Ratio

WILEY

ORIGINAL ARTICLE OPEN ACCESS

Generics as Expectations: Typicality and Diagnosticity

Peter Gärdenfors[1] | Matías Osta-Vélez[2,3]

[1]Department of Philosophy and Cognitive Science, Lund University, Lund, Sweden | [2]Department of Philosophy, Heinrich Heine University Düsseldorf, Düsseldorf, Germany | [3]Institute of Philosophy, FHCE, University of the Republic, Montevideo, Uruguay

Correspondence: Peter Gärdenfors (peter.gardenfors@lucs.lu.se)

Received: 23 February 2024 | **Revised:** 8 June 2024 | **Accepted:** 14 August 2024

Funding: During the period of writing of this paper, the author Matías Osta-Vélez was supported by the Deutsche Forschungsgemeinschaft (DFG), grant SCHU 566/17.1 'Parameterised Frames and Conceptual Spaces' and by the Comisión Sectorial de Investigación Científica (CSIC), University of the Republic, Uruguay, grant 22520220100257UD.

Keywords: conceptual spaces | diagnosticity | generics | reasoning | typicality

ABSTRACT

Generic statements play a crucial role in concept learning, communication and education. Despite many efforts, the semantics of generics remain a controversial issue, as they do not seem to fit our standard theories of meaning. In this article, we attempt to shed light on this problem by focusing on how these sentences function in reasoning. Drawing on a distinction between *property* and *diagnostic* generics, we defend three theses: First, property generics are not about facts but express relations between concepts. Second, generics play an important role in everyday reasoning by interacting with our expectations about the world. Third, diagnostic generics emphasise properties that separate the category in the generics from other categories in the same contrast class. We use the theory of conceptual spaces to advance measures of typicality and diagnosticity capable of modelling different aspects of generics and apply them to the modifier effect and the inverse conjunction fallacy. Finally, we discuss the pragmatics of generics.

1 | Introduction

Natural languages typically have a rich inventory of quantifiers to accurately express the distribution of properties over classes of individuals. Despite this, speakers have a strong predisposition to convey this kind of information using generic sentences like 'ducks lay eggs' or 'Frenchmen like wine'. Remarkably, even if generics are *prima facie* ambiguous regarding the statistical information they convey, they prove highly effective in transmitting conceptual knowledge during both learning and everyday communication (see Krifka et al. 1995; Leslie and Lerner 2022; Gelman 2010).

Within analytic philosophy, one finds several attempts to analyse generics through the study of their truth conditions (e.g., Pelletier and Asher 1997; Cohen 2004; Rooij and Schulz 2020). For instance, it is often claimed that generics are not real universals but implicitly something like 'Typically, birds fly' or 'Usually penguins do not fly'. It has therefore been suggested that the 'deep structure' of generic statements combines two predicates with a hidden operator (named '*Gen*') that refers to some adverb of quantification like 'typically', 'generally' or 'usually'. If *Gen* is what lays behind the surface structure of generics, it is assumed that they have a logical form, like $Gen(x)(Bird(x) \rightarrow Fly(x))$, and that they are analysable in terms of truth-functional structure (see Chierchia 1995; Pelletier and Asher 1997).[1]

This approach, however, faces multiple technical problems as well as limitations in explaining the role of generics in everyday communication and reasoning (see Leslie 2008; Leslie and Lerner 2022). To mention one, let us assume that a generic is true if each individual of the kind would have the property if things were as normal as possible. Then, a generic like 'ducks lay eggs' would come out as false since, in the most common circumstances, it is very easy to find an individual that does not meet the property (most 'normal' duck specimens do not in fact lay eggs). Pelletier and Asher (1997) tried to amend this problem by arguing that some generics quantify over some subset of a category as long as that subset denotes some 'genuine' subkind of the category. In our case, 'ducks lay eggs' would actually quantify over the subkind 'female ducks' and would therefore be true. However, as Leslie (2008, 8) pointed out, this strategy would turn obviously false generics into true ones: What prevents 'ducks are female' from being true if the generic can refer to the genuine sub-kind 'female ducks'?

In this paper, we defend the view that generics are more related to the *structure* of concepts than to their extensions (cf. Cimpian, Gelman, and Brandone 2010; Hampton 2012; Leslie 2012). More specifically, generic sentences express information about relations between properties and categories. We shall argue that they cannot be evaluated as isolated sentences but only in association with clusters of sentences encoding conceptual knowledge. Instead of looking for truth conditions, we focus on how generics are used, that is, their role in reasoning, learning and communication. Our analysis will highlight their cognitive and pragmatic aspects, particularly their interaction with *expectations* in reasoning (see also Leslie 2008; Prasada et al. 2013).

There has been an extensive discussion concerning how many types of generics should be distinguished. For example, Leslie, Khemlani, and Glucksberg (2011) and Prasada et al. (2013) suggest at least five types.[2] We propose a new classification that distinguishes *property generics* dealing with characteristic properties of objects and *diagnostic generics* dealing with properties that are diagnostic for categories.

In the literature, particular attention has been devoted to a type of generics that highlight properties deemed to be 'striking', such as 'mosquitoes carry the West Nile virus' and 'sharks attack bathers' (Leslie 2008; Prasada et al. 2013; Sterken 2015; Rooij and Schulz 2020). We will show that striking generics can be analysed in terms of diagnosticity and affective valence. We argue that such generics do not constitute an additional subtype of generic sentence.

The article continues as follows. Section 2 looks into property generics and defends the thesis that they express relations between concepts. We propose a Gricean strategy to explain that a generic sentence is evaluated in the light of the background knowledge that agents have and the context of use. In Section 3, we use the theory of conceptual spaces for explaining how generics reflect structural features of the organisation of conceptual knowledge and how they can be understood as playing a crucial role in expectation-based reasoning. Section 4 shows that our model can explain the modifier effect and the inverse conjunction fallacy (Connolly et al. 2007; Jönsson and Hampton 2006), which involve compositions of concepts.

In Section 5, we turn to diagnostic generics, that is, generics that convey diagnostic properties of categories. We argue that they are indeed of a different nature than property generics and we propose a formal measure of diagnosticity. Section 6 discusses the pragmatics of generics, focusing on their role in teaching.

2 | Generics as Expectations

2.1 | Background Knowledge and Expectations

We view the meaning of generics as determined by their use in reasoning and communication, aligning with Greenberg and Harman's functional perspective (2006). Thus, a generic's meaning cannot be determined by taking it in isolation, but arises from its interplay with the user's background knowledge. This idea is our primary thesis.

One way in which background knowledge is manifested is through expectations about the world, that is, clusters of propositions that arise in response to new information. For instance, if someone exclaims, 'There is a squirrel!' I expect to see an animal of a specific shape, size and colour. If the squirrel turns out to be white, I would be surprised, as this fact contradicts my expectations. Expectations play a pivotal role in reasoning, serving as implicit premises used for making inferences under uncertainty (see Gärdenfors and Makinson 1994).

