
Modal Semantics for Knowledge Bases Dealing with Vague Concepts

Brandon Bennett

Division of Artificial Intelligence
School of Computer Studies
University of Leeds, Leeds LS2 9JT, England
brandon@scs.leeds.ac.uk

Abstract

The paper investigates the characterisation of vague concepts within the framework of modal logic. This work builds on the *supervaluation* approach of Fine and exploits the idea of a *precisification space*. A simple language is presented with two modalities: a necessity operator and an operator ‘it is unequivocal that’ which is used to articulate the logic of vagueness. Both these operators obey the schemas of the logic *S5*. I show how this language can be used to represent logical properties of vague predicates which have a variety of possible precise interpretations. I consider the use within KR systems of number of different entailment relations that can be specified for this language. Certain vague predicates (such as ‘tall’) may be indefinite even when there is no ambiguity in meaning. These can be accounted for by means of a three-valued logic, incorporating a definiteness operator. I also show the relationship between observable quantities (such as height) and vague predicates (such as ‘tall’) can be represented *via* axioms involving precise comparative relations (such as ‘taller’). I consider how Williamson’s ‘logic of clarity’ can be combined with the semantics for unequivocality and how the clarity operator can be related to observables.

1 INTRODUCTION

Natural language descriptions are pervaded by the use of vague concepts of a variety of different kinds. Perhaps the most often cited cases are predicates such as ‘tall’, where it is not clear exactly how high a person or object should be to count as being tall. Here the vagueness is associated with the problem of reconciling a qualitative concept with a continuous range of possible measurements. In more subtle cases it may not even be clear what objective facts are at stake. (Consider the meaning of ‘disaster’ in the sentences ‘The eruption of Krakatoa was a disaster’ and ‘The best man’s speech was a complete disaster’.) Often the use of a seemingly simple concept may involve different kinds of vagueness at the same time. For instance, in a sentence such as ‘This frog is green’ the word ‘green’ is vague not only in that there may be borderline cases of green (a yellowy green or a dark brownish green) but also because the frog may have a yellow belly, black eyes and a pink tongue. Hence, ‘green’ in some circumstances means predominantly green rather than green all over. Although the different types of vagueness share certain properties, it is implausible that all these phenomena could be accounted for by a completely uniform logical analysis.

Vagueness is often thought of as a deficiency of natural language: it is imagined that if vague concepts were replaced by clearly defined concepts the resulting precise language would be superior to existing natural languages. Against this view it may be argued that vagueness actually increases the utility of language because it enables one to make true statements without having to use the extremely cumbersome descriptions that would be needed to make them completely precise.

In constructing AI knowledge bases, researchers have typically attempted to eliminate vagueness from concepts by precise definitions. However, vagueness poses

an extremely serious threat to this enterprise. Since ordinary language concepts are vague it is very easy for contradictions to occur because of slight differences in the intended meaning of different occurrences of predicates in a set of axioms. In the construction of large axiomatised ontologies (Guha and Lenat 1990) this problem is especially acute. I believe that a logical treatment of vagueness can be particularly useful in preventing inconsistency arising in this way. Similar problems also occur when one attempts to combine information from two or more different databases (Lehmann and Cohn 1994).

In this paper I shall show that by a suitable logical treatment of vagueness it is possible to safely combine data and axioms where the meanings of predicates may vary from instance to instance. Moreover, I shall show that it is possible to do this without completely enfeebling inference capabilities. Whether or not vagueness is a defect or an asset, the ability to manipulate vague concepts would certainly be very useful in a wide range of KR&R systems.

1.1 RELATION OF THE PRESENT APPROACH TO PREVIOUS WORK

In computer science the approaches to vagueness that have received the most attention are those based on *multi-valued* logics (Łukasiewicz and Tarski 1930). *Fuzzy* logic (Zadeh 1965, Zadeh 1975, Goguen 1969, Dubois and Prade 1988) can be regarded as a generalised multi-valued logic. An extensive discussion of the pros and cons of fuzzy logic can be found in (Elkan 1993). In accord with Elkan, I suggest that fuzzy logic may be an appropriate formalism for modelling the relation between continuous valued observables and the meanings of vague qualitative predicates; however, it is not suitable as a formalism for carrying out logical reasoning. This is primarily because propositional operators whose values are determined purely in terms of the fuzzy truth values of their arguments cannot take account of either logical or domain-specific constraints holding among the argument propositions.

Another approach to vagueness was given by Fine (1975) in terms of *supervaluations*. The idea is that vague expressions are associated not with a single extension but with a set of different possible extensions. This can be modelled by means of a Kripke-style possible worlds semantics (Kripke 1959). In a recent book by Williamson (1994), it is argued that vagueness should be regarded as an epistemic phenomenon and proposes a modal treatment of vagueness rather different from that of Fine. This paper also takes a model-theoretic approach incorporating and extending

a number of previous ideas. However, whereas previous models were proposed mainly in support of philosophical thesis about vagueness, I shall be concerned with the management of vagueness within AI knowledge bases. Consequently, my concern is to specify a formal language within which logical dependencies between vague expressions can be represented and which provides a complete inference mechanism.

1.2 VAGUENESS, UNCERTAINTY AND GENERALITY

I distinguish sharply between vagueness and uncertainty.¹ Although they are both associated with lack of precision of information and may both be regarded as properties of propositions, the two properties are complementary rather than parallel. The more uncertain a proposition is, the less likely it is to be true; but the more vague it is the more likely it is to be in some sense true. In other words by stating information in terms of vague concepts we may be more certain about our claims. Another common misunderstanding is that generality and disjunctive information are a form of vagueness or of uncertainty. For instance one might (wrongly) regard the statement ‘I am in my twenties’ as more vague than ‘I am 29’ but although the first sentence is less specific it is not at all vague (and both are — or were — equally certain): being vague (but just as certain) I would say ‘I am approaching 30’. Generality and disjunction are similar to vagueness in that they are ways of making certain (or at least more certain) claims regarding situations about which we do not know all the precise details.

