

What is a forest?

On the vagueness of certain geographic concepts*

Brandon Bennett[†]
School of Computing, University of Leeds, Leeds LS2 9JT, UK
brandon@comp.leeds.ac.uk

February 21, 2002

Abstract

The paper examines ways in which the meanings of geographical concepts are affected by the phenomenon of vagueness. A logical analysis based on the theory of supervaluation semantics is developed and employed to describe differences and logical dependencies between different senses of vague concepts. Particular attention is given to analysing the concept of ‘forest’ which exhibits many kinds of vagueness.

1 Introduction

Vagueness is ubiquitous in spatial and geographical concepts and tends to persist even where steps are taken to give precise definitions. For example, in the guide book to the Ordnance Survey’s *Land-Line* data-set (Ordnance Survey 2000), a ‘road’ is defined as: “A metalled way for vehicles.” This does tell us something about what is meant by ‘road’ but the definition is still vague in many respects. We may be unsure about what surfaces count as ‘metalled’ — neither condition of the surface nor any restrictions on its spatial extent is specified. And the term ‘way’ could be understood in many more or less general senses. ‘Vehicle’ is also a very general term. The OS definition of road would seem to apply to bicycle paths, which may not be intended. It also seems to rule out cobbled or paved streets which one might expect to be classified as roads.

An understanding of the meanings of vague geographical concepts is relevant to many practical problems that involve determining and allocating land types. For instance, if we want to answer a question such as ‘How rapidly is the forested area of the earth shrinking?’ the problem of demarcating forest areas is central. Similar problems apply to the identification and classification

*A slightly modified version of this paper was published in *Topoi* **20**: 189–201, 2001.
©2001 *Kluwer Academic Publishers*.

[†]This work was supported by the EPSRC under grant GR/M56807.

of ‘deserts’ (and the problem of measuring and monitoring the progress of desertification). Information sources made available by the USGCRP (n.d.) include multiple data sets on topics such as soil, precipitation, vegetation, temperature, land cover, etc. Several of these data sets have ‘desert’ as a specific class, each with its own method of compilation and concept of what actually constitutes a desert — absence of vegetation, annual rainfall below a particular (and varying) threshold, number of months per year exceeding a precipitation threshold, type of soil, ecosystem characteristics, etc. The result is a set of maps that produces a very different distribution of deserts according to which classification you choose to use at any one time. Moreover, the concept of desert can itself be variously classified into sub-types (e.g. ‘desert, mostly bare’, ‘sand desert, partly blowing’, ‘other desert and semi-desert’, ‘polar desert’, ‘tropical desert’). A further example of the importance in environmental modelling of clarifying vague terms is provided by Alker, Joy, Roberts and Smith (2000) who consider issues in defining the concept of a ‘Brown-field’ which is often used in formulating development policies.

In this paper I shall explore a possible approach to the logical and semantical analysis of vagueness and apply this to specific problems in the definition of geographical concepts. I focus in particular on the example of the concept ‘forest’, looking at different ways in which the term can be interpreted and how these effect the determination of the spatial extensions associated with features classified as forests. Before colouring the reader’s perception of the issues by suggesting a particular theoretical approach, I end the introduction by enumerating a list of questions, each of which addresses one of the main aspects of vagueness associated with the term ‘forest’, and hence has no clear-cut answer. I shall return to these later and examine them in the light of a *supervaluationist* account of vagueness.

1. Is a forest a natural feature or one determined by convention and legality?
2. Does ‘forest’ refer to an integral feature or can it be applied to an arbitrary region of land?
3. What type of vegetation can constitute a forest? (i.e. what species and how big must they be?)
4. How dense must the vegetation be?
5. How large an area must a forest occupy?
6. Are there any constraints on its shape?
7. Must a forest be self connected, or can it consist of several disjoint parts?
8. Must it be maximal or could it share a border with another region of forest?
9. Is a clearing a part of or a hole in a forest?
10. Are roads and paths going through a forest parts of the forest?
11. How should seasonal and other temporal variations be taken into account?
12. If part of a forest is felled and subsequently re-grown, does it remain part of the forest throughout?¹

¹Accompanying its ‘Land Usage of the World’ data the web site www.ecoworld.com gives

2 The Nature of Vagueness

To get a purchase on the multiplicity of ways in which one might bend or stretch the meaning of a concept such as ‘forest’, we will need a theoretical framework within which properties of the variable meanings of vague expressions can be clearly articulated. Unfortunately, the concept of vagueness itself is not completely clear, so I shall first try to be precise about the phenomenon I wish to analyse.

