A Theory of Vague Adjectives Grounded in Relevant Observables

Brandon Bennett
School of Computing,
University of Leeds,
Leeds LS2 9JT, England
brandon@comp.leeds.ac.uk

Abstract

The ambiguity and vagueness of natural language vocabulary is one of the biggest obstacles to formalising reasoning systems that can operate with natural language concepts. The paper presents a formal theory of the logic of vague adjectives, which is based on two key ideas: 1) a sharp distinction between count noun concepts and adjectival concepts is build into both syntax and semantics; 2) valid inferences with vague adjectives are explained by means of a theory of the relevance of particular observables to the applicability of any given adjective.

The theory is presented as a first-order formalism in which conceptual elements are combined by means of a small number of primitive logical predicates, which are axiomatised and given a set-theoretic semantics. The relationship of the proposed formalism to existing approaches using fuzzy logic or supervaluationistic modal semantics is also considered. A possible solution to the *sorites* paradox is also outlined.

Introduction

The ambiguity and vagueness of natural language vocabulary is one of the biggest obstacles to the construction of robust general-purpose ontologies. Because of these phenomena, it is extremely difficult to give precise specifications of concept meanings that are acceptable to a wide community of users.

In the current paper I develop a logical language which I call Vague Adjective Logic (**VAL**). The purpose of **VAL** is not to model natural language in all its peculiar details, but rather to provide a formal theory that captures the essential logic and semantic function of vague adjectives. A theory of how these expressions work is important to many AI applications, especially those that must interface in some way with human descriptions of the world.

The Nature of Vagueness

Many natural language terms are vague because of an indeterminate threshold of applicability in one or more relevant properties. For example: tall, red, cup, mountain. We know that 'tallness' of a person is dependent on their height, but there is no definite height measure at which someone becomes tall. This kind of vagueness is known as *sorites* vagueness, from the Greek word meaning 'heap'. The question of how many grains of sand, make a heap gives rise some deep philosophical puzzles. In particular, the seemingly true statement that "if one removes just one grain of sand from a heap, one still has a heap," leads to the absurd conclusion that even if we remove all the sand we still have a heap. In the geographic domain vague senses of concepts lead to indeterminism in the extensions of features such as forests and mountains (Bennett 2001; Cohn & Gotts 1996; Burrough & Frank 1996; McGee 1997; Roy & Stell 2001).

Existing Approaches to Modelling Vagueness

A general-purpose knowledge representation language must provide apparatus for dealing with sorites vagueness. The best known approach is *Fuzzy Logic* (Goguen 1969; Zadeh 1965; 1975; Dubois & Prade 1988), which takes a kind of probabilistic approach to the applicability of concepts. Fuzzy Logic has been very successful in certain kinds of application but it does have some problems. One is that it does not provide an objective way to assign numbers indicating the degree to which a given object satisfies a given vague predicate; the numbers have to be assigned on a rather *ad hoc* basis. A second problem is that it is inferentially quite weak: fuzzy inference rules make it impossible to derive definite conclusions on the basis of vague information, whereas humans do this all the time. For instance from 'The mountain is far from the sea; and my house is beside the mountain' we would infer that 'My house is far from the sea'. But in fuzzy logic the entailment would only hold up to some degree. Several articles attacking or defending fuzzy logic inference are collected in (Elkan 1994).

An alternative but less well know alternative is *supervaluation semantics* (Fine 1975), a framework within which one can accommodate multiple senses of a concept within a single formal specification. This approach also enables one to specify logical constraints which apply either to all senses or to a restricted range of senses of a concept. One way of bringing the semantical richness of supervaluation semantics into the object language is to employ a modal logic (Bennett 1998).

Semantic approaches of this kind have so far been little applied to Knowledge Representation for AI. Notable exceptions are (Bennett 2002), which considers how vagueness affects the identity conditions of physical objects and (Halpern 2004) which considers how vagueness affects perceptual observations of continuously varying properties.

Parameterised Meanings and Relevant Observables

Supervaluation theory can be seen as modelling sorites vague concepts as having a meaning that varies according to one or more parameters, whose values are not specifically stated. In a given context, assertions which involve semantically related concepts are taken as sharing the same parameter-space and hence varying in a mutually consistent way. Valid inference can then be modelled in terms of entailments that hold whatever the choice of these underlying indeterminate parameters. A fundamental idea of the proposed theory is that the ascription of a vague adjective a judgement that sets arbitrary thresholds for a number of *relevant objective measures*.

Within any given context I assume that a language user employs a choice of thresholds that is consistent across usages of all concepts. Thus, where two or more concepts have some semantic interdependence, this will be maintained by consistent usage of thresholds. For example 'tall' and 'short' are (in a particular context) mutually exclusive and are dependent on the objective observable of height. Thus the height threshold above which a person to be considered 'tall' must be greater than the height threshold below which a person is considered 'short'. Note that as well as depending on a particular context in which the terms are used, there is also an explicit dependence of these thresholds on the type of thing to which the terms are applied; i.e. in this example they are applied to a person.

To describe the collection of threshold values that affect a judgement involving vague concepts, I adopt the word *standpoint* as a theoretical term. A standpoint is a semantic object associated with a complete assignment of threshold values to parameters associated with the vague vocabulary of a formal language. Standpoints, can be regarded as *indices* relative to which a vague assertion is interpreted. Thus, they are similar to the possible worlds in Kripke semantics for modal logics. In fact they can be seen as an elaboration of the idea of a *precisification* that is often taken as indexing propositions in supervaluation semantics. The difference is that standpoints have an internal structure, which explicitly describes the dependence vague propositions on a particular choice of threshold judgements.

Contextual Relativity of Vague Adjectives

A widely recognised problem that affects many vague adjectives is that the criteria for their ascription may strongly depend upon the type of thing to which they are applied. For instance a 'large man' may be much smaller in absolute terms than a 'large elephant' (or even a relatively small elephant). Because of this we cannot simply treat both adjectives and count nouns as predicates. Hence, the present theory incorporates explicit representation of the combination of adjectives with count nouns.

