Generics and Mental Representations

Ariel Cohen* (arikc@bgumail.bgu.ac.il)

Ben-Gurion University, Israel

Linguistics and Philosophy, 27:5, 2004, 529–556

Abstract. It is widely agreed that generics tolerate exceptions. It turns out, however, that exceptions are tolerated only so long as they do not violate *homogeneity*: when the exceptions are not concentrated in a salient "chunk" of the domain of the generic. The criterion for salience of a chunk is cognitive: it is dependent on the way in which the domain is mentally represented. Findings of psychological experiments about the ways in which different domains are represented, and the factors affecting such representations, account for judgments of generic sentences, facts which cannot be explained by linguistics alone.

The reason for the homogeneity requirement itself is, in turn, also dependent on cognitive considerations. Generics express default rules, and psychological findings have shown that, the more homogeneous the domain, the easier it is for subjects to infer rules about it.

Thus, cognitive results form a crucial part of a comprehensive account of the meaning of a linguistic expression.

Keywords: mental representations,generics,probability

* I would like the following people for helpful and insightful discussions about issues relating to this paper: Chris Gauker, Susan Gelman, Michiel van Lambalgen, Michael Morreau, and Dan Osherson.

1. Introduction

What is the meaning of a sentence? If you ask a semanticist, you are likely to get an answer in terms of truth conditions (and entailments) and felicity conditions (and presuppositions and implicatures). If, on the other hand, you ask a psychologist, you are likely to get an answer in terms of mental representations; the psychologist will talk about the way the sentence or parts of it are represented in the mind/brain, and how they may be related to other concepts.

Obviously, the two types of answer are somehow related. After all, when one judges the truth or felicity of a sentence, one presumably uses some mental representations. However, there appears to be no necessity to combine the two in the actual practice of research. Psychologists can and do make considerable progress analyzing mental representations, without having to rely on any linguistic theory explaining truth and felicity; and semanticists can and do make considerable progress analyzing the logical and pragmatic properties of sentences without having to rely on any psychological theory explaining how the meanings of these sentences are mentally represented.

There is no *a priori* reason why there shouldn't be a case where the nature of mental representation turns out to be crucial for a formal theory of meaning; a case where the predictions of the theory rely crucially not only on whether certain concepts are represented in the mind/brain, but on *how* they are represented. Indeed, a widely used textbook of formal semantics states: "If such [mental] representations are crucial in mediating between symbols and their content, we must not exclude them from semantics" (Chierchia and McConnell-Ginet, 1990, 47). In this paper I propose that the interpretation of generics constitutes precisely such a case.

2. Generics

Generics present a thorny problem for semantics, and few claims about generics are widely agreed upon.[1] One such claim that *is* uncontroversial is the empirical observation that generics tolerate exceptions. For example, (1) is true and felicitous despite the existence of albino ravens.

(1) Ravens are black.

Yet, sometimes the toleration of a generic to exceptions appears very low: often, a property may hold of the vast majority of individuals in the domain of a generic, and yet the sentence is rejected:[2]

(2) a. ?Mammals are placental mammals.

b. ?Books are paperbacks.

c. ?Israelis live on the coastal plain.

d. ?Chinese speak Mandarin.

e. ?People are over three years old.

f. ?Crocodiles die before they attain an age of two weeks.

[1] See Krifka et al. (1995), Cohen (1996), and Pelletier and Asher (1997) for reviews of current theories.

[2] Some may claim that such generics are false, though in my judgment they are odd, just like sentences whose presuppositions fail. At any rate, I think it is clear that the sentences in (2) are not unproblematically true.

g. ?Primary school teachers are female.

h. ?Bees are sexually sterile.[3]

The majority of mammals are placental mammals, and yet (2.a) is bad; the majority of books are paperbacks, and yet (2.b) is bad; the majority of Israelis live on the coastal plain, and yet (2.c) is bad; and so on. In all the sentences in (2), the vast majority of instances do satisfy the predicated property, and yet the generic sentence is rejected.

Interestingly, when we change these sentences slightly, by adding an adverb of quantification such as *usually* or *generally*, the sentences become perfectly acceptable, in fact true:

(3) a. Mammals are usually placental mammals.

b. Books are usually paperbacks.

c. Israelis usually live on the coastal plain.

d. Chinese usually speak Mandarin.

e. People are usually over three years old.

f. Crocodiles usually die before they attain an age of two weeks.

g. Primary school teachers are usually female.

h. Bees are usually sexually sterile.

In earlier work (Cohen, 1999) I have proposed that the interpretation of a generic is a probability judgment: $\mathbf{gen}(\psi, \phi)$ is true iff the conditional probability of ϕ given ψ is high (specifically, greater than

[3] Examples (2.b), (2.f), and (2.h) are from Carlson (1977); example (2.e) is due to Henk Zeevat (pc).

0.5). In addition, generics (but not adverbs of quantification) carry a *homogeneity* presupposition. The generic $\mathbf{gen}(\psi, \phi)$ presupposes that its domain, ψ, is homogeneous, in the following sense: for any psychologically salient criterion by which ψ may be partitioned into subsets, the conditional probability of ϕ ought to be roughly the same given every such subset of ψ. That is to say, the domain of a generic may not have "chunks" where there are significantly more ϕs or significantly fewer ϕs than there are in the rest of ψ. In this paper I intend to demonstrate that the sentences in (2) violate homogeneity, hence their badness.

More formally, we can define the homogeneity presupposition as follows:

DEFINITION 1 (Homogeneity).
The generic $\mathbf{gen}(\psi, \phi)$ *presupposes that exactly one of the following holds:*

1. *for every psychologically salient partition* Ω *on* ψ*, and for every* $\psi' \in \Omega$, $P(\phi|\psi')$ *is high*

2. *for every psychologically salient partition* Ω *on* ψ*, and for every* $\psi' \in \Omega$, $P(\phi|\psi')$ *is low*

Note that since we define homogeneity to be a presupposition, generics that violate it ought to be considered infelicitous because of presupposition failure. Instead of presupposition, we could have made the homogeneity requirement an entailment, in which case sentences that violate it would be false, or an implicature, in which case sentences violating it may be true but odd. Nothing in this paper relies on homogeneity being a presupposition—one may choose the formulation that best fits one's judgment, so long as all agree that sentences such as the ones in (2) are not unproblematically true.

Homogeneity corresponds rather well to the pre-theoretical notion of what a generic sentence means. For example, suppose a friend of mine is coming to Israel for a visit, and is worried about whether she will be able to manage, speaking only English. I reassure her by saying:

(4) Israelis speak English.

Observe that (4) means more than simply that if you meet an Israeli, he or she is likely to speak English; in addition, the sentence requires that, wherever you go in Israel, whichever group of Israeli society you associate with, a member of this community will be likely to speak English. Indeed, suppose my friend spent all her visit in a town where nobody speaks English, or with members of some group of Israeli society where English is rarely spoken. In such a case, she would be justified in accusing me of misleading her.

Note that if, instead of the generic (4), I used an adverb of quantification, there would be no presupposition of homogeneity; if my promise to my friend were (5) instead of (4), she would not be able to blame me for her predicament.

