Generics in Defeasible Reasoning: Exceptionality, Gradability and Content Sensitivity

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

The problem of representing defeasible information is a long-standing topic of discussion in Knowledge Representation: for example, considering logic-based ontology representation languages, in Description Logics many proposals for defining defeasibility and typicality have been formalised, mostly emerging from existing approaches from the non-monotonic logic literature. On the other hand, little attention has been devoted to study the capability of these approaches in capturing the interpretation of typicality and exceptions from a formal ontological and cognitive point of view.

To address this, we here consider the notion of generics as discussed in the linguistic and cognitive literature, i.e. a category of sentences about classes of individuals and admitting exceptions. We then use our analysis as a possible guide for understanding the important features of defeasible information in commonsense reasoning. In this paper we analyze different aspects of generics and we provide some initial desiderata for formalizing defeasible reasoning in DLs.

Keywords

Generics, Defeasible Reasoning, Description Logics

1. Introduction

Representing and reasoning with *defeasible* information is a long-standing topic of discussion in Artificial Intelligence (AI), dating back to the origins of the field of Knowledge Representation (KR): in the presence of stronger conflicting information (or exceptions) with such defeasible information, one wants to retract what we would have inferred in view of new information. In its formalisation in different *non-monotonic logics* [1], this form of reasoning has been considered since the earliest days of KR as one of the parts of the common-sense that artificial systems should have to be considered actually intelligent [2, 3]. The classical example in the non-monotonic logics literature is the *Penguin example* (see, e.g., [1]): if we know that Tweety is a bird and we also know that birds fly, then we are willing to infer that Tweety flies. However, if we come to know that Tweety is in fact a penguin and this makes us retract the previous conclusion: we are more inclined to say that Tweety does not fly instead.

Considering logic-based ontology representation languages, in Description Logics (DLs) many proposals for defining defeasibility and typicality have been formalised: as a matter of fact,

CAOS VII: Cognition and Ontologies, 9th Joint Ontology Workshops (JOWO 2023), co-located with FOIS 2023, 19-20 July, 2023, Sherbrooke, Québec, Canada.
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)

most of them emerge from existing approaches in non-monotonic logics [4, 5]. On the other hand, little attention has been devoted to study the capability of these approaches in capturing the interpretation of typicality and exceptions from the point of view of formal ontology and cognitive aspects. Thus, often a lack of discussion about the philosophical and cognitive assumptions of this kind of reasoning is noted. Namely, when looking at non-monotonic logic only as a tool, evaluating the systems with the sole criterion of evaluation to check the functionality of the particular formal system proposed w.r.t. a particular reasoning problem, we end up with a fragmented set of approaches. This, clearly, increases the difficulty to compare and therefore properly evaluate comparatively such systems from a more general point of view. Moreover, since in the end these tools should be used to model knowledge and reason in real world scenarios, we also need criteria that allow us to decide if the ontological and cognitive assumptions that the formal systems imply are justified or not. For these reasons, we are here interested in discussing these foundational aspects of defeasible reasoning in DLs with the goal of developing a DL system based on an ontologically and cognitively well-justified foundation.

A particularly interesting take on defeasibility, and defining our focus in this paper, comes from the literature on *generics* [6, 7]. Generics are defined as sentences reporting a regularity regarding particular facts that can be generalised even if they admit exceptions. In the penguin example, for instance, the sentence "birds fly" can be considered a generic since it states something about the individuals of the class "bird", but still allowing exceptions like the members of the sub-classes "penguins", "ostriches" and "emus".

Studying defeasible reasoning based on generics appears to be promising for a number of reasons: first of all, according to [7], the central examples that are taken as reference points in the literature on non-monotonic logic can all be considered as involving generics, and therefore we can put forward the idea that other kinds of defeasible reasoning can all be traced back to an account of generics; secondly, discussing generics allows to directly address the issue of exceptions, which is a distinctive feature of both generics and defeasible reasoning [8, 9]; lastly, it provides an excellent starting point to reflect on the ontological, metaphysical, and cognitive assumptions explicitly or implicitly made by the various accounts of defeasible reasoning [10].