Although dependence on count nouns is a key feature of many uses of vague adjectives, there are also types of proposition where properties of different kinds of thing are compared in an absolute way. For example, we may say that 'This elephant is taller than that man'. In such a case we do not normally relativise the idea of tallness to the different types of animal involved. The formal language presented here also supports representation of assertions involving this kind of absolute comparison.

Overview of VAL Syntax

Count Nouns and Adjectives

Although the language that will be presented is essentially a variant of first-order logic, it does have quite significant differences in its syntax. In particular, I make a clear syntactic distinction between count nouns ('woman', 'giraffe', 'table', 'pen', etc.) and adjectives ('tall', 'happy', 'heavy', etc.). In the following paragraphs I will explain why this distinction is necessary. In my opinion, treating these two types of concept as essentially the same is a fundamental shortcoming in the way that logical languages have been applied to modelling natural language concepts.

In standard first-order logic, both count nouns and adjectives are treated as predicates. Thus, 'Jane is a tall woman' would be represented by something line $\mathsf{Tall}(jane) \wedge \mathsf{Woman}(jane)$. and similarly 'Gerald is a tall giraffe' would be written as $\mathsf{Tall}(gerald) \wedge \mathsf{Giraffe}(gerald)$. But this leads to a problem: 'Tall' being a predicate will, in the usual first-order semantics, correspond to a set of objects. In order for both of these propositions to be true, both Jane and Gerald must be members of the set of 'tall things'. But of course even a small giraffe will be much taller than a tall woman. Thus, the criterion for membership of the set of tall things cannot depend merely on height, it must also depend on the type of thing in question — in other words, it will depend on the count noun of which the object is an instance. But the difficulty comes when we realised that any given object can be considered as an instance of many different count nouns. For example, Gerald is certainly an animal as well as a giraffe. Now the real problem comes when we want to make sense of a statement 'Gerald is a tall animal but he is not a tall giraffe'. This cannot be consistently represented using the standard first-order treatment of adjectives, since it can only be true if Gerald is both a member of the set of tall things and not a member of this set.

A solution to this problem can be obtained by treating adjectives not as classical predicates, but rather as operators modifying the meaning of count nouns. Let a be a vague adjective and $\mathbf{c}$ a category of object (which would normally correspond to a natural language count noun). The ascription of a to x, considered as an instance of $\mathbf{c}$ can be written as a complex predication of the form $[\mathsf{a} \circ \mathbf{c}](\mathbf{x})$. Here the adjective a is relativised to the category $\mathbf{c}$. For example 'John is a tall man' would be represented as $[\mathsf{tall} \circ \mathbf{man}](john)$. This notation indicates the compositional semantics of adjectives. According to this analysis, the meaning of the adjective operates on the count noun in order to produce a complex count noun, which is a restricted or otherwise modified version of the original. It is this modified count noun which is then applied as a predicate to a named object.

However, for computational manipulation (e.g. use of theorem provers) it is convenient reduce this notation to a standard first-order form. In order to do this I reify both count nouns and adjectives, so that one can write

$$\mathsf{Isa}(x, \mathsf{a}, \mathbf{c}) \ .$$

For instance, $\mathsf{Isa}(john, \mathsf{tall}, \mathbf{man})$. Here, the '$\mathsf{Isa}$' predicate is just used to glue the various reified components together.

Since, count nouns have been reified, their use as predicates will also need to be modified. Instead of writing $\mathbf{c}(\mathbf{x})$, to mean x is an instance of count noun $\mathbf{c}$, I write $\mathsf{Isa}(x, \mathbf{c})$.

Restrictive *vs* Non-Restrictive Adjectives The meaning of most adjectives is such that, when applied to a count noun, the resulting modified count noun applies to a subset of the objects to which the original count noun applied. For instance, the set of 'tall men' is a subset of the set of 'men'. These will be called restrictive adjectives, and I will use the logical predicate $\mathsf{Restrictive}(\mathsf{a})$, to assert that adjective 'a' is restrictive. The meaning of this predicate can be defined as follows:

D1) $\mathsf{Restrictive}(\mathsf{a}) \ \equiv_{\text{def}} \ \forall xc[\mathsf{Isa}(x, \mathsf{a}, c) \rightarrow \mathsf{Isa}(x, c)]$

By contrast, non-restrictive adjectives are those such that $\mathsf{Isa}(x, \mathsf{a}, \mathbf{c})$ can hold, when $\mathsf{Isa}(x, \mathbf{c})$ is not true. For instance, we can have

$\mathsf{Isa}(quacky, \mathsf{rubber}, \mathbf{duck}) \wedge \neg\mathsf{Isa}(quacky, \mathbf{duck})$.

In the current work, the focus is on restrictive adjectives, since these are the ones for which relatively strong inferential rules hold. However, the distinction is made explicit so that non-restrictive adjectives can be employed without fear of incorrect inferences. Inferences involving non-restrictive adjectives would require special axioms tailored to their specific meanings.

Absolute *vs* Graded Adjectives A second classification of types of adjectives is the distinction between those that are either true or false for a given object (relative to the count noun by which it is referred) and those that hold to varying degrees for different objects. The latter may be called *graded* (or simply *vague*). The classical semantics for predicates is well suited to absolute (restrictive) adjectives but provides no mechanism for interpreting graded adjectives. The focus of this paper is on adjectives that are restrictive and graded and these will be specified using the logical predicate $\mathsf{ResGrad}(a)$.

Adjectives and Relevant Observables

A key idea of the present theory is that the use of any given vague adjective can be explained in terms of a number (usually a small number) of relevant observables. For example, in deciding whether a person is 'tall', the relevant observable is their height. In **VAL** the relevance between an observable and an adjective is formalised as a primitive logical relation. Thus we may write $\mathsf{Rel}(\mathbf{height}, \mathsf{tall})$.

In some cases we may want to specify that the applicability of a particular vague adjective depends only on a single observable. For instance we might state $(\forall o)[\mathsf{Rel}(o, \mathsf{tall}) \leftrightarrow (o = \mathbf{height})]$. More generally, there may be a number of observables relevant to a given adjective.

For natural language adjectives it may not always be completely clear which observables are relevant. This is because of a deep conceptual ambiguity in certain natural terms. Considering this in detail would go beyond the scope of the present work. However, there are plenty of examples where the relevant observables are fairly clear. And moreover, the formalism presented here will allow one to do useful reasoning without assuming full knowledge of all instances of the Rel relation.

