Non-Monotonic Generalisation of an Ontology

Gabriele Sacco[1,2,*], Loris Bozzato[3], Oliver Kutz[2] and Michael Grüninger[4]

[1]Fondazione Bruno Kessler, Via Sommarive 18, 38123 Trento, Italy
[2]Faculty of Engineering - Free University of Bozen-Bolzano, NOI Techpark - Via Bruno Buozzi 1, 39100, Bolzano, Italy
[3]DiSTA - Università dell'Insubria, Via O. Rossi 9, 21100 Varese, Italy
[4]Department of Mechanical and Industrial Engineering, University of Toronto, Toronto, Ontario, Canada

Abstract

Human beings constantly deal with situations where their model of the world does not perfectly correspond to reality. In these cases, exceptions emerge, and so it is necessary to include them in the model and reason with them. If the latter phenomenon has a long history as a research topic in the field of artificial intelligence, the modelling one has received less attention. Therefore, in this paper, we start exploring this new perspective by discussing a non-monotonic generalisation of the ontology of damaged solid physical objects. By non-monotonic generalisation we mean a generalisation of the ontology which allows to represent exceptions. If one aspect of this consists in applying a non-monotonic logic, the other, more original, is how to actually manage the exceptions at the level of the axioms of the ontology. The present work concentrates on this latter aspect. In particular, we define what we mean by non-monotonic generalisation and why the notion of damage is suited for exploring the representation of exceptions. Then we outline the methodology to generalise the ontology of solid physical objects. Since this is intended as a case study which will constitute the first step of a more general line of research, we finally discuss some important future directions.

Keywords

Exceptions, Ontology, Damage, Defeasible Reasoning, Non-Monotonic Logics, Knowledge Representation

1. Introduction

We as humans are constantly in situations where our model of the world does not exactly match reality. For example, imagine to order online a book knowing it has 200 pages, but then, when delivered, yours has 199 pages because it misses page 25. In cases like this, normally we would not conclude that the knowledge that the book has 200 pages is false, but rather consider our copy an unlucky exception.[1] Studying these cases has been the interest of the research in *common-sense reasoning* since the early days of Artificial Intelligence (AI). The goal was understanding and so modelling such situations, in order to make artificial agents able to deal with them in a human-like way. The capability of humans involved has been called *defeasible reasoning* or *non-monotonic reasoning* or also *default reasoning*, and has been studied in philosophy [1, 2], computer science [3, 4] and cognitive sciences [5].

As it may be suggested by the terminology used, the focus of the investigations has been *reasoning*, so here we want to advance a different perspective, which sees defeasibility and exceptions under the lens of ontological modelling. In this sense the main question transforms from how to reason with defeasibility to how to model defeasible knowledge. Consider the example about the book used above, the classic line of research would consider it as the same of knowing that birds fly and discovering that Pingu the penguin does not or that months have either 30 or 31 days except for February, because

Proceedings of the Joint Ontology Workshops (JOWO) - Episode XI: The Sicilian Summer under the Etna, co-located with the 15th International Conference on Formal Ontology in Information Systems (FOIS 2025), September 8-9, 2025, Catania, Italy
*Corresponding author.
gsacco@fbk.eu (G. Sacco); loris.bozzato@uninsubria.it (L. Bozzato); oliver.kutz@unibz.it (O. Kutz); gruninger@mie.utoronto.ca (M. Grüninger)
0000-0001-5613-5068 (G. Sacco); 0000-0003-1757-9859 (L. Bozzato); 0000-0003-1517-7354 (O. Kutz); 0000-0003-3925-2398 (M. Grüninger)

[1]Note that here the concept of "knowledge" is used in a non-strict and non-technical way. In fact, from a philosophical point of view, knowledge cannot be false by the very definition of the term, e.g. in the tradition explicating knowledge as variants of "justified true belief". Therefore, we should speak of justified belief rather than full-fledged knowledge: however, for the sake of simplicity, we will continue to use the term knowledge with this loose sense also in the rest of the paper.

from a formal reasoning perspective they are equivalent. Our general interest is rather in investigating how this different cases should be ontologically understood and so coherently modelled. For instance, the 199 pages book can be considered an exception because it is a case of damaged object, and this is definitely not the case for Pingu and February.

This inquiry is especially relevant for the modelling methodologies. Indeed, the question addressed, as hinted above, is exactly how we should model those portions of the world where exceptions arise. As the examples already show, from an ontological perspective the phenomena which allows or cause such exceptional cases to appear are quite eterogeneous and it is not trivial to decide how they should be axiomatised and their relation with the exceptions should be captured in the ontology. Moreover, there is also a more practical case where a non-monotonic generalisation may be needed, that of ontologies where there are core axioms which are agreed by most and axioms which extend the theory in different arguable ways and so may be considered defeasible.

Therefore, in this paper we propose to explore how to *non-monotonically generalise* an ontology by considering the particular case of *damage*. We start from the ontology of damaged solid physical objects developed in [6] and we discuss why it would make sense to generalise it non-monotonically, what does it mean to make such a generalisation and how it can be done. The discussion will be conducted in the theoretical framework for defeasibility and exceptions proposed in [7], which provides important clarifications and interpretations of key points. We propose this investigation as the starting point for the more general research on the non-monotonic generalisation of ontologies. Consequently, one of the goals of the paper is also to use this specific case to individuate and formulate the crucial issues which need to be addressed for the non-monotonic generalisation of ontologies.

We start by introducing the ontology of Damaged Solid Physical Objects (SoPhOs) in Section 2. Then, in Section 3 we explain why it can be non-monotonically generalised and so what is a non-monotonic generalisation of an ontology. At this point, we describe the framework for defeasibility and exceptions in Section 4. In Section 5 we discuss the notion of *dimension of damage*, which lays the basis for the non-monotonic extension of SoPhOs through *exceptional predicates* proposed in Section 6. Finally, in Section 7 we conclude with a discussion some related issues and the future directions we envision for generalising this approach beyond SoPhOs.