We aim to demonstrate that generics can be viewed as expectations utilised in reasoning and that they play a significant role in communicating our conceptual commitments. Instead of discussing generics as statements that can be either true or false, we will discuss them as having different 'degrees of strength'. Following Osta-Vélez and Gärdenfors (2022), this will be done in terms of the *typicality* and *diagnosticity* of concepts.

2.2 | The Internal Structure of Generics

In their paradigmatic form (bare plurals), generics relate a property to a category or kind. A central point of our analysis is that the property involved, and our knowledge about it, determines (for the most part) the content of a generic (cf. Prasada et al. 2013). Our approach requires a theoretical framework that allows us to analyse the relations between categories and properties.

Such a framework is presented in Gärdenfors and Stephens' (2018) analysis of inductive reasoning. They propose that, in addition to the traditional epistemological distinction between *knowing that* and *knowing how* (Ryle 1949), *knowing what* should also be considered. *Knowing that* pertains to knowledge concerning relations between agents and propositions, while *knowing how* relates to an agent's abilities, dispositions and actions. In contrast, *knowing what* involves the ability to categorise and understand the relation between categories and properties. Gärdenfors and Stephens (2018) argue that induction concerns *knowing what* rather than *knowing that* and we submit that the same applies to generics.

The typical form of a generic is 'Category C has property P'; and its evaluation depends crucially on our knowledge of the interaction between C and P.[3] Surprisingly, few theories of generics focus on this. Those that do engage with this issue (e.g., Declerck 1991; Heyer 1990) do not build upon a substantive theory of background knowledge or category structure.

We claim that property generics (as sentences expressing *knowledge what*) are evaluated by considering different compatibility relations between the semantic domain of the predicated property and the category in the sentence. In brief, generics concern more or less prototypical or diagnostic properties of concepts. In what follows, we will explain how these ideas can be made more precise by using conceptual spaces as a model of the structure of concepts.

3 | Conceptual Spaces and Prototypical Structure

3.1 | An Outline of the Theory of Conceptual Spaces

The theory of *conceptual spaces* (Gärdenfors 2000, 2014) builds on two central ideas about the structure of concepts and properties: (i) they are composed of clusters of *quality dimensions*, for example *height*, *temperature*, or *pitch*; (ii) they have a geometric structure that is the result of the integration of the specific geometrical structures of the dimensions. The dimensions are often equipped with a distance function (a metric).

The notion of a *domain* is defined as a set of dimensions that are separable from all other dimensions. For instance, colour properties are composed of three fundamental dimensions of colour perception: *hue*, *saturation* and *brightness*. A central thesis in the theory of conceptual spaces is that natural properties (like colours) correspond to convex regions of a single domain. This is called 'Criterion P' in Gärdenfors (2000). A region is convex when for every pair of points x and y in the region, all points *between* them are also in the region.

We now define a *conceptual space* as a collection of one or more domains with a distance function that represents similarity relations between objects. The closer two points are in the space, the more similar the objects they represent are. Within this framework, objects are seen as instances of concepts and are mapped into points of the space, and concepts (categories and properties) are represented as regions of the space.

While adjectives typically refer to single domains, the meaning of most nouns in natural language is more complex, defined by clusters of properties across several domains. These domains often exhibit correlations. For instance, in the concept of FRUIT, properties in domains such as *size*, *weight*, *ripeness*, *colour*, *texture* and *taste* tend to co-vary. These co-variations generate expectations that are crucial for inferential processes that exploit semantic properties.

To illustrate this, consider the concept APPLE as a subregion within the broader conceptual space of FRUIT. This 'Apple space' can be represented using dimensions such as *colour*, *taste*, *ripeness*, *texture*, *size* and *shape*.[4] The 'fruit space' itself is a subset of the Cartesian product of these dimensions, with the 'apple space' occupying specific subregions corresponding to the possible properties of apples, as depicted in Figure 1.

An advantage of representing concepts in this way is that it allows us to account for the *prototypical* structure of categories in a natural way (Rosch 1975; Gärdenfors 2000). Defined as convex regions within n-dimensional spaces, the central point in each region can be interpreted as the prototype for the property or concept.[5] In other words, a prototype can be seen as a list of properties that are most typical for a concept. This allows to represent graded membership and degrees of *typicality* (Hampton 2007).

Even if several domains are involved in representing a concept, not all of them have the same *salience*. For example, the shape of an apple is more salient than its colour. To model this, we assume that the representation of a concept also contains information about the weights of domains (see Gärdenfors 2000, Section 4.2). These weights are used to determine the distance function in the space and the degree of similarity between different objects. Since a property is assigned a particular domain, the weight of the domain can now be used as a measure of how characteristic a property for a particular concept is. This construction reflects that there is an intimate connection between similarity judgments and how characteristic properties are.

In the literature, references to 'essential properties' of categories are often made (e.g., Gelman 2003; Haslanger 2014; Noyes and Keil 2019). Even though one need not subscribe to Aristotelian essentialism, it is natural for humans to judge certain properties of categories as more essential than others (Gärdenfors 2000, section 4.2.2). In psychological theories (e.g., Osherson and Smith 1981; Rips 1995), 'core' properties are supposed to determine the meaning of concepts while other properties are considered 'peripheral' even though some may have diagnostic value and thus be useful for some cognitive tasks.[6] The core is considered to be those characteristic (sometimes called essential) properties of concepts, while peripheral properties may or may not be present in the instances.[7]

3.2 | Generics Express Expectations

The way generics interact with prior knowledge can vary and, based on the available information, thereby have different

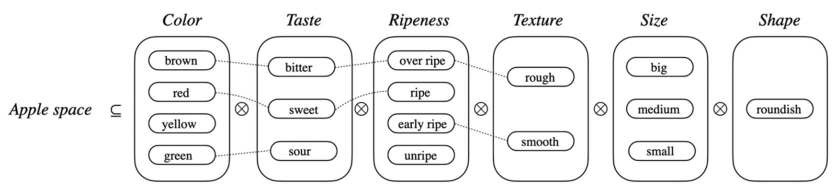

FIGURE 1 | 'Apple space' as a subregion of a 'Fruit space'. The dotted lines represent some of the correlations among properties of different domains.

influence on how they are used in reasoning. For example, 'birds fly' is generally accepted as a robust generic. When combined with the defining generic 'penguins are birds', 'birds fly' entails 'penguins fly' and will come into conflict with the generic 'penguins do not fly', which would typically be part of background knowledge. However, 'penguins fly' has less expectation strength than 'penguins do not fly' since a flying penguin is further from the penguin prototype than a non-flying one. Thus, given the background information presented here, the generic 'birds fly' would be eliminated in reasoning about penguins to avoid inconsistencies (see Gärdenfors and Makinson 1994; Osta-Vélez and Gärdenfors 2020, 2022; Gärdenfors and Osta-Vélez 2023).