Observables and their Relative Magnitudes The logic of observables and adjectives will be grounded by associating each observable o with a relationship describing the relative magnitude of an observable in relation to a pair of individuals. This will be represented by the primitive relation $\mathsf{Obgeq}(o, x, y)$, which states that individual x exemplifies the observable o to a degree greater than or equal to individual y.

Thus, for any given observable, the relation $\mathsf{Obgeq}(o, x, y)$ specifies a preorder over the domain of individuals — i.e. it satisfies the following axioms:

A1) $\forall ox[\mathsf{Obgeq}(o, x, x)]$
A2) $\forall oxyz[(\mathsf{Obgeq}(o, x, y) \wedge \mathsf{Obgeq}(o, y, z)) \rightarrow \mathsf{Obgeq}(o, x, z)]$

For many common observables (those that may be regarded as one dimensional) the corresponding ordering is total. Hence, I add an auxiliary axiom

A3) $\forall oxy[\mathsf{Total}(o) \rightarrow (\mathsf{Obgeq}(o, x, y) \vee \mathsf{Obgeq}(o, y, x))]$

It will often be useful to employ the relations asserting that an individual exemplifies an observable to an equal degree or to a strictly greater degree than another. Thus I define:

D2) $\mathsf{Obeq}(o,x,y) \equiv_{\mathrm{def}} \mathsf{Obgeq}(o,x,y) \wedge \mathsf{Obgeq}(o,y,x)$

D3) $\mathsf{Obg}(o,x,y) \equiv_{\mathrm{def}} \mathsf{Obgeq}(o,x,y) \wedge \neg\mathsf{Obgeq}(o,y,x)$

Comparatives An important aspect of the meaning of adjectives is their relationship to associated comparative relations. Thus, the adjective 'tall' is intimately connected with the comparative 'taller'. It is interesting to note that even when we are dealing with vague adjectives, such as 'tall', the corresponding comparative 'taller' is more or less precise (leaving aside certain possible quibbles about how height should be measured). Hence, the logic of comparatives is very significant in reasoning about vague adjectives.

The essence of the theory presented here is that adjectival comparative relations are determined by the orderings of the values of relevant observables. In the case of adjectives that have just a single relevant observable, the correspondence is simple, as the comparative is determined directly by the observable ordering. But, more generally, several observables may be involved. For instance, speed, agility, strength and stamina might all be considered relevant to being athletic. In such cases we may have pairs of individuals for which their relative ordering with respect to one relevant observable is opposite to their ordering with respect to another (one person may be faster than another but have less stamina). Here the adjectival comparative is not straightforwardly determined by the observables. In terms of abstract semantics, the ordering can be modelled by an arbitrary compound function of the different observables, satisfying appropriate monotonicity conditions (which will be specified later in the section on Semantics). These conditions correspond to axiomatic constraints, which determine the dependence of the adjectival comparative relations on the observable orderings, via the relevance relation, Rel.

The primitive comparative relation of the theory is $\mathsf{Atleastas}(a,x,y)$, which means that object x exemplifies the adjective a to at least as great a degree as object y. This relation satisfies the following axioms:

A4) $\forall axy[\forall o[\mathsf{Rel}(o,a) \rightarrow \mathsf{Obgeq}(o,x,y)] \rightarrow \mathsf{Atleastas}(a,x,y)]$

A5) $\forall axy[(\forall o[\mathsf{Rel}(o,a) \rightarrow \mathsf{Obgeq}(o,x,y)] \wedge \exists o[\mathsf{Rel}(o,a) \wedge \mathsf{Obg}(o,x,y)]) \rightarrow \neg\mathsf{Atleastas}(a,y,x)]$

A6) $\forall axy[(\mathsf{Atleastas}(a,x,y) \wedge \neg\mathsf{Atleastas}(a,y,x)) \rightarrow \exists o[\mathsf{Rel}(o,a) \wedge \mathsf{Obg}(o,x,y)]]$

A7) $\forall axy[(\mathsf{Atleastas}(a,x,y) \wedge \mathsf{Atleastas}(a,y,z)) \rightarrow \mathsf{Atleastas}(a,x,z)]$

A4 ensures that if an object's observable values for all observables relevant to a given adjective are as least as great the relevant values of another object, then the first object satisfies the adjective to at least the same degree as the second. **A5** ensures that if an object's relevant observables are at least as great as those of another and in at least one case are greater than the other, then the other object cannot satisfy the object to at least as great a degree as the second.[1] (This means that the first object satisfies the adjective to a strictly greater degree than the second — see definition of $\mathsf{More}(a,x,y)$ below.) **A6** requires that if one thing exhibits an adjective to a strictly greater degree than another it must exhibit at least one relevant observable to a greater degree than the other. Finally **A7** imposes the condition of transitivity on the Atleastas relation.

It should be noted that although these axioms place tight constraints on the Atleastas relation, in terms of the relative magnitudes of relevant observables, these relative magnitudes do not fully determine the extension of Atleastas.

It is easy to define additional relations to express the condition that one objects exhibits an adjective to a greater degree than another, or that two objects equally satisfy the adjective.

D4) $\forall axy[\mathsf{Equally}(a,x,y) \leftrightarrow (\mathsf{Atleastas}(a,x,y) \wedge \mathsf{Atleastas}(a,y,x))]$

D5) $\forall axy[\mathsf{More}(a,x,y) \leftrightarrow (\mathsf{Atleastas}(a,x,y) \wedge \neg\mathsf{Atleastas}(a,y,x))]$

The 'More' relation gives us a convenient means to define relations that correspond closely to natural language comparatives. For instance,

$$\mathsf{Taller}(x,y) \equiv_{\mathrm{def}} \mathsf{More}(\mathsf{tall},x,y) .$$

Axioms Relating Adjectives and Comparatives Now that we have introduced the adjectival comparative relations, we can use these to formulate some key inferential principles that are used in reasoning about restrictive graded adjectives.