(5) Israelis generally speak English.

Lowe (1991) makes a similar observation, regarding the following generic sentence:

(6) Fido chases cats.

Lowe writes:

> The sort of empirical evidence which *would* support or undermine the truth of [(6)] is a matter that is broadly familiar to us all (though it is by no means a *simple* matter). If Fido is found to chase a good many different individual cats, of varying sizes, colours, ages and temperaments, then, *ceteris paribus*, we shall consider [(6)] to be empirically well supported; if not, then not (p. 295, original emphases).

According to Lowe, then, it is not sufficient for Fido to chase many cats; they must also be cats of many varieties. That is to say, the domain of cats must be homogeneous with respect to the psychologically salient criteria used to partition it, such as size, color, etc.

What sort of arguments are there to support such a theory? There are arguments of two types: empirical and theoretical. One possible approach is to concentrate on the notion of psychologically salient partitions, to determine empirically what constitutes a psychologically salient criterion, and then to use this information to predict the acceptability of various generic sentences, hence provide an empirical test for the theory. Another route is the theoretical one: to demonstrate that the homogeneity presupposition follows from some more general principles, e.g. cognitive principles.

Both approaches call for information outside the realm of linguistics proper. Determining the salience of a partition or discovering the cognitive principles that give rise to homogeneity is not something that can be done using the methods of formal semantics or pragmatics; one must go to psychology to find out the answers. In this paper, I intend to do just that, and argue for my proposed account of genericity using the data and theory of psychological research.

3. What is a salient partition?

3.1. Theories of conceptual representation

The theory that generics require their domain to be homogeneous cannot yet be directly tested, unless we can provide independent evidence that, for a given category, in a given context, a certain partition is, or is not, salient. The place to look for this sort of evidence is in psychological theories of how categories are mentally represented.

How do we determine that certain individuals belong to a certain category? The most widely held view is that this judgment is based on similarity; the members of some category are similar to one another, and different from members of other categories. Hence, categorization is based on judging the similarity between two mental representations. There are two general theories explaining how people arrive at a similarity judgment: the featural and the geometric approach (see Smith, 1995 for an overview). The main difference between the theories is that they assume two different representations of concepts.

According to the featural approach, concepts are represented as sets of features; according to the geometric approach, they are represented as points in a multidimensional space. To illustrate, consider the category *mammal*. We could represent the concept of each species of mammal using a set of features. For example, *lion* will have, among others, the features HAS A MANE, CARNIVOROUS, HAS A PLACENTA, LACTATES, WARM BLOODED, etc.

Alternatively, we could represent mammals in terms of their values on a number of scales, or dimensions. Consider, for example, the dimension SIZE. Then *mouse* would be near one end and *elephant* near the other end, with *lion* somewhere in between. In this way, we can consider each concept to be a point in a multidimensional space, where each dimension represents some scale.

On the basis of their different representations of concepts, the two theories explain similarity judgments in different ways. According to the featural theory, the judgment of similarity between two concepts is a function of the number of features they share; according to the geometric approach, the similarity between two concepts is a function of the distance between them in the relevant multidimensional space.

One difference between the featural and the geometric approaches is that the former, but not the latter, can directly be transformed into a hierarchical, treelike representation. There is a simple way to derive a tree from a featural representation: concept A is a subordinate of concept B just in case A has all the features of B (and possibly additional features). For example, *lion*, being a subordinate of *mammal*, has all the features associated with the concept *mammal*, (e.g. LACTATES, WARM BLOODED) in addition to other features (e.g. HAS A MANE).

Things are different with a geometric representation. Unlike the featural approach, where a subordinate may have more features than a superordinate, in the geometric representation all concepts, at all levels, have values along the same dimensions. The value of a superordinate along some dimension is the (subjective) average of the values of all its subordinates along this dimension (Smith and Medin, 1981, p. 102). For example, each mammal has some value along the SIZE dimension, and the superordinate *mammal* itself also has a value, the average of the sizes of all mammals. This fact means that, unlike the featural approach, a superordinate and a subordinate may conflict on the value of a certain dimension: hence, one cannot derive a hierarchical structure from the geometric representation.

Let us now consider a concrete example of a category represented using each of the two methods. Here is what a tree representation of *mammal* and some of its subordinates might look like:

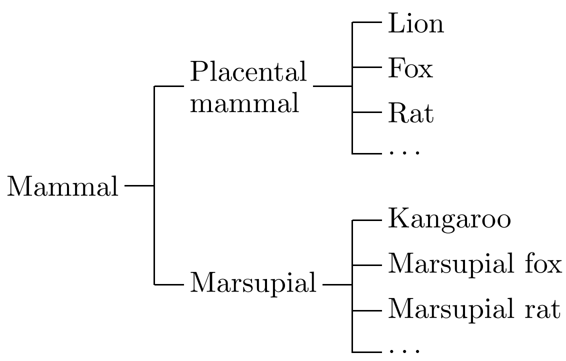

Compare the treelike representation of mammals with a geometric representation, taken from Henley (1969). She characterizes dimension 1 as SIZE, and dimension 2 as FEROCITY.