Consequently, the goal of this paper is to report on the works on generics with an ontological and semantical lens and to extract some characteristics that are relevant for a comparison of the existing approaches to non-monotonic logic. In particular, we provide the following contributions: we analyse different aspects of generics (Sec. 2–4) and we provide some initial desiderata for formalising defeasible reasoning in DLs (Sec. 5).

2. On Generics

First of all, let us give a more detailed presentation of generics. As said above, generics are general propositions that admit exceptions [6]. Examples of generics found in the literature include

1. *"Lions have manes"*,
2. *"Dogs have four legs"* and
3. *"Mosquitoes carry the West Nile virus"*.

As we can see from these examples, an intuitive way to formalise these sentences, the kind of straightforward reading, suggests to consider them as universally quantified propositions. However, under some further scrutiny, it becomes evident that they are not: in fact, we consider them 'globally true' even if there are individuals in the class under consideration that do not have the predicated property. In the cases above, exceptions could be (1) female or cub lions, (2) unfortunate dogs that encountered some accident and lost a leg and (3) non-infected mosquitoes. Therefore, when considering the sentences as universally quantified statements, such exceptions would render them simply as false.

Some important observations have been made on generics in [6]. A first remark is that the number of exceptions that the three generics above admit is different. In the first case, we can imagine that the number of lions without a mane is probably more than 50% of all the existing lions; in the second, exceptions would be just few individuals; in the third case, the exceptions would be almost all the members of the class. We will discuss this in more detail below, but this observation already suggests that generics can clearly not be treated as simple quantified sentences.

A second important point is the difference between generics and *direct kind predications* [6]. Even if their form in natural language is essentially the same, the latter predicate a property directly to *the kind*, whereas generics predicate a property to *the individuals belonging to the kind*. For instance, sentences like *"Dinosaurs are extinct"* or *"Dogs are widespread"* are direct kind predications. Thus, you can say that Bob the dog has four legs, following Example 2, but it does not make sense to say that Bob the dog is widespread, being a direct kind predication.

These two initial observations show that not all the sentences that in natural language have this general form are generics. In particular, this means that there are other kinds of generalisations. However, as we have seen, the peculiar characteristic of generics is that they tolerate exceptions, that is there are individuals that belong to the class and fail to have the predicated property. Therefore, the question that needs to be tackled is in what semantical sense generalisation can be true and at the same time admit exceptions.

This issue has been addressed mainly in two distinct fields: linguistics and cognitive psychology [6, 8]. In the former, the focus has been on giving a semantics to generics in order to model them formally; in the latter, the attention has been directed towards what generics can say about our cognitive system of generalisation and therefore about concepts and kinds. Of course, the two research interests are related and influence each other, however, it is important to remember that they are different in both their goals and thus in their results. In the following sections, we will briefly present some of the core research from both these fields and discuss if and and in what way they may help with our goal of giving a theoretical and ontological grounding to non-monotonic reasoning in KR languages.

2.1. Generics and Cognition

The main points of the analysis of generics from the perspective of cognitive psychology are referred to by [6] as the *generic-as-default hypothesis* and *sensitivity to content-based factors* [11, 8].

The generic-as-default hypothesis affirms that generic sentences correspond to the default generalising mechanism of the mind, that is the fundamental one, from which the other systems

derive. This is very interesting and arguably counter-intuitive if we consider that from a theoretical and logical point of view the standard generalisation is the universal one based on the standard semantics. Moreover, as we will see, this will be in conflict with some semantic approaches to generics, where generics are indeed reduced in some way to universal quantified proposition.