A8) $\forall acxy[(\mathsf{ResGrad}(a) \wedge \mathsf{Isa}(x,a,c) \wedge \mathsf{Isa}(y,c) \wedge \mathsf{Atleastas}(a,y,x)) \rightarrow \mathsf{Isa}(y,a,c)]$

A9) $\forall acxy[(\mathsf{ResGrad}(a) \wedge \mathsf{Isa}(x,a,c) \wedge \mathsf{Isa}(y,c) \wedge \neg\mathsf{Isa}(y,a,c)) \rightarrow \mathsf{More}(a,x,y)]$

Informally **A8** says that if object x is considered as satisfying a given (graded) adjective and a second object y is of the same kind and also satisfies that adjective to the same or greater degree, then y also satisfies that adjectives. (Note that the ascription of the adjective is — as always in this theory — relative to the count noun under which the objects are described. So the axiom requires that both objects are described in terms of the same count noun.)

A9 concerns the case where we have two objects of the same kind and one satisfies a graded adjective, whereas the other does not. This implies that the first object exemplifies the adjective to a greater degree than the second.

[1] It is possible that one may wish to drop this axiom in certain cases. For instance if a two objects are equal in all observables relevant to an adjective apart from small differences, one may wish to allow them to be considered as equally exemplifying the adjective.

Converse Observables Adjectives often come in pairs of polar opposites (e.g. 'tall' and 'short'). Here the same observable(s) is (are) relevant to both adjectives; however, the sense in which the magnitude of the observable(s) contributes to the justification for ascribing the adjective is reversed. Thus to state semantic relationships between adjectives we will often want to refer to this *converse* relationship. The converse of an observable is expressed by the function $\mathsf{conv}(x)$.[2] Thus we can write

$$\mathsf{conv}(\mathsf{tall}) = \mathsf{short}\ .$$

The converse relationship between adjectives corresponds to an analogous relationship between observables. Hence, I overload the conv function an allow it to also operate on observables; so that $\mathsf{conv}(\mathbf{flexibility})$ is an observable whose order relation is the inverse of the order relation associated with **flexibility**. Hence, we might write

$$\mathsf{conv}(\mathbf{flexibility}) = \mathbf{rigidity}\ .$$

When an observable gives rise to an ordering, it seems that we often have a natural preference to associate a particular direction to that ordering. For instance, in ordering objects in terms of their vertical extension, the natural ordering is in terms of increasing height, not 'shortness'. However, for the purpose of formalising the theory, it is much more convenient if all observables have a converse, even if this doesn't necessarily correspond to a common natural language term.

The logic of the converse operator is specified by the following axioms:

A10) $\forall x[\mathsf{conv}(\mathsf{conv}(x)) = x]$
A11) $\forall oa[\mathsf{Rel}(o, a) \leftrightarrow \mathsf{Rel}(\mathsf{conv}(o), \mathsf{conv}(a)]$
A12) $\forall axy[\mathsf{Obgeq}(a, x, y) \rightarrow \mathsf{Obgeq}(\mathsf{conv}(a), y, x)]$
A13) $\forall xac[\neg(\mathsf{Isa}(x, a, c) \wedge \mathsf{Isa}(x, \mathsf{conv}(a), c))]$

A10 makes the conv function self-inverse. **A11** means that the observables relevant to an adjective are exactly the converse observables of those relevant to the converse adjective. **A12** ensures that converse observables induce opposite orderings on the domain of objects. From **A11**, **A12** and the axioms for $\mathsf{Atleastas}$ it follows that the $\mathsf{Atleastas}$ relation runs in opposite directions for converse adjectives:

$$\forall axy[\mathsf{Atleastas}(a, x, y) \rightarrow \mathsf{Atleastas}(\mathsf{conv}(a), y, x)]$$

A13 is a further important axiom, which ensures that no object can exemplify both an adjective and its converse.

Adjectival Modifiers The proposed language also enables the modifiers 'very' and 'quite' to be applied to adjectives as functional operators. Thus we can write, for example, $\mathsf{Isa}(\mathit{gerald}, \mathsf{very}(\mathsf{tall}), \mathbf{giraffe})$ or $\mathsf{Isa}(\mathit{john}, \mathsf{quite}(\mathsf{clever}), \mathbf{man})$. The very operation is governed by the following axioms:

A14) $\forall acx[\mathsf{Isa}(x, \mathsf{very}(a), c) \rightarrow \mathsf{Isa}(x, a, c)]$
A15) $\forall acxy[(\mathsf{ResGrad}(a) \wedge \mathsf{Isa}(x, \mathsf{very}(a), c) \wedge \mathsf{Isa}(y, c) \wedge \neg\mathsf{Isa}(y, \mathsf{very}(a), c)) \rightarrow \mathsf{More}(a, x, y)]$

The natural language modifier 'quite' is somewhat more subtle and seems to be ambiguous. For instance 'John is quite tall' could be interpreted as meaning either:

- John is not really 'tall' although he is almost high enough to be 'tall';
- John is tall but not very tall;

or it could be taken as equivalent to the disjunction of these two interpretations.

In order to cater for this range of meanings I first introduce a the adjectival operator $\mathsf{als}(a)$, to be read as 'at least somewhat'. Thus, the predicate $\mathsf{Isa}(x, \mathsf{als}(\mathsf{tall}), \mathbf{man})$ is true of any man, x, of at least somewhat tall stature.

A16) $\forall acx[\mathsf{Isa}(x, a, c) \rightarrow \mathsf{Isa}(x, \mathsf{als}(a), c)]$
A17) $\forall acxy[(\mathsf{ResGrad}(a) \wedge \mathsf{Isa}(x, \mathsf{als}(a), c) \wedge \mathsf{Isa}(y, c) \wedge \neg\mathsf{Isa}(y, \mathsf{als}(a), c)) \rightarrow \mathsf{More}(a, x, y)]$

Although als does not itself correspond to a basic adjectival modifier of natural language, it does enable one to define modifiers whose meaning corresponds well with the different senses of 'quite' mentioned above.

Additional Non-Logical Relations As well as representing information relating to vague adjectives and observables, we may sometimes want to reason with additional relationships between objects. Since the theory is encoded within ordinary 1st-order logic, we can simply employ any additional predications $\mathsf{R}(x_1, \ldots, x_n)$ we wish to use.