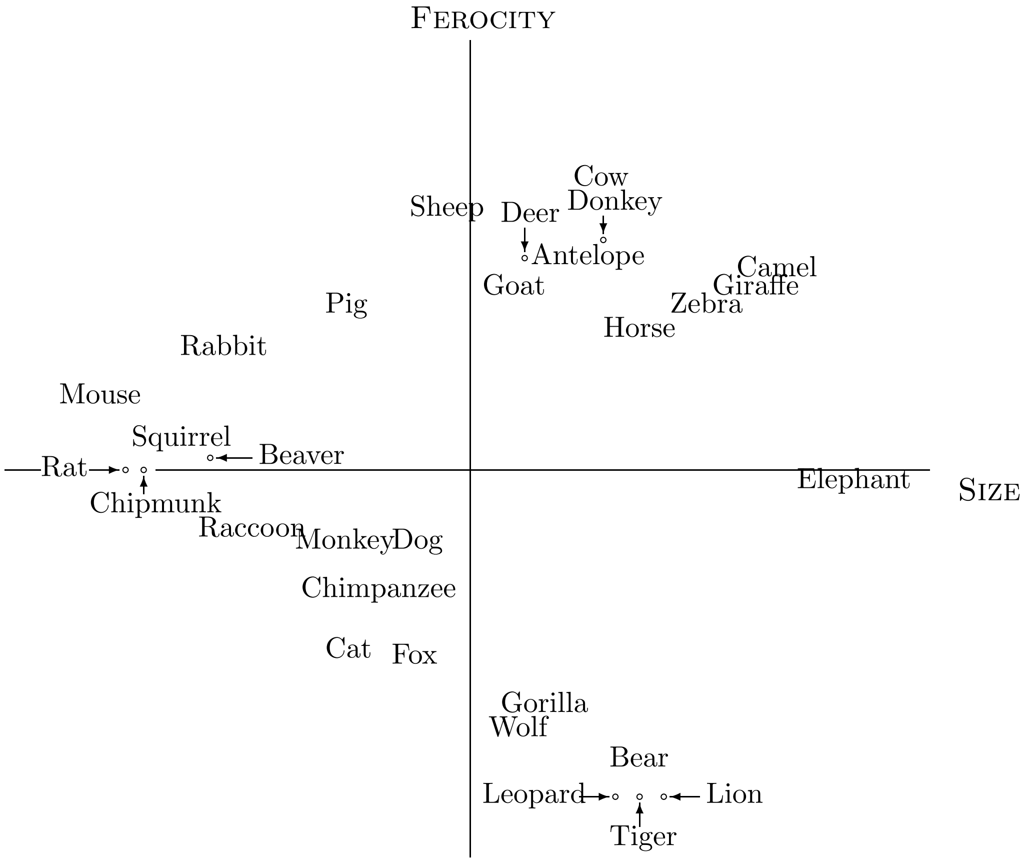

3.2. Homogeneity and conceptual representation

Using these two theories, we can be more specific about what it means for the domain of the generic $\mathbf{gen}(\psi, \phi)$ to be homogeneous. If ψ is a node in a tree, the generic presupposes that $P(\phi|\psi')$ is roughly the same for every node ψ' immediately lower than ψ in the tree.

If, on the other hand, ψ is an n-dimensional space, then $\mathbf{gen}(\psi, \phi)$ presupposes that $P(\phi|\psi')$ is roughly the same for every ψ' that is a "slice" along one of the dimensions of ψ. That is to say, for all ψ's that share the same value along one of the dimensions of ψ, $P(\phi|\psi')$ is more or less the same.

Thus, if *mammal* is represented as a tree, to say that this concept is homogeneous with respect to a property P is to say that P holds of roughly the same proportion of placental mammals and marsupials. If, on the other hand, *mammal* is represented as a multidimensional space, its homogeneity requires that for any given size, or any given level of ferocity, the property P holds of roughly the same proportion of mammals of that size or ferocity. For example, the probability that a large mammal, or a small mammal, or a domesticated mammal has property P should be more or less the same as the probability that a randomly selected mammal has property P.

Which criterion of homogeneity should we choose? The answer depends, of course, on which theory of conceptual representation we choose. Is *mammal* represented as a tree or as a multidimensional space? In a sense, both theories are correct; there are cases where one is more suitable, and cases where the other fits the data better, but no theory is best in absolute terms. Sattath and Tversky (1977) put this in the following way:

> The appropriateness of tree vs. spatial representation depends on the nature of the task and the structure of the stimuli…the two models may be appropriate for different data and may capture different aspects of the same data (pp. 337–338).

So, two factors affect the choice of representation: the nature of the task and the structure of the stimuli. Let us first consider the nature of the task, i.e. the setting of the experiment. One important factor about the nature of the experiment turns out to account for a difference between the interpretation of generics expressing nominal predication, and those expressing verbal or adjectival predication.

3.3. Nominal vs. verbal/adjectival predication

In similarity judgment tasks, subjects are usually presented with pairs of names of concepts, and are asked to judge (using a numerical scale) how similar the members of each pair are. In some tasks, the pairs only consist of concepts at the same level in the hierarchy (e.g. *cat* and *cow*); in others, some of the pairs contain both a concept and its superordinate (e.g. *cat* and *mammal*).

This difference in the task may significantly affect conceptual representation, because it is impossible to represent concepts and their superordinate category in the same multidimensional space. The explanation for this is as follows. Subjects usually judge a concept to be very similar to its superordinate. So, all species of mammals are similar to the concept *mammal*, although they themselves may be very different from one another. For example, *cat* and *cow* are judged very dissimilar, although both are judged very similar to *mammal*. However, if the degree of dissimilarity between two concepts is judged on the basis of the distance between them in the multidimensional space, it is impossible to construct a space where the points corresponding to species of mammals must all be close to one point (corresponding to the superordinate) yet may be very far from one another. Thus, we expect the tree representation to provide a better account of the data when the task calls for similarity judgments of concepts as well as their common superordinate. This, indeed, has been found to be the case (Tversky and Hutchinson, 1986; Shoben and Ross, 1987).

Now consider (2.a), repeated below:

(7) ?Mammals are placental mammals.

This sentence presupposes that *mammal* is homogeneous with respect to the property of being a placental mammal. The sentence contains two terms, referring to a concept (*placental mammal*) and its superordinate (*mammal*). Based on the consideration above, we would expect the tree model to be triggered. This means that (7) presupposes that placental mammals and marsupials have roughly the same proportion of placental mammals. This, of course, does not hold: all placental mammals have this property, but no marsupial does. Hence, homogeneity is not satisfied, and (7) is bad.

Now note what happens if we omit the subordinate category name, and a verbal property takes its place, as in (8).

(8) Mammals have a placenta.

Since being a placental mammal is the same as having a placenta, (8) ought to be judged the same as (7); in fact, however, it is significantly better. The reason is that here, since the name of the concept *placental mammal* is not explicitly mentioned, the tree representation is less likely to be triggered. If the geometric representation is triggered instead, homogeneity is satisfied. This is because, as is well known, the marsupials, while different from placental mammals in not having a placenta, are similar to them in most other respects. There is, for example, a marsupial version of a fox, a marsupial version of a rat, etc.

The marsupial fox and the placental fox are similar in terms of their size and ferocity; so are the marsupial rat and the placental rat, and so on. Therefore, mammals of any given size or level of ferocity will have roughly the same proportion of placental mammals among them. Hence, (8), unlike (7), satisfies homogeneity.

It might be argued that the problem with (7) is not lack of homogeneity, but the fact the same word, *mammals*, occurs in both subject and predicate.[4] Perhaps this is what makes the sentence awkward, and homogeneity is irrelevant. To counter this objection, consider the following example:

(9) a. ?Bees are workers.

b. Bees collect pollen.

Bees are divided into workers, who collect pollen, the queen, who doesn't, and the drones, who don't collect pollen either. Almost all bees are workers, yet (9.a) is odd. Sentence (9.b), in contrast, is much better, in fact true. Note that (9.a) does not contain the same word in both subject and predicate, and yet it is awkward. The homogeneity requirement provides a simple explanation for the facts exemplified in (9). Sentence (9.a), which refers to concepts at two levels of representation, triggers a tree representation of *bee*:

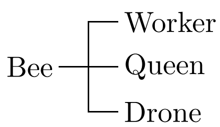

Since the queen and drones do not collect pollen, homogeneity fails. In contrast, (9.b) does not trigger the tree model, but rather a geometric

[4] I am indebted to Michael Morreau for this comment.

representation. In section 5 below we will see why (9.b), represented in a multidimensional space, satisfies homogeneity.

Note that this type of explanation is compatible with Montague's (1973) theory, according to which nominal and verbal predication have the same semantics: the two types of predication differ only in their pragmatics, namely in their effects on homogeneity.

At this point I would like to consider whether it is possible to provide an alternative account of the facts exemplified by (7), (8), and (9). Could one argue instead that Montague's generalization does not hold in the case of generics, and propose a difference in the logical forms of nominal and verbal predication that would account for the phenomena under discussion?