The sensitivity to "content-based factors hypothesis" supposes that the generalisations made by generics are sensible to elements related to what the generalisations are actually about. [8] individuates these factors as:

> "The counterinstances are negative, and:
>
> If F lies along a characteristic dimension for the Ks, then some Ks are F, unless K is an artifact or social kind, in which case F is the function or purpose of the kind K;
>
> If F is striking, then some Ks are F and the others are disposed to be F;
>
> Otherwise, almost all Ks are F." [8, p. 43]

The first factor is that the members of the kind that do not have the predicated property are exceptions if they only lack the property, that is, if they do not also have an alternative property. This makes these members exceptions, instead of counter-examples that falsify the statements. For example, "Lions have manes" is true because the lions that fail to have a mane lack this property and do not have something else in place of it.

The second factor corresponds to the observation that the acceptability of a generic depends on "characteristic properties" of the kind. Examples of characteristic properties are peculiar physical attributes, as for example manes for lions, or reproductive methods for animals, like in the case of laying eggs for ducks. Similarly, with respect to artefacts or social kinds what is relevant is whether the property corresponds to the function or the purpose of the kind. For example an orange crushing machine should crush oranges, and this makes the proposition "orange crushers crush oranges" true even for completely new orange crushers that never crushed an orange.

The third factor is the poignancy of a property, as for example whether it is a deadly property. This is the case for generics like "sharks attack bathers" and "mosquitoes carry the West Nile virus". However, it is relevant here to note that the individuals that do not have the property in question have the disposition to have it, that is, they could potentially have it.

The last element is the more intuitive one, where most of individuals have the predicated property. Examples are "tigers are striped" and "dogs have four legs". This factor seems to suggest that quantified propositions using *most* or equivalent adverbs of quantification are closely related to generics. However, a precise discussion of this point is beyond the scope of this paper, and we leave it for future research.

This cognitive analysis of generics is not directly connected with a specific formal characterisation of generics, but it gives some insights on what should be taken into account in any such formalisation. In particular, it also guides us in extracting some relevant elements for our analysis of defeasible reasoning.

Regarding the first element, that is, the generic-as-default hypothesis, it is interesting because it would suggest that from a cognitive point of view, defeasible reasoning is more basic and foundational than classical deductive reasoning, a point also defended by [12]. In fact, generics seem to be what actually introduces exceptions and therefore defeasibility in human reasoning: if they are the natural mode of generalisation then also defeasible reasoning is the natural mode of reasoning. This, for the moment, is only a conjecture that needs to be further investigated. However, it seems to accord with the intuitions of the first generation of researchers in defeasible reasoning, especially in the field of Artificial Intelligence. These researchers showed a strong interest in this reasoning because they considered it the type of reasoning that is used in everyday situations. Indeed, with the suggestion of [11], we can consider generics as judgements corresponding to Kahneman and colleagues' System 1.

2.2. Generics and Ontology

The second aspect, namely the sensitivity of generics to content-based factors, gives us something more concrete to discuss with respect to what can be elicited for a discussion of the theoretical desiderata for defeasible reasoning. In fact, as [8, 11] calls them, the factors quoted above are "worldly truth makers". In other words, that list is the description of how the world has to be to make a generic true. This means that they are not semantic conditions, but rather ontological ones, because they explain how reality has to be, not how the language is related to reality.

However, the content-dependency should not be completely overlooked. In fact, an important observation that has been remarked in the literature about generics is that they are not quantifiers, or not purely extensional, in the sense that they tell nothing about how many individuals are instances of the generalisation [11, 7, 9]. Consequently, we hold that in developing a semantic for generics it is needed to refer to the content of the particular generic sentence.

A last interesting observation about these worldly truth conditions for generics, related to the previous one, is that they are essentially conditions on the generalised property, rather than on the individuals of the kind to which the property is predicated. That is, we cannot simply look at the individuals instantiating or not the property, but we need to look at the property itself and see if it is striking, for example. The only exception is the last point, which in fact is stated as a classical quantified statement. In fact, even if the first point refers to "counterinstances", it actually says something about the property, namely if it has *positive alternatives*, as we have seen in the case of the mane of the lions.