Semantics

Since **VAL** is formulated within standard 1st-order logic, it comes already equipped with the standard set-theoretic semantics. However, this is a very general semantics, which does not place any particular structural restrictions on the semantic interpretation of the fundamental primitives of the **VAL** theory. To give a better idea that the theory captures the intuitive meanings of the notions it seeks to formalise, it will be useful to provide a more specific interpretation, within which the fundamental primitives ($\mathsf{Obgeq}(o, x, y)$, $\mathsf{Rel}(o, a)$, $\mathsf{Isa}(x, c)$, $\mathsf{Isa}(x, c)$ and $\mathsf{Atleastas}(a, x, y)$) are associated with semantic structures that capture intuitive properties of their meaning.

I start by considering suitable denotations and structures corresponding to each of the different syntactic elements of **VAL**. Following this, I shall combine these to define formal model structures, which can function as interpretations of any consistent **VAL** formula.

Individual Constants and Count Nouns For the purpose of the present analysis, constants and count

[2] For theorem proving purposes, it may be better to represent converses by means of a symmetric binary relation, $\mathsf{Conv}(\mathsf{a}_1, \mathsf{a}_2)$; however, it seems to be axiomatically simpler to employ the functional representation.

nouns are given the semantics of classical constants and predicates; that is, $\delta(x)$ is an element of the domain of entities $\mathcal{D}$, and the extension of a count noun $\mathbf{c}$ is a subset $\delta(\mathbf{c})$ of $\mathcal{D}$. As usual I then specify that $\mathsf{Isa}(x, \mathbf{c})$ is true if and only if $\delta(x) \in \delta(\mathbf{c})$.

Non-Vague Adjectives. The operation of a non-vague adjective on a count noun predicate to give a modified count noun predicate can be modelled as a function $v : 2^{\mathcal{D}} \rightarrow 2^{\mathcal{D}}$.

Vague Adjectives We wish to model the possible extensions of a vague predicate according to weaker or stricter interpretations of its meaning. And we model these degrees of strictness in terms variations in the choice thresholds for relevant observables corresponding to different standpoints. To formalise this I use structures that I call Nests. These are simply sequences of subsets of the domain, such that each successive element of a Nest is a subset of the preceding one, and also such that nesting corresponds to successive increases in threshold variables. A nest that is appropriate for interpreting a given vague adjective will be one whose ordering corresponds to a succession of threshold increases on the observables relevant to that adjective.

Another important part of the **VAL** models is the structure Comp, which specifies the semantics of adjectival comparatives. Intuitively, one individual must be a-er than another if all relevant observables of adjective a are possessed to a greater degree in the one individual than the other. Comp is made to respect this condition.

I allow the ampliative and comparative structure of adjectives to vary according to the type of thing that is being considered. Thus, in the semantics, both Amp and Comp specify interpretations for adjective/count noun pairs, rather than simply giving an interpretation for each adjective.

Notation for Describing Orderings. Specifying of the semantics for ResGrad adjectives involves the construction of ordering structures which describe the way that objects are ranked relative to relevant observables.

To say that the object pair $\langle x, y\rangle$ satisfies the order relation ω, I will normally write $x \succeq_\omega y$. Where the order is referred to by a complex term or induced by a Nest structure (see below), I also write this in the form $(\tau : x \succeq y)$.

A *complemented order structure* is a tuple: $\langle \mathcal{D}, \Omega, inv\rangle$, where:

- $\mathcal{D}$ is a non-empty set of entities,
- Ω is a set $\{\ldots, \succeq_i, \ldots\}$ of ordering relations on $\mathcal{D}$, such that if $\omega \in \Omega$ then $\omega^{-1} \in \Omega$ (where ω^{-1} is the inverse of ordering ω).
- $inv : \Omega \rightarrow \Omega$ is the function mapping each $\omega \in \Omega$ to its inverse.

We need a notation to state that two entities satisfy each of a set of order relations. Thus, where Ω is a set of orderings on $\mathcal{D}$ and $x, y \in \mathcal{D}$, I define a *joint ordering* relation

- $x \succeq_\Omega y \;\equiv_{\text{def}}\; (\forall \omega \in \Omega)[x \succeq_\omega y]$

I use these joint orderings to define ordered nestings of the domain relative to sets of observable orderings. Let $S \subseteq \mathcal{D}$ and $\mathcal{S} = [S_0, S_1, \ldots, S_i, \ldots]$ be a a denumerable sequence of subsets of $\mathcal{D}$. I define:

- $\mathsf{Subseq}(\mathcal{S}_1, \mathcal{S}_2)$ iff $\mathcal{S}_2$ contains a subset of the elements of $\mathcal{S}_1$ occurring in the same order as in $\mathcal{S}_1$
- $\mathsf{UpInc}(S, \omega, \mathcal{D}) \;\equiv_{\text{def}}$
  $(\forall x, y \in \mathcal{D})\; [(x \in S \wedge y \succeq_\omega x) \rightarrow y \in S]$
- $\mathsf{Nest}(\mathcal{S}, \omega, \mathcal{D}) \;\equiv_{\text{def}}$
  $(\forall S_i \in \mathcal{S})\; [\mathsf{UpInc}(S_i, \omega, \mathcal{D}) \;\wedge\; (S_i \supseteq S_{i+1})]$
- $\mathsf{NESTS}(\omega, \mathcal{D}) \;=_{\text{def}}\; \{\mathcal{S} \mid \mathsf{Nest}(\mathcal{S}, \omega, \mathcal{D})\}$

$\mathsf{UpInc}(S, \omega, \mathcal{D})$ can be read as S *is upward inclusive over* $\mathcal{D}$ *relative to* ω. If Ω is a set of orderings $\mathsf{UpInc}(S, \Omega, \mathcal{D})$ is defined via the definition of joint ordering. For any Nest sequence $\mathcal{S}$ over $\mathcal{D}$, an induced ordering on $\mathcal{D}$ is defined as follows:

- $(\mathcal{S} : x \succeq y) \;\equiv_{\text{def}}\; (\exists S_i \in \mathcal{S})[(x \in S_i) \wedge \neg(y \in S_i)]$

Semantics of Adjective Comparison. The relation $\mathsf{Atleastas}(a, x, y)$, will be interpreted by a comparison relation $\mathsf{Comp}(a, x, y)$. In order to accord with its intuitive meaning, the Comp relation must be heavily constrained to accord with the orderings on relevant observables of the objects that are compared. In specifying the semantics this is done by reference to a complemented order structure.