I am not aware of any proposal exploiting such an idea, but one line of approach suggests itself. It is well known (Krifka et al., 1995) that there are generics, such as (10.a), which predicate some property directly of a kind, whereas characterizing generics, e.g. (10.b), predicate some property of instances of the kind.

(10) a. Dinosaurs are extinct.

b. Dinosaurs were very large.

One might suggest that while verbal/adjectival predication may result in a characterizing generic, generic nominal predication is restricted to direct kind predication. Thus, (7) and (8) would have different logical forms: (11.a) and (11.b), respectively.

(11) a. $\mathbf{IS\text{-}A}(\uparrow\mathbf{mammal}, \uparrow\mathbf{placental\text{-}mammal})$

b. $\mathbf{gen}_x[\mathbf{mammal}(x)][\mathbf{have\text{-}placenta}(x)]$

The formula (11.a) says that the kind ↑**mammal** is subsumed by the kind ↑**placental-mammal**. Since this is not the case, sentence (7) is predicted to be false.[5] In contrast, (11.b) says that, in general, if x is a mammal, x has a placenta. Since this is, indeed, the case, (8) is predicted to be true, as desired.

Similarly, one could suggest that (9.a) says that the kind *bee* is subsumed by the kind *worker*, which is false; while (9.b) says that generally, if x is a bee, x collects pollen, hence its truth.

While the approach sketched here might appear attractive, it suffers from some grave difficulties, which ultimately render it untenable. One difficulty involves indefinite singulars. It is well known that direct kind predication is impossible with indefinite singular generics; unlike (10.a), (12) is bad (unless it is read taxonomically).

(12) *A dinosaur is extinct.

However, it is quite possible to form a nominal predication generic with an indefinite singular:

(13) A dog is a mammal.

Therefore, it would seem that nominal predication generics cannot be cases of direct kind predication.

This difficulty is not insurmountable, however. There are reasons to believe that the generic use of indefinite singulars is fundamentally different from that of "real" generics, i.e. sentences that involve generic quantification (Cohen, 2001). If this is so, one cannot conclude from

[5] Of course, (7) is not simply false, but sounds quite odd. But let us assume, for the sake of argument, that this difficulty can be resolved.

the behavior of indefinite singulars anything regarding the behavior of generics like (7), (8), and (9).

A more serious problem involves adverbs of quantification. Like indefinite singulars, they, too, are impossible in sentences expressing direct kind predication:

(14) *Dinosaurs are always extinct.

And, again like indefinite singulars, adverbs of quantification are quite acceptable with nominal predication:

(15) a. Mammals are usually placental mammals.

b. Bees are generally workers.

One could, perhaps, salvage the direct kind predication view by postulating that the adverb somehow changes the logical form, so that sentences like the ones in (15) are not cases of direct kind predication, yet sentence like (7) and (9.a) still are. Such an explanation, while conceivable, would be unmotivated and would complicate the proposal further. In contrast, under the proposal made in this paper, the answer is quite simple: the sentences in (15) are fine because adverbs of quantification, unlike generics, do not presuppose homogeneity.

Yet a third objection to the idea that nominal predication generics express direct kind predication is the following. If nominal predication generics do not involve quantification, they should not exhibit scope ambiguities. But they do exhibit such ambiguities, just like regular characterizing generics. For example, in the most plausible reading of (16.a), the disjunction takes wide scope, and the sentence can be paraphrased as (16.b).

(16) a. Whales are mammals or fish (from Schubert and Pelletier, 1987).

b. =Whales are mammals or whales are fish.

In contrast, (17.a) is most easily interpreted as (17.b), where the disjunction takes narrow scope.

(17) a. Pets are cats or dogs.

b. =In general, if x is a pet, x is a cat or x is a dog.

This objection is quite serious. In my opinion it proves conclusively that nominal predication generics and verbal/adjectival generics have the same logical form: both are characterizing generics and involve the generic quantifier. Thus, nominal and verbal/adjectival predication do have the same semantics, and the only way to account for the difference between their interpretations is pragmatically, via the homogeneity presupposition.

4. Tree representations

Some concepts, such as *mammal*, apparently can be represented either as a tree or as a multidimensional space, depending on the task. But recall that Sattath and Tversky (1977) point out that the representation may also depend on the nature of the stimuli. There are concepts that are naturally represented as trees, and others that are naturally represented geometrically. In this section we will consider more closely concepts that are represented as trees, and generic sentences involving such concepts. In the next section we will consider concepts that receive geometric representations.

Take, for example, the concept *movie*. Rifkin (1985) has found that it is mentally represented as a tree:

- Movie
  - Comedy
  - Horror
  - Drive-in movie
  - Thriller
  - Love story
  - Drama
  - Science fiction
  - Western movie

With the possible exception of *drive-in movie*, the criterion used to partition the concept appears to be based on the content of the movie. Let us assume that the concept *book* is similarly represented, perhaps by the following tree:

- Book
  - Encyclopedia
  - Mystery
  - Adventure
  - Science fiction
  - . . .

Now consider (2.b), repeated below:

(18) ?Books are paperbacks

Assuming the representation of *book* as above, is the domain homogeneous with respect to the property of being a paperback? The answer is no: mystery books are typically paperbacks, whereas encyclopedias typically have a hard cover. Thus, the homogeneity requirement is not satisfied, and this is why (18) is bad.

Rosch et al. (1976) have identified in hierarchical representations a special level, which they call the *basic level.* For example, *dog* is a basic level concept, while *mammal* is a superordinate and *cocker spaniel* is a subordinate. Members of basic level categories have a significant number of properties in common, and are relatively easily distinguished from members of other basic categories. This means that when we observe the environment, the first categorization is made at the basic level (e.g., *This is a dog*, rather than *This is a mammal* or *This is a cocker Spaniel*). Rosch et al., 1976 have discovered a number of properties that distinguish basic level categories with respect to superordinate and subordinate categories.

Of particular relevance for our purposes here is the fact that basic level concepts have different criteria for decomposing them into concepts at a lower level.

Specifically, Rosch et al. demonstrate that basic level concepts may be distinguished from one another by their shapes, while subordinates may not. Thus, if we partition a superordinate concept into basic level concepts, shape is a possible partitioning criterion; but not if we further partition a basic level concept into its subordinates. Here, for example, is part of the representation of *fruit* (basic level concepts are in **bold**):[6]

[6] The concept *fruit* apparently has no interpretable geometric representation (Shoben, 1976).

