wedge PPwx) \vee (w = y \wedge w = x)). \qquad \text{Overlap}$$

To understand what the problem is with respect to our approach, consider again the schema in Figure 5. Let us assume that φ corresponds to $being\_a\_frame \vee being\_a\_handlebar\_grip$. For Unrestricted Fusion′, there needs to exist a fusion corresponding to the things that satisfy $(being\_a\_frame \vee being\_a\_handlebar\_grip)$, that in our case are the frame and the handlebar grip. But we do not have such a fusion, because the second conjunct of the definition of a fusion, namely $\forall y(Pyz \rightarrow \exists x(\varphi x \wedge Oyx))$, is not respected. In fact, the only whole composed by the frame and by the handlebar grip is the bicycle, but it has two parts, namely wheels and spokes, which do not overlap with the parts satisfying φ.

These are the main problems we identified for the merging of classical mereology and perceptron operators for classification tasks. We note that there is an easy way to avoid the difficulties with the decomposition and the composition axioms: one can assume that the mereological structure we are trying to classify already respects classical mereology. In such a case, we are excluding models that could be classified as bicycles according to the conceptual structure described by the operator, but which are not compatible with classical mereology. Therefore, the true difficulty for our goal is the transitivity of parthood.

4. Encoding mereology in weighted description logics

The problem we are interested in to solve is classification by *mereological similarity*: that is, given a prototypical mereological structure defining a concept, such as (the concept of) a bicycle, we want to describe how to define such a concept via the perceptron operator in DLs by computing the weights to be assigned to each kind of part.

In this section, we will sketch a formalisation of this setting that can work as a starting point for a more involved technical discussion.

Given a DL signature $\Sigma = \mathrm{NC} \uplus \mathrm{NR} \uplus \mathrm{NI}$, we consider the role `hasImmediatePart` $\in \mathrm{NR}$ to represent the *immediate parthood* relation and its inverse `isImmediatePartOf` $\in \mathrm{NR}$. Here, the idea is to use *immediate parthood* instead of the classical *parthood* or *proper parthood* in order to have already an initial coarse "control" over the hierarchies of parts.

Our definitions will be based on immediate parthood, but for convenience we also assume the role `hasPart` $\in \mathrm{NR}$ (with inverse `isPartOf` $\in \mathrm{NR}$) to represent the *(proper) parthood* relation. Thus, we assume that `hasImmediatePart` $\sqsubseteq$ `hasPart` (and `isImmediatePartOf` $\sqsubseteq$ `isPartOf`) and that `hasPart` is transitive. As noted above, transitivity of parthood is one of the more problematic elements: here we maintain it for the sake of simplicity.

In our setting, we assume to have a given DL knowledge base defining objects (among other properties) via the immediate parthood relation. Thus, we assume to consider a knowledge base of the form $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$ where elements of the ABox $\mathcal{A}$ are statements about individuals using roles and concepts. I.e. they may be defined in terms of `hasImmediatePart` and `isImmediatePartOf` as for example in the assertions of the kind $\mathtt{hasImmediatePart}(\mathtt{bike}, \mathtt{wheel})$ and $\mathtt{hasImmediatePart}(\mathtt{wheel}, \mathtt{spoke})$, describing the structure of an individual `bike`. But also, we may have other descriptions using concepts, such as $\mathtt{Wheel}(\mathtt{wheel})$, $\mathtt{Spoke}(\mathtt{spoke})$, $\mathtt{Bike}(\mathtt{bike})$ and so on. The TBox $\mathcal{T}$ contains, as usual, terminological knowledge about the domain in the form of general concept inclusions, such as $\mathtt{Bike} \sqsubseteq \mathtt{Vehicle}$, etc.

