Cognition 151 (2016) 80–95

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/COGNIT

Vagueness, graded membership, and conceptual spaces

Igor Douven

Sciences, Normes, Décision (CNRS), Paris-Sorbonne University, France

ARTICLE INFO

Article history:
Received 8 September 2014
Revised 5 March 2016
Accepted 8 March 2016
Available online 17 March 2016

Keywords:
Similarity
Conceptualisation
Vagueness
Graded membership
Conceptual spaces

ABSTRACT

This paper is concerned with a version of Kamp and Partee's account of graded membership that relies on the conceptual spaces framework. Three studies are reported, one to construct a particular shape space, one to detect which shapes representable in that space are typical for certain sorts of objects, and one to elicit degrees of category membership for the various shapes from which the shape space was constructed. Taking Kamp and Partee's proposal as given, the first two studies allowed us to predict the degrees to which people would judge shapes representable in the space to be members of certain categories. These predictions were compared with the degrees that were measured in the third study. The comparison yielded a test of the account of graded membership at issue. The outcome of this test was found to support the conceptual spaces version of Kamp and Partee's account of graded membership.

1. Introduction

Many if not most predicates in natural language are vague, meaning that they admit of borderline cases, that is, cases to which they neither fully apply nor fully fail to apply. This common characterization of vagueness raises the question of what it is for a predicate to apply only to *some extent*. Researchers have been tempted to answer this question in terms of a graded membership relation, for instance, as developed in fuzzy set theory (Zadeh, 1965). But it is widely acknowledged that such formal models must leave something to be desired as long as the notion of graded membership itself stands in need of clarification (Lindley, 2004).

Kamp and Partee (1995) propose to model graded membership in the context of prototype theory and to thereby explicate graded membership in psychologically realistic terms, specifically, in terms of similarity to prototypical instances of a predicate. However ingenious their account, it is by their own admission incomplete, inasmuch as the constraints that it seems reasonable to impose on a graded membership function are insufficient to guarantee uniqueness of that function.

Decock and Douven (2014) and Douven and Decock (in press) show how uniqueness can be obtained by embedding Kamp and Partee's proposal in the conceptual spaces framework (Gärdenfors, 2000, 2014; Nosofsky, 1987, 1988, 1989). As they also show, the resulting account of graded membership has clear empirical content. Three studies were conducted to put this account to the test. The first study aimed to determine the structure of a particular conceptual space, specifically a space for the representation of container-like objects such as bowls, vases, and pots. The second study was meant to gain insight into which elements of this space (if any) represent *typical* bowls, vases, and so on. In the conceptual spaces version of Kamp and Partee's account, this information is sufficient to calculate the degree to which any shape represented in the space belongs to a category of shapes whose typical instances are also represented in the space. The third study was aimed at eliciting people's actually held degrees of category membership for these shapes in order to compare those degrees to the predicted degrees obtained from the first two studies.

Kamp and Partee's proposal is, first and foremost, a formal semantics for languages with vague predicates. The results of our studies constitute evidence for the material adequacy of this semantics, which according to Tarski—the founding father of formal semantics—is one of two desiderata to be satisfied by any semantics.[1] It is to be noted, however, that as an account of graded membership, Kamp and Partee's proposal has a number of well-established rivals. From a psychological perspective, it will be especially interesting to know how the proposal fares in comparison with those rivals. This question will be taken up in the general discussion.

E-mail address: igor.douven@paris-sorbonne.fr

[1] For a semantics of a given language to be materially adequate, it must accord with the linguistic behavior of competent speakers of that language. Tarski's other desideratum is formal correctness, meaning that the semantics must be free of paradoxes. That the conceptual spaces version of Kamp and Partee's semantics is formally correct is proven in Douven and Decock (in press).

http://dx.doi.org/10.1016/j.cognition.2016.03.007
0010-0277/

2. Theoretical background

Philosophers have been concerned with both the metaphysics and the logic of vagueness, the main metaphysical question being what the nature of borderline cases is, and the main logical question being whether classical logic needs to be modified—and if so to what extent—if it is to elucidate inferential relationships among sentences of a language some of whose predicates are vague (see, e.g., Fine, 1975; Shapiro, 2006; Smith, 2008; Williamson, 1994).

By contrast, and not surprisingly, psychologists have been more concerned with studying the *cognitive mechanisms* putatively underlying vagueness. Because it is primarily *concepts* that are vague—vagueness of propositions or assertions can be understood as deriving from their involving vague concepts—psychological research on vagueness must start from some account of concepts. While the nature of concepts is a controversial issue, much of the best psychological work on vagueness has assumed some form of prototype theory (Hampton, 1998, 2007; Kamp & Partee, 1995). According to this theory, not all instances of a concept need be equally characteristic of it; some instances may be more characteristic of the concept than others, and the most characteristic instance(s) count(s) as the concept's prototype(s).