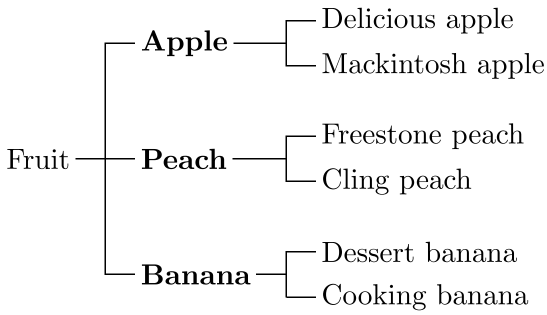

Note that the basic level concepts, *apple*, *peach*, and *banana*, differ in their shapes; in contrast, the subordinates of any one of them (e.g. *delicious apple* and *mackintosh apple*) have very similar shapes. We would, therefore, predict the following: if a property involving shape is predicated of a superordinate (e.g. *fruit*) in a generic sentence, then the existence of even one basic level concept that does not satisfy the property would suffice to violate homogeneity. On the other hand, if some shape is predicated of a basic level concept, homogeneity would be satisfied, and the sentence would be acceptable. Anything subordinate to basic level concepts, of course, also cannot be partitioned by shape, so if a shape is predicated of a subordinate, the sentence would also be fine.

This prediction is, indeed, borne out:

(19) a. ?Fruits are round.

b. Apples are round.

c. Delicious apples are round.

Sentence (19.a) is rejected, although most fruits probably are, indeed, round. This is because the superordinate concept may be partitioned according to shape, resulting in some fruits (e.g. *apple*) that are round, but some (e.g. *banana*) that are not.

In contrast, (19.b) is fine, although not all apples are round (some are oblong). The reason is that, being a basic level concept, *apple* may not be partitioned according to shape. Thus, homogeneity is satisfied, and (19.b) is acceptable (and true). Of course, *delicious apple*, too, being a subordinate level, may not be partitioned according to shape, hence (19.c) is fine.

It is reasonable to assume that not only shape, but other perceptually salient properties behave in the same way: basic level concepts may not be partitioned according to perceptually salient properties. We would predict, then, that generics that predicate a perceptually salient property of a basic level concept will not violate homogeneity. This does, indeed, appear to be the case:

(20) a. Roses are red.

b. Tigers have stripes.[7]

Although there are certainly white and yellow roses, (20.a) is fine. Since *rose* is a basic level concept, if we assume that basic level concepts cannot be partitioned into their subordinates based on color, this judgment would thereby be explained. Indeed, our first example, sentence (1), is explained in exactly the same way, since *raven* is a basic level concept.

Sentence (20.b) is also acceptable (and true), despite the existence of albino tigers. Since *tiger* is a basic level concept, and having stripes is a perceptually salient property, this judgment is precisely as predicted.

[7] I am indebted to Dan Osherson for pointing out the significance of this example to me.

And, of course, (21) is also fine, since *Bengal tiger* is a subordinate of *tiger*.

(21) Bengal tigers have stripes.

An additional property of basic level concepts has been proposed by Tversky and Hemenway (1984). They claim that superordinates may be partitioned according to their parts, but basic level concepts may not. Take, for example, the following representation of *furniture*, from Rosch et al. (1976):

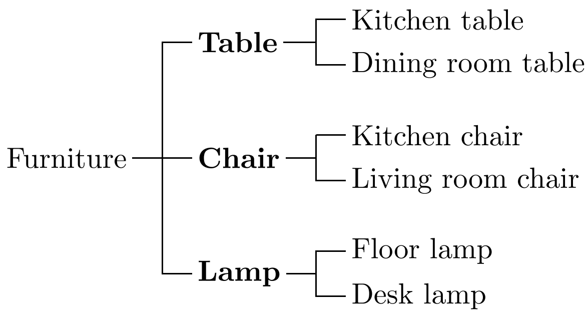

The basic level concepts *table*, *chair*, and *lamp* differ in the parts they have. For example, tables have tops, chairs have backs, and lamps have light bulbs. But the subordinates do not differ in the same way: kitchen tables and dining room tables have essentially the same parts.

Thus, we predict that if a generic predicates of a superordinate level concept the property of having a certain part p, any basic level concept that does not have p would suffice to violate homogeneity. In contrast, if having p is predicated of a basic level concept (or its subordinate), homogeneity ought to be satisfied. This does, indeed, seem to be the case:

(22) a. ?Furniture has four legs.

b. Tables have four legs.

c. Coffee tables have four legs.

Furniture, being a superordinate level concept, may be partitioned according to its parts. Since some furniture (e.g. *lamp*) do not have four legs, homogeneity is not satisfied and the sentence is bad. In contrast, *table* may not be partitioned according to the parts it has. Although some tables do not have four legs, they do not form a class into which the basic level concept may be decomposed. Rather, these tables are spread more or less evenly among the subordinates of *table*. Hence, homogeneity is satisfied, and the sentence is fine. The same holds for the subordinate *coffee table*, accounting for the acceptability of (22.c).

The same phenomenon applies in biological taxonomies, where the basic level is higher in the hierarchy; for example, Rosch et al. have discovered that it is not species of birds (e.g. *robin*, *duck*, etc.) that are at the basic level, but rather the category *bird* itself. And, indeed, there is a difference between (23.a) and (23.b).

(23) a. ?Animals have legs.

b. Birds have wings.

c. Ducks have flat bills.

Since *animal* is a superordinate level concept, it may be partitioned with regard to its parts. Many animals have legs, but some do not; hence, homogeneity is violated, and (23.a) is bad. In contrast, (23.b) is fine, although not all birds have wings (kiwis, moas). The reason is that partitioning the domain of birds according to their parts is disallowed, since this is a basic level concepts, hence the sentence is fine.

Duck, being a subordinate concept, may of course also be partitioned according to its parts, hence the acceptability of (23.c), despite the existence of exceptions such as the merganser, which has a narrow bill.

Murphy and Wisniewski (1989) have suggested another way in which superordinates and basic level concepts differ. According to them, superordinates may be partitioned according to the *scene* they are in, but basic level concepts may not. For example, the following is a possible mental representation:

Animal ─┬─ Animal in the wild
        └─ Animal in the zoo

This representation is fine, because *animal*, being a superordinate category, may be partitioned according to scenes: wild nature and of the zoo. In contrast, the following representation is *not* possible:

Tiger ─┬─ Tiger in the jungle
           └─ Tiger in the zoo

Here, it is impossible to partition *tiger*, a basic level concept, according to scenes.

Now, consider the following pair of sentences:

(24) a. ?Animals live in the wild.

b. The tiger lives in the jungle.

Sentence (24.a) is odd; this is because the concept *animal* may be decomposed according to the scenes in which various animals appear. In particular, as we have seen, it may be partitioned into wild animals and zoo animals. Since zoo animals do not live in the wild, homogeneity is violated and the sentence is bad.

In contrast, (24.b) is fine. This is because the basic level category *tiger* may not be partitioned according to scenes. It may, of course, be partitioned in other ways, but presumably in every subordinate of *tiger*, the proportion of jungle tigers would be roughly the same. Hence, homogeneity is satisfied, and (24.b) is fine. Of course, the same considerations hold for the subordinate *Bengal tiger*, hence (25) is fine.

(25) The Bengal tiger lives in the jungle.

5. Geometric representations

We have seen above that some concepts are naturally represented as trees. There are other concepts that are naturally represented as a multidimensional space.