Semantics of Adjective Predications. Given an adjective and a count noun we wish to determine a set of individuals satisfying that adjective/count noun combination. Moreover, we wish to be able to apply modifiers (very and als) to the adjective in order to make the requirement imposed by the adjective more or less strict.

To model this, each adjective/count noun compound is associated with a list of sets, which represent successively stricter interpretations of the adjective as applied to that count noun. Thus, the interpretation function $\mathsf{Amp}(\mathsf{a}, \mathbf{c})$ returns a list of sets, such that the first element is the set of those entities, of the kind $\mathbf{c}$, that satisfy the adjective to a limited degree. The second element is the set of entities of this kind that satisfy the adjective to at least a normal degree.

So, $\mathsf{Amp}(\mathsf{tall}, \mathbf{giraffe})$ denotes the list of sets $[G_0, G_1, G_2, G_3, \ldots]$, where:

$G_1 =$ the set of giraffes that are 'quite tall' or taller,
$G_1 =$ the set of giraffes that are 'tall' or taller,
$G_1 =$ the set of giraffes that are 'very tall' or taller,
$G_1 =$ the set of giraffes that are 'very very tall' or taller, ...

Using the head and tail operations we can extract from the nest, a set corresponding to either the unqualified adjective, or the adjective qualified by als or any number of very operators.

VAL Models

A **VAL** model is a tuple $\langle \mathsf{Voc}, \mathsf{Ord}, \mathsf{Ass} \rangle$, where:

- Voc is a VAL vocabulary $\langle N, O, C, A, R \rangle$,
- Ord is a complemented order structure, $\langle \mathcal{D}, \Omega, inv \rangle$
- Ass is an assignment structure

$$\langle \delta_N, \delta_O, \delta_C, \rho, \mathsf{Amp}, \mathsf{Comp}, \delta_R \rangle \ ,$$

which determines the denotations of each name variable, observable constant and count noun, and also the observables relevant to each adjective:

- δ_N is a function from N into $\mathcal{D}$,
- δ_O is a function from O into Ω,
- δ_C is a function from C to subsets of $\mathcal{D}$,
- ρ is a function from A to subsets of Ω.
- Amp is a function $(A \times C) \rightarrow \mathsf{NESTS}(\rho(\mathsf{a}), \mathcal{D})$, such that for every a and $\mathbf{c}$ we have $\mathsf{head}(\mathsf{Amp}(\mathsf{a}, \mathbf{c})) \subseteq \delta_{\mathbf{C}}(\mathbf{c})$
- Comp is a subset of $A \times \mathcal{D} \times \mathcal{D}$, such that:
  $\langle \mathsf{a}, x, y \rangle \in \mathsf{Comp}$ only if $x, y \in \delta_C(\mathbf{c})$;
  and, for every $x, y \in \delta_C(\mathbf{c})$,
  if $x \succeq_{\rho(\mathsf{a})} y$ then $\langle \mathsf{a}, x, y \rangle \in \mathsf{Comp}$,
  and if $x \succeq_{\rho(\mathsf{a})} y$ and for some $o \in \rho(\mathsf{a})$ $y \not\succeq_o x$ then $\langle \mathsf{a}, y, x \rangle \notin \mathsf{Comp}$.
- δ_R is a mapping from R to 2^{Tuples}, where Tuples is the set of all finite tuples over $\mathcal{D}$,

Interpretation

Given a **VAL** model
$\langle \langle \mathcal{D}, \Omega, inv \rangle, \langle N, O, C, A, R \rangle, \langle \delta_N, \delta_O, \delta_C, \rho, \mathsf{Amp}, \mathsf{Comp} \rangle \rangle$
the interpretation of a **VAL** formula is given by the following specification.

Preliminary to characterising the relational symbols, it will be convenient to specify the interpretations of the nominal symbols:

- $[\![x]\!] = \delta_N(x)$, where $x \in N$;
- $[\![o]\!] = \delta_o(o)$, where $o \in O$;
- $[\![c]\!] = \delta_c(c)$, where $c \in C$;
- $[\![\mathsf{conv}(\omega)]\!] \ = \ inv([\![\omega]\!])$;

Semantics of the primitive relations can now be specified as follows:

- $[\![\tau_1 = \tau_2]\!] = \mathbf{t}$ if $[\![\tau_1]\!] = [\![\tau_2]\!]$, else $= \mathbf{f}$,
- $[\![\mathsf{Rel}(\omega, \mathsf{a})]\!] = \mathbf{t}$ if $[\![\omega]\!] \in \rho(\mathsf{a})$, else $= \mathbf{f}$,
- $[\![\mathsf{Obgeq}(\omega, x, y)]\!] = \mathbf{t}$ if $\delta_N(x) \succeq_\omega \delta_N(y)$, else $= \mathbf{f}$, where $\succ_\omega$ is the order relation $[\![\omega]\!]$,
- $[\![\mathsf{Atleastas}(\mathsf{a}, x, y)]\!] = \mathbf{t}$ if $\langle \mathsf{a}, [\![x]\!], [\![y]\!] \rangle \in \mathsf{Comp}$, else $= \mathbf{f}$,
- $[\![\mathsf{Isa}(\mathbf{c}, x)]\!] = \mathbf{t}$ if $[\![x]\!] \in [\![\mathbf{c}]\!]$, else $= \mathbf{f}$,
- $[\![\mathsf{Isa}(\alpha, \mathbf{c}, x)]\!] = \mathbf{t}$ if $[\![x]\!] \in \mathsf{head}([\![\alpha \circ \mathbf{c}]\!])$,
- $[\![\mathsf{R}(x_1, \ldots, x_n)]\!] = \mathbf{t}$ if $\langle [\![x_1]\!], \ldots, [\![x_n]\!] \rangle \in \delta_R(\mathsf{R})$, else $= \mathbf{f}$,

In the clause for interpreting $\mathsf{Isa}(\alpha, \mathbf{c}, x)$ it will be seen that this is defined in terms of $[\![\alpha \circ \mathbf{c}]\!]$. This gives the interpretation of an adjective/count noun combination in terms of a NEST returned by the Amp function.