The mereological prototype we want to consider can be specified at a conceptual level in another TBox $\mathcal{P}$. Concepts of $\mathcal{P}$ are described exclusively in terms of the immediate part relations: thus for example we have, $\mathtt{Wheel} \sqsubseteq \exists \mathtt{isImmediatePartOf}.\mathtt{Bike}$, $\mathtt{Spoke} \sqsubseteq \exists \mathtt{isImmediatePartOf}.\mathtt{Wheel}$ etc. In other words, $\mathcal{P}$ provides the representation of the perfect mereological match for the prototype: our goal now is to provide, using the perceptron operator, a way to relax the strict requirements of this DL description and allow for the classification of elements that "satisfy enough" of the mereological structure of the prototype. In terms of the game in Section 2, we can see this procedure as fixing the ideal bicycle that will be used to compare to the actual instances we are checking: then, we need to set how we can decide if the instances resemble enough the prototype to be classified as a bicycle, even if not necessary a perfect one.

The "approximate" version of the concept can be expressed via the perceptron operator, as shown intuitively in Section 2: the question then is how, starting from the structure of the mereological prototype, we can assign weights to the components in a way that is compatible with the mereological properties. For example, a naive bottom-up method for the computation of the weights in the tooth expression can be defined as follows. Let us denote with $\mathrm{NC}_{\mathcal{P}} \subseteq \mathrm{NC}$ the set of atomic concepts appearing in $\mathcal{P}$. For every concept $C \in \mathrm{NC}_{\mathcal{P}}$, we consider the set:

$$rel(C) = \{\, D \in \mathrm{NC}_{\mathcal{P}} \mid D \sqsubseteq \exists \mathtt{isImmediatePartOf}.C \in \mathcal{P} \,\}$$

We also use $rel^{*}(C)$ to denote the closure of this set (i.e. all the concepts that are reachable via `isImmediatePartOf` by C). For each $C \in \mathrm{NC}_{\mathcal{P}}$, we define its weight as:

$$w(C) = \sum_{D \in rel(C)} w(D) + 1$$

Note that the weight of each part is computed "bottom-up" from the leaves of the mereological structure (the concepts describing atomic parts) to the top concept (the concept describing the whole object). Of course, this measure is quite rough: for example, it provides more weight to the parts that are more detailed in terms of sub-parts. A more refined computation, for example, could proceed by *levels* of the part-hood hierarchy, so that composite parts (e.g. a bicycle wheel) have always more weight than the their single components (e.g. the spokes and tire of the wheel).

If $E \in \mathrm{NC}_{\mathcal{P}}$ is the top concept of the mereology (i.e. the one defining the whole prototypical object we are describing), the total cost of E should be $w(E)$. Then the "similarity degree" can be defined as $C \in \mathrm{NC}_{\mathcal{P}}$ is $deg(C) = w(C)/w(E)$, with the limit case $deg(E) = w(E)/w(E) = 1$.

Thus we can use these measures to express the "approximate" definition of the prototypical object as:

$$\nabla\!\!\!\!\nabla^t(\,\exists\texttt{hasPart}.C_1 : deg(C_1),\,\ldots,\,\exists\texttt{hasPart}.C_n : deg(C_n)\,)$$

with $\{C_1, \ldots, C_n\} = rel^*(E)$ and $t \in [0, 1]$ a threshold defining the level of similarity.

Referring again to our simple game in Section 2, now we are able to define our ideal bicycle (the mereological prototype in $\mathcal{P}$) and we have a way to decide (by the encoding in a tooth expression) if what we have is a bicycle or not for very easy cases. However, for the situations where the issues discussed in Section 3 emerge, this simple solution is still not enough.

5. Conclusions

In this paper, we introduced the possibilities and issues of using mereological definitions for classification via perceptron operators. We first described an example to show informally how the perceptron operator works. Then, by moving towards the limits of the application of the operator to mereological wholes, we discussed some of the problems that emerge. In particular, these limits include counting instances of the same concept, computing the weights of parts of parts (i.e., considering the transitivity of the parthood relation) and of overlapping parts, and managing situations where the parts are scattered or re-arranged. We then discussed some issues due to the adoption of classical mereology as the reference formalisation for parthood: the issue of transitivity emerged also in this case, but we noted also new issues caused by the decomposition and composition axioms. Finally, we proposed a DL formalisation for the basic scenarios of the approach, which can be used as a starting point for future discussions in order to find more general or more appropriate technical solutions for the emerged problems.