One such concept is what Wish (1970) and Wish et al. (1972) call *nation*. They have found that nations are partitioned along four dimensions: Cold War political alignment, economic development, geography, and culture. The first dimension is irrelevant today, but the three others certainly are. Thus, nations that have a similar level of economic development (e.g. USA, Germany, Britain), nations that share a geographical region (e.g. India, Indonesia, and China), or nations that are perceived to share a similar culture (e.g. Spain and South American countries), are grouped together as psychologically relevant subsets of the domain.

Hence, we expect nationality based domains not be homogeneous with respect to properties relating to economic development, geography, or culture. This prediction is, indeed, borne out. For example, the vast majority of Americans are classified as middle class, and two thirds of them own homes, yet (26) is bad.

(26) ?Americans $\left\{ \begin{array}{l} \text{are middle class.} \\ \text{own homes.} \end{array} \right\}$

The other two dimensions can explain the unacceptability of (2.c) and (2.d), repeated below:

(27) a. ?Israelis live on the coastal plain.

b. ?Chinese speak Mandarin.

Although the majority of Israelis do live on the coastal plain, if we decompose the domain by geographical locations, we will find subclasses of Israelis who do not—e.g. those who live in Jerusalem. Therefore, the domain of Israelis is not homogeneous with respect to the property of living on the coastal plain, and (2.c) is bad. Similarly, while the majority of Chinese speak Mandarin, there are regions and cultural groups where Mandarin is not spoken; hence, the domain is not homogeneous, and the sentence is bad.

Some of the dimensions are salient in any context, whereas the salience of others is context-dependent. In particular, if the instructions given to the experimental subjects refer to some scale, it becomes a salient dimension, however unnatural it may be regarded initially. For example, Sadler and Shoben (1993) elicited similarity judgments between pairs of professions. A control group was given no further instructions. The results of this group could be represented as a three-dimensional space: one dimension corresponded to whether the job involved mental or manual work (with *priest* at one end and *butcher* at the other end); the second dimension corresponded to whether the job took place outdoors or indoors (with *ranger* at one end and *librarian* at the other end); the third dimension expressed the extent to which the job involved working with people (with *teller* at one end and *welder* at the other end).

Two groups of subjects were given additional instructions. One group was asked to adopt the point of view of an IRS tax auditor who was trying to determine which occupations were more likely to be involved in tax fraud. The other group was asked to imagine being a political campaign manager who was trying to predict the typical attitudes of various occupations toward the issue of regulating agricultural chemicals. The representation of the similarity judgments of each of these two groups retained the mental vs. manual dimension. However, instead of the other two, each contained a new dimension: likelihood to commit tax fraud (with *physician* at one end and *priest* at the other end) for one group, and level of concern about agricultural chemicals (with *farmer* at one end and *chauffeur* at the other end) for the other. Thus, when the context calls for the consideration of some scale, it becomes one of the dimensions of the representation, and, consequently, can induce a salient partition.

In Sadler and Shoben's experiments, a property (e.g. likelihood to commit tax fraud) was made salient by asking subjects to consider it. But context can highlight a property simply by alluding to it in a sentence. For example, Barclay et al. (1974) presented subjects with sentences, and then asked them to recall nouns appearing in these sentences with the aid of verbal cues. They found that for different sentences, different recall cues were better. For example, if "piano" was recalled from (28.a), the cue "something heavy" was better than "something that makes a nice sound." The reverse was true if the word was recalled from (28.b).

(28) a. The man lifted the piano.

b. The man tuned the piano.

Barclay et al. concluded that when the word is encountered in a sentence, it is represented differently, depending on the property predicated of the concept in the sentence. When (28.a) is interpreted by the hearer, the heaviness of the piano becomes a salient property, whereas when (28.b) is heard, the sound of the piano becomes salient.

Given these results, I suggest the following: whenever a generic predicates a property of some concept, and the predicate contains reference to a value on a scale, the concept is represented as a multidimensional space, with this scale as one of its dimensions. The prototypical case of a value on a scale is, of course, a number word (Horn, 1972). Hence, any sentence whose predicate contains an explicit number word will be ruled out, by failure of homogeneity, so long as there are *any* exceptions to the predicated property.

This does, indeed, appear to be the case:

(29) a. ?Buildings are less than 1000 feet tall.

b. ?Animals weigh less than two tons.

c. ?Shoes are size 7 and above.

When (29.a) is encountered, the domain of buildings is represented with HEIGHT as one of the dimensions. Since a few buildings taller than 1000 feet do exist (e.g. Sears Tower), homogeneity fails. With (29.b), the domain of animals is represented as a multidimensional space, with WEIGHT as one of its dimensions. Since some animals (e.g. the whale) do weigh over two tons, homogeneity is violated. And when (29.c) is encountered, the domain of shoes is represented with SIZE as one of its dimensions. Since some shoes are smaller than size 7, homogeneity fails and the sentence is bad.

In this way we can account for (2.e) and (2.f), repeated below:

(30) a. ?People are over three years old.

b. ?Crocodiles die before they attain an age of two weeks.

In both sentences, the predication of age involves a number word. Therefore, in this context we can assume that age is one of the dimensions used in the representation of these domains. Clearly, given such a partition, the domains are not homogeneous: although most people are over three years old, babies are not; and although the majority of crocodiles do die before they attain an age of two weeks, adult crocodiles obviously did not. Hence, the homogeneity requirement is not satisfied, and the sentences are bad.

Note that when age is not mentioned explicitly, it is not a salient partition:

(31) a. People want to have a lot of money.

b. Crocodiles eat fish.

Sentence (31.a) is fine, although babies do not want to have a lot of money; and (31.b) is fine, although very young crocodiles do not eat fish. This is because age is not one of the dimensions that are used in the representations, and hence babies and young crocodiles, although they are exceptions, do not constitute a salient “chunk” that could violate homogeneity.

Even when age *is* mentioned, but does not trigger a scale (e.g. because the predicate does not contain a number word), it is not used as a dimension in the partition. Consider (32), suggested to me by an anonymous reviewer.

(32) Dogs live longer than rabbits.

This sentence refers to age, but its predicate does not contain any value on the age scale. Hence, the domain is not partitioned according to age, and the sentence is fine, although some rabbits live longer than some dogs.

The SEX dimension, in particular, is an important scale, and may be used to partition a domain (see, for example, Rosenberg and Kim, 1975). This scale contains few values: *female*, *male*, and, in some species, *neuter*. When the predicate mentions one of these values, sex becomes a salient partition.

We can therefore account for (2.g) and (2.h), repeated below:

(33) a. ?Primary school teachers are female.

b. ?Bees are sexually sterile.

While the majority of primary school teachers are, indeed, female, if we partition the domain according to sex, there will be a subclass of teachers who are not female—the male teachers. The existence of this subclass is enough to rule (33.a) out as a violation of the homogeneity constraint. Similarly, although the majority of bees are, in fact, sexually sterile, if we partition them according to sex, we will get subclasses (the drones and queens) that are not.

When the sex dimension is not mentioned, the domain is not partitioned with respect to it. This is exemplified by (9.b), repeated below:

(34) Bees collect pollen.

Sentence (34) is acceptable (and true), although the queens and drones do not collect pollen; however, since the sex dimension is not mentioned in the sentence, it does not induce a partition on the domain, which is therefore homogeneous.