Now we can move on to the proposals of how to formalise generics: we first report on the discussion of the syntax used for them and then illustrate and discuss some approaches used for the semantics.

3. The Formalisation of Generics

3.1. The Syntax of Generics

In most presentations, the syntactic way of representing generics is quite uncontroversial and it is inspired by the framework [13] developed for adverbs of quantification. It consists in

a tripartite structure of the form $Q[R][S]$, where Q corresponds to the quantifier, R is the *restrictor*, which sets the domain of Q, and S is the *scope*, which expresses the properties that $Q\ Rs$ have.
Using the example in [6], we can represent the generic "*Typhoons arise in this part of the Pacific*" as

$$Gen\ x\ [Typhoon(x)][\text{Arises-in-this-part-of-the-Pacific}(x)]$$

or as

$$Gen\ s\ [\text{In-this-part-of-the-Pacific}(s)]\exists y[Typhoon(y) \wedge \text{Arises-in}(y, s)].$$

The former corresponds to the interpretation of the generic statements as "*Typhoons in general have a common origin in this part of the Pacific*", whilst the latter corresponds to the alternative interpretation "*There arise typhoons in this part of the Pacific*".
At this point, we need to remark that even if the logical syntax for generic statements is the same as that for sentences with adverbs of quantification, the former cannot be reduced to the latter. Actually, generic statements cannot be reduced to other quantified statements, for a deeper argument see [11]. In fact, the main debated issue about generics is how to interpret them semantically.

3.2. The Semantic Approaches to Generics

According to [6], we can recognise five main approaches in giving the semantics to generics. Here, we will briefly address each of them in order to discuss if and eventually what they can suggest about the theoretical grounding of defeasible reasoning.

3.2.1. Relevant Quantification

One way to try to explain the meaning of generics is by considering a generic as a universal quantification only over *relevant individuals*. This makes Gen in some way context-dependent, since what is relevant should be determined by the context. The main problem of this approach is that as it is, that is without criteria to decide what is a relevant individual, it can justify the truth of *any* generic sentence.
For our purpose, this makes explicit the fact that we need to be able to decide over which elements of the domain we are reasoning on. That is, we need to find a criterion that is able to discriminate between the individuals we are considering and those outside of the scope. Conversely, we can try to see this the other way around: trying to define the characteristics according to which certain individuals are regarded as exceptional, i.e. those individuals which, using the standard quantifiers, should satisfy the generic/inference, but actually do not.

3.2.2. Indexical Approach

Another approach is to consider Gen not only context-dependent, but also indexical. That is, what counts is not only the semantic context in which the generic is found, but also the context of utterance. This means that what counts to determine the truth value of a generic includes the intention of the person uttering it.
This approach has a similar problem to the previous one, because again this explanation seems

to be incomplete since it lacks clear criteria regarding what makes a specific generic true or false.
The intuition of evaluating a generic according to its context of utterance does not seem helpful for our goal. Since we are interested in using the analysis on generics to gain insights to use in formalising defeasible reasoning, the context of utterance is something that could not be easily considered. Rather, if there was strong evidence supporting this approach, it could be an important objection to our purpose. However, since this is not the case, we will simply maintain a careful eye on further developments from this direction.