In this paper I shall assume that the state of the world is completely determinate in terms of certain precise *objective measurements* that can always be obtained for any given objects in any given situation. These determine the truth of all precise propositions and also indirectly determine the validity of inferences which involve vague concepts. The content of this assumption will be made precise by means of a formal (possible worlds) semantics. I also assume that vagueness is confined to the predicates of the language. This means that there are no vague objects² or nominal terms and the logical operators are also taken to be completely definite.

¹The subject of uncertainty has also received much attention from AI researchers (see e.g. (Shafer and Pearl 1990)).

²The logic of vague objects is likely to be very different from that of vague predicates. A theory of vague spatial objects can be found in (Cohn and Gotts 1996).

2 PRECISIFICATION SPACES AND SUPERVALUATIONS

The supervaluation theory of vagueness proposed by Fine (1975) is based on the idea that expressions of a language which includes vague expressions can be assigned different semantic values according to different possible ways in which the vague expressions can be interpreted.³ Each such interpretation is called a *precisification*. The simplest model of this is to identify each precisification with a classical two-valued truth-assignment to the atomic propositions of the language. Precisifications are then treated just like possible worlds in a Kripke semantics. Fine's semantics has been elaborated and clarified by Pinkal (1995).

2.1 MODELS FOR NECESSITY AND UNEQUIVOCALITY OPERATORS

Fine's analysis of vagueness assumes we can specify a set of *admissible* precisifications. These are associated with valuations of atomic propositions which are taken to reasonably accord with natural intuitions about the meaning of vague concepts. Of course, from a purely formal point of view one cannot distinguish between reasonable and unreasonable valuations but Fine takes admissibility as a primitive notion restricting the range of possible models. Whilst this is perfectly coherent it presents methodological problems in that it obscures the relationship between valuations and analytic constraints on predicate meanings. Such constraints are assumed to be dealt with before the model theoretic analysis begins. By contrast I impose no constraints on precisifications — they are associated with arbitrary valuations. This enables necessary constraints to be modelled axiomatically in the logic rather than treated as restrictions imposed from outside. Thus, I identify validity with truth which is guaranteed in arbitrary valuations, not just in those that are admissible. Within this more general framework the concept of an 'admissible' precisification can easily be introduced as being one satisfying some necessary theory, Θ : we say that ' ϕ is true for all admissible precisifications' corresponds to the condition that $\Theta \rightarrow \phi$ is true for arbitrary valuations.

A formal language for knowledge representation will typically be used to specify theoretical as well as factual relationships. In my analysis of vagueness, the difference between these two types of proposition will

³The term 'supervaluation' appears to originate with van Frassen, who used this kind of semantic structure to model the phenomenon of *presupposition* which occurs in natural languages (van Frassen 1968, van Frassen 1969).

be important and it will be useful to explicitly distinguish them both syntactically and semantically. To do this I employ the modal logic *S5* (see e.g. (Hughes and Cresswell 1969)). This enables theoretical information to be represented by formulae of the form $\Box\phi$ (factual formulae will not contain modal operators). In order to articulate the logic of vagueness I introduce into the language the modality \mathbf{U} meaning 'unequivocally' and its dual \mathbf{S} (where $\mathbf{S}\phi \equiv \neg\mathbf{U}\neg\phi$) meaning 'in some sense'. Thus, for example, $\mathbf{U}P \rightarrow \mathbf{S}Q$ asserts that if P is unequivocally true (i.e. true under any reasonable interpretation) then Q is in some sense true (i.e. true under some reasonable interpretation). This extended language will be called *S5U*.

As a model for an intensional logic of vagueness, within which the necessity modality can be used to specify analytic truths, I employ a Kripke semantics in which each 'world' is indexed by a pair $\langle w, p \rangle$, w being a member of a set of possible worlds W and p being a member of a set of precisifications P . I shall refer to the semantic entity referred to by a pair $\langle w, p \rangle$ as a *point*. Each point is associated with a classical valuation of formulae which do not contain modal operators.

An *S5U* model is a structure $\langle W, P, I, \mathcal{R}, \delta \rangle$, where I is a set of individual constants, \mathcal{R} is a set of relation symbols and δ is a function from $W \times P \times \mathcal{R}$ to tuples of elements of I . I write $\langle w, p \rangle \Vdash \phi$ to mean that ϕ is true at $\langle w, p \rangle$. For an atomic proposition $R(x_1, \dots, x_n)$ I specify that

$$\langle w, p \rangle \Vdash R(x_1, \dots, x_n) \quad \text{iff} \quad \langle x_1, \dots, x_n \rangle \in \delta(w, p, R)$$

The truth-conditions of the Boolean truth functions and quantifiers can then be specified in the usual way:

$$\begin{aligned} \langle w, p \rangle \Vdash (\alpha \wedge \beta) & \quad \text{iff} \quad \langle w, p \rangle \Vdash \alpha \quad \text{and} \quad \langle w, p \rangle \Vdash \beta \\ \langle w, p \rangle \Vdash (\alpha \vee \beta) & \quad \text{iff} \quad \langle w, p \rangle \Vdash \alpha \quad \text{or} \quad \langle w, p \rangle \Vdash \beta \\ \langle w, p \rangle \Vdash \neg\phi & \quad \text{iff} \quad \langle w, p \rangle \not\Vdash \phi \\ \langle w, p \rangle \Vdash \forall x\phi(x) & \quad \text{iff} \quad (\forall i \in I)\{\langle w, p \rangle \Vdash \phi(i)\} \end{aligned}$$

The semantics for quantification is based on the simplifying assumption that the domain of quantification (I) is constant for every world/precisification point $\langle w, p \rangle$. This is probably not realistic. A more adequate theory would need a more complex semantics with variable domains and would need to address problems of identifying counterpart individuals at different worlds. Such considerations are beyond the scope of the present paper.