6. Where does the homogeneity requirement come from?

We have seen some arguments that generics presuppose that their domain is homogeneous. One question arises naturally: why? Why do generics have such a requirement? If it can be shown that homogeneity arises from general cognitive principles, the proposal made in this paper may be further strengthened.

Before answering this question, let us first consider a different, but related one. I have claimed that the *meaning* of a generic is probability judgment; but what is its *use*? Why would anyone want to utter a probability judgment? Fisher (1959) considers

> a gambler, who, for example, stands to gain or lose money, in the event of an ace being thrown with a single die. To such a man the information supplied by a [probability judgment]. . . will seem not merely remote, but also incomplete and lacking in definiteness in its application to the particular throw in which he is interested. Indeed, by itself it says nothing about that throw (p. 32).

What Fisher's gambler needs is not a probability judgment, but a *rule* dictating what to do in any particular gambling situation.

Elsewhere (Cohen 1996; 1997) I have proposed that while the *meaning* of a generic is a statement of high probability, its *use* is to state a default rule. Someone who utters (1), repeated below, suggests a rule of reasoning, according to which, if we know that Nevermore is a raven, we may conclude that it is black, unless we know something to the contrary.

(35) Ravens are black.

The meaning and use are not unrelated: the default rule is *sound* (or "useful") iff the generic sentence is acceptable and true.

Indeed, a rule can be inferred from a probability judgment: we can take a probability judgment and apply it to a given situation. For example, if we know that the probability that a raven is black is high, we can use this information to conclude that Nevermore is black. But we can only do this if we are allowed to assume that the situation we are in is not an exceptional one, a situation where there are many more albino ravens than usual. For example, if we happen to be in a sanctuary for albino ravens, where they are kept to protect them from the attacks of their black brethren, it will be inappropriate to apply this rule. In most other situations, however, the rule is a good one to follow.

In other words, when we decide to use the probability that a random raven is black to infer a rule that ravens are, by default, black, we assume that the population of ravens is homogeneous with respect to the property of being black. We assume that in whatever situation we expect to find ourselves, the percentage of black ravens would be more or less the same. It is precisely because this assumption makes sense, that the rule is sound, and, therefore, the generic (35) is acceptable.

More generally, when we want to infer a rule from a probability judgment, we have to assume that the reference class is homogeneous, or, in Fisher's (1959) words, "subjectively homogeneous and without recognizable stratification" (p. 33). That is to say, for any sample that we may reasonably wish to apply the rule to, the probability judgment ought to be roughly the same; every such sample must be a representative one.

Do people, indeed, require that a domain be homogeneous when they infer a rule about it? Suppose we wanted to find out if a domain ψ is homogeneous with respect to property ϕ, by observing properties of its instances. If we pick individuals that are similar to one another, and find out that they have property ϕ, this would not be good enough; the specimen might all be subsumed by one subclass, which has properties different from other subclasses of the domain, which may not satisfy ϕ. If, on the other hand, the instances we examine are dissimilar, they are likely to belong to different subclasses, and the domain is more likely to be homogeneous. Therefore, if, indeed, when people conclude that a rule applies to a domain, they require the domain to be homogeneous, the following ought to be the case: when one notes that some subordinates, say s_1 and s_2, of a domain ψ, satisfy property ϕ, and uses this fact to argue that ϕ applies to ψ as a whole, the strength of this argument will be judged to be proportional to the perceived dissimilarity of s_1 and s_2.

Exactly this behavior has, indeed, been observed by Osherson et al. (1990). They consider inductive arguments, for example:

(36) Sparrows have sesamoid bones.

Eagles have sesamoid bones.

All birds have sesamoid bones.

In this argument, information about subclasses (sparrows and eagles) of a class (birds) is used to derive a rule about all the members of the class.[8] Note that here, as in the rest of the arguments used in Osherson et al.'s experiments, the sentences use detailed biological information, which normal subjects are not expected to know beforehand. Thus, their judgments are expected to rely solely on the perceived strength of the formal argument.

Not surprisingly, Osherson et al. have found that if we have information about more subclasses, the argument is judged to be stronger. For example, the argument in (37) is stronger than that in (36).

(37) Hawks have sesamoid bones.

Sparrows have sesamoid bones.

Eagles have sesamoid bones.

All birds have sesamoid bones.

This is because in (37) we have evidence coming from an additional subclass, hawks, which was absent in (36). Knowing about more members of a class that satisfy the property (with none that violate it)

[8] The conclusion here is not a generic, but a universal, so the rule is not a default rule; but this does not affect the point I am making here.

increases the probability that all members satisfy it, hence strengthens the argument.

It turns out that not all supporting evidence strengthens an argument in the same way. Osherson et al. have tested the following argument:

(38) Hippopotamuses have a higher sodium concentration in their blood than humans.

Hamsters have a higher sodium concentration in their blood than humans.

All mammals have a higher sodium concentration in their blood than humans.

Information about hamsters adds to information about hippopotamuses to provide an argument about all mammals. This argument is judged by subjects to be quite strong. Interestingly, if instead of hamsters, we add information about rhinoceroses, the argument is judged to be considerably weaker:

(39) Hippopotamuses have a higher sodium concentration in their blood than humans.

Rhinoceroses have a higher sodium concentration in their blood than humans.

All mammals have a higher sodium concentration in their blood than humans.

Now note that hippopotamuses and rhinoceroses are rather similar; they are both large herbivorous African mammals. But hippopotamuses and hamsters are rather different, in their size, habitat, behavior, etc. Thus, knowing that hippopotamuses and rhinoceroses share some property leaves open the possibility that this property is only characteristic of some subclass of mammals, say that of the large herbivorous African mammals; if this is the case, the domain of mammals would not be homogeneous with respect to this property, and hence the argument is weak. In contrast, it is hard to think of a natural subclass of mammals that includes both hippopotamuses and hamsters but leaves other mammals out. Therefore, if we find that both hippopotamuses and hamsters share a property, the domain is more likely to be homogeneous, hence the argument is judged stronger. These results, then, indicate that, given an inductive argument, the more evidence it provides that the domain is homogeneous, the stronger it is judged.

7. Conclusion and Future Work

In this paper I have proposed that generics presuppose their domain to be homogeneous. Homogeneity follows from the fact that generics are interpreted as probability judgments, and are used to express default rules. Homogeneity is defined relative to conceptual representations. Hence, a complete account of the meaning of generics crucially requires input from cognitive psychology. Generics, then, are a case where empirical results and theoretical considerations from psychology may be used as evidence for or against a formal account of a linguistic phenomenon.

The evidence I have brought in support of this claim is of the following sort: find a domain that has been studied by psychologists; find out what these studies tell us about the possibility of salient partitions of this domain; demonstrate that a generic that violates homogeneity with respect to these partitions is bad, and one that does not—is good.

This type of evidence is, by its nature, limited to the domains whose representations have been investigated empirically; it is, in fact, further restricted by the fact that it is hard to construct natural and relevant generic sentences for some of these domains, e.g. action verbs (Hemeren, 1996), adjectives denoting feelings (Bush, 1973), or theoretical concepts in computer science (Adelson, 1985).