2. Damaged Solid Physical Objects

In this section we will present the *Ontology of Solid Physical Objects* and its extension developed in [6] for damaged objects. Then, we will argue how damage can be understood as exceptionality.

2.1. Ontology of Solid Physical Objects

The design of the Ontology of Solid Physical Objects (SoPhOs) follows the principle that each module axiomatizes a different parthood relation for solid physical objects. A solid physical object is defined to be *an object that is self-connected, has some shape and boundary, is made of some material, and occupies some space.* The ontology consists of four modules (see Figure 1). Each module features a characteristic ontology in the upper ontology and gives rise to corresponding distinct parthood relations following the approach of mereological pluralism.

A solid physical object is a material object. Matter constitutes solid physical objects and is one of the prime characteristics we determine for solid physical objects. These theories within the Matter Module axiomatize the constitution relation *constitutes* between *mat* (matter) and *material_object*, defines *chunkOf* as the parthood relation within matter and *portionOf* as the parthood relation within material objects. .

The basic intuition of the Structure Module is that a physical object is self-connected, so the external connection axioms and corresponding parthood ontology are axiomatized by an ontology that is synonymous with the mereotopology T_{cisco} [8].

The Shape Module reuses the Box World Ontology [9, 10], which is an ontology for shapes composed of surfaces. As such, it can represent aspects of an object that have different dimensionality e.g. the 2D

	Characteristic	Parthood on Characteristic	Parthood on Characteristic with Spatial Relationship	Characteristic vs. Object	Defined Parthood on Characteristic Object	Parthood on Characteristic Object with Location	Connection on Characteristic	Damage Related
Theories	**Matter Module**	T_{Matter}	T_{occupy} (adjusted)	$T_{Constitution}$	$T_{Material\_Object}$	$T_{confinement}$	T_{Attach}	Removal of material
Relations	*Mat* (continuous)	*chunkOf*	*region_part*	*constitutes*	*portionOf* Damage: *ab_portion*	*confinedIn* Damage: *ab_confinement*	*attach*	
Theories	***Shape Module***	$T_{BoxWorld\_feature}$	$T_{MT\ Multidimensional\ Object\ Mereotopology}$	T_{Bounds}	T_{Piece}	$T_{Containment}$	T_{Affix}	Handle of mug; Dent in chair leg; Egg in box;
Relations	*ShapeFeature* (discrete)	*featureOf* $T_{MWorld\_component}$ *bSCmpOf* *sCmpOf*	*surface_part* *box_part*	*bounds*	*pieceOf* Damage: *ab_piece* $T_{component}$ *componentOf* Damage: *ab_component*	*containedIn* (convex hull) Damage: *ab_containment*	*affix* T_{Joint} *joint*	

Figure 1: Module Ontologies in Ontology of Solid Physical Objects (SoPhOs)

surface of a 3D object together with one-dimensional features such as edges and ridges. Ultimately, we want to categorize sufficient number of types of shape to identify common atomic shapes (i.e. cylinder, box, curved cylinder, ball) and differentiate the pieces of one object to say that each piece is a atomic shape. We name the placeholder for parthood ontology of the shape module as Boundary Ontology, to include axioms that separate physical object into pieces by its physical boundary according to convexity.

The Spatial Module represents the location of physical objects within abstract space together with the spatial relationship between the regions occupied by objects. Incorporating with the representation of enclosure in the Occupy Ontology [11], the Containment Ontology defines the containment parthood relation.

2.2. Damage as Abnormality

SoPhOs is used in [6] as the basis for an extension able to model also damaged solid physical objects. The approach adopted is that of considering damage as some sort of abnormality of some of the parthood relations used in SoPhOs. Specifically, variants of these parthood relations are introduced which correspond to the abnormal versions of the originals (see Figure 1).

The choice of using abnormality for representing damaged solid physical objects comes from understanding damage as a deviation from the ideal definition of the class to which the object is an instance. That is, we consider the definition as specifying the "*intended properties of an (ideal) object, so that any divergence from the specification of these properties (e.g. missing or spurious parts) are indications of damage.*" [6]

Moreover, in [6] the authors 'use the notion $\neg Ab$ to describe the condition of an object complying with its design, that is, of being not abnormal, or "ideal" '. Consequently, the notion of *not abnormal*, or in other terms of *normality*, is introduced for describing the condition of 'respecting' the ideal or design specifications. For this reason damage, as a deviation from the ideal, corresponds to a kind of abnormality. Interestingly, in the paper, there is only one kind of normality, but many abnormalities, corresponding to different parthood relationships. Considering the account just described, this is not surprising since there is only one way of complying to the ideal specifications, namely having those requirements, while there may be many different ways of deviating from them. To stress this interesting difference, we will talk of 'normality' at the singular and of 'abnormalities' in the following. As we will discuss later in Section 5, this is one aspects of what we call *dimensions* of damage, where a dimension is, broadly, a way in which an object may be damaged, that is a way in which it can differ from the ideal specifications.

Here, we will not explore systematically the relationship between damage and normality, however we will make some points which needs to be kept in mind during the rest of the work. Firstly, damage may be considered as some kind of abnormality, but abnormality taken in a broader sense does not reduce to damage, even if we consider only solid physical objects. In fact, there may be different reasons or modes in which an object diverge from the ideal, in fact, consider these examples:

1. A chair that had lost a leg.
2. A drawer assembled with the bottom part upside down.

These two cases arguably represents two cases of abnormality for solid physical objects, but only 1 is a case of damage, while 2 can be considered a case of mis-assembly. This distinction is clear if we have in mind, for instance, the scenario motivating the work in [6], that is a marketplace for second-hand furniture: surely one would have different attitudes with respect of the two items listed above.