3.2.3. Probabilities

The main proponent of this approach identified by [6] is Ariel Cohen, who suggests to understand generics in a probabilistic way. In this approach, there are two main types of generics: *absolute generics* and *relative generics*. We can use two examples to show what they are: an instance of the first type is "*Tigers are striped*" and it can be paraphrased as "*A randomly chosen tiger is more likely to be striped than not*"; an instance of the second type is instead "*Mosquitoes carry the West Nile virus*", which can be understood as "*In choosing a mosquito and another insect, it is more likely that the mosquito carries the West Nile virus rather than the other insect*". The difference, as we can see in the examples, is that absolute generics refer to the probability of satisfying the property among the individuals of the kind; whereas in the case of relative generics we are considering also alternative kinds which could satisfy that property. Therefore, in the first case the generic is absolute because we use a threshold that tells us if the probability is high enough to make the generic true. While in the second case it is relative because we compare two probabilities and we look at which of those is the greater one.
There are some counterexamples that strike at the very structure of this approach [6], however, what seems to be the most relevant one for us is that for reasoning purposes it is simply not very helpful. In fact, it seems that it would be very difficult to use this approach to reason on individuals, since it does not tell us when a specific individual satisfies or does not satisfy the generic.
Nonetheless, we found interesting the idea used in the case of relative generics of comparing alternatives to the kind and the property to evaluate the generic. On one hand, it allows to deal with complicated cases like that of mosquitoes carrying the West Nile virus, where the individuals satisfying the generic are a minority; on the other hand, it allows to avoid to decide arbitrarily a threshold to reach for the satisfaction of the generic.

3.2.4. Modal Interpretations and the Notion of Normality

Many researchers, with [7, 9] among them, rely on the possible worlds semantics to interpret generics. The intuition is that generics say something about what normally holds for the individuals of a kind, and so they refer to normal possible worlds to evaluate if a generic is true or not. For example, possible worlds are ordered with respect to normality and a generic statement is considered true if and only if it is true in the most normal worlds for the kind we are considering [7].
According to [6], a shortcoming of these approaches is the fact that they do not seem to be able

to explain why generics like "*Mosquitoes carry the West Nile virus*" or like "*Sharks attack bathers*" are accepted as true. In fact, it does not sound correct to interpret them as "*Normally, mosquitoes carry the West Nile virus*" and "*Normally, sharks attack bathers*". Difficulties arise also for the interpretation of generics like "*Ducks lay eggs*", since the solution proposed of restricting the domain to the relevant sub-kind, in this case female ducks, gives rise to new problems and is experimentally adverse.

A more fundamental issue is that the problem of explaining generics is not really solved, but rather moved to the background. In fact, without a discussion and a definition of what '*normally*' means we are essentially begging the question. Even if the definition of what is normal in each application domain is left to the experts of that domain, if there is not a common definition of normality available, this could easily lead to a fragmentation of the approaches, having a different treatment of generics in each field. This would mean that generics are not actually a unified phenomenon, but rather different phenomena that are explained differently in each specific domain.

However, this approach gives us interesting elements to discuss for our purpose. Firstly, the proponents of this account emphasise the fact that generics have an intensional character, since they cannot be explained by a purely extensional approach. That is, they cannot be simply treated as quantified sentences. According to [7] and [9], this is due to the *normative* connotation of generics. This means that accidental generalisations are not generics. To be a generic there should be a tie between what is meant by the subject and the property that is predicated. This is why, for instance, there are generics that make sense and can be evaluated true or false even if there are no instances of them. Examples taken from [7] include "*This machine crushes oranges*" and "*Kim handles the mail from Antarctica*". Both of them can be considered true even if the machine has never been used for crushing oranges, because it is the machine's function to crush oranges, and because there is no mail from Antarctica, but that geographical area is Kim's responsibility. As we have seen, also [11, 8] agrees that generics cannot be treated as quantified statements because they cannot have a purely extensional representation. Nevertheless, [11] argues that referring to intensionality is still too coarse-grained. Consequently, we have to refer to the actual content of the generic as we discussed above in the context of the content-sensitivity factor.

Secondly, if discussed properly, normality could be an explicit criterion to look at to distinguish what is relevant and what is not for the evaluation of the truth of generics. Moreover, discussing normality would be interesting also because it allows to address the problem from the other way around, that is, to discuss what are exceptions to a generic. A step in this direction has been done, for instance, by [9]. Even if her aim is to account for the similarities and differences of two linguistic variants of generics, the analysis is useful for more general purposes. Therefore, further work on the analysis of this notion could be of great help in the understanding of generics.

