THE INHABITANTS OF VAGUENESS MODELS

Galit W. Sassoon
ILLC/ University of Amsterdam

Abstract

Gradable predicates apply to individuals to different extents. The ontology underlying this intuition consists of at least a domain of individuals, a domain of graded-properties (mappings of individuals to values), and a domain of 'property values' – measurable subparts of individuals – their heights, weights, happiness extents, etc. This paper explores the hypothesis that individuals are identified by the sum of their property values. Intuitions regarding the (im)plausibility of cross-index identity in different cases directly follow, as well as an intuitive distinction between vagueness and accidental ignorance.

1 Introduction

As part of the common practice of experimental research in cognitive psychology, objects are systematically represented by clusters of property values. Individuals correspond to points in an n-dimensional space, for some number n, where each dimension (axis) is some scalar property. For example, the set of possible individuals presented as stimuli in a given experiment can be represented with the 2-dimensional space generated from the scalar properties denoted by *red* and *long*, or from the dimensions *color* and *shape*, the latter seen as nominal-scale properties, assigning to entities values such as 'red', 'blue', 'square' and 'circle'. The result is a set of individuals including a red square, a blue square, etc. Many other examples can be found in, for instance, Murphy (2002).

The ontology underlying these experimental designs consists of different types of instances:

(1) a. A domain of *time points*, D_p
b. A domain of *individuals*, D_x (including 'stages', i.e. individuals at a given time p)
c. A domain of *property values* $D_r \subseteq D_x$ ('satisfaction extents'): measurable subparts of individuals – their heights, lengths, weights, colors, happiness extents, etc.[1]
d. A domain of *graded-properties*, $D_{<x,r>}$ (mappings of individuals to values)

Now let us consider the ontology underlying language and thought more broadly, i.e. the ontology which is the core of the interpretation of natural language expressions. Members of the domain of possible individuals within this ontology can also be characterized as points in an n-dimensional space, with intrinsic properties such as those listed in (1c-d) above forming the generating coordinates (height, length, temperature, color, happiness extents and so on). On this proposal the domain D_x is a theoretical construct, consisting of all the possible combinations of property-values which we humans can conceive of. Hence, the domain consists of 'possible' individuals. Upon encountering actual, real-world entities, we attempt to create a link between

[1] Back to Russell (1905), linguists assume that adjectives like *tall* map individuals to degrees. Clearly, *tall*'s degrees reflect the ordering between entities' heights, but little is agreed about them beyond this fact (Klein, 1991). The notion of degrees as property values (heights, lengths, etc.) resembles Moltmann's (2006) notion of 'tropes' and Kennedy's (2001) 'degree sorts'.

them and the corresponding elements of D_x; but no empirical observation of an actual, real-world entity can pinpoint one single individual in D_x, because only partial, proper subsets of the property values of worldly entities are accessible to us. Thus, observations only isolate a subset of the domain consisting of all those possible individuals that, for all that we know, might still be the observed entity, i.e. all those that share with the observed entity all the property values accessible to us.

This paper explores the predictive fruitfulness of this hypothesis, namely the hypothesis that:

(2) The elements of the domain of possible individuals D_x are identified by and only by their property values.

Methodologically speaking, this study purports to contribute to the general enterprise of bridging the gap between researches in linguistics and psychology, by aiming towards a unified conception of individuals. However, more significantly, this study purports to shed new light on central problems much discussed in the philosophical literature, problems pertaining to cross world identity and to vagueness versus ignorance.

In a nutshell, an important consequence of the proposal that individuals are identified by the sum of their intrinsic property values is that it directly predicts a certain range of our ontological or metaphysical intuitions concerning cross world identity.

On the one hand, intuitively, an individual who is, for instance, six feet tall in this world is not the same person in an alternate world (bound to the same time) in which his height is different. This intuition directly follows. The property values of an individual per a time point are index invariant (they do not vary across worlds), for one individual in D_x cannot be identified with two different sums of property values.

On the other hand, as discussed in part 2, if the property values of two individuals are the same across alternative worlds (bound to the same time) the individuals are the same. They are the same regardless of how these property values categorize under vague predicates, e.g. even if six feet tall counts as *tall*, in this world, but not in the alternate realities, and so the given individual is *tall* in this world and is *not tall* in the alternate realities. This can be viewed as an instance of Leibniz's law; two individuals with all the same property values must be identical. The crux of this proposal pertains to vague predicates; membership in their denotations plays no role at all in distinguishing between individuals (Sassoon, 2007; Frank Veltman, pc). It follows that their only use is to blur differences between many different property values so as to form secondary, coarse-grained and therefore general categories, e.g. the whole dense set of possible heights can collapse into general categories such as 'tall', 'short', and 'borderline'.

As discussed in part 3, the proposal in (2) implies a departure from the Kripkean perspective on the reference of singular terms in that, according to (2), the referent of a noun like *Aristotle* is uniquely specified in a context c if and only if Aristotle's sum of property values is entirely specified in c. This never happens. Speakers, as well as language communities, never have access to the entire set of property values of a proper name's referent. For example, if all except the referent's height is specified, many possible individuals, exactly alike except differing in height, might form the unique referent of, e.g., *Aristotle*. Thus, the referent varies across the worlds consistent with c.

A benefit of this departure concerns how to represent truth value gaps caused by vagueness versus those caused by accidental ignorance within a supervaluationist semantics. In practice, only a subset of the property values of the referents of singular terms are known in actual

contexts and therefore we do not have a unique individual to refer to, but instead have sets of possible individuals. Consequently, terms may create truth value gaps in otherwise non-vague statements like, for instance, *Aristotle is taller than Plato*. This is a clear case of ignorance about property values, for we do not know the heights of these famous people. While the ordering of individuals by their heights is fixed (world-invariant), we do not know exactly which possible individuals in D_x are being compared. For all that we know, many possible individuals exactly alike except in height may form the unique referent of, e.g., *Aristotle*.