Consequently, a second point to make is that damage may interact with other forms of abnormality. Using the example above, it is clear that the chair mentioned in (1) can be abnormal also due to a mis-assembled part, e.g. another leg assembled in the wrong position. This would make the chair abnormal in two different ways, maybe also at different times: damaged in the case of (1) and mis-assembled in the other.

A third element to keep in mind consists of the fact that different forms of abnormality can require different ontological treatments and so, as we will see, different ways of generalising the ontology making it non-monotonic. We will discuss more in detail the case of damage in the following, but consider, for instance, a case where abnormality is interpreted as departing from typicality, as *Typical chairs have four legs*, but this chair here has only three. In this case, we may argue that typicality affects expectations or beliefs rather than objects and so this kind of abnormality should not impact on the modelling of the ontology, but only the reasoning. We will discuss a bit more in detail this topic in Section 4.

3. What is a Non-Monotonic Generalisation of an Ontology?

Keeping these considerations in mind, we will now explain why the interpretation of damage as a kind of abnormality is what makes this ontology a good starting point for the exploration of what a *non-monotonic generalisation* of an ontology could be and what do we mean by it.

A first hint that this ontology of damaged solid physical objects is particularly suited for introducing and exploring such a generalisation comes from the fact that abnormality has been a phenomenon linked to defeasibility since the first investigations on it, being the key notion of one of the first formal approaches, namely *Circumscription* [12, 13].

This connection has already been acknowledged in [6], even if for taking the distances from it. Their claim is that in this case abnormality is not to be understood as typicality, and so it does not correspond to the use it has in non-monotonic reasoning. We agree that here abnormality is not typicality, however, we think that this does not disqualify the generalisation towards non-monotonicity of the ontology. In fact, the proposed interpretation of damage as deviating from the intended properties of an object fits perfectly with a non-monotonic reading: "objects O have property P, unless they are *damaged*" or "objects O have property P, except for the *damaged* ones". This reading correspond to considering damaged objects as exceptions to the general knowledge of how the non-damaged, or normal, objects should be.

Therefore, we consider the ontology of damaged solid physical objects a good use case for starting the investigation of the non-monotonic generalisation of an ontology. However, we want to point out that there are other relevant uses of ontologies which would benefit from being able to accept exceptions.

Consider, for example, the health and medical domain: even if there are standard or normal descriptions which constitute the basic notions of the ontology, we have a lot of irregularities or exceptions. Think, for example, of some normal information about the heart like "The heart is on the left side of the body" or "There are no *venae cavae* in the left side of the heart". Now, both of them admit exceptions, which in the health domain are important at least as the regularities. You may have a person with *situs* *inversus*, and so with the heart on the right side of their body, or another one with a persistent left superior *vena cava*.

Another case is the merging of different ontologies. In this case, a novel approach may be that of deciding for each ontology which conflicting axioms we want to keep as strict and which instead we are willing to make defeasible, that is admitting exceptions. Consider, for instance, [14] which is a case of axiom weakening: the idea would be that instead of weakening axioms in the sense of making them more general and so restoring the consistency, we can instead weaken them by allowing exceptions and making the final ontology non-monotonic.

In both these situations, it would be useful to *generalise* the ontology or the ontologies non-monotonically, that is, such that they are able to model exceptional entities.

A first objection which can be made is that a natural way to understand this generalisation would be to use one of the non-monotonic logics developed so far as the formal language used for the ontology. This is true, yet, we think it would not be enough. Even if exploiting these logics is a necessary part of this non-monotonic generalisation, our comprehension is that we need to address *ontologically* the matter of exceptionality and the related notions involved, like defeasibility, normality, typicality, ideal, according to the domain we are modelling. More concretely, we want to address these issues at the level of the axiomatisation and not only at the level of the language, which is mainly focused on reasoning rather than modelling.

Returning to the case of damaged objects: consider the case of a damaged book, say a copy of the *Tractatus logico-philosophicus* with a page with a ripped corner. We consider it an exception with respect of the ideal *Tractatus logico-philosophicus*. In order to model this scenario we will generalise non-monotonically a relevant ontology, in our case the ontology of solid physical objects. The first step is to choose a suitable non-monotonic logic as our language, but the other fundamental step is to adjust the axiomatisation in order to capture the ontological behaviour of the exceptional entities. For instance, we want that our ontology model damage as something transitive with respect of parthood, in order to conclude that if the page of a book is damaged then also the book is damaged. More generally, if a part is damaged, then the whole of which it is part is damaged too.

Consequently, with non-monotonic generalisation of an ontology we intend the application of a non-monotonic language to an ontology, but especially introducing axioms which allow to model also the ontological features of exceptions.

4. Defeasibility and Exceptions

In this section, we will briefly introduce the approach we endorse with respect to the exceptionality and defeasibility, which is the one presented in [7]. The goal of the work in [7] is to provide a characterization of the notion of exception and then use such notion to recognize the different representation of exception and ontological commitments in existing non-monotonic systems: intutively, the aim is to allow the modeller of a specific domain to choose the system that best captures the idea of exception that is needed for the specific knowledge of interest.

The first step in this direction is to recognize what is considered as "exception" in common sense terms. Thus, the authors first consider a collection of common-sense definitions of exception: what emerges is that such definitions all have in common *(i)* a notion of generalisation and *(ii)* an instance that is excluded from such general statement. Considering the notion of generalisation, the paper then shows the distinction between actual exceptions with *counter-examples* and *errors*: the case of exceptionality is recognized in the case where both the general statement and the particular instance are considered to be valid. These initial considerations already allow us to provide a first intuitive explanation of the notion of exception: *an exception is an individual justifiably excluded from a generalisation, without causing a contradiction.*

This notion of exception is then refined by studying the notion of *generalisation*, considering its characterization in the literature of linguistics, philosophy and cognitive science, in particular in the concepts of *generics* and *ceteris paribus laws*. What emerges is that in generic propositions, there can be instances that do not instantiate the predicate property and the notion of a defeasible generalisation can be distinguished from an absolute (universal) generalisation.