The truth-conditions for the necessity and unequivocality operators are as follows:

$$\begin{aligned} \langle w, p \rangle \Vdash \Box\phi & \quad \text{iff} \quad (\forall u \in W)\{\langle u, p \rangle \Vdash \phi\} \\ \langle w, p \rangle \Vdash \mathbf{U}\phi & \quad \text{iff} \quad (\forall q \in P)\{\langle w, q \rangle \Vdash \phi\} \end{aligned}$$

A formula ϕ is valid in this semantics if it is forced at every world/precisification in every model. I write $\models_{S5U} \phi$.

The proof theory of this language can be specified in at least two different ways. One method is to first translate formulae of $S5U$ into purely 1st-order formulae. By reifying propositions and the indices of the semantic points and propositions, a condition $\langle w, p \rangle \Vdash \phi$ can be represented by a formula $\text{holds}(w, p, \phi)$. The semantics conditions for the connectives in $S5U$ can easily be specified as axioms constraining the holds relation. Given that I am assuming a constant domain of quantification for every point $\langle w, p \rangle$, quantifiers occurring in ϕ can be moved outwards to operate on the holds relation. The meaning of the modal operators can be expressed directly by quantification over the index variables. For unmodalised formulae the indices will be fixed constants w_0 and p_0 denoting the ‘actual’ world/precisification.

Alternatively, we may want to formalise the proof theory within the object language of $S5U$. To do this we add to a proof procedure for 1st-order logic the usual $S5$ schemata and the rule of necessitation for each of the two operators \Box and U . For completeness we will also need additional schemata expressing logical interactions between the two different modalities. The most obvious of these is the commutation schema,

$$\Box U\phi \leftrightarrow U\Box\phi . \quad (\text{Com})$$

Seegerberg (1973) gave a ‘2-dimensional’ modal logic which is very similar, although rather more expressive than my logic $S5U$. This includes an axiom equivalent to

$$\Box U(\Box\alpha \vee U\beta) \leftrightarrow (\Box U\alpha \vee \Box U\beta) \quad (\text{Orth})$$

which enforces a property pertaining to the orthogonal nature of the two modalities. The completeness proof given by Seegerberg strongly suggests that these schemata provide a complete axiomatisation of $S5U$. However, Seegerberg’s system contains additional schemata involving other modal operators and it is possible that these induce further properties of \Box and U , which must be axiomatised for completeness.

By means of well-known theorems of $S5$ together with the commutation schema it is easy to show that there is a finite number of ‘modalities’ in $S5U$ — i.e. a finite number of logically distinct strings of the symbols \neg , \Box , \Diamond , U , and S . For example, it is fairly easy to show that

$$\Box S\Box\phi \leftrightarrow S\Box\phi .$$

All the distinct modalities which do not include negation are shown in Figure 1 (arrows indicate entail-

ment). Each has a negated counterpart, so that the total number of modalities is 22. $S5$ also has the property that any formula can be systematically reduced to an equivalent formula within which no modal operator occurs within the scope of an other modal operator. This reduction allows all $S5$ formulae to be transformed into a relatively simple normal form (MCNF — see (Hughes and Cresswell 1969)) and forms the basis of decision procedure for $S5$ theoremhood. It may be possible to establish a similar though considerably more complex normal form for $S5U$. This would be very useful for automated theorem proving.




Figure 1: Distinct modalities in $S5U$

2.2 CONTRAST WITH MULTI-VALUED APPROACHES

The $S5U$ semantics means that every classical theorem is unequivocally true even if it contains vague predicates. Thus

$$\models_{S5U} U(\text{tall}(\text{John}) \vee \neg\text{tall}(\text{John}))$$

and

$$\models_{S5U} U\neg(\text{tall}(\text{John}) \wedge \neg\text{tall}(\text{John})) .$$

This is because each precisification/world is associated with a purely classical valuation, which makes every classical theorem true. By contrast, classical theorems involving vague predicates are generally not valid in many-valued or truth-functional logics, where the truth-value of a conjunction or disjunction is determined solely by the semantic value of the argument propositions. In such a logic, if $\text{tall}(\text{John})$ is not definitely true or false then it will have some indefinite truth value and so will $\neg\text{tall}(\text{John})$. It may also be

that $\text{thin}(\text{John})$ is also indefinite and has the same semantic value as $\text{tall}(\text{John})$. Hence, in such a case the tautology $\neg(\text{tall}(\text{John}) \wedge \neg\text{tall}(\text{John}))$ must have the same truth value as the contingent proposition $\neg(\text{thin}(\text{John}) \wedge \neg\text{tall}(\text{John}))$.

2.3 PRECISIFICATIONS AND PREDICATE SENSES

A precisification space is an abstract structure exhibiting the ways in which a language can be made precise in different ways. In the case of a real learnable language this space must be determined by the different possible interpretations of specific expressions. Since I am assuming that object names and logical operators are perfectly precise, the possible interpretations of any expression must be determined by the possible interpretations of its constituent predicates.

Consider the natural language predicate ‘ x is a murderer’. Although some necessary and sufficient properties of a murderer are uncontroversial, other properties are open to debate. Thus there is a wide variety of possible precise definitions of murderer. For instance one definition would be that: a murderer is a human being who has intentionally killed another human being. However, some people (e.g. certain strict vegetarians) might contend that a human being who intentionally kills any kind of animal is also a murderer.