Our future work in this direction will be mainly focused on overcoming the difficulties discussed above. This has to be done both from the point of view of the definition of the operator, for example as relating to the options to faithfully compute the weights, and regarding the open issue which mereological theory to adopt. In this respect, a deeper consideration of the declination of classical mereology in DLs is needed (see e.g. [6]), a theory which in its more expressive versions can only be fully formulated in first- and/or second-order logic. Moreover, some research on alternatives of classical mereology that, for example, do not accept transitivity of parthood would be useful. As a related direction, we are also currently studying how some notion of *exception* in DLs (like the one used in [7, 8]) can be used in combination with perceptron operators, in particular in relation to graded or multi-relational definitions of defeasibility [9, 10].

References

[1] D. Porello, O. Kutz, G. Righetti, N. Troquard, P. Galliani, C. Masolo, A toothful of concepts: Towards a theory of weighted concept combination, in: M. Simkus, G. E. Weddell (Eds.), Proceedings of the 32nd International Workshop on Description Logics, Oslo, Norway, June 18-21, 2019, 2019.

[2] P. Galliani, O. Kutz, D. Porello, G. Righetti, N. Troquard, On knowledge dependence in weighted description logic, in: D. Calvanese, L. Iocchi (Eds.), GCAI 2019. Proceedings of

the 5th Global Conference on Artificial Intelligence, volume 65 of *EPiC Series in Computing*, EasyChair, 2019, pp. 68–80. URL: https://easychair.org/publications/paper/s17p. doi:10.29007/hjt1.

[3] P. Galliani, G. Righetti, O. Kutz, D. Porello, N. Troquard, Perceptron connectives in knowledge representation, in: C. M. Keet, M. Dumontier (Eds.), Knowledge Engineering and Knowledge Management, Springer International Publishing, Cham, 2020, pp. 183–193.

[4] A. C. Varzi, A. J. Cotnoir, Mereology, Oxford: Oxford University Press, 2021.

[5] A. Varzi, Mereology, in: E. N. Zalta (Ed.), The Stanford Encyclopedia of Philosophy, Spring 2019 ed., Metaphysics Research Lab, Stanford University, 2019.

[6] C. M. Keet, O. Kutz, Orchestrating a network of mereo(topo)logical theories, in: Proceedings of the Knowledge Capture Conference, K-CAP 2017, ACM, New York, NY, USA, 2017, pp. 11:1–11:8. URL: http://doi.acm.org/10.1145/3148011.3148013. doi:10.1145/3148011.3148013.

[7] L. Bozzato, T. Eiter, L. Serafini, Reasoning on *DL-Lite*$_{\mathcal{R}}$ with defeasibility in ASP, Theory Pract. Log. Program. 22 (2022) 254–304. URL: https://doi.org/10.1017/S1471068421000132. doi:10.1017/S1471068421000132.

[8] L. Bozzato, T. Eiter, L. Serafini, Reasoning with justifiable exceptions in $\mathcal{EL}_{\bot}$ contextualized knowledge repositories, in: Description Logic, Theory Combination, and All That, volume 11560 of *Lecture Notes in Computer Science*, Springer, 2019, pp. 110–134.

[9] L. Giordano, D. T. Dupré, An ASP approach for reasoning in a concept-aware multipreferential lightweight DL, Theory Pract. Log. Program. 20 (2020) 751–766. URL: https://doi.org/10.1017/S1471068420000381. doi:10.1017/S1471068420000381.

[10] L. Bozzato, T. Eiter, R. Kiesel, Reasoning on multirelational contextual hierarchies via answer set programming with algebraic measures, Theory Pract. Log. Program. 21 (2021) 593–609. URL: https://doi.org/10.1017/S1471068421000284. doi:10.1017/S1471068421000284.