This leads to a distinction of *universal* and *defeasible generalisations*: while an universal generalisation can be explained as an universal quantification, "all of the Ps are Qs" ($\forall x(Px \rightarrow Qx)$), a defeasible generalisation is explained as a conditional universal quantification, "a strict subset of Ps is also a subset of Qs". In [7], universal generalisation is called $\forall$-generalisation and defeasible generalisation is called $\widetilde{\forall}$-generalisation. The different non-monotonic systems formally characterize the meaning of $\widetilde{\forall}$-generalisation and thus the meaning of "strict subset": those elements that do not belong to such set, intuitively characterize *exceptions*. Thus, we can see a defeasible generalisation as a universal generalisation where we remove from the range of universal quantifier the exceptional individuals. In this light, the notion of exception in [7] is presented as: *An exception is an individual belonging to a justified subset of the domain which if explicitly excluded from the scope of the quantifier of a false $\forall$-generalisation transforms it in a true $\widetilde{\forall}$-generalisation.* The definition above leaves open the notion of *justification*: the notion of exception in the different formal systems is actually dependent on their interpretation of justification.

Starting from these distinctions, four known formal systems for representation of defeasibility are compared, showing the differences that may occur from an ontological perspective. In particular, the non-monotonic systems in [4] are systematically compared with respect to such framework, namely *Closed World Assumption*, *Circumscription*, *Default logic* and *Autoepistemic logic*. Given the explanation of exception provided above and the understanding of the key aspects of the formal systems, the systems are compared along four different lines: these can be seen as features of the characterization of exceptions in the different systems that the modeller can choose to better suit the modelling problem at hand. In particular, systems are compared with respect to:

(i). *Syntactic or Semantic approach:* whether the defeasibility is represented through a syntactic or semantic representation;

(ii). *Epistemic or Ontological level:* whether defeasibility is represented as ontological object (ontological level) or it emerges at the level of knowledge (epistemic level);

(iii). *Explicit or implicit representation of Exceptions:* whether the concept of exception is clearly represented as a class of objects or it emerges implicitly from the representation of the defeasibility;

(iv). *Logical or meta-logical:* whether defeasible generalisations are modelled at the logical level (like classical FOL formulas) or if these are outside the language, and thus meta-logical.

The work in [7] is intended to provide a way to choose the right formalism for the representation problem at hand: thus the current paper can be seen as an example of such an application to the case of the SoPhOs ontology.

In fact, the understanding of damage as abnormality aligns perfectly to this account of defeasibility and exception, where an exception is an instance of a generalisation excluded from it:

Concept Explication 1. *An* ***exception*** *is an individual belonging to a* justified *subset of the domain which if explicitly excluded from the scope of the quantifier of a false $\forall$-generalisation transforms it in a true $\widetilde{\forall}$-generalisation.*

In this case the $\forall$-generalisation is the description of the class of objects in terms of their ideal specification, namely their designed features. The justification, as we will see in the next section, correspond to the specific dimension of damage involved. That is, the type of damage explain and so justify the exclusion of the specific object from the class as it is ideally described.

Finally, note that we do not need a precise account of what damage is, instead we only commit to an understanding of damage as making an object deviate from the designed specifications of the class it belongs to. Consequently, the definition of the class of the objects becomes a true $\widetilde{\forall}$-generalisation. Therefore, the source of these defeasible generalisations is their being ideal because they regard the intended properties, but not necessarily the actual ones. Therefore, for some specific instances, there can be a discrepancy with the concrete reality those generalisations refer to, and these instances can be the damaged ones.

5. Dimensions of Damage

Now we can start to discuss how to actually generalise the ontology non-monotonically. As we have argued in Section 3, the reason for applying this generalisation to SoPhOs is the interest to represent damaged solid physical objects. Consequently, it is in representing damage that defeasibility is involved, where damage is understood as diverging from the ideal. In this sense, the description of the ideal object can be considered the defeasible generalisation which is subject to the exceptions, which in this case are the damaged objects. Therefore, in this section we explore how an object can diverge from its ideal definition and so be considered a damaged object. We call these different ways *dimensions of damage* and they will give us the theoretical basis for the formal non-monotonic generalisation of SoPhOs discussed in the next section.

In order to explore these dimensions, we can start from the examples used in [6] (see Figure 2 and Figure 3). We have the intended properties of a three legged chair `MyChair`, which has exactly three legs, each made of three pieces, and five damaged instances of that chair. The standard design of `MyChair` with three legs is represented in Axioms 1-3 as in Figure 2 below, showing the intended properties of components, pieces and portions respectively. Axioms 4-8 in Figure 3 show an example specification of damaged chairs. C_1 denotes a replaced leg resulting an abnormal component; C_2 denotes a dent in one leg of a three-leg chair resulting an abnormal piece; C_3 denotes one leg missing; C_4 denotes when there is an extra leg; last but not least, C_5 denotes some additional matter is added to the chair.