There is, however, an alternative way to test the theory directly. A research project is currently underway,[9] along the following lines. Suppose we created a completely novel domain—say, the nonexistent animal, *zarg*. Suppose we drew pictures of zargs on a page, and we draw them already partitioned into groups, by visual cues (such as location on the page, color, etc.). Now, we could draw the zargs so that some of them have a certain property (say, horns) whereas others do not. We could then test the subjects' judgments of generic sentences about this novel domain, and see whether the sentences that violate homogeneity are rejected, and the ones that do not are accepted. Such an experiment would provide a direct test of the theory, and would allow us maximal freedom in manipulating the structure of the domain, so as to affect its homogeneity.

[9] In joint work with Susan Gelman.

References

Adelson, B.: 1985, 'Comparing Natural and Abstract Categories: A Case Study from Computer Science'. *Cognitive Science* **9**, 417–430.

Barclay, J. R., J. D. Bransford, J. J. Franks, N. S. McCarrell, and K. Nitsch: 1974, 'Comprehension and Semantic Flexibility'. *Journal of Verbal Learning and Verbal Behavior* **13**, 471–481.

Bush, L. E.: 1973, 'Individual Differences Multidimensional Scaling of Adjectives Denoting Feelings'. *Journal of Personality and Social Psychology* **25**(1), 50–57.

Carlson, G.: 1977, 'Reference to Kinds in English'. Ph.D. thesis, University of Massachusetts at Amherst. Also published 1980, New York: Garland.

Chierchia, G. and S. McConnell-Ginet: 1990, *Meaning and Grammar*. Cambridge, Massachussetts: MIT Press.

Cohen, A.: 1996, 'Think Generic: The Meaning and Use of Generic Sentences'. Ph.D. thesis, Carnegie Mellon University. Published 1999, Stanford: CSLI.

Cohen, A.: 1997, 'Generics and Default Reasoning'. *Computational Intelligence* **13**(4), 506–533.

Cohen, A.: 1999, 'Generics, Frequency Adverbs and Probability'. *Linguistics and Philosophy* **22**, 221–253.

Cohen, A.: 2001, 'On the Generic Use of Indefinite Singulars'. *Journal of Semantics* **18**(3), 183–209.

Fisher, R. A.: 1959, *Statistical Methods and Scientific Inference*. Edinburgh: Oliver and Boyd, second edition.

Hemeren, P. E.: 1996, 'Frequency, Ordinal Position and Semantic Distance as Measures of Cross-Cultural Stability and Hierarchies for Action Verbs'. *Acta Psychologica* **91**, 39–66.

Henley, N. M.: 1969, 'A Psychological Study of the Semantics of Animal Terms'. *Journal of Verbal Learning and Verbal Behavior* **8**, 176–184.

Horn, L. R.: 1972, 'On the Semantic Properties of Logical Operators in English'. Ph.D. thesis, University of California, Los Angeles.

Krifka, M., F. J. Pelletier, G. Carlson, A. ter Meulen, G. Link, and G. Chierchia: 1995, 'Genericity: an Introduction'. In: G. Carlson and F. J. Pelletier (eds.): *The Generic Book*. University of Chicago Press, pp. 1–124.

Lowe, E.: 1991, 'Noun Phrases, Quantifiers and Generic Names'. *The Philosophical Quarterly* **41**(164), 287–300.

Montague, R.: 1973, 'The Proper Treatment of Quantification in Ordinary English'. In: J. Hintikka, J. Moravcsik, and P. Suppes (eds.): *Approaches to Natural Language*. Dordrecht: Reidel.

Murphy, G. L. and E. J. Wisniewski: 1989, 'Categorizing Objects in Isolation and in Scenes: What a Superordinate Is Good For'. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **15**(4), 572–586.

Osherson, D. N., E. E. Smith, O. Wilkie, A. López, and E. Shafir: 1990, 'Category-Based Induction'. *Psychological Review* **97**(2), 185–200.

Pelletier, F. J. and N. Asher: 1997, 'Generics and Defaults'. In: J. van Benthem and A. ter Meulen (eds.): *Handbook of Logic and Language*. Amsterdam: Elsevier, pp. 1125–1177.

Rifkin, A.: 1985, 'Evidence for a Basic Level in Event Taxonomies'. *Memory and Cognition* **13**(6), 538–556.

Rosch, E., C. B. Mervis, W. D. Gray, D. M. Johnson, and P. Boyes-Braem: 1976, 'Basic Objects in Natural Categories'. *Cognitive Psychology* **8**, 382–439.

Rosenberg, S. and M. J. Kim: 1975, 'The Method of Sorting as a Data-gathering Procedure in Multivariate Research'. *Multivariate Behavioral Research* **10**, 489–502.

Sadler, D. D. and E. J. Shoben: 1993, 'Context Effects on Semantic Domains as Seen in Analogy Solution'. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **19**(1), 128–147.

Sattath, S. and A. Tversky: 1977, 'Additive Similarity Trees'. *Psychometrika* **42**(3), 319–345.

Schubert, L. K. and F. J. Pelletier: 1987, 'Problems in the Representation of the Logical Form of Generics, Plurals, and Mass Nouns'. In: E. LePore (ed.): *New Directions in Semantics*. London: Academic Press, pp. 385–451.

Shoben, E. J.: 1976, 'The Verification of Semantic Relations in a Same-Different Paradigm: An Asymmetry in Semantic Memory'. *Journal of Verbal Learning and Verbal Behavior* **15**, 365–379.

Shoben, E. J. and B. H. Ross: 1987, 'Structure and Process in Cognitive Psychology Using Multidimensional Scaling and Related Techniques'. In: R. R. Ronning, J. A. Glover, J. C. Conoley, and J. C. Witt (eds.): *The Influence of Cognitive Psychology on Testing*. Hillsdale: Lawrence Erlbaum Associates, pp. 229–266.

Smith, E. E.: 1995, 'Concepts and Categorization'. In: E. E. Smith and D. N. Osherson (eds.): *Thinking*, Vol. 3 of *An Invitation to Cognitive Science*. Cambridge, Massachusetts: MIT Press, 2nd edition, pp. 3–33.

Smith, E. E. and D. L. Medin: 1981, *Categories and Concepts*. Cambridge, MA: Harvard University Press.

Tversky, A. and J. W. Hutchinson: 1986, 'Nearest Neighbor Analysis of Psychological Spaces'. *Psychological Review* **93**(1), 3–22.

Tversky, B. and K. Hemenway: 1984, 'Objects, PArts, and Categories'. *Journal of Experimental Psychology: General* **113**(2), 169–197.

Wish, M.: 1970, 'Comparisons among Multidimensional Structures of Nations based on Different Measures of Subjective Similarity'. *General Systems* **15**, 55–65.

Wish, M., M. Deutsch, and L. Biener: 1972, 'Differences in Perceived Similarity of Nations'. In: A. K. Romney, R. N. Shepard, and S. B. Nerlove (eds.): *Multidimensional Scaling*, Vol. 2: Applications. New York: Seminar Press, pp. 289–313.