To be more specific about the relationship between ‘killer’ and ‘murderer’ we could stipulate that if someone is a murderer in any sense then they are unequivocally a killer:

$$\Box \forall x [(\mathbf{S} \text{ murderer}(x)) \rightarrow \mathbf{U} \text{ killer}(x)] .$$

but in the opposite direction we might have the weaker implication

$$\Box \forall x [\text{ killer}(x) \rightarrow (\mathbf{S} \text{ murderer}(x))] .$$

If we then give axioms specifying the meaning of these concepts there will be certain analytic facts that we are sure should be true whatever precise sense they are given. For example:

$$\Box \mathbf{U}(\text{killer}(x) \leftrightarrow \exists y[\text{killed}(x, y)])$$

$$\Box \mathbf{U}(\text{murderer}(x) \rightarrow \exists y[\text{human}(y) \wedge \text{killed}(x, y)])$$

Other axioms will only apply to some but not all interpretations of vague predicates:

$$\Box \mathbf{S}\forall x[\text{murderer}(x) \leftrightarrow \exists y[\text{animal}(y) \wedge \text{killed}(x, y)]] .$$

Sometimes to specify properties of vague predicates it may be useful to introduce artificially sharpened versions of these predicates. Thus, if we wanted to use ‘murderer1’ in a sense which is sharper than the vague predicate ‘murderer’, we might use the following axiom:

$$\Box \forall x[\text{murderer1}(x) \rightarrow (\mathbf{S} \text{ murderer}(x))] \wedge$$

$$\Box \forall x[(\mathbf{U} \text{ murderer}(x)) \rightarrow \text{murderer1}(x)]$$

and define murderer1 by

$$\Box \mathbf{U}\forall x[\text{murderer1}(x) \leftrightarrow \exists y[\text{human}(y) \wedge \text{killed}(x, y)]] .$$

Equivocal predicates may include those that are obviously ambiguous, having two or more quite separate meanings rather than a cluster of similar meanings. One way to distinguish certain ambiguous concepts is to say that there is no case to which it applies in every sense. Thus we could write

$$\neg \diamond \exists x[\mathbf{U} \text{ nut}(x)]$$

to say that there is no possible situation in which there is something that is unequivocally a ‘nut’ (nothing is both a nut of the edible variety and of the mechanical variety).

As well as vague and ambiguous predicates, most knowledge bases will contain a large number of precise predicates, whose meaning is not in question. The sharpness of certain predicates can be attested by axioms such as

$$\Box \forall x[(\mathbf{S} \text{ integer}(x)) \rightarrow (\mathbf{U} \text{ integer}(x))] .$$

It should be clear from the various examples given in this section that the $S5\mathbf{U}$ language provides an expressive representation for describing the logical properties of vague predicates. Especially those whose vagueness consists in them having a cluster of related meanings.

2.4 VALIDITY DEFINED USING THE MODAL DEFINITELY OPERATOR

Within a precisification semantics such as $S5\mathbf{U}$ it is possible to define the relation of entailment among vague propositions in a number of different ways. Philosophical examination of this semantics has often sought to demonstrate that some particular definition is correct. However, from the point of view of describing possible inferences within a database containing vague information, it is not necessary to adopt a single definition of entailment. Many kinds of entailment can be defined in terms of the unambiguous property of validity and each of these tells us something different about the logical relationship between propositions.

Some of the more useful possible definitions of entailment are the following:

- $\phi_1, \dots, \phi_n \models_{\text{local}} \psi$ iff
 $\models_{S5U} \mathbf{U}((\phi_1 \wedge \dots \wedge \phi_n) \rightarrow \psi)$
- $\phi_1, \dots, \phi_n \models_{\text{global}} \psi$ iff
 $\models_{S5U} (\mathbf{U}\phi_1 \wedge \dots \wedge \mathbf{U}\phi_n) \rightarrow \mathbf{U}\psi$
- $\phi_1, \dots, \phi_n \models_{\text{arguable}} \psi$ iff
 $\models_{S5U} \mathbf{S}((\phi_1 \wedge \dots \wedge \phi_n) \rightarrow \psi)$
- $\phi_1, \dots, \phi_n \models_{\text{reliable}} \psi$ iff
 $\models_{S5U} (\mathbf{S}\phi_1 \wedge \dots \wedge \mathbf{S}\phi_n) \rightarrow \mathbf{U}\psi$
- $\phi_1, \dots, \phi_n \models_{\text{s-reliable}} \psi$ iff
 $\models_{S5U} (\mathbf{S}\phi_1 \wedge \dots \wedge \mathbf{S}\phi_n) \rightarrow \mathbf{S}\psi$
- $\phi_1, \dots, \phi_n \models_{\text{co-reliable}} \psi$ iff
 $\models_{S5U} \mathbf{S}(\phi_1 \wedge \dots \wedge \phi_n) \rightarrow \mathbf{U}\psi$
- $\phi_1, \dots, \phi_n \models_{\text{co-s-reliable}} \psi$ iff
 $\models_{S5U} \mathbf{S}(\phi_1 \wedge \dots \wedge \phi_n) \rightarrow \mathbf{S}\psi$

Philosophers (Fine 1975, Williamson 1994) have been most concerned with the first two definitions of validity. So-called ‘local’ validity means that under-any interpretation where the premisses are true the conclusion is also true. However, because \mathbf{U} satisfies the rule of necessitation and the T schema, $\models_{S5U} \mathbf{U}\phi$ holds just in case $\models_{S5U} \phi$; so, for formulae not containing modalities, local validity is equivalent to classical validity.

To avoid this reduction of entailment in vague languages to classical validity, it has been argued that an appropriate characterisation of valid entailment in vague languages is that if the premisses are unequivocally true the conclusion is also unequivocally true. This has been called ‘global’ validity and is strictly weaker than local validity (the former is derivable from the latter by the K schema).

What I call ‘arguable’ validity is perhaps the weakest useful entailment relation. A sequent is arguably valid if the corresponding implication is valid under some interpretation. If we move the \mathbf{S} operator inwards in the schematic theorem corresponding to arguable entailment we get $((\mathbf{U}\phi_1 \wedge \dots \wedge \mathbf{U}\phi_n) \rightarrow \mathbf{S}\psi)$. Thus to show that a sequent is arguably valid we only need to show that if all the premisses are ‘unequivocally’ true then the conclusion is ‘in some sense’ true. This form of entailment is suitable for query applications, where we are interested in whether a vague query might possibly be a consequence of some vague data. Ideally a procedure which tests for arguable validity would also give further information about what senses of the

query and database facts make the entailment valid. The user could then assess the reliability of the inference.

‘Reliable’ entailment on the other hand is very strong. It makes valid only those arguments such that however the premisses are interpreted, the conclusion is unequivocally true. This notion of entailment might be useful in certain safety critical applications where we want to take some action only if an entailment must hold whatever interpretations are given to vague concepts in the premisses and conclusion.