$$\begin{gathered} MyChair(x) \supset (\neg Ab(x) \\ \equiv (\exists y_1, y_2, y_3) Leg(y_1) \wedge Leg(y_2) \wedge Leg(y_3) \wedge (y_1 \neq y_2) \wedge (y_1 \neq y_3) \wedge (y_2 \neq y_3) \\ \wedge componentOf(y_1, x) \wedge componentOf(y_2, x) \wedge componentOf(y_3, x) \\ \wedge ((\forall z) componentOf(z, x) \supset z = y_1 \vee z = y_2 \vee z = y_3)) \end{gathered} \tag{1}$$

$$\begin{gathered} Leg(x) \supset (\neg Ab(x) \\ \equiv (\exists y_1, y_2, y_3) bottom(y_1) \wedge side(y_2) \wedge top(y_3) \wedge (y_1 \neq y_2) \wedge (y_1 \neq y_3) \wedge (y_2 \neq y_3) \\ \wedge pieceOf(y_1, x) \wedge pieceOf(y_2, x) \wedge pieceOf(y_3, x) \\ \wedge ((\forall z) pieceOf(z, x) \supset z = y_1 \vee z = y_2 \vee z = y_3)) \end{gathered} \tag{2}$$

$$\begin{gathered} MyChair(x) \supset (\neg Ab(x) \\ \equiv (\exists y) Mat(y) \wedge constitutes(y, x) \wedge ((\forall z) portionOf(z, x) \equiv chunkOf(z, y))) \end{gathered} \tag{3}$$

Figure 2: $T_{mychair}$: Representation of Standard MyChair Design

Each class of solid physical objects is axiomatized by sentences of the form seen in Figure 2. By using the abnormality predicate Ab, we allow the existence of objects in a class even if they do not satisfy the conditions for the ideal object in that class; inconsistency is avoided since such an object is simply an abnormal instance of the class. Furthermore, we can use the parthood relations in SoPhOs to identify the nature of the abnormality – missing matter vs. spurious matter, missing shape feature vs. unintended shape features, missing components vs. extra components.

$$MyChair(C_1) \wedge Leg(L_1) \wedge Leg(L_2) \wedge Leg(L_3)$$
$$\wedge ab\_component(L_1, C_1) \wedge componentOf(L_2, C_1) \wedge componentOf(L_3, C_1) \supset Ab(C_1) \quad (4)$$

$$MyChair(C_2) \wedge Leg(L_4) \wedge Leg(L_5) \wedge Leg(L_6)$$
$$\wedge component(L_4, C_2) \wedge componentOf(L_5, C_2) \wedge componentOf(L_6, C_2)$$
$$\wedge Dent(P_1) \wedge Bottom(P_2) \wedge Side(P_3) \wedge Top(P_4)$$
$$\wedge ab\_piece(P_1, L_4) \wedge pieceOf(P_2, L_4) \wedge pieceOf(P_3, L_4) \wedge pieceOf(P_4, L_4) \supset Ab(C_2) \quad (5)$$

$$MyChair(C_3) \wedge Leg(L_7) \wedge Leg(L_8) \wedge componentOf(L_7, C_3) \wedge componentOf(L_8, C_3)$$
$$\supset Ab(C_3) \quad (6)$$

$$MyChair(C_4) \wedge Leg(L_9) \wedge Leg(L_{10}) \wedge Leg(L_{11}) \wedge Leg(L_{12})$$
$$\wedge ab\_component(L_9, C_4) \wedge componentOf(L_{10}, C_4) \wedge componentOf(L_{11}, C_4) \wedge componentOf(L_{12}, C_4)$$
$$\supset Ab(C_4) \quad (7)$$

$$MyChair(C_5) \wedge constitutes(M_1, C_5) \wedge chunkOf(M_2, M_1) \wedge ab\_portion(M_2, C_5) \supset Ab(C_5) \quad (8)$$

Figure 3: $T_{chairexample}$: Representation of Damaged MyChair Examples

5.1. Dimensions as a Subset of the Signature

The first thing we can notice is that damage may be related to different relations of the ontology, as it emerges by the introduction of the abnormal variants of different mereological relations of SoPhOs. This can be considered a first understanding of "dimension" of damage, since it is a first specification of the divergence from the ideal and it is the sense already captured in [6] by introducing the abnormal version of some parthood predicates. Therefore, we can consider a dimension of damage the relation or property of the ontology affected by the damage. Consequently, the set of all the dimensions of damage understood in this way corresponds to a subset of the signature used in the ontology. For instance, in SoPhOs, one dimension of damage may be being damaged with respect of the relation `componentOf`, while another dimension may be being damaged relatively to the `pieceOf` relation, and the set of all these dimensions corresponds to the darker cells of Figure 1.

5.2. Dimensions as Justification of the Exceptionality

However, there is a different sense in which we can talk of dimensions of damage, which specify further how the damaged object diverges from the intended properties. Indeed, in the examples, this is shown by the fact that both C_1 and C_4 are damaged because of the `componentOf` relation, which is substituted by `ab_component`, but the intended damage to be represented is different, respectively a replaced leg and an extra leg. For this reason we propose here a refinement of the treatment of the damaged objects used in [6], obtained through a discussion of these examples. This refinement will be the theoretical basis for the formal non-monotonic extension of the ontology.

What emerges from the examples, which exploit only the first sense of dimension of damage, is that there are some dimensions of damage which are very similar, but there are treated differently, and others which are treated similarly but are quite different. For example, the former are the cases of the missing leg of C_3 and of the extra leg of C_4. In C_3 the abnormality, that is the damage, emerge from the absence of a part, namely the third leg. Therefore, the divergence from the ideal consists in violating the condition about the number of parts, but all the other parts are not abnormal. Also in the case of C_4, which has an extra leg, the divergence from the intended properties of `MyChair` consists in a different number of parts, particularly, a different number of components. However, in this case we have the extra leg which is related to MyChair through the ab_componentOf relation.

If from one side this allows to immediately identify what is responsible for the divergence from the ideal descriptions, on the other we are treating two symmetric cases of the same divergence, that is the violation of the condition about the exact number of components, in different ways.

The second situation, that of different kinds of divergence treated similarly, emerges comparing C_1, which has a replaced leg and C_4: the replaced leg of C_1 is represented as $\texttt{ab\_component}(\texttt{L}_1, \texttt{C}_1)$, that is in the same way of the extra leg of C_4. However, the divergence from the ideal, that is the reason these are cases of abnormality, seems to be different in nature: if the case of the extra leg concerns the number of parts, that is the quantity of the parts, the case of the replaced leg regards rather a qualitative aspect. Consequently, the divergence of C_1 from the ideal specification of MyChair regards some properties of the leg itself instead of a "meta" property like the number of legs of the chair. This suggests also that the leg should diverge from the ideal specifications of the legs, at least those used for MyChair, and this divergence is transmitted to the whole.