Another source of truth value gaps is the absence of linguistic conventions regarding the spatial borders of referents of terms. Consider, for example, statements of the form *Mount Etna occupies n square meters* (cf., Lewis, 1988, 1993). While the size of an individual is world-invariant, *Mount Etna* may refer to many possible partially overlapping regions; thus, no convention tells us precisely which individual in D_x is to be measured. Finally, in some cases, when the existence of a linguistic convention uniquely specifying the referent is debatable, the question remains open, whether personal ignorance or semantic indeterminacy is the case. Different conventions for defining the region of a given city or state x may be held by people with equal authority regarding the interpretation of x in the given language. Different conventions may prevail in court, in novels, and in other daily life situations. On the present proposal, vagueness and accidental ignorance are represented in one and the same model, as discussed in part 3. This helps dealing with the given examples. At the same time, the differences between the two phenomena, when such exist, are captured.

The main consequences of the present study are therefore tightly related to vagueness. Let us call models for the interpretation of vague predicates *vagueness models*, including in particular the kind of models used by *supervaluationist* theories (to be described shortly). Vague predicates (like *tall, heavy,* and *happy*) are characterized by the absence of sharp boundaries. Some individuals exist, for which we cannot tell whether they fall under the predicate or not. They form a denotation gap. Vagueness models, then, include *partial information states* ('contexts') c, representing the common knowledge of possible communities of speakers (Stalnaker 1978). The interpretation of linguistic expressions in contexts c is modeled via a set of indices T_c, the worlds (Stalnaker 1978) or completions (also called *precisifications*; van Fraassen, 1969; Fine, 1975; Kamp, 1975) consistent with c. Worlds are alternative universes or realities, while completions are information states, which are 'classical' in the sense that every statement is either true or false. Completions may include more information than worlds do; for instance, each completion determines cutoff points for vague predicates, although this type of information may well not be part of the actual world (Kamp, 1975). This paper refers mainly to completions, but where differences between completions and worlds are relevant, they are explicitly discussed. Truth of a statement S in c is defined based on these indices:

(3) a. S is true in c if and only if S is true in every $t \in T_c$;
b. S is false in c if and only if S is false in every $t \in T_c$, and
c. S's truth value is undetermined otherwise.

Vague predicates are associated with graded properties (Kennedy, 1999). Intuitively, individuals can satisfy these properties to different extents, and they fall under the predicates if and only if their satisfaction extents (height, health, happiness extent, etc.) exceed the predicates' contextual cutoff point ('*standard for membership'*). Though satisfaction extents usually vary

with time, this paper does not aim to deal with issues specifically pertaining to changes over time. Thus, reference to time indices is omitted.

The graded property of some predicates, especially multidimensional ones, is indeterminate, i.e. completion-dependent (Klein, 1980). This is manifested by indeterminacy with regard to the interpretation of the derived comparative of the adjective in question. For example, unlike the interpretation of the comparative *taller*, the interpretation of, for example, *tastier* or *more similar to Dan* is indeterminate. A predicate like *tall* always assigns a given individual, x_1, the same value per a time point (x_1's height). By contrast, a predicate like *similar to Dan* may assign x_1 different values in different completions. Crucially, the reason for this is not that x_1 changes from one completion to another, for it doesn't. Rather, it is the graded property associated with *similar to Dan* that changes. The value of x_1 in *similar to Dan* in a given completion depends on what parts of x_1 (which of x_1's property values) are taken into account in similarity comparisons (general look, hair color, facial expressions, character, etc.) Hence, formally:

(4) a. Let f be a function associating each $t \in T_c$ and predicate P with a *graded property*, f(P,t), i.e., a function from individuals $x \in D_x$ to values $r \in D_r$.
b. P is true of an individual $x \in D_x$ in t if and only if x's value in P, f(P,t)(x), exceeds P's standard in t.
c. x's value in P is determined in c if and only if it is completion-invariant: $\forall t_1,t_2 \in T_c$, $f(P,t_1)(x) = f(P,t_2)(x)$.

For example, two completions t_1 and t_2 may differ along the graded property of a predicate like *beautiful*, $f(\text{beautiful},t_1) \neq (f(\text{beautiful},t_2)$, representing different interpretations. The values of individuals in *beautiful* are determined (index-invariant) if and only if the context uniquely selects one possible interpretation (one measure of beauty) for this domain of individuals. Again, the individuals' sums of property values remain fixed across completions, for, as stated in (1c), property values are simply subparts of individuals; two individuals in two different worlds cannot have different subparts yet be one and the same (cf., Lewis, 1986). It is the interpretation (graded property) of *beautiful* that varies, taking different aspects of an individual into account across different completions.

These notions generalize to expressions other than one-place predicates. For example, a context c may fail to determine whether some individuals, x_1 and x_2 stand in a relation like *love*, perhaps because they do not possess much love to one another, or perhaps because in c we do not yet know what *love* is – that is, which property values the graded property of *love* represents. Then in c the pair $<x_1,x_2>$ falls in the gap of the two-place predicate *love*. Thus, in a vagueness model, a vague predicate's graded property, as well as its denotation (set of instances) may vary across indices of evaluation.

To recapitulate, this paper studies the consequences of the proposal that individuals are identified by and only by their property values. Part 2 explores our intuitions concerning whether, in different cases, individuals inhabiting different indices can be identified. Part 3 aims to capture the distinction between vagueness and accidental ignorance while representing the two phenomena within one and the same model. Part 4 briefly addresses additional implications of the model developed in this paper, pertaining to demonstratives, identity statements and the interaction of terms with modals.