The remaining three entailment relations are weaker than ‘reliable’ but are similar in that we only assume that the premisses are ‘in some sense’ true. For ‘s-reliable’ entailment we only require that some sense of the conclusion is implied whatever sense is given to the premisses. In ‘co-reliable’ and ‘co-s-reliable’ we assume that, although all the premisses are only true ‘in some sense’, they are all interpreted consistently. Thus the \mathbf{S} operator acts on the conjunction of premisses rather than each premiss separately.

Notice that, when we are concerned with ‘reliable’ and ‘s-reliable’ entailment, two separate database facts α and β are not equivalent to the conjunction $\alpha \wedge \beta$. This is because we assume that uses of a predicate within a single formula are consistent; whereas, two distinct formulae which share one or more vague predicates may be true for different interpretations of these predicates.

2.5 THE SORITES ARGUMENT

I now consider the famous *sorites* paradox which is perhaps the best known example of the way that vagueness seems to defy logical analysis. Detailed accounts may be found in (Black 1970, Burns 1991, Williamson 1994). The original example of this paradox concerns the question of how many grains of sand constitute a ‘heap’ (‘sorites’ is Greek for heap) but I shall consider the directly analogous case of the bald man (which will be of more relevance to many academics).

Assume we have a theory in which arithmetic addition is defined. The following argument is then classically valid:

1. $\forall x[(\text{num_hairs}(x) = 0) \rightarrow \text{bald}(x)]$
2. $\forall x\forall y[(\text{bald}(x) \wedge (\text{num_hairs}(y) = \text{num_hairs}(x) + 1)) \rightarrow \text{bald}(y)]$
3. $\therefore \forall x[(\text{num_hairs}(x) = 10,000,000) \rightarrow \text{bald}(x)]$

Precisification semantics enables one to avoid this

paradox because in border-line cases we can say that a person is ‘bald’ under some interpretations of this concept but not under others. One can then weaken the second premiss to say: if an n -haired man is bald under all reasonable interpretations of ‘bald’ then an $(n + 1)$ -haired man is bald under at least one reasonable interpretation of ‘bald’:

$$\forall x \forall y [(\mathbf{U} \text{ bald}(x) \wedge (\text{num_hairs}(y) = \text{num_hairs}(x) + 1)) \rightarrow \mathbf{S} \text{ bald}(y)]$$

Under this interpretation, the argument is not valid according to the semantics for \mathbf{U} given above.

This analysis is perhaps not a complete solution of the paradox. One might argue that if someone with n hairs is bald ‘in some sense’ then someone with $n + 1$ hairs must be bald in some (slightly weaker) sense. So the paradox re-emerges with respect to the predicate ‘... is bald in some sense’. It may be argued that fuzzy logic gives a better account of sorites-style paradoxes because it allows the degree to which someone is bald to increase monotonically as hairs are lost (see e.g. (Goguen 1969). However, other problems with fuzzy logic were noted above (section 2.2).

3 ANALYTIC AND OBSERVATIONAL VAGUENESS

The $S5\mathbf{U}$ semantics given in section 2.1 assumes that every precisification assigns classical truth values (true or false) to every atomic predicate. This model gives a good account of vague concepts that correspond to a cluster of precise concepts, each having definite truth conditions. However, there are also cases where there seems to be no ambiguity in the meaning of the word but nevertheless there is indeterminacy as to when the concept should be applied. For instance, the predicate $\text{tall}(x)$ does not seem to have a variety of senses, although there is no specific height which qualifies a person as being tall. In judging the truth of a proposition such as $\text{tall}(\text{John})$ it seems perfectly reasonable to say in certain cases that its truth is indefinite. Of course different judges might still quibble about what was an indefinite case.

In $S5\mathbf{U}$ we can say that a proposition is indefinite in that it is neither true in every sense nor false in every sense: $\neg \mathbf{U} \phi \wedge \mathbf{S} \phi$. However we cannot say that a proposition is indefinite in every sense. This motivates the construction of a three-valued precisification semantics.

3.1 THREE-VALUED PRECISIFICATION SPACES

In a three-valued precisification model (Fine 1975, Pinkal 1995), a precisification is associated with a function mapping each atomic predicate and each tuple of arguments to a value in the set $\{\mathbf{t}, \mathbf{f}, \mathbf{i}\}$ (true, false and indefinite). A precisification which only assigns the values \mathbf{t} and \mathbf{f} is called *complete*. If a precisification p assigns the value \mathbf{t} whenever q assigns \mathbf{t} and \mathbf{f} whenever q assigns \mathbf{f} , then we write $p \geq q$ and say that p *extends* q .

Fine suggests that one should constrain the space of precisifications \mathcal{P} by what he calls the Resolution Principle (RP) requiring that whenever an atom indefinite at a precisification there is an extension of that precisification which makes it true and another that makes it false. In formal terms this means that for any precisification $p \in \mathcal{P}$ and any atomic proposition A , if $\llbracket A \rrbracket_p = \mathbf{i}$, then there must exist precisifications $q, r \in \mathcal{P}$ such that $q \geq p$, $r \geq p$, $\llbracket A \rrbracket_q = \mathbf{t}$ and $\llbracket A \rrbracket_r = \mathbf{f}$. RP means that all maximal elements of \mathcal{P} are complete. A further suggested constraint is that there should be a minimal precisification p_0 of which every world is an extension.

If RP is dropped and we consider models which have arbitrary sets of valuations we get a very general semantics within which a wide variety of operators can be defined. Such models will not be explored here but they may well prove useful in modelling aspects of vagueness, lack of knowledge or uncertainty.