- $[\![\ \mathsf{als}(\mathsf{a}) \circ \mathbf{c}]\!] = \mathsf{Amp}(\mathsf{a}, \mathbf{c})$, where $\mathsf{a} \in A$,
- $[\![\mathsf{a} \circ \mathbf{c}]\!] = \mathsf{tail}(\mathsf{Amp}(\mathsf{a}, \mathbf{c}))$, where $\mathsf{a} \in A$,
- $[\![\mathsf{very}(\alpha) \circ \mathbf{c}]\!] = \mathsf{tail}([\![\alpha \circ \mathbf{c}]\!])$,

The clauses for the interpretation function dealing with the Boolean connectives and quantification as those of standard 1st-order logic, so will not be presented here.

Extensions Further Work

Reasoning and its Implementation

A series of theorem proving experiments have been carried out in order to check whether the proposed axioms do indeed entail the inferences involving vague adjectives that we would intuitively consider to be valid. These were carried out using the Otter theorem prover (McCune 1990) using axioms that are essentially equivalent to those given in this paper.[3]

It was found that all expected inferences could be proved as long as the number of individuals involved in the chain of reasoning was relatively small (up to 5 was usually tractable). However, it is clear that when dealing with larger domains the theory could not be successfully managed by a general purpose theorem prover. The main problem seems to stem from certain existential commitments in axioms (such as **A6**) and the fact that equality sometimes plays a significant part in reasoning. Nevertheless, considering the size of the theory it seems to be relatively effective when used for inferring implications among ground facts. However, it is likely that a special purpose prover designed specifically to handle this theory would be much more effective.

Adding Value Entities

In the theory presented so far I have considered observables only as determining relative orderings on the domain of individuals. However, to represent certain modes of description and reasoning, it would be useful to be able to refer to actual values of observables for particular individuals and to consider orderings in terms of relations over domains of values.

To this end one can extend the theory by introducing value functions of the form $\mathsf{v}(o, x)$ which gives the value of observable o for the individual x. Then, one can introduce a further primitive relation $\mathsf{Vgeq}(o, v_1, v_2)$

[3] In fac,t additional predicates representing sort restrictions on different kinds of entity were added, and these were used as pre-conditions for many of the axioms. This does not affect provability, but prevents the theorem prover from generating vast numbers of ill-sorted formulae that cannot contribute to a proof.

representing the condition that $v_1 \geq v_2$ relative to observable o.

A18) $\forall ox[\mathsf{Vgeq}(o, x, x)]$
A19) $\forall oxyz[(\mathsf{Vgeq}(o, x, y) \wedge \mathsf{Vgeq}(o, y, z)) \rightarrow \mathsf{Vgeq}(o, x, z)]$
A20) $\forall oxy[(\mathsf{Vgeq}(o, x, y) \wedge \mathsf{Vgeq}(o, y, x)) \rightarrow x = y]$
A21) $\forall oxy[\mathsf{Total}(o) \rightarrow (\mathsf{Vgeq}(o, x, y) \vee \mathsf{Vgeq}(o, y, x))]$

The reason we need to relativise the ordering to a particular observable is that we may need to handle different orderings over the same values corresponding to different ways of comparing these values. In particular observables o and $\mathsf{conv}(o)$ will share the same domain of values, but will be associated with opposite orderings over this domain. Nevertheless, identity between values will hold independently of the ordering, which enables one to refer to specific values whose identity will be fixed whatever ordering they are related to (thus if someone is 6' tall they are also 6' short). Hence the use of logical equality in **A20**.

Having introduced value entities, the observable ordering relation Obgeq over individuals would no longer be taken as primitive and the axioms **A1** and **A2** would not be required. Instead Obgeq can be defined in term of the ordering of observable values. Thus we would have:

D6) $\mathsf{Obgeq}(o, x, y) \equiv_{\mathrm{def}} \mathsf{Vgeq}(o, \mathsf{v}(o, x), \mathsf{v}(o, y))$

Moreover, **A12** should be replaced by the following constraint on Vgeq:

A22) $\forall ov_1v_2[\mathsf{Vgeq}(o, v_1, v_2) \leftrightarrow \mathsf{Vgeq}(\mathsf{conv}(o), v_2, v_1)]$

Embedding within a Supervaluation Semantics

So far we have modelled the logic of vague adjectives within an essentially classical reasoning system. However, since the thresholds on application of vague adjectives are essentially subjective, there are many possible interpretations corresponding to any given state of the world. The reasoning of the system so far covers inferences that hold in any consistent interpretation that satisfies the axioms on the logical predicates. However, in certain cases (e.g. the presence of multiple agents with different judgements) it is useful to distinguish subjective from 'unequivocal' assertions. This can be done by generalising to a supervaluation where each interpretation corresponds to a particular *standpoint*. Within this more general framework we can introduce a modal *unequivocality* operator $\blacksquare\phi$ asserts that ϕ is true for any standpoint (this is just an $S5$ modality). This is used to assert non-subjective semantic constraints constraining the non-logical vocabulary.

A Non-Monotonic Solution of the *Sorites* Paradox

Though the formalism I have presented seeks to model the variability in the meaning of vague adjectives it does not directly tackle the famous *sorites* paradox (see e.g. (Keefe & Smith 1996) for many good articles and references this, and see also (Halpern 2004)).

Having spent considerable time trying to formulate the semantics of vague concepts, I have come to the conclusion that the solution to the *sorites* requires not only semantics but also an analysis of a particular mode of commonsense inference. Specifically, the sorites seems to involve a kind of default reasoning. If I imagine a heap (e.g. of sand) then I imagine a typical, and therefore quite large heap. Clearly, removing one grain will not significantly alter the size of the heap. Thus it will still be very close to a typical heap and therefor must be considered a heap. However, the 'heap-minus-a-grain' is *not* a default heap. It is a default heap with a modification. Thus we cannot apply default reasoning directly; we have to take account of the extra information, which qualifies our understanding of the modified heap.