Thirdly, this approach is particularly relevant for us, since [7], who argued and discussed the relation between generics and defeasible reasoning, endorse this approach. In fact, a classic approach in non-monotonic logic, known as *preferential* or *selection* semantics [1], is very similar to this way of understanding generics. Where the latter uses an order among possible world, using normality as a criterion, the former uses an order among models, using normality as a criterion too.

A last point that is interesting, but seems to be mostly overlooked, is the interpretation of "normally" when used to paraphrase a generic. In fact, with a generic of the form "$Ks\ \phi s$", can be rephrased as "*Normally*, $Ks\ \phi s$" or as "*Normal* $Ks\ \phi s$". The difference between the two variants is substantial: in the first case, we are saying that what is normal is the state of affairs corresponding to the statement; whereas, in the second case, we mean that we are speaking only of normal instances. However, these interpretations may not be completely independent. For example, one could argue that the general interpretation is the former, but that it corresponds to the latter.

3.2.5. Prototypes

Researchers in prototype theory approaches often endorse a similar position. In fact, even if they do not use the notion of normality, they affirm that generics should be interpreted as statements about typical individuals. Using the example above, they would say that a generic can be rephrased as "*Typical* $Ks\ \phi s$".
There are two main ways reported by [6] to understand the notion of *typicality*: the first is to refer to a stereotype, the second to refer to a prototype. The distinction between the two is not deeply discussed in the text, however it appears that the core difference is that stereotypes emerge from culture, whereas prototypes are related to the cognitive sphere. Nonetheless, [6] report an objection that they apply to both. The concern is that in this approach it seems that for someone to hold a belief is regarded a sufficient condition to make it a true generic. However, whilst we think that this can apply to the stereotype view, since, as we have said, stereotypes emerge from cultural background and therefore they are not necessarily grounded in reality, it does not easily apply to the prototype view. In fact, the latter approach is inspired by a theory in cognitive science, which tries to describe what structure concepts have. Therefore, how the prototype is actually built does not necessarily depend simply on the beliefs of people.
The account that we will propose can be considered as belonging to this general approach. In fact, discussing briefly the prototype theory about concepts coming from the cognitive science side will make evident its link to generics and therefore to defeasible reasoning.

4. The Prototype Theory About Concepts

According to the endorsers of the so-called *prototype theory* about concepts, to fall under a concept does not mean to satisfy a precise definition, but rather to satisfy enough features or constituents of that concept [14]. Another way of describing this account is to compare the individuals and see if they are similar enough to fall under the same concept. Ideally, then, it could be possible to have an individual which is the best example of that concept, that is the *prototype* of that concept. However, it is better to think of the prototype as a more abstract object, that is, as a "concept that was constituted from the different ways in which the category members resembled each other and differed from nonmembers" [15, p. 80].
For instance, we classify Cluedo, tag, football and Pac-Man all as games, because each of them satisfies enough features of the concept game. In fact, developing a definition of game which is able to encompass all four of them is very complicated, to be generous, and even if we were able to make one, finding an instance of game which would be left out would not be so difficult.

[15] individuates four phenomena about categories on which prototype theory relies: *vagueness*, *typicality*, *genericity* and *opacity*.

- *Vagueness* is about the 'boundaries' of the concept, it means that it is not always clear if an individual lies inside or outside the scope of the concept we are considering. For example, it is not clear for the biologist if viruses should be considered living organisms or not.
  In the literature on generics, [9] relies on a similar idea of vagueness to account for their exception tolerance mechanism. In this case, what is vague is which individuals are actually considered as satisfying the generic. Therefore, this first property of the prototype theory about concepts can be seen also in generics with small differences, at least according to [9].