Note that this is a general principle when reasoning with dynamic information. If someone starts with the sentence '*x* is *C*', and later receives the new information '*x* is *P*', compatible with *C*, her uncertainty regarding the expected properties of *x* will be reduced, and she will be able to make more precise inferences with her background knowledge. A paradigmatic case of this is when *P* is a subordinate of *C* (like in PENGUIN and BIRD). In the conceptual spaces framework, reducing uncertainty about an object *x* implies to reduce the *volume* of the subregion of the space in which *x* is located.[8] Such a reduction, however, may make new properties prototypical (*not flying* is prototypical for PENGUIN, but not for BIRD). This analysis shows that reasoning with generics has non-monotonic properties since adding more information to the given background knowledge can lead to some generics no longer being accepted.[9]

The most important consequence of this analysis is that it does not require any additional (hidden) logical operators like *Gen*. In support of this, it can be noted that there is no known language that has a dedicated, articulated generic operator.[10] Instead, we assume that the acceptability of generics can be determined by the strength of their expectations within the appropriate conceptual space.

Our analysis thereby explains why it is easier for children to learn to use generics than to use sentences involving quantifiers. There is evidence that generics are the default form of generalisation used by children when reasoning about kinds (Gelman 2010; Gelman et al. 2015). According to our interpretation, these generalisations rely solely on the organisation of conceptual knowledge, present from early development stages (Mattos and Hinzen 2015). In contrast, mastering quantifiers seems to depend not on conceptual knowledge but on extensional reasoning forms, which are less intuitive for children (Gelman 2010). Thus, interpreting generics as 'knowledge-what' offers a more natural explanation for their prevalence in infant cognition.

4 | Representing Property Generics in Conceptual Spaces

4.1 | Property Generics and Typicality

Our analysis of how property generics express expectations focuses on the agent's representation of categories as associated with clusters of properties. The main idea is that when one categorises an object *x* as *C*, expectations about properties that *x* is supposed to have can be used to generate generics. A crucial point is that these expectations will respect an ordering that follows from the *prototypical* structure of the concept. For instance, if I am told that *x* is an apple, I will form expectations about *x* having the prototypical properties of apples, like being red, sweet, round and smooth. To a lesser extent, I will expect that *x* is green, sour, or rough. As we will see later, the fact that our expectations have different strengths (that is, they show different degrees of defeasibility) implies that generics are also graded in this sense.

The underlying rationale for this procedure of generating an expectation ordering is a version of the Gricean principle of *maximal informativeness* (Grice 1975). If you are informed that an object *x* should be categorised as, for example, a bird, but you do not know more about what kind of bird *x* is, then you expect that relevant generics can be applied to *x*: 'Birds fly', 'birds sing', 'birds have wings', 'birds have beaks', 'birds build nests', etc. The principle of maximal informativeness says that your informant should have communicated something more specific if these generics do not apply to *x*. However, when new information about *x* is added, for example, that *x* is a penguin, some characteristic properties of birds will no longer be applicable. For example, the generic 'birds fly' will no longer apply, but it will be replaced by 'penguins do not fly'. Further generics may be applicable, for example, 'penguins eat fish'.

It is well-known that generics and factual statements behave in different ways linguistically. One test is proposed by Lawler (1973):

1. Blue whales eat plankton.
2. A blue whale eats plankton.
3. Blue whales can be seen around the Cape of Good Hope.
4. A blue whale can be seen around the Cape of Good Hope.

(1) describes a characteristic property of BLUE WHALE. It can be exchanged for the indefinite singular version in (2). The generic expresses a relation between the concept of BLUE WHALE and the property of *feeding on plankton*. In contrast, (3) is a factual statement. A test for this is that it cannot be exchanged for the indefinite singular version in (4) (see Carlson 2009; Krifka 2013). As we will show in Section 5, this test only applies to property generics but not to diagnostic generics.

A second test for the difference between generics and factual sentences is that generics are not 'upward entailing'. This means that in (1), BLUE WHALES cannot be replaced by the superordinate MAMMAL since 'Mammals eat plankton' is not acceptable. On the other hand, factual sentences are upward entailing: (3) entails 'Mammals can be seen around the Cape of Good Hope'.

The fact that sentences (1) and (2) express the same content despite their different logical forms further indicates that generics form a special class of sentences. The upshot is that although a generic is a sentence, it expresses a different kind of knowledge than factual sentences.

The pragmatic principle above, plus the prototypical organisation of the expectation order, tells us which set of generics fits better with some piece of information. Now, within the set of generics, some are more expected than others. In other words, it is possible to establish an internal ordering of prototypical properties that will translate into an ordering of generics according to how expected they are. Surprisingly, this aspect of generics does not seem to have been investigated experimentally. For example, compare the following two generics:

5. Elephants have trunks.
6. Elephants are grey.

The property *having a trunk* is more characteristic of elephants than *being grey*. One can quite easily accommodate the occurrence of a white or a black elephant and perhaps also a pink one. However, an elephant without a trunk is an injured elephant and is much less expected than a non-grey elephant. In terms of prototypes, a non-grey elephant is more similar to the prototypical elephant than an elephant without a trunk.

4.2 | A Measure of Typicality

The main prediction of our approach is that the degree of expectation of a property generic will be a positive function of the *typicality* of the property within a conceptual space. The general rule is that the more characteristic a property is for a particular concept, the more expected is the generic. This rule explains why the proportion of instances is not decisive for how useful a generic is in arguments. For example, 'books are paperbacks' describes a highly prevalent feature of books but does not concern a characteristic property. From our perspective, a way to determine such an expectation ordering is to consider the relative typicality of each property in its conceptual space. For instance, *round* and *smooth* are more characteristic properties of APPLES within a fruit space than *red* since instances of non-red apples are more typical than non-round or rough apples. This means that the former two properties will have priority over the latter in an expectation ordering (see Figure 2).

The key idea is to interpret a property generic '*M*s are *R*' as '*M*s that are similar to the prototype have the property *R*'. The strength of such a generic can be determined by measuring the distance to the closest point where the property *R* is *not* satisfied. We denote the complement or the property *R* as $\neg R$. We can use the distance function from the conceptual space to obtain this kind of information. For example, the distance from an apple that is not flattened to a prototypical apple is larger than the distance from an apple that is not red to a prototypical apple (green apples are not far from the prototype).