Prototype theorists regard prototypes as one of the two basic elements involved in how humans categorize the world; the second basic element is a function measuring similarity between items. Whether or not something is categorized as falling under a given concept depends on how similar the thing is to the concept's prototype(s). Because similarity is a graded notion—items can be more or less similar to each other—prototype theory might seem promising as a basis for developing an account of graded membership that permits items to fall under a concept to differing degrees, which would naturally explain borderline cases as cases with less-than-perfect membership. Osherson and Smith (1981) may have been the first to observe that, even assuming these elements, it is still far from straightforward to arrive at an account of graded membership. As they note, it cannot simply be that an item falls under a concept to a degree equal or proportional to its similarity to the concept's prototype(s). By way of counterexample, just consider that a robin is more similar to the bird prototype than a pelican is, yet there can be no doubt that both robins and pelicans are birds to the fullest degree.

2.1. Kamp and Partee on graded membership

Important progress on this matter was made in Kamp and Partee's (1995) attempt to formulate a semantics for vague terms on the basis of prototype theory. Their attempt starts by considering a language with possibly vague "simple predicates," that is, predicates with monolexemic expressions in English. They define a partial model $\mathfrak{M}$ for this language, consisting of a set of individuals (the universe of discourse) and an interpretation function defined on that set. This partial model lets each predicate divide the universe into three parts: a part containing the clear instances of the predicate—its *positive* extension—a part containing the clear non-instances of the predicate—its *negative* extension—and a part containing the remaining ("indeterminate") cases. In principle, each of the three parts may be empty; if the third part is empty, the predicate is crisp. This model serves as a basis for defining a partial truth predicate, where Fa is true/false/neither if and only if the interpretation function locates a among the clear instances of F/the clear non-instances of F/the remaining objects.

More relevant to our present concerns, the model is used as a basis for defining a notion of graded membership. Central to this definition is the so-called supermodel $\mathfrak{M}^*$ that Kamp and Partee construct by endowing $\mathfrak{M}$ with a class of completions, where for each simple predicate, a completion splits up the set of indeterminate cases of the predicate by grouping some with the predicate's clear instances and some with its clear non-instances. Importantly, Kamp and Partee do not consider all possible ways of splitting up the set of indeterminate cases but only those that respect similarity orderings, meaning that if a completion groups an indeterminate instance a of F with the clear instances of that predicate, then it should also group anything that is at least as similar as a to the F prototype(s) with the clear F instances.[2] The basic idea of their proposal is, then, to let the degree to which an item a falls under a concept F be given by the proportion of completions that group a with the clear F instances.

Apart from relying on similarity orderings, Kamp and Partee also argue for imposing some formal constraints on any graded membership function, such as that it should take values only in the interval $[0,1]$, and that the function's value for a given case should reflect the size of the set of completions that make that case come out as falling under the relevant concept (the larger the set, the higher the degree of membership should be). While, as Kamp and Partee show, any membership function that satisfies their constraints will be a normalized measure in the technical sense of measure theory, they admit that the constraints are not enough to fix a unique graded membership function: "[T]he constraints do not determine the [graded membership] function μ completely. Indeed, it is far from clear on what sorts of criteria a particular μ could or should be selected" (Kamp and Partee, 1995:153). Their paper leaves this "non-uniqueness problem" open.

Hampton (2007:367) is surely right to note, in commenting on Kamp and Partee's proposal, that for some concepts we should have no difficulty in fixing a unique membership function μ. As an example, he gives the concept of tallness (as pertaining to adults). For the sake of argument, suppose this has the interval $[170\text{ cm}, 180\text{ cm}]$ as its boundary region, in the sense that everyone shorter than 170 cm is clearly not tall, and everyone taller than 180 cm is clearly tall, while everyone whose height is in between constitutes a borderline case of tallness. Then a natural suggestion is that the relevant sets of completions are related to the ratios of this height scale, whence μ would take the value $x/10$ for any person who is $(170 + x)$ cm tall, with $x \in [0, 10]$. Hampton does note as a possible drawback of this suggestion that it yields a linear membership function, where on the grounds of previous empirical work (e.g., McCloskey & Glucksberg, 1978) one would expect an S-shaped membership function instead. But an arguably deeper problem with the suggestion is that it seems rather limited in its scope, given that it is unclear how it generalizes to concepts that fail to linearly order their domain of application.

2.2. The conceptual spaces framework

Decock and Douven (2014) solve the non-uniqueness problem that besets Kamp and Partee's account of graded membership by embedding it in the conceptual spaces framework that cognitive psychologists have developed over the past twenty years. In this framework, concepts are thought of geometrically, as regions of metrical spaces, which are one- or multi-dimensional structures with a distance metric defined on them. What makes these mathematical objects *conceptual* spaces is that their dimensions are supposed to represent fundamental qualities that objects may have to differing degrees. Depending on the values it assumes with respect to these dimensions, an object is mapped onto a specific point in the space. The distance between that point and a second point in the space onto which a different object is mapped is supposed to

[2] This approach builds on Fine's (1975) influential treatment of vagueness.