Another element which may suggest that the case of having an extra leg should be treated in a way more similar to that of a missing leg then to that of the replaced leg is that while in the latter case it is clear which leg is the abnormal component, in the former it is not necessarily so. In other terms, it is not necessarily evident which of the legs is the extra one. So it may be better to represent C_4 simply as

$$\texttt{Chair}(C_4) \wedge \texttt{Leg}(l_1) \wedge \texttt{Leg}(l_2) \wedge \texttt{Leg}(l_3) \wedge \texttt{Leg}(l_4) \wedge$$
$$\texttt{componentOf}(l_1,\ C_4) \wedge \texttt{componentOf}(l_2,\ C_4) \wedge \texttt{componentOf}(l_3,\ C_4) \wedge \texttt{componentOf}(l_4,\ C_4)$$

This description makes C_4 an abnormal chair even without the presence of ab_component because it violates the condition that it must have exactly three legs.

This latter consideration leads to a more general point: C_1, C_2 and C_5 do not explain what is the divergence from the ideal properties. In other words, they use the general technique of introducing abnormal version of the mereological relations involved in the damage as all encompassing for the relative damages. However, the justification for the divergence, or the exceptionality, is not ontologically explicit.

Consider, for example, C_1: the reason for one of its leg, the replaced one, to be considered abnormal is not present. This could be something like being of a different material or having a different shape, that is there is a divergence already at the level of the ideal specifications of the legs of the three-legged chair. Therefore, this would be something similar to the case of C_2, where $\texttt{Dent}(\texttt{P}_1)$ violates the ideal description of the leg. Even so, the formal description of C_2 does not show the divergence from the ideal three-legged chair, but only the divergence of L_4 from the ideal specifications of the leg. What is lacking is the fact that the dent makes the leg abnormal and so also the chair results abnormal.

This is exactly the central point. In order to being able to distinguish the different dimensions of damage, a first step is surely to identify the affected relation or property, but then we need also to distinguish how, and so why, this relation or property is the subject of the divergence from the ideal specifications. In other words, what is needed is to make clear what justifies the divergence from the ideal description of the damaged object. Or also, what justifies the exceptionality of the damaged object.

Therefore, we can distinguish two different senses of 'dimension' of damage: (i) a broad one, according to which relations or classes are those responsible for the divergence from the ideal, we may have different dimensions of damage in the sense that damage may regard different relations or properties of our ontology; (ii) a specific one, where 'dimensions of damage' refer to a deeper understanding of the nature of damage or of how something may be damaged. In the latter sense, the dimension correspond to the kind of justification for the diversion from the ideal specification, that is for the justification of the exceptionality.

6. A Non-Monotonic Generalisation of the Ontology of Damaged Solid Physical Objects

In this section we explore an actual generalisation of the ontology of damaged solid physical objects. We start by individuating a suitable formalism for achieving non-monotonicity and then we proceed by introducing axioms which exploit what we call *exceptional predicates*.

6.1. Choosing the Right Formalism

As said in Section 4, the distinctions made in [7] help in choosing the right formalism for the targeted problem, so we can use them also in this case. Here we are starting from an ontology which has already introduced a notion of abnormality through new relations and properties, and we want to proceed along these lines. Therefore, a decision between the distinctions has already been made. Defeasibility here is emerging at the ontological level, since damage is not something dependent from the knowledge of a subject, but it regards rather properties or relations of objects, as the discussion on the dimensions of damage above shows. This already individuates among the formalisms taken in considerations in [7], unsurprisingly, *circumscription*, since it is the only one with addressing defeasibility at an *ontological* level. However, we can check if the other commitments of circumscription with respect of the other distinctions comply also with the present goal.

Circumscription is a *semantic* approach: this correspond to the treatment of damage adopted here, since we are using predicates and we want to individuate the damaged objects instantiating the corresponding properties and relations. Therefore, also its *explicit* representation of exceptions is a desirable commitment, since in this case the exceptions are exactly the damaged objects. Finally, we are clearly using formulas of the language for representing them and so also modelling defeasibility at the logical level is welcomed.

This means that a circumscription-like formalism is fitting for the non-monotonic generalisation of the ontology of damaged objects. Of course, we do not claim here that this is the only possible choice, but for the present purpose of introducing and exploring the non-monotonic generalisation of an ontology through the case of the ontology of damaged solid physical objects, the fact that circumscription is a good option is enough.

6.2. Generalising through Exceptional Predicates

We have seen that to understand damage as exceptionality it means that the damaged object is considered an exception, which means that this object is still an instance of the class corresponding to the undamaged objects, but it does not respect its definition. Therefore, we can individuate the deeper sense of dimensions of damage by specifying how these definitions may be violated.[2]

Therefore, to generalise the ontology we can start by individuating the axioms which will be made defeasible. For example, consider Axiom 1 a damaged $MyChair$ would falsify it in some way, therefore the first step is to negate it. To do so we make explicit the universal quantification and remove the reference to abnormality, since we will reintroduce this notion later to make the axiom defeasible,:

$$\neg(\forall x)MyChair(x) \supset (\exists y_1, y_2, y_3)Leg(y_1) \wedge Leg(y_2) \wedge Leg(y_3) \wedge (y_1 \neq y_2) \wedge (y_1 \neq y_3) \wedge (y_2 \neq y_3)$$
$$\wedge\ componentOf(y_1, x) \wedge componentOf(y_2, x) \wedge componentOf(y_3, x)$$
$$\wedge\ ((\forall z)componentOf(z, x) \supset z = y_1 \vee z = y_2 \vee z = y_3))$$