To develop this idea more formally, let us assume that a concept *M* is represented by a region in a conceptual space $C(M)$ with a prototypical point p^M. For the definition of typicality, it turns out that one must distinguish between *prototypical* properties *R* for the concept *M*,[11] that is, properties for which it holds that $R(p^M)$ and *non-prototypical* properties (for which it does not hold). For a non-prototypical property, it is the closest point to the prototype that determines its typicality. We therefore define the typicality of a property *R* in $C(M)$—written $T_M(R)$—as a measure of its expectedness level as follows:

Typicality measure:

i. For any prototypical property R_i, in $C(M)$, $T_M(R_i) = \min_{x \in \neg Ri(x)} d(x, p^M)$.
ii. For any non-prototypical property R_k in a conceptual space $C(M)$, $T_M(R_k) = -\min_{x \in R_k(x)} d(x, p^M)$.

It is important to note that in our typicality measure, we do not count the numbers of instances, but the criterion is based

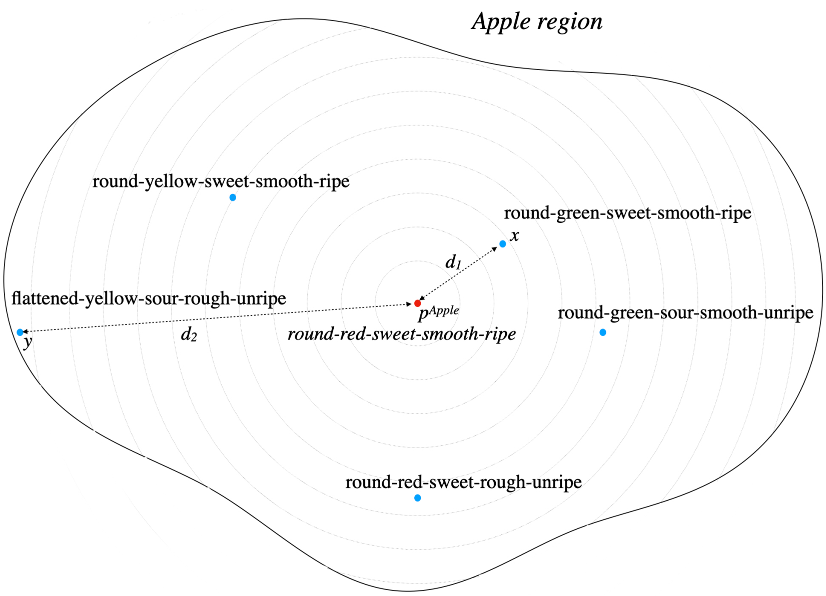

FIGURE 2 | A representation of property distances in the apple region of a fruit space. The central point represents the apple-prototype with the most expected properties in each domain. *Green* is more typical than *non-round* (and thus, more expected) because the distance between the prototype and the closest instance not satisfying the property *round* (leftmost point) is bigger than the distance from the prototype to the closest instance with the property *green* ($d_1 < d_2$).

on similarity to the prototype. In other words, our model is not probabilistic. Furthermore, since non-prototypical properties are always assigned negative values, they will be less typical than all the prototypical properties.

To be sure, there exist recent probabilistic (Bayesian) models of generics, for example, Cohen (2001), Tessler and Goodman (2019) and Rooij and Schulz (2020). Tessler and Goodman (2019) are to some extent similar to our approach since they also consider the role of background knowledge in evaluating generics. However, they analyse such sentences in terms of prevalence (probability) rather than in terms of typicality. A common problem with probabilistic models (with the exception of Goodman, Tenenbaum, and Gerstenberg (2015)) is that they do not account for the nature of concepts or their role in reasoning, while the framework presented here builds on a theory of concepts. Furthermore, probabilistic models sometimes give wrong results since some property that is improbable may be judged to be typical. For example, the prototypical turtle is an adult, so 'Turtles become 30 years old' is an acceptable generic, since turtles who reach adulthood become old, while people in general do not take into account that it is very probable that a turtle dies before adulthood.

4.3 | The Modifier Effect and the Inverse Conjunction Fallacy

We next present some further applications of our model. We show that the modifier effect and the inverse conjunction fallacy, both dealing with compositions of concepts, can be explained with the aid of our typicality measure.

The 'modifier effect', as identified by Connolly et al. (2007), is a cognitive bias that demonstrates a deviation from logical set-inclusion principles in the evaluation of generic statements. Specifically, individuals tend to perceive a generic statement about a concept (e.g., 'Ravens are black') as more likely to be acceptable compared to a similar statement where the concept is specified with a modifier (e.g., 'Jungle ravens are black').

Jönsson and Hampton (Jönsson and Hampton 2012, see also Hampton, Passanisi, and Jönsson 2011) replicated the effect and also showed that the strength of the effect is dependent on the typicality of the modifier: the less typical the modifier, the stronger the effect. For example, in comparison to the generic 'Lambs are white', the modified generic 'Norwegian lambs are white' was, on average, judged to be less acceptable than the modified generic 'Fluffy lambs are white'. The reason is that Norwegian lambs are less typical examples of lambs than fluffy lambs.

This result is predicted by our model in the cases when the prototype of the non-typical concept (NORWEGIAN LAMBS) is closer to the area of the conceptual space covering the property *non-white* than is the prototype of the more typical property (*fluffy lambs*). The property of the latter is presumably close to the prototype of LAMB.

Jönsson and Hampton (2012, 101) show that also for typical modifiers such as 'fluffy', the modified generic 'Fluffy lambs are white' on average is judged to be less acceptable than the unmodified 'Lambs are white' (although the effect is not as strong as for non-typical categories). Our model cannot directly explain this effect since the prototype for FLUFFY LAMBS presumably is very close to that of LAMBS. However, as noted by Jönsson and Hampton, pragmatic factors such as Gricean maxims may explain this phenomenon.[12]

A related phenomenon, called the inverse conjunction fallacy, has been studied by Jönsson and Hampton (2006) in a series of experiments. Again, a generic (like 'All candles are made of wax') is compared to a modified generic containing a conjunction of properties ('All expensive purple candles are made of wax'). The experiments show that, on average, participants judged the original generic to be acceptable more often than they judged the modified generic to be acceptable. Again, our model predicts this for the case where the prototype of the modified concept (EXPENSIVE PURPLE CANDLES) is closer to the area of the conceptual space covering the property *not made of wax* than is the prototype of the unmodified concept (CANDLE).[13]

5 | Diagnostic Generics

The typicality measure provides a solid platform for understanding the acceptability conditions of a property generic. However, generics can also be evaluated from an informational perspective. To do so, we must shift our attention from the internal to the external structure of categories, that is, to superordinates and contrast classes.