Fine investigates a number of ways in which the semantics of complex propositions should be determined by a 3-valued precisification space. The simplest proposal is that the truth value of a complex proposition at a precisification p is determined by the truth values assigned according to 2-valued classical logic to all complete extensions of p . Specifically, ϕ is true at p if it is true at all complete extensions of p ; ϕ is false at p if it is false at all complete extensions of p ; otherwise ϕ is indefinite. This is very plausible semantics for propositions involving only classical connectives. Fine also suggests that a definiteness operator \mathbf{D} can be added to the language. He gives two possible semantic specifications for this operator. One is that the truth of a formula $\mathbf{D}\phi$ depends only on the value of ϕ at the minimal (i.e. most indefinite) world p_0 :⁴

$$p \Vdash \mathbf{D}\phi \quad \text{iff} \quad p_0 \Vdash \phi$$

According to this specification, the truth of $\mathbf{D}\phi$ is con-

⁴Under this definition, \mathbf{D} is essentially the same as Pinkal’s ‘ \square ’ operator.

stant over all precisifications.

Fine also suggests a more complex semantics, within which a proposition may be definite at some precisifications and not at others. This involves treating the **D** operator as a 3-valued truth functional operator whose truth-table is given in Table 1. The table also shows how the Boolean connectives can be treated as truth functional except in the case of a conjunction $\alpha \wedge \beta$ of two indefinite propositions. Here the truth value of the conjunction is not a simple truth-function of the values of the conjuncts because, even though α and β may be indefinite, it may be impossible for them to be true together (e.g. α might be ‘John is tall’ and β ‘John is not tall’). To make this precise, we stipulate that when $\llbracket \alpha \rrbracket_p = \mathbf{i}$ and $\llbracket \beta \rrbracket_p = \mathbf{i}$, then: $\llbracket \alpha \wedge \beta \rrbracket_p = \mathbf{i}$ iff there is some precisification $q \geq p$ (q is an extension of p) such that $\llbracket \alpha \rrbracket_q = \mathbf{t}$ and $\llbracket \beta \rrbracket_q = \mathbf{t}$; otherwise $\llbracket \alpha \wedge \beta \rrbracket_p = \mathbf{f}$.

ϕ	$\neg\phi$
t	f
i	i
f	t

$\alpha \wedge \beta$	t	i	f
t	t	i	f
i	i	i/f	f
f	f	f	f

ϕ	D ϕ
t	t
i	f
f	f

Table 1: Truth tables for \neg and \wedge and **D**.

It should be noted that the **D** operator is semantically very different from the **U** operator. **D** ϕ is true at a world/precisification just in case ϕ is true at that world/precisification, whereas **U** ϕ is true at a world if ϕ is true at that world for every precisification. At the semantic level it is easy to combine both operators into a single language. We just replace the 2-valued precisification space of **S5U** with the 3-valued space. In the combined language we can say things like **S**(\neg **D** tall(John) \wedge \neg **D** \neg tall(John)), meaning that ‘in some sense John is neither definitely tall nor definitely not tall’.

The object-level proof theory of this language is somewhat problematic. However, the semantics can easily be described in classical 1st-order logic in a similar fashion to the method suggested in section 2.1. This provides an indirect way of testing validity.

4 OBSERVABLES AND VAGUENESS

The three-valued precisification semantics allows one to model predicates which are vague even when their interpretation is unequivocal. However, it does not provide any means for specifying the way in which the

justification for applying a vague predicate depends on the state of the world. In this section I examine how the truth of a vague qualitative concept may depend on the value of some continuously variable observable parameter. To do this I consider a language which contains, as well as predicates (which may be more or less vague) functions corresponding to the values of observable properties.

4.1 OBJECTIVE MEASURES AND POSSIBLE WORLDS

It is plausible that the meanings of vague concepts depend to some extent on objective measurements. To formalise these I introduce a class of entities which will be called *observables*. These can be regarded as types of measurement which can be applied to objects. Each observable has a fixed arity, which is the number of objects to which it applies. To make the model theory simpler I shall assume that we have only monadic and dyadic observables. Examples of observables are **height** and **weight** (monadic) and **distance** (dyadic).

I consider each possible world to be associated with a complete assignment of real values to every n -ary observable for each n -tuple of arguments. Thus each world will assign a height and weight to each object and a distance to each pair of objects. This semantics for precise observables can readily be combined with a precisification semantics for vague predicates. A model for a vague language with observable functions is then a structure

$$\mathcal{M}_{VLO} = \langle W, P, I, \mathcal{R}, \delta, O_1, \sigma_1, O_2, \sigma_2, \rangle,$$

where: $\langle W, P, I, \mathcal{R}, \delta \rangle$ is a precisification model of the kind described in section 2.1; O_1 is a set of names of monadic observables; σ_1 is a function from $W \times O \times I$ to real numbers; O_2 is a set of names of dyadic observables and σ_2 is a function from $W \times O \times I \times I$ to real numbers.

Note that unlike the valuation of relations, the valuation of observable functions does not depend on the precisification. Thus these functions are objective (non-vague) properties of the world.⁵ The possible values of observables will normally be constrained by a theory Θ_{obs} , which may be regarded as a formalisation of physical laws. Thus the observable **distance** might be axiomatised to be a Euclidean metric.

⁵One could easily add vague functions to the language also. These would have values that could vary between different precisifications for the same world.

4.2 PREDICATE MEANINGS

In a setting where there is no uncertainty about the state of the world and no controversy regarding the sense of a concept, the justification as to whether a tuple of objects satisfies some predicate will appeal to the values of observables evaluated with respect to those objects. To model these justification conditions one may axiomatise precise predicates by specifying their necessary connections to observables.

For vague predicates the situation is more complex. Some vague predicates are uncontroversially linked with a single observable. In such a case the predicate may still be vague because the cut-off value for the predicate may be debatable. An example of this is the predicate **tall**. We can assume that this is determined purely by the observable **height**. (Of course some argument might arise in comparing heights when one person has a stoop or curly hair. However this kind of vagueness is best modelled with the **U** operator and can be separated from that which occurs even when the height measurement is uncontroversial.) Many vague predicates will depend upon more than one observable. For example whether a person is ‘big’ may depend on both their height and weight. A tall person might be described as big even if lighter than a shorter person. On the other hand if the tall person were only moderately tall but extremely thin one would be reluctant to call him ‘big’.