- *Typicality* means that the instances of a concept can be more or less 'good' examples of that concept. For instance, lions are better examples of Carnivora than pandas.
  Here, there is a link with generics, too. In fact, as we have seen, there are accounts of generics that resort to typicality. However, here there is an important distinction that could lead to a development of this kind of approach. In this case, typicality is a graded notion, that is, you can compare different individuals and see which ones are more typical instances of the concepts. The tendency in the literature on generics, instead, is to use typicality in a *Boolean* way, so to say: either you are a typical instance, or you are not. Trying to use the idea of typicality of the prototype theory could improve an interpretation of generics by making it more flexible.

- *Genericity* is the phenomenon which actually corresponds to the formulation of what we have called generics so far, that is, statements that seem to be about the entire class, but that admit exceptions.
  This is the most evident tie between generics and the prototype theory, which strongly supports the attempt to take into consideration also this theory about the structure of concepts for a formal account of generics and, therefore, for a formalisation of defeasible reasoning.

- *Opacity* means that the criteria of categorisation are not clear to the one doing the categorisation. This is directly connected with the problem of formulating a precise definition for the concept, in fact the case of games above is an example of this phenomenon: it is not clear according to which criteria those games are games.
  In a certain sense, we can state that this property can be found in generics, too, and it is witnessed precisely by the difficulty of giving a semantic characterisation of them. Basically, we can consider the struggle in developing approaches that represent and maybe explain generics as a struggle to make this opacity a bit more transparent.

These four main properties show the deep connection of the prototype theory developed in the cognitive research on the structure of concepts to the research on generics. Consequently, we think that they can be usefully related to generics and therefore to defeasible reasoning.

5. The Desiderata for a Non-monotonic Logic Aware of Generics

So far we have discussed generics and the prototype theory of concepts. Now we can sum up what emerged as a useful theoretical foundation and discussion for defeasible reasoning approaches.

1. *Exceptionality*: this is actually more a general presupposition to the purpose. As we have seen, we need a criterion that is able to discriminate and somehow explain when a generic is true and when it is not. That is, which manages to distinguish the relevant individuals from the irrelevant ones. Consequently, it should "recognise" when we have an exception and when we do not. In terms of defeasible reasoning, it means that this criterion is able to tell us when the inference applies and when it does not.
   In the approaches discussed above, two of these criteria were normality and typicality. But in the first case we do not have a proper theory of normality, and so this is not an effective solution. Whereas, in the case of typicality, we can take advantage of the prototype theory of concepts, developed in cognitive science, to have a more precise notion of this criterion.

2. *Gradability*: as we have seen in the case of relative generics in the probabilistic approach and in the definition of typicality in the case of the prototype theory for concepts, it could, and we think it is, useful to apply a relative or comparative approach, rather than an absolute one. This is more in line also with the vagueness that characterises prototype theory, since we do not have a clear boundary that allows us to say this lays inside and this outside, conversely we need to compare the elements and decide if they are similar enough to fall in the same class. This means, in the case of typicality for example, that instead of typical individuals and atypical ones of some concept, we have *more or less* typical individuals.

3. *Content sensitivity*: an important point that emerged at different times during the above discussion is the irreducibility of generics to a purely extensional explanation. Their content sensitive character has been understood mainly in two different ways: in terms of a semantics that resorts to possible worlds and the notion of normality, or alternatively as a semantics that resorts to relations among concepts, as in the case of the prototype approach. This last interpretation is more in line with the property of content-sensitivity formulated in the analysis of generics from the point of view of cognitive psychology.

6. Conclusions

In this paper we provided an overview on generics and discussed the ways in which they can be of use for giving a principled foundation to non-monotonic reasoning methods. We derived three desiderata from this discussion, exceptionality, gradability, and content sensitivity.