5.1 | Hierarchies of Concepts

Categories can be understood as organised in hierarchical structures with a 'horizontal' and a 'vertical' dimension (Rosch 1988). The horizontal dimension concerns contrast relations between categories at the same abstraction level. For example, DOG, CAT, ELEPHANT, HORSE are in a contrasting relation since any object falling under one of them is automatically excluded from the others. This type of relation occurs when categories are included in the same partitioning of some other concept of a higher level of abstraction, called superordinate (in our example, MAMMAL). The subordinate-superordinate relation is encoded in the 'vertical' dimension. We will refer to the contrast class of concept X as the set $CC(X)$ containing the concepts that are different from X but share its immediate superordinate concept with it. For example, the contrast class of LION is the set of concepts that refer to other big cats that are not lions, that is, CC(LION) = {TIGER, CHEETAH, JAGUAR, LEOPARD, COUGAR, etc.}.

A significant part of our expectations about a concept is inherited from knowledge of its superordinate. For example, one may have minimal knowledge of what a lynx is, but the mere fact of knowing that it is a feline will generate expectations (of varying strength) that are derived from the defining properties and characteristics of such a superordinate concept. We submit that inherited expectations play an important role when evaluating generics with minority characteristic properties.

Consider the following generic sentences:

7. Birds lay eggs.
8. Bees are sterile.
9. Lions have manes.

From an extensional point of view, it seems puzzling that (7) is true while the property only applies to a minority of the individuals of the class (adult females). On the other hand, (8) is false but the specified property applies to a majority of the individuals of the class. The confusion disappears when we consider these generics in interaction with inherited expectations.

Since *egg-laying* is prototypical for BIRD, the typicality criterion offers a straightforward explanation of why (7) is accepted as true. However, the informational relevance of (7) depends on its interaction with the set of inherited expectations {'Birds are animals', 'animals reproduce'}, which consequently entails 'birds reproduce'. By specifying a mode of reproduction for BIRD, (7) adds relevant information to *x* and thus reduces the degree of uncertainty in the (inferentially available generic) 'birds reproduce'.

Similarly, we can explain why (8) is false by appealing to the fact that the property in it is not prototypical of BEE. However, this explanation can be complemented from an informational point of view if we note that (8) can be interpreted as implying 'bees cannot reproduce', a generic that conflicts with the set of inherited expectations {'Bees are animals', 'animals reproduce'}, which have as a consequence that bees reproduce.[14]

5.2 | A Measure of Diagnosticity

Many generics turn out to express properties that are diagnostic for a concept. A property is diagnostic when it maximises the dissimilarity between a concept and the concepts in its contrast class, speeding up categorisation (Tversky 1977).[15] For example, in (9) the property of having a mane is diagnostic for lions since having manes is what distinguishes them most (at least their adult males) from other concepts of felines. In other words, *having a mane* is diagnostic for LION in *CC*(LION) = {TIGER, CHEETAH, JAGUAR, LEOPARD, CAT, etc.}.

In our model, we address the influence of the contrast class by considering the prototype of the immediate superordinate concept. In the case of (9), explaining why it is an informationally relevant generic involves considering its relationship to the contrast class *CC*(LION). More specifically, the informational value of (9) is related to the fact that the property 'has a mane' is not present in any other of the categories in *CC*(LION) and therefore is highly diagnostic for LION. Using our typicality measure, we can define a measure of diagnosticity as follows:

Given a category M, its immediate superordinate concept N, and a property R_i, the diagnosticity of R_i is equal to $\frac{T_M(R_i)}{T_N(R_i)}$.

Recall that $T_M(R_i)$ expresses the typicality of property R_i in relation to the concept M.

We predict that generics with diagnostic properties are easier to endorse than generics with non-diagnostic ones because more information is added to agents' conceptual systems. For instance, (9) should be seen as more interesting than 'lions have whiskers' because *having manes* contains information that is specific to the concept and helps to differentiate it from other categories, while *having whiskers* is common to all members in *CC*(LION) and can be inferred from characteristic knowledge of the superordinate FELINE.

Diagnosticity is important since diagnostic properties are those that carry the most information about a concept when compared to similar categories. So, for example, when children learn about a new concept, being told about such properties will speed up their learning (Gelman 2003).

Now, diagnosticity comes in degrees. Some properties can have diagnostic value for a concept and still be non-prototypical or even rare. Consider the generic:

10. Frenchmen eat horse meat.

Horse meat is atypical in the diet of most French people (although it is common in a few regions of the country); why then is (10) considered an acceptable generic? The answer is that the diagnostic value of the property in (10) in relation to *CC*(FRENCHMEN) compensates for its low degree of typicality. In other words, the property *eat horse meat* will be more expected for FRENCHMEN than for any other concept in *CC*(FRENCHMEN).

The above measure assures that for any two similarly distinctive properties of a concept (with respect to its contrast class), the one with the highest typicality degree will be the one with more diagnostic value. This would allow us to capture the intuition that generics with prototypical properties with diagnostic value are stronger than generics with atypical properties that also have diagnostic value: For instance, 'Frenchmen speak French' is stronger than 'Frenchmen eat horse meat'.

5.3 | Generics With Striking Properties

In the literature (e.g., Leslie 2008; Prasada et al. 2013), so-called 'striking generics' have been brought forward as a special type of generics. Two examples are the following:

11. Ticks carry Lyme disease.
12. Sharks kill people.

Such generics are problematic for any quantificational approach since only a small minority of ticks carry Lyme disease and shark attacks are very rare.

On our analysis, striking generics are special cases of diagnostic generics. (11) expresses that even though very few ticks carry the Lyme disease, this is more frequent in ticks than in any other concept in CC(TICKS) (the superordinate may perhaps be BUGS—ticks are not insects but arachnids). Ticks may, in fact, be the only animals that carry the disease. Similarly, even though sharks' attacks are rare, encounters with sharks are more dangerous than encounters with other animals.

The linguistic form of a striking generic is generally a combination of a noun phrase and a verb phrase that is not based on 'is' or 'has'. Indeed, all of the striking generics in the tests of Prasada et al. (2013) contain such verbs (see their table A, 420).

Although they are ultimately diagnostic generics, striking generics could form a special subclass of generics if we take into account their affective dimension. Since striking generics refer to properties or dispositions that are 'dangerous and to be avoided' (Prasada et al. 2013, 409), they must bear negative valence, that is, they elicit negative emotions, and as a consequence draw more attention. This phenomenon aligns with the 'automatic vigilance hypothesis', which posits that negative stimuli command more attention than neutral or positive stimuli (Pratto and John 1991; Öhman, Flykt, and Esteves 2001). This automatic attentional bias towards negative information makes such statements more memorable and, consequently, more likely to be acceptable, despite their rareness.

5.4 | Comparing Property and Diagnostic Generics

Our analysis of property generics and diagnostic generics suggests that there are two factors at work in the evaluation of a generic.[16] The first has to do with the degree of typicality of the property in the generic, and the second with its informational contribution, that is, how diagnostic the property expressed in the generic is. To illustrate the difference between them, compare (9) with the generic 'Lions have bones'. Although very strong, the latter generic may seem obvious because the information in it is inferentially available from the fact that lions are mammals. That is, minimal knowledge about the concept LION allows you to infer that generic. On the other hand, (9) contains information that cannot be inferred from the superordinate of LION and which contributes to the specification of the concept in relation to its contrast class.