Another problem is that concepts such as tall are relative to a class. Thus a tall child is not as tall as a tall woman. This phenomenon is undoubtedly a significant feature of natural vague concepts however it will not be tackled in the current paper.

4.3 COMPARATIVES AND THE PROPERTIES OF VAGUE CONCEPTS

Despite the fact that the applicability of a vague predicate such as **tall** is indeterminate, comparative relations which seem to be derived from these predicates may be completely determined by observables. For instance the relation **taller** could be defined in terms of **height** as follows:

$$\Box(\forall x\forall y[\mathbf{taller}(x, y) \leftrightarrow (\mathbf{height}(x) > \mathbf{height}(y))])$$

Here the truth of the mathematical comparative relation ‘>’ is determined directly by a model of the form of \mathcal{M}_{VLO} .

It is also clear that willingness to ascribe the predicate tall to a person increases the greater their height; so, if any reasonable judge calls one person tall, then she

must regard all taller people as tall:

$$\mathbf{U}\Box(\forall x\forall y[(\mathbf{tall}(x) \wedge \mathbf{taller}(y, x)) \rightarrow \mathbf{tall}(y)])$$

Axioms involving comparatives provide a straightforward and very general way of constraining the meanings of vague concepts.

5 THE LOGIC OF CLARITY AND THE ‘DISTANCE’ BETWEEN WORLDS

Williamson (1994) proposes that vagueness should be described by means of a propositional operator, **C**, where **C** ϕ is read ‘It is clear that ϕ ’. The meaning of this operator is defined relative to a possible worlds semantics in which there is a metric over the set of possible worlds giving a ‘distance’ between any two worlds. We then say that a proposition **C** ϕ is true at a world w (i.e. ϕ is *clearly* true at w) just in case ϕ is true at all worlds within a certain distance of w . This distance is taken as providing a *margin of error* which ensures that ϕ is clearly true since it is true not only at the actual world w but at all worlds which differ from w by some small amount. In this section I look at the relationship between Williamsons’s logic of clarity, precisification spaces and observables.

It seems reasonable to suppose that the distance between possible worlds is determined by observables and that the margin of error required for an atomic proposition to be ‘clearly’ true will vary according to the meaning of its predicate. Thus each predicate can be associated with a metrical function of observables and a margin. Together these constitute a relevant *measure* for that predicate. It is implausible that the precise specification of this measure is a fixed property of the predicate. Rather these will vary according to a spectrum of ways in which the predicate could be interpreted. Moreover, one might expect that each distinct predicate would be associated with its own space of measures. However, Williamson considers an idealisation of this picture in which, within a given model, the certainty of all predicates is determined by a single metric and margin on the possible worlds.

The fact that we do not need separate measures for each predicate may seem odd; but the notion of a metric is very general and can encompass distance within a multi-dimensional space of possible variations of atomic predicates.

Williamson considers specifying the semantics of **C** in terms of *fixed margin models*. Such a model is a

quadruple $\langle W, d, \alpha, \llbracket \cdot \rrbracket \rangle$, where W is a set, d a metric⁶ on W , α a non-negative real number, and $\llbracket \cdot \rrbracket$ is a mapping of formulas to subsets of W such that for all formulas ϕ, ψ :

- $\llbracket \neg\phi \rrbracket = W - \llbracket \phi \rrbracket$
- $\llbracket \phi \wedge \psi \rrbracket = \llbracket \phi \rrbracket \cap \llbracket \psi \rrbracket$
- $\llbracket \mathbf{C}\phi \rrbracket = \{w \in W : (\forall x \in W)[(d(w, x) \leq \alpha) \rightarrow x \in \llbracket \phi \rrbracket]\}$

Under these conditions it can be shown that \mathbf{C} behaves exactly as the \Box operator of the modal logic $S4$. However, Williamson argues that this operator is too strong in that it obeys the schema $\Box\phi \leftrightarrow \Box\Box\phi$. This would rule out the possibility of second-order vagueness as exemplified by a claim such as ‘John is clearly tall but not clearly clearly tall’.⁷ In order to allow for non-redundant iteration of \mathbf{C} , this operator can be interpreted in terms of a *variable margin model* where

$$\llbracket \mathbf{C}\phi \rrbracket = \{w \in W : (\exists\delta > \alpha)(\forall x \in W)[(d(w, x) \leq \delta) \rightarrow x \in \llbracket \phi \rrbracket]\}.$$

If validity is identified with truth in all such models then \mathbf{C} is constrained by exactly the same schemata as the \Box operator of the modal logic KT . Since the schemata K and T are intuitively correct properties of any plausible clarity operator and any refinement of the idealised semantics could only weaken the logical properties of \mathbf{C} , Williams argues that KT provides an adequate specification of the logic of clarity.

The semantics of clarity is orthogonal to the idea of precisification in that to evaluate $\mathbf{C}\phi$ at a some world/precisification, we look at the valuations for a predicate at a neighbouring worlds with respect to the same precisification. To take account of this model-theoretically, one could generalise the notion of precisification to include a measure on possible worlds as well as valuation on predicates. Since models include arbitrary precisifications, every possible combination of predicate valuations and measures will correspond to a point in the precisification space. The reason that measures are considered part of precisifications is that I take them to be part of the sense ascribed to a predicate. Hence if we combine \mathbf{C} with the unequivocal operator \mathbf{U} , the \mathbf{U} operator asserts that a proposition is true for every assignment *and* every measure (not just every assignment for a fixed measure).

By adding measures to a precisification model (either the simple $S5\mathbf{U}$ model or the 3-valued version) one can construct a semantics within which \mathbf{U} and \mathbf{C} (and in the 3-valued case also \mathbf{D}) are definable. This we can

say things like ‘The eruption of Krakatoa was clearly a disaster in every sense’. This means that for every precisification the proposition ‘The eruption of Krakatoa was a disaster’ is true for every possible world which is within the required margin of similarity of the actual world (the appropriate margin being dictated by the precisification).