A wider discussion on the relationship between generics and the literature on non-monotonic reasoning from the point of view of cognitive science and philosophy is still needed in order to further complete the discussion on the theoretical foundations of defeasible reasoning and

for a complete review of the existing approaches in DLs. For example, [11] observes that inferences involving generics are not “formal", meaning that they depend on the particular generic considered. Moreover, [6] affirm that the main assumption of the non-monotonic reasoning literature seems to be reflected in the experiments: that is, normally, the property predicated by an accepted generic is also attributed to an arbitrary exemplar of the kind involved in the generic; however, it is observed that the literature considers only non-troublesome cases corresponding to high-prevalence generics.

A full technical development of our account, however, is also needed in order to verify the feasibility from a formal and computational point of view and to provide a precise comparison with existing approaches (see [16] for an initial proposal). Finally, it would be interesting to perform an evaluation of the desiderata we identified from a cognitive point of view.

References

[1] C. Strasser, G. A. Antonelli, Non-monotonic Logic, in: E. N. Zalta (Ed.), The Stanford Encyclopedia of Philosophy, Summer 2019 ed., Metaphysics Research Lab, Stanford University, 2019.

[2] J. McCarthy, Programs with common sense, in: Proceedings of the Teddington Conference on the Mechanization of Thought Processes, Her Majesty’s Stationary Office, London, 1959, pp. 75–91.

[3] J. McCarthy, P. J. Hayes, Some Philosophical Problems from the Standpoint of Artificial Intelligence, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1987, p. 26–45.

[4] L. Giordano, V. Gliozzi, A. Lieto, N. Olivetti, G. L. Pozzato, Reasoning about typicality and probabilities in preferential description logics, 2020. URL: https://arxiv.org/abs/2004.09507. doi:10.48550/ARXIV.2004.09507.

[5] K. Britz, J. Heidema, T. Meyer, Modelling object typicality in description logics, in: A. Nicholson, X. Li (Eds.), AI 2009: Advances in Artificial Intelligence, Springer Berlin Heidelberg, Berlin, Heidelberg, 2009, pp. 506–516.

[6] S.-J. Leslie, A. Lerner, Generic Generalizations, in: E. N. Zalta, U. Nodelman (Eds.), The Stanford Encyclopedia of Philosophy, Fall 2022 ed., Metaphysics Research Lab, Stanford University, 2022.

[7] F. J. Pelletier, N. Asher, Generics and defaults*, in: J. van Benthem, A. ter Meulen (Eds.), Handbook of Logic and Language, North-Holland, Amsterdam, 1997, pp. 1125–1177.

[8] S.-J. Leslie, Generics: Cognition and acquisition, Philosophical Review 117 (2008) 1–47.

[9] Y. Greenberg, Exceptions to Generics: Where Vagueness, Context Dependence and Modality Interact, Journal of Semantics 24 (2007) 131–167.

[10] M. Johnston, S.-J. Leslie, Cognitive Psychology and the Metaphysics of Meaning, in: Metaphysics and Cognitive Science, Oxford University Press, 2019.

[11] S.-J. Leslie, Generics and the structure of the mind, Philosophical Perspectives 21 (2007) 375–403.

[12] K. Stenning, M. van Lambalgen, Human Reasoning and Cognitive Science, MIT Press, 2012.

[13] D. Lewis, Adverbs of quantification, Formal semantics-the essential readings 178 (1975) 188.
[14] E. Margolis, S. Laurence, Concepts, in: E. N. Zalta, U. Nodelman (Eds.), The Stanford Encyclopedia of Philosophy, Fall 2022 ed., Metaphysics Research Lab, Stanford University, 2022.
[15] J. A. Hampton, Concepts as prototypes, volume 46 of *Psychology of Learning and Motivation*, Academic Press, 2006, pp. 79–113.
[16] G. Sacco, L. Bozzato, O. Kutz, Introducing weighted prototypes in description logics for defeasible reasoning, in: A. Dovier, A. Formisano (Eds.), Proceedings of the 38th Italian Conference on Computational Logic, Udine, Italy, June 21-23, 2023, volume 3428 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2023. URL: https://ceur-ws.org/Vol-3428/short3.pdf.