This informational factor in the evaluation of generics can be explained in terms of Wilson and Sperber's notion of 'relevance':

> When is an input relevant? Intuitively, an input (a sight, a sound, an utterance, a memory) is relevant to an individual when it connects with background information he has available to yield conclusions that matter to him [...]. In relevance-theoretic terms, an input is relevant to an individual when its processing in a context of available assumptions yields a positive cognitive effect. (Wilson and Sperber 2004, 608)

Generics that are coherent with our expectations and provide information that reduce their uncertainty have a positive cognitive effect and therefore are relevant in communication. Generics that provide concept-specific information that help to separate a concept from the categories in the contrast class are also relevant in a similar sense. On the other hand, generics that potentially contradict our expectations or induce false conclusions are irrelevant and therefore often rejected.

Like property generics, diagnostic generics are not upward entailing: 'Lions have manes' does not entail 'Mammals have manes' and 'Turtles have shells' does not entail 'Reptiles have shells'. Hence, this property is a test for determining whether a sentence is a generic or not, also for causal generics. In contrast, factual sentences such as 'Lions live in Africa' entail 'Mammals live in Africa', so such sentences are upward entailing.

In contrast to property generics, however, diagnostic generics do not pass Lawler's (1973) test: The generic 'Frenchmen eat horse meat' does not express the same as 'A Frenchman eats horse meat' and similarly 'Lions have manes' is different from 'A lion has a mane'. This observation supports that the two types of generics are indeed different. The differences between property generics, diagnostic generics and factual sentences are summarised in Table 1.

6 | The Pragmatics of Generics

So far, we have only considered generics from the perspective of an agent that evaluates its acceptability. In brief, our answer is that the acceptability of a generic is determined from prototypical properties and distances in a conceptual space. In this section, we turn to a different type of agents, that is, somebody who hears or reads a generic. This perspective has not received much attention in the literature, since the focus has been on the truth conditions for generics. We, however, think the perspective is important since it concerns the role of generics in human communication and how this role differs from that of factual statements.

The question is how a new generic will affect the cognitive state of the agent, that is, how it will change the agent's state of knowledge. We thereby turn to further pragmatic aspects of generics, focusing on the question: What is the *use* of generics in social interactions? (Gelman 2021)

Several studies argue that young children understand generics and can distinguish them from non-generics on the basis of

TABLE 1 | Types of generics and their differences.

	Property generics	Diagnostic generics	Factual sentences
Passes Lawler's test	Yes	No	No
Upward entailing	No	No	Yes

several types of cues. In particular, Gelman (2010) shows that already by about the age of two and a half, children start producing generics, and at the age of four, generics constitute 3% of the sentences produced. This may not seem to be a high rate, but it should be considered in the context of all other sorts of utterances that children can make. They rely on generics when they learn about categories and they use them when drawing inferences. Children also interpret generics as different from quantificational sentences containing 'all', 'some' or 'most'.

We speculate that children at an early age learn to reason with the expectations that are generated by the generics they hear, but we know of no empirical investigations related to this position. Novel nouns provided by care-givers are interpreted as inference-rich and thus denoting coherent concepts. This assumption helps children figure out which categories are kinds (Gelman 2003, 184; Cimpian and Scott 2012). Children expect that the first count nouns they learn will refer to basic level concepts.

The nouns then appear in generics (both property generics and diagnostic generics) together with an adjective that denotes a typical or diagnostic property of the noun. Gelman (2003, 174) writes that generics have two main functions for children: Firstly, generics teach children particular concept-wide generalisations. Secondly, they indicate that instances of a concept are alike in important ways.

Generics also have a central role in *teaching*, in particular in what is called 'natural pedagogy', that is, teaching by parents and others in everyday circumstances (Csibra and Gergely 2009). We tell our children, already when they are small, things like: 'cats say meow, dogs say woof, and cows say moo'.[17] Later in school, they learn generics like 'tigers have stripes', 'copper conducts electricity' and 'democracies have freedom of speech'. Such property generics are a way of presenting characteristic properties of various categories (Leslie 2008). Learning about categories is primarily done via their characteristic properties.[18] And when it comes to diagnostic generics containing striking properties such as 'dogs bite people', they function as guidelines for caution in actions (Sterken 2015). In brief, generics play an important role in learning about the world from a young age.

7 | Conclusion

This article proposes an analysis of generics that privileges their cognitive and pragmatic dimensions over their logical structure. We have argued that property generics should be distinguished from diagnostic generics. We have put forward three theses building on this distinction. The first thesis is that both property and diagnostic generics function as expectations that are added to other expectations when reasoning. The second thesis is that property generics do not deal with facts about the world but express typicality relations between concepts and properties. In other words, they express knowledge-what rather than knowledge-that. Based on representing concepts in conceptual spaces, we have proposed a measure of typicality. The third thesis is that diagnostic generics emphasise properties that separate the concept in the generics from other categories in the same contrast class. We have also shown that property generics and diagnostic generics do not fulfill upward entailment. In addition, Lawler's test can be used to distinguish property and diagnostic generics.

Finally, we have claimed that understanding generics requires an account of their pragmatic dimension. We believe that the peculiarity of this class of sentences is that they provide an efficient way of conveying prototypical and diagnostic information about kinds. Generics can be understood as a linguistic counterpart of Rosch's principle of cognitive economy in categorisation (Rosch 1988, 312), that is, the disposition to reduce the informational complexity of stimuli for having categories that are easier to use in cognitive tasks. This aspect is also mirrored by our use of Grice's principle of maximal informativeness. Therefore, principles of cognitive economy could be part of the explanation of why generic statements have such a central role in communication and learning.

An advantage of our approach is that it opens up for new kinds of empirical investigations. Since the measures of typicality and diagnosticity that we have proposed can be applied to determining how acceptable a generic is, the predictions that are generated can be tested experimentally. As far as we are aware, there have not been any tests of how subjects judge the strength of generics. Of course, the testing depends on having a conceptual space with an appropriate distance measure established, but there exist such spaces (Douven 2016). And other spaces and distances can be estimated by using similarity judgments and multi-dimensional scaling techniques.

Acknowledgements

Open Access funding enabled and organized by Projekt DEAL.

Conflicts of Interest

The authors declare no conflicts of interest.

Endnotes

[1] For instance, for Pelletier and Asher (1997), generics of the form '*Fs* are *Gs*' are true if and only if for each individual in *F*, in the most normal possible worlds for *F* (ordered according to some contextual information), that individual has property *G*.