This only represent the existence of a damaged objects, in fact we can rewrite the above formula as

[2] This approach can also, in principle, open the possibility of establishing a grade of severity of the damage understood as how many conditions given in the definition are violated. This could be used to decide when a damaged object ceases to be a valid instance, even if an exceptional one, and when it starts to be something else. However, we will keep this line of research for future work.

$$(\exists x)MyChair(x) \land \neg((\exists y_1, y_2, y_3)Leg(y_1) \land Leg(y_2) \land Leg(y_3) \land (y_1 \neq y_2) \land (y_1 \neq y_3) \land (y_2 \neq y_3)$$
$$\land\ componentOf(y_1, x) \land componentOf(y_2, x) \land componentOf(y_3, x)$$
$$\land\ ((\forall z)componentOf(z, x) \supset z = y_1 \lor z = y_2 \lor z = y_3)))$$

Now we need to specify the dimensions of damage we want to consider and so individuate which part of the axiom is violated. We can consider, for instance, the missing leg chair. This would mean that the violation regards the absence of one of the components y_1, y_2 or y_3:

$$(\exists x)MyChair(x) \land (\exists y_1, y_2)Leg(y_1) \land Leg(y_2) \land (y_1 \neq y_2)$$
$$\land\ componentOf(y_1, x) \land componentOf(y_2, x) \land ((\forall z)componentOf(z, x) \supset \neg(z = y_1 \lor z = y_2))$$

Then, we can use this description of the damage missing leg for $MyChair$ as the definition of the predicate $MissingLeg_{MyChair}$, which represent this dimension of damage. To do so, we remove the existential quantifier, since we do not want to commit with the necessary existence of such a chair and instead introduce a universal quantifier binding the predicate with the definition:

$$(\forall x)MissingLeg_{MyChair}(x) \supset$$
$$(MyChair(x) \land (\exists y_1, y_2)Leg(y_1) \land Leg(y_2) \land (y_1 \neq y_2)$$
$$\land\ componentOf(y_1, x) \land componentOf(y_2, x)$$
$$\land\ ((\forall z)componentOf(z, x) \supset \neg(z = y_1 \lor z = y_2)))$$

Next, we impose that $MissingLeg_{MyChair}$ is a predicate to be minimised, that is which has the smaller extension possible according to the knowledge we have. We call these predicates, representing the dimensions of damage and which are minimised *exceptional predicates* and they are the main means for the non-monotonic generalisation of the ontology.

Finally, to make the starting Axiom 1 defeasible, we will collect all the exceptional predicates, which in this case represent the different dimensions of damage, in a more general exceptional predicate, like $\forall(x)Damaged_{MyChair}(x) \supset (MissingLeg_{MyChair}(x) \lor ExtraLeg_{MyChair}(x) \lor ...)$ and insert it in Axiom 1 as

$$(\forall(x)(MyChair(x) \land \neg Damaged_{MyChair}(x)) \supset$$
$$(\exists y_1, y_2, y_3)Leg(y_1) \land Leg(y_2) \land Leg(y_3) \land (y_1 \neq y_2) \land (y_1 \neq y_3) \land (y_2 \neq y_3)$$
$$\land\ componentOf(y_1, x) \land componentOf(y_2, x) \land componentOf(y_3, x)$$
$$\land\ ((\forall z)componentOf(z, x) \supset z = y_1 \lor z = y_2 \lor z = y_3))$$

The steps described here are specific of the damaged solid physical objects ontology we are considering, however the methodology used is generalisable beyond this specific scenario, at least when the chosen formalism is circumscription. Therefore, when we want to non-monotonically generalise an ontology, the procedure to follow is:

1. Individuate the axioms which should be made defeasible, that is which may have exceptions.
2. Negate them in general.
3. Individuate the relevant ways in which the axioms can be falsified.
4. Introduce these modes of falsification in the ontology as exceptional predicates.
5. Reformulate the original axioms by posing the condition of not being abnormal.

In case another formalism is preferred to circumscription, further work is needed to properly generalise the above procedure. Notwithstanding, here we setted the starting point in order to develop such a generalisation.

7. Conclusion and Future Works

In this paper we explored how to non-monotonically generalise an ontology by considering the particular case of *damage*. We started from the ontology of damaged solid physical objects developed in [6] and we discussed why it would make sense to generalise it non-monotonically, what does it mean to make such a generalisation and how it can be done. The discussion will be conduct in the theoretical framework for defeasibility and exceptions proposed in [7], which provide important clarifications and interpretations of key points.

As an exploratory work interested mainly in the non-monotonic generalisation of an ontology, we did not enter in the detailed ontological analysis of damage, although this is an interesting topic to investigate in the future. A first intuition which can be explored is that damaged objects need to have underwent an event or an activity which changed them from undamaged to damaged. That is, to be damaged, something had to be previously not damaged.

Secondly, as we said above, it would be worth looking into the possibility of establishing a graded notion of damage, thus being able to distinguish between damages which simply affect, ruining, objects and damages which destroy the interested object, completely changing its previous classification under a specific class.

Thirdly, also the notion of dimension of damage may be further investigated. In this sense, one aspect which definitely is worth exploring is the relations between the different dimensions and how the exceptional predicates modelling them may interact. Moreover, for instance, [15] introduce a taxonomy of part-whole relations which may establish another framework for individuating other dimensions.

Finally, a thorough comparison between the original ontology of damaged solid physical objects of [6] and its fully defined non-monotonic generalisation in terms of models entailed would be useful both as a possible evaluation of the latter and also as a source of more insights on the general process of non-monotonic generalisation of an ontology.

In this regard, the present investigation had the aim of serving as the starting point for the more general research on the non-monotonic generalisation of ontologies. Therefore, it also allowed to raise some elements which needs to be taken into consideration in the future work.