An augmentation of the nesting semantics for vague adjective/count noun pairs allows one to add default extensions giving upper and lower bounds on the set of 'typical' individuals of a given kind qualified by a particular vague adjective — e.g. the set of 'typical tall giraffes'. These additional bounds are located in the nesting structure somewhere between the set of 'tall giraffes' and 'very tall giraffes'. Individuals in the sets between these bounds do not contain any borderline cases of 'tallish' giraffes. A default inference rule is then added, such that, in the absence of information to the contrary, one can assume that a given individual described as say a 'tall giraffe' lies within the typical bounds of the nesting, rather than at the extremes of the the limiting bands. Given some assertions about the height differences between giraffes at different levels of the 'tallness' nesting, it is then be possible to reason about giraffes who differ in some way from a typical tall giraffe (e.g. by being a whisker's width shorter), and thereby to make a defeasible (non-monotonic) inference about whether such giraffes are likely to count as 'tall'. Slippery slope reasoning can be blocked because we can infer that a series of small differences can add up to a difference that takes us beyond the distance between typical cases and limiting cases.

Relationship to Fuzzy and Probabilistic Approaches

The reader may be wondering whether any accommodation is possible between the kind of approach presented here and the much more established fuzzy and probabilistic accounts of vagueness. The two approaches are generally considered as rivals, but nevertheless I believe that some integration may be possible.

In the supervaluation inspired approaches one treats assertions as true according to some point of view and seeks to capture necessary inferences that hold despite the subjectivity of these assertions. In the fuzzy approaches both the assertions and the modes of inference are treated as quantitative rather than absolute. In my opinion the quantitative/statistical treatment of inference is problematic as it does not account for the fact that definite inferences can be made even on the basis of subjective assertions (i.e. inferences that are necessary relative to the given subjective interpretation). However, a statistical treatment of the subjective assertions does seem useful, since we may want to have some measure of how closely an assertion corresponds to the typical use of the vocabulary involved. This would give us some way of grounding these assertions in reality by saying 'how true' they are when applied to a given situation.

I believe this may be achieved in the following way. In the supervaluationistic semantics we have a space of possible interpretations of the whole language each of which is self-consistent but differs from other interpretations. Each of these interpretations are equally valid, however remote they are from normal usage of the concepts involved. But suppose we introduce a probability distribution over the space of possible interpretations. Then a quantitative measure of the truth of an assertion can be obtained by summing the probabilities over all interpretations which make it true. Similarly, a measure of the degree of validity of an entailment can be obtained as a kind of conditional probability: we take the subspace of interpretations, for which the premisses of the entailment are true, and normalise the probability distribution to 1 over this subspace; then we sum the probabilities of interpretations where the consequent is true over the normalised subspace.

Conclusion

I have presented a formal language for reasoning with vague adjectives and describing their dependence on relevant observables. This dependency explains how, even though the extensions of adjectives are often not well-defined it is still possible to carry out useful reasoning with information expressed using such adjectives. I believe that this kind of analysis of logical structure hidden behind vagueness is a novel contribution to Knowledge Representation and is potentially useful in many AI applications.

Ongoing work at Leeds is investigating the formal characterisation of geographic features, which are classified by means of vague adjectives (Santos, Bennett, & Sakellariou 2005).

References

Bennett, B. 1998. Modal semantics for knowledge bases dealing with vague concepts. In Cohn, A. G.; Schubert, L.; and Shapiro, S., eds., *Principles of Knowledge Representation and Reasoning: Proceedings of the 6th International Conference (KR-98)*, 234–244. Morgan Kaufman.

Bennett, B. 2001. What is a forest? on the vagueness of certain geographic concepts. *Topoi* 20(2):189–201.

Bennett, B. 2002. Physical objects, identity and vagueness. In Fensel, D.; McGuinness, D.; and Williams, M.-A., eds., *Principles of Knowledge Representation and Reasoning: Proceedings of the Eighth International Conference (KR2002)*, 395–406. San Francisco, CA: Morgan Kaufmann.

Burrough, P., and Frank, A., eds. 1996. *Geographical Objects with Undetermined Boundaries*, GISDATA. Taylor and Francis.

Cohn, A. G., and Gotts, N. M. 1996. The 'egg-yolk' representation of regions with indeterminate boundaries. In Burrough, P., and Frank, A. M., eds., *Proceedings, GISDATA Specialist Meeting on Geographical Objects with Undetermined Boundaries*, 171–187. Francis Taylor.

Dubois, D., and Prade, H. 1988. An introduction to possibilistic and fuzzy logics. In P.Smets; E.H.Mamdani; D.Dubois; and H.Prade., eds., *Non-Standard Logics for Automated Reasoning*. Academic Press.

Elkan, C. 1994. The paradoxical success of fuzzy logic. *IEEE Expert* 9(4):3–8. Followed by responses and a reply.

Fine, K. 1975. Vagueness, truth and logic. *Synthèse* 30:263–300.

Goguen, J. 1969. The logic of inexact concepts. *Synthese* 19:325–373.

Halpern, J. Y. 2004. Intransitivity and vagueness. In *Principles of Knowledge Representation: Proceedings of the Ninth International Conference (KR-2004)*, 121–129.

Keefe, R., and Smith, P. 1996. *Vagueness: a Reader*. MIT Press.

McCune, W. 1990. Otter 2.0 users guide. Technical report, Argonne National Laboratory, Argonne, Illinois.

McGee, V. 1997. Kilimanjaro. *Canadian Journal of Philosophy* 141–195. supplementary volume 23.

Roy, A. J. O., and Stell, J. G. 2001. Spatial relations between indeterminate regions. *International Journal of Approximate Reasoning* 27:205–234.

Santos, P.; Bennett, B.; and Sakellariou, G. 2005. Supervaluation semantics for an inland water feature ontology. In Kaelbling, L. P., and Saffiotti, A., eds., *Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI-05)*, 564–569. Edinburgh: Professional Book Center.

Zadeh, L. 1965. Fuzzy sets. *Information and Control* 8:338–353.

Zadeh, L. A. 1975. Fuzzy logic and approximate reasoning. *Synthese* 30:407–428.