[2] Prasada et al. (2013) distinguish between majority characteristic, minority characteristic, majority statistical, striking, and quasi-definitional generics. Even though the mapping is not precise, we would sort majority characteristic and quasi-definitional generics as property generics and the rest as diagnostic generics.

[3] We use the term 'category' to denote the set of entities that fall under a concept.

[4] These domains are not an exhaustive representation of APPLE but are chosen for illustrative purposes.

[5] For bounded convex regions, the central point is well-defined. For unbounded regions, however, this does not hold.

[6] A property is said to be diagnostic when its presence or absence provides useful information for determining whether an entity belongs to a particular category. A property can be diagnostic without being considered 'essential'. Consider, for instance, the properties *having a mane* and *being carnivore* for the concept LION. The presence of a mane will allow us to more quickly identify an animal as a lion compared to knowing that the animal is a carnivore, even though it is

more difficult to conceive of cases of vegetarian lions than of lions without manes. In other words, although both properties are closely associated with the concept, being a carnivore is more central to LION than having a mane (for an explanation of the notion of centrality of a property, see Sloman, Love, and Ahn 1998).

[7] We do not believe that a sharp distinction between essential and non-essential properties can be made. However, a weak form of the distinction can be maintained by saying that the dimensions with the highest weights are those that are most essential for a concept. A definitional generic, such as 'a kangaroo is a marsupial', picks out a property with the highest weight for the concept.

[8] Osta-Velez and Gärdenfors (2022) propose various ways in which our expectations about object-properties are updated in the light of new information.

[9] This mechanism has been formalized in Gärdenfors and Makinson (1994). The key idea is that *p non-monotonically entails q* means that *q* follows from *p*, *together with* all the propositions in the background information that are 'sufficiently well' expected in the light of *p*. The technical specification of what is meant by 'sufficiently well' is to require that any proposition *p* that is added to the reasoning be more expected than not-*p* in the expectation ordering. In the example above, 'penguins do not fly' is more expected than 'penguins fly'. The formalities of this procedure and its relations to typicality are developed in greater detail in Osta-Vélez and Gärdenfors (2022).

[10] Furthermore, generic sentences are always syntactically unmarked, unlike sentences with quantifiers (Leslie 2012). This is another sign that generics are basic in our communication.

[11] For any point $x \in C(M)$ and property $R \subseteq D_i \in C(M)$, $R(x)$ expresses that the coordinate corresponding to domain D_i falls under the subregion corresponding to R.

[12] Another possible explanation takes the possible vagueness of prototypes into account. It may be that the prototype for 'Norwegian lambs' and similarly 'fluffy lambs' is more difficult to identify than the prototype for 'lambs'. Such a vagueness will affect the distance measures (Douven et al. 2013).

[13] Again, another explanation takes the possible vagueness of prototypes into account. It may be that the prototype for EXPENSIVE PURPLE CANDLE is more difficult to identify than the prototype for CANDLES and thereby, the closest point where expensive purples candles are not made of wax will be closer.

[14] Note that the previous explanations cohere with Cimpian and Markman's (2009) hypothesis that generics about biological kinds are interpreted in the light of naïve biological theories implicitly held by cognitive agents.

[15] Note that we define diagnosticity in terms of similarities. This makes it different from cue validity which is defined in terms of probabilities.

[16] Cohen (2001, 60–64) makes a distinction between 'absolute' and 'relative' generics which is seemingly similar to ours. A main difference, however, is that he analyses generics in terms of probabilities rather than our use of typicality and diagnosticity. Leslie and Learner (2022, section 4) argue that there are cases that Cohen's probabilistic approach cannot handle.

[17] Children's picture books of animals and other object categories highlight the diagnostic properties of the categories.

[18] Rooij and Schulz (2020) emphasize the role of learning generics.

References

Carlson, G. 2009. "Generics and Concepts." In *Kinds, Things and Stuff: Mass Terms and Generics*, edited by F. Pelletier, 16–36. New York: Oxford University Press.

Chierchia, G. 1995. "Individual Level Predicates as Inherent Generics." In *The Generic Book*, edited by F. Pelletier and G. Carlson, 176–223. Chicago: University of Chicago Press.

Cimpian, A., S. A. Gelman, and A. C. Brandone. 2010. "Theory-Based Considerations Influence the Interpretation of Generic Sentences." *Language and Cognitive Processes* 25, no. 2: 261–276.

Cimpian, A., and E. M. Markman. 2009. "Information Learned From Generic Language Becomes Central to Children's Biological Concepts: Evidence From Their Open-Ended Explanations." *Cognition* 113, no. 1: 14–25.

Cimpian, A., and R. M. Scott. 2012. "Children Expect Generic Knowledge to Be Widely Shared." *Cognition* 123, no. 3: 419–433.

Cohen, A. 2001. "Relative Readings of Many, Often, and Generics." *Natural Language Semantics* 9: 41–67.

Cohen, A. 2004. "Generics and Mental Representations." *Linguistics and Philosophy* 27: 529–556.

Connolly, A. C., J. A. Fodor, L. R. Gleitman, and H. Gleitman. 2007. "Why Stereotypes Don't Even Make Good Defaults." *Cognition* 103, no. 1: 1–22.

Csibra, G., and G. Gergely. 2009. "Natural Pedagogy." *Trends in Cognitive Sciences* 13, no. 4: 148–153.

Declerck, R. 1991. "The Origins of Genericity." *Linguistics* 29, no. 1: 79–102.

Douven, I. 2016. "Vagueness, Graded Membership, and Conceptual Spaces." *Cognition* 151: 80–95.

Douven, I., L. Decock, R. Dietz, and P. Égré. 2013. "Vagueness: A Conceptual Spaces Approach." *Journal of Philosophical Logic* 42: 137–160.

Gärdenfors, P. 2000. *Conceptual Spaces. The Geometry of Thought*. Cambridge, MA: MIT Press.

Gärdenfors, P. 2014. *Geometry of Meaning. Semantics Based on Conceptual Spaces*. Cambridge, MA: MIT Press.

Gärdenfors, P., and D. Makinson. 1994. "Nonmonotonic Inference Based on Expectations." *Artificial Intelligence* 65, no. 2: 197–245.

Gärdenfors, P., and M. Osta-Vélez. 2023. "Reasoning With Concepts: A Unifying Framework." *Minds and Machines* 33, no. 3: 451–485.

Gärdenfors, P., and A. Stephens. 2018. "Induction and Knowledge-What." *European Journal for Philosophy of Science* 8: 471–491.

Gelman, S. A. 2003. *The Essential Child: Origins of Essentialism in Everyday Thought*. New York: Oxford University Press.

Gelman, S. A. 2010. "Generics as a Window Onto Young Children's Concepts." In *Kinds, Things, and Stuff*, edited by J. Pelletier, 100–120. New York: Oxford University Press.