The logic of \mathbf{C} can also be related to the values of observables. Since the distance measure between possible worlds is determined by observables, in a language including observable functions it is possible to constrain the behaviour of the clarity operator to take account of this dependence. One way to do this is in terms of comparatives. Analogously to the example given in section 4.3 one could state the following axiom ensuring that if someone is clearly tall according to some precisification then every taller person is also clearly tall:

$$\mathbf{U}\Box(\forall x\forall y[(\mathbf{C}\text{tall}(x) \wedge \text{taller}(y, x)) \rightarrow \mathbf{C}\text{tall}(y)])$$

6 CONCLUSION

This paper has drawn on previous work in philosophical analysis of vagueness in order to construct formalisms which may prove useful for handling vague concepts within AI knowledge bases. Although the semantic models I have presented are not completely original, they seem to be little known to AI researchers. This is perhaps because they have hitherto been developed principally for the exposition of philosophical arguments rather than for practical purposes. Consequently, in order to specify useful representation languages for dealing with vague information, certain modifications are required.

By the seemingly trivial generalisation of precisification spaces to include all arbitrary valuations of relations rather than just ‘admissible’ ones I have provided a logically neutral model structure within which analytic (i.e. necessary) properties of both vague and sharp predicates can easily be stated. I have suggested that philosophical debates concerning the correct definition of entailment within a vague language are irrelevant to KR. Moreover each of variety of possible entailment relations can yield useful information about possible inferences among propositions involving vague concepts.

I have distinguished the vagueness that occurs where a predicate has a cluster of different senses from the vagueness associated with the lack of a clear cut off point for a qualitative predicate with respect to continuously variable objective measurement. I believe that these different kinds of vagueness require quite

⁶i.e. it satisfies the usual definition of a metric.

⁷I am not really sure how much sense this makes.

distinct model theoretic treatment but can nevertheless be incorporated within a single, albeit rather complex, semantics. Further elaboration of the details of this semantics is the subject of ongoing work. My approach contrasts with that of other treatments (such as fuzzy logic) where different types of vagueness (and sometimes also uncertainty) are treated uniformly. I think that my analysis suggests that such an assimilation is unlikely to give a satisfactory of the various aspects of vagueness.

Vagueness is an extremely subtle phenomenon which is very hard to get hold of within a logical framework. Consequently there are likely to be many deficiencies in the semantical models I have proposed. Nevertheless I believe that my semantic analysis is in many respects correct and can provide a framework which will enable AI systems to take account of certain important aspects of vagueness. I also believe that a satisfactory account of vagueness is essential to the solution of many crucial problems in AI. Vagueness is a fundamental property of natural languages and consequently must be taken seriously in any KR system that attempts to manipulate natural concepts.

Acknowledgements

This work was supported by the EPSRC under grant GR/K65041.

I would like to thank professor A.G. Cohn for helpful discussions on the subject matter of this paper.

References

- Black, M.: 1970, Reasoning with loose concepts, *Margins of Precision*, Cornell University Press, Ithaca, chapter 1, pp. 1–13.
- Burns, L. C.: 1991, *Vagueness: An Investigation Into Natural Languages And The Sorites Paradox*, Vol. 4 of *Reason and Argument*, Kluwer.
- Cohn, A. G. and Gotts, N. M.: 1996, A mereological approach to representing spatial vagueness, in J. D. L C Aiello and S. Shapiro (eds), *Principles of Knowledge Representation and Reasoning*, Proc. 5th Conference, Morgan Kaufmann, pp. 230–241.
- 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, London.
- Elkan, C.: 1993, The paradoxical success of fuzzy logic, *Proceedings of the National Conference on Artificial Intelligence (AAAI-93)*, pp. 698–703.
- Fine, K.: 1975, Vagueness, truth and logic, *Synthese* **30**, 263–300.
- Goguen, J.: 1969, The logic of inexact concepts, *Synthese* **19**, 325–373.
- Guha, R. V. and Lenat, D. B.: 1990, Cyc: a mid-term report, *AI Magazine* **11**(3), 32–59.
- Hughes, G. E. and Cresswell, M. J.: 1969, *An Introduction to Modal Logic*, Methuen, London.
- Kripke, S.: 1959, A completeness theorem in modal logic, *Journal of Symbolic Logic* **24**, 1–14.
- Lehmann, F. and Cohn, A. G.: 1994, The EGG/YOLK reliability hierarchy: Semantic data integration using sorts with prototypes, *Proc. Conf. on Information Knowledge Management*, ACM Press, pp. 272–279.
- Lukasiewicz, J. and Tarski, A.: 1930, Untersuchungen über den aussagenkalkül, *Comptes rendus des séances de la société des sciences et des lettres de Varsovie* **3**(23), 1–21. Reprinted in A. Tarski, *Logic, Semantics, Metamathematics*, 1956, Clarendon Press, trans. J.H. Woodger.
- Pinkal, M.: 1995, *Logic and Lexicon: the semantics of the indefinite*, Vol. 56 of *Studies in linguistics and philosophy*, Kluwer, Dordrecht. translated by Geoffrey Simmons.
- Seegerberg, K.: 1973, Two-dimensional modal logic, *Journal of Philosophical Logic* **2**, 77–96.
- Shafer, G. and Pearl, J. (eds): 1990, *Readings in Uncertain Reasoning*, Morgan Kaufmann, Palo Alto.
- van Frassen, B. C.: 1968, Presupposition, implication and self-reference, *Journal of Philosophy* **65**, 136–52.
- van Frassen, B. C.: 1969, Presupposition, supervaluations and free logic, in K. Lambert (ed.), *The Logical Way of Doing Things*, Yale University Press, New Haven, chapter 4, pp. 67–91.
- Williamson, T.: 1994, *Vagueness*, The problems of philosophy, Routledge, London.
- Zadeh, L.: 1965, Fuzzy sets, *Information and Control* **8**, 338–353.
- Zadeh, L. A.: 1975, Fuzzy logic and approximate reasoning, *Synthese* **30**, 407–428.