The first and more pressing one is how the methodology to develop a non-monotonic generalisation can be fully generalised, also beyond the formalism chosen.

Another important element to investigate in the future works is abnormality, or exceptionality in general, with respect of relations. In this paper we have not dedicated much space to this, but this issue was implicitly present in the discussion of Section 5: taking, for instance, the case of `ab_component`, do we understand this as "the *relation* `componentOf` is abnormal"? Maybe because it is violating its defining axioms. Or rather as "the *component* we are referring to is abnormal"? As, for example, in the case of the replaced leg? Or even like "this component is making the *whole* abnormal"? Like in the case of the extra leg?

Finally, it would be interesting to investigate other cases where a non-monotonic generalisation of an ontology seems to be a promising approach, like those mentioned in Section 3. This would mean both an application to other domains, like the medical one, but also other scenarios, like the merging of two ontologies.

Acknowledgments

We acknowledge the financial support through the "Abstractron" project funded by the Autonome Provinz Bozen - Südtirol (Autonomous Province of Bolzano-Bozen) through the Research Südtirol/Alto Adige 2022 Call and by the European Commission funded projects "Next Generation Internet Transatlantic Fellowship Program NGI Enricher" (grant #101070125).

Declaration on Generative AI.

The authors have not employed any Generative AI tools.

References

[1] R. Koons, Defeasible Reasoning, in: E. N. Zalta (Ed.), The Stanford Encyclopedia of Philosophy, Summer 2022 ed., Metaphysics Research Lab, Stanford University, 2022.

[2] C. Strasser, G. A. Antonelli, Non-monotonic Logic, in: E. N. Zalta, U. Nodelman (Eds.), The Stanford Encyclopedia of Philosophy, Winter 2024 ed., Metaphysics Research Lab, Stanford University, 2024.

[3] D. Makinson, Bridges from classical to nonmonotonic logic, volume 5 of *Texts in computing*, College Publications, 2005.

[4] R. J. Brachman, H. J. Levesque, Chapter 11 - Defaults, in: Knowledge Representation and Reasoning, The Morgan Kaufmann Series in Artificial Intelligence, Morgan Kaufmann, 2004, pp. 205–235. doi:https://doi.org/10.1016/B978-155860932-7/50096-0.

[5] L. E. G. Castañeda, M. Knauff, Defeasible reasoning and belief revision in psychology, in: The Handbook of Rationality, The MIT Press, 2021. URL: https://doi.org/10.7551/mitpress/11252.003.0035. doi:10.7551/mitpress/11252.003.0035. arXiv:https://direct.mit.edu/book/chapter-pdf/2370191/s026100_9780262366175.pdf.

[6] Y. Ru, M. Grüninger, What's the damage? abnormality in solid physical objects, in: L. Jansen, D. P. Radicioni, D. Gromann (Eds.), Proceedings of the Joint Ontology Workshops 2018 Episode IV: The South African Spring co-located with the 10th International Conference on Formal Ontology in Information Systems (FOIS 2018), Cape Town, South Africa, September 17-18, 2018, volume 2205 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2018.

[7] G. Sacco, L. Bozzato, O. Kutz, Know your exceptions, in: C. Trojahn, D. Porello, P. P. F. Barcelos (Eds.), Formal Ontology in Information Systems - Proceedings of the 14th International Conference (FOIS 2024), Enschede, The Netherlands, July 2024, volume 394 of *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2024, pp. 254–267. URL: https://doi.org/10.3233/FAIA241312. doi:10.3233/FAIA241312.

[8] M. Grüninger, C. Chui, Y. Ru, J. Thai, A mereology for connected structures, in: Formal Ontology in Information Systems, IOS Press, 2020, pp. 171–185.

[9] M. Grüninger, A. Delaval, A first order cutting process ontology for sheet metal parts, in: Proceedings of the 2009 conference on Formal Ontologies Meet Industry, 2009, pp. 22–33.

[10] M. Grüninger, S. Bouafoud, Thinking Outside (and Inside) the Box, in: J. Hastings, O. Kutz, M. Bhatt, S. Borgo (Eds.), Proceedings of the First Interdisciplinary Workshop on SHAPES, Karlsruhe, Germany, September 27, 2011, volume 812 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2011.

[11] B. Aameri, M. Grüninger, Location ontologies based on mereotopological pluralism, Appl. Ontology 15 (2020) 135–184.

[12] J. McCarthy, Circumscription - A form of non-monotonic reasoning, Artif. Intell. 13 (1980) 27–39. URL: https://doi.org/10.1016/0004-3702(80)90011-9. doi:10.1016/0004-3702(80)90011-9.

[13] J. McCarthy, Applications of circumscription to formalizing common-sense knowledge, Artif. Intell. 28 (1986) 89–116. URL: https://doi.org/10.1016/0004-3702(86)90032-9. doi:10.1016/0004-3702(86)90032-9.

[14] G. Righetti, D. Porello, N. Troquard, O. Kutz, M. M. Hedblom, P. Galliani, Asymmetric hybrids: Dialogues for computational concept combination, in: F. Neuhaus, B. Brodaric (Eds.), Formal Ontology in Information Systems - Proceedings of the Twelfth International Conference, FOIS 2021, Bozen-Bolzano, Italy, September 11-18, 2021, volume 344 of *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2021, pp. 81–96. URL: https://doi.org/10.3233/FAIA210373. doi:10.3233/FAIA210373.

[15] A. Galton, Yet another taxonomy of part-whole relations, in: L. Jansen, D. P. Radicioni, D. Gromann (Eds.), Proceedings of the Joint Ontology Workshops 2018 Episode IV: The South African Spring

co-located with the 10th International Conference on Formal Ontology in Information Systems (FOIS 2018), Cape Town, South Africa, September 17-18, 2018, volume 2205 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2018. URL: https://ceur-ws.org/Vol-2205/paper4_caos1.pdf.