Gelman, S. A. 2021. "Generics in Society." *Language in Society* 50, no. 4: 517–532.

Gelman, S. A., S. J. Leslie, A. M. Was, and C. M. Koch. 2015. "Children's Interpretations of General Quantifiers, Specific Quantifiers and Generics." *Language, Cognition and Neuroscience* 30, no. 4: 448–461.

Goodman, N., J. Tenenbaum, and T. Gerstenberg. 2015. "Concepts in a Probabilistic Language of Thought." In *The Conceptual Mind: New Directions in the Study of Concepts*, edited by E. Margolis and S. Laurence, 623–654. Cambridge, MA: MIT Press.

Greenberg, M., and G. Harman. 2006. "Conceptual Role Semantics." In *The Oxford Handbook of Philosophy of Language*, edited by B. C. Smith, 295. Oxford: Oxford University Press.

Grice, H. P. 1975. "Logic and Conversation." In *Syntax and Semantics, Vol. 3, Speech Acts*, edited by P. Cole and J. Morgan, 41–58. New York: Academic Press.

Hampton, J. A. 2007. "Typicality, Graded Membership, and Vagueness." *Cognitive Science* 31, no. 3: 355–384.

Hampton, J. A. 2012. "Generics as Reflecting Conceptual Knowledge." *Recherches Linguistiques de Vincennes* 41: 9–24.

Hampton, J. A., A. Passanisi, and M. L. Jönsson. 2011. "The Modifier Effect and Property Mutability." *Journal of Memory and Language* 64, no. 3: 233–248.

Haslanger, S. 2014. "The Normal, the Natural and the Good: Generics and Ideology." *Politica & Società* 3, no. 3: 365–392.

Heyer, G. 1990. "Semantics and Knowledge Representation in the Analysis of Generic Descriptions." *Journal of Semantics* 7, no. 1: 93–110.

Jönsson, M. L., and J. A. Hampton. 2006. "The Inverse Conjunction Fallacy." *Journal of Memory and Language* 55, no. 3: 317–334.

Jönsson, M. L., and J. A. Hampton. 2012. "The Modifier Effect in Within-Category Induction: Default Inheritance in Complex Noun Phrases." *Language and Cognitive Processes* 27, no. 1: 90–116.

Krifka, M. 2013. "Definitional Generics." In *Genericity*, edited by A. Mari, C. Beyssade, and F. Del Prete, 372–389. New York: Oxford University Press.

Krifka, M., F. Pelletier, G. Carlson, A. ter Meulen, G. Chierchia, and G. Link. 1995. "Genericity: An Introduction." In *The Generic Book*, edited by G. Carlson and F. Pelletier, 1–124. Chicago: University of Chicago Press.

Lawler, J. M. 1973. "Studies in English Generics." Doctoral diss., University of Michigan.

Leslie, S. J. 2008. "Generics: Cognition and Acquisition." *Philosophical Review* 117, no. 1: 1–47.

Leslie, S. J. 2012. "Generics articulate default generalizations." *Recherches Linguistiques de Vincennes* 41: 25–44.

Leslie, S. J., S. Khemlani, and S. Glucksberg. 2011. "Do all Ducks Lay Eggs? The Generic Overgeneralization Effect." *Journal of Memory and Language* 65, no. 1: 15–31.

Leslie, S. J., and A. Lerner. 2022. "Generic Generalizations." In *The Stanford Encyclopedia of Philosophy* (2023 Edition), edited by E. N. Zalta and U. Nodelman. https://plato.stanford.edu/archives/fall2022/entries/generics/.

Mattos, O., and W. Hinzen. 2015. "The Linguistic Roots of Natural Pedagogy." *Frontiers in Psychology* 6: 1424.

Noyes, A., and F. C. Keil. 2019. "Generics Designate Kinds but Not Always Essences." *Proceedings of the National Academy of Sciences of the United States of America* 116, no. 41: 20354–20359.

Öhman, A., A. Flykt, and F. Esteves. 2001. "Emotion Drives Attention: Detecting the Snake in the Grass." *Journal of Experimental Psychology: General* 130, no. 3: 466–478.

Osherson, D. N., and E. E. Smith. 1981. "On the Adequacy of Prototype Theory as a Theory of Concepts." *Cognition* 9, no. 1: 35–58.

Osta-Vélez, M., and P. Gärdenfors. 2020. "Category-Based Induction in Conceptual Spaces." *Journal of Mathematical Psychology* 96: 102357.

Osta-Vélez, M., and P. Gärdenfors. 2022. "Nonmonotonic Reasoning, Expectations Orderings, and Conceptual Spaces." *Journal of Logic, Language and Information* 31: 77–97.

Pelletier, F. J., and N. Asher. 1997. "Generics and Defaults." In *Handbook of Logic and Language*, edited by J. van Benthem and A. ter Meulen, 1125–1179. Cambridge, MA: MIT Press.

Prasada, S., S. Khemlani, S. J. Leslie, and S. Glucksberg. 2013. "Conceptual Distinctions Amongst Generics." *Cognition* 126, no. 3: 405–422.

Pratto, F., and O. P. John. 1991. "Automatic Vigilance: The Attention-Grabbing Power of Negative Social Information." *Journal of Personality and Social Psychology* 61, no. 3: 380–391.

Rips, L. J. 1995. "The Current Status of Research on Concept Combination." *Mind and Language* 10: 72–104.

Rooij, R. V., and K. Schulz. 2020. "Generics and Typicality: A Bounded Rationality Approach." *Linguistics and Philosophy* 43, no. 1: 83–117.

Rosch, E. 1975. "Cognitive Representations of Semantic Categories." *Journal of Experimental Psychology: General* 104, no. 3: 192–233.

Rosch, E. 1988. "Principles of Categorization." In *Readings in Cognitive Science, a Perspective From Psychology and Artificial Intelligence*, edited by A. Collins and E. E. Smith, 312–322. San Mateo, CA: Morgan Kaufmann Publishers.

Ryle, G. 1949. *The Concept of Mind*. London: Hutchinson.

Sloman, S. A., B. C. Love, and W. K. Ahn. 1998. "Feature Centrality and Conceptual Coherence." *Cognitive Science* 22, no. 2: 189–228.

Sterken, R. 2015. "Generics, Content and Cognitive Bias." *Analytic Philosophy* 56, no. 1: 75–93.

Tessler, M. H., and N. D. Goodman. 2019. "The Language of Generalization." *Psychological Review* 126, no. 3: 395.

Tversky, A. 1977. "Features of Similarity." *Psychological Review* 84, no. 4: 327–352.

Wilson, D., and D. Sperber. 2004. "Relevance Theory." In *The Handbook of Pragmatics*, edited by L. R. Horn and G. Ward, 607–632. Oxford: Blackwell.