I regard vagueness as a lack of clearly defined criteria for the applicability of a concept. Thus, it is a property of language not of the world itself. Typical examples of vague propositions in the geographical domain are: ‘All *mountains* are *very high*’; and ‘*Near* the *marsh* is a *dense thicket*’. The words given in italics are the principal sources of vagueness. ‘Mountains’, ‘marsh’ and ‘thicket’ are vague feature classifiers: they do not have precise, universally acknowledged definitions; and consequently the spatial extent occupied by such features is potentially controversial. ‘High’ and ‘dense’ are adjectives, which give some indication of physical properties of a feature but do not specify any definite measurable requirement. ‘Very’ accentuates vague adjectives but does not make them any more definite.

Vagueness should be sharply differentiated from *uncertainty*, which is a distinct (though interacting) phenomenon. Uncertainty arises from lack of exact knowledge about an object or situation, and is thus an epistemic state rather than a feature of language. Although modelling of uncertainty is extremely important in the processing and interpretation of spatial information, it will not be considered in the current paper. I shall assume we are dealing with idealised data which is completely certain and accurate. Although philosophers generally have little trouble separating the notions of vagueness and uncertainty this is not always the case in certain branches of science.

Vagueness can often lead to uncertainty in that where a concept such as ‘forest’, ‘desert’ or ‘swamp’ is vague we will in many cases be uncertain how to demarcate the spatial extension of entities to which these concepts apply. On the other hand, if we are not completely certain of the exact details of some information we want to report, we may employ vagueness as a means of increasing the certainty of what we say, while at the same time conveying a sense of imprecision. For example, a statement such as ‘The chair is *in the corner of* the room’ is vague but can often be said with certainty, whereas an exact specification of the location of a chair (or even a range of possible locations) will typically be uncertain. This example also illustrates the fact that vagueness is not merely a defect of language; it also often facilitates communication without the cumbersome language required to achieve precision.

Vagueness is also sometimes confused with generality. If I say ‘I shall see you again later this month’, this is an example of generality, since the claim can be fulfilled in many alternative ways. However, it is not vague, since ‘later this month’ refers to a precise period of time (I assume all relevant events take

the following definition of forest: “Forest: Land under natural forests or planted stands of trees. Also includes logged areas to be replanted in the near future, after logging.”

place within a single time-zone). But if I say ‘I shall see you in a few weeks’, this is vague (as well as general) since there is no hard and fast definition what periods of time can be described as ‘a few weeks’. Generality *per se* poses no real problems for the logician, since it is handled perfectly well by classical logic.

Having, separated it from some closely related issues we must ask whether vagueness *qua* ‘lack of definite criteria of application’ is a single uniform property of linguistic expressions or whether it comes in different varieties. I argue that it is useful to distinguish between at least two kinds, or modes, of vagueness, which require somewhat different logical description. I call these modes *conceptual vagueness* and *sorites vagueness*.

Conceptual vagueness occurs where there is no single completely adequate definition of a conceptual term. Certain requirements may be clearly identifiable, whereas for other conditions it is arguable whether or not they are necessary. Certain combinations of these conditions may capture typical senses of the term but none is representative of all possible senses. Thus if we take the intersection of plausible definitions we get a concept that is much too strict (perhaps even unsatisfiable), whereas if we take their disjunction we get a concept that is overly general. This kind of vagueness is closely related to *ambiguity*. If a word is ambiguous, it has two or more distinct senses that are clearly distinguishable. However, a conceptually vague term corresponds to a complex cluster of many overlapping senses, such that we cannot say exactly what senses make up the cluster. Moreover, one can meaningfully use a conceptually vague term without being committed to any one of its possible precise interpretations.

Sorites vagueness is the kind of indeterminacy that affects the thresholds at which we assert properties such as ‘tall’, ‘large’ or ‘heavy’. Sorites vague predicates divide entities with respect to some continuously varying quantity, without being committed to any specific boundary value. In such cases I suggest that there is nothing vague about the conceptual content of the predicate — it is just that the content is unspecific about the precise boundaries of application. Indeed these boundaries could not be made precise without changing the meaning of the concept.

The two types of vagueness are not mutually exclusive. Indeed many natural terms are affected by both. In a pure case of sorites vagueness it is uncontroversial which factors are relevant to the ascription of a term or how those factors should be measured — it is only the threshold that is at issue. However, in many cases the relevant factors *and* suitable thresholds are at issue. For example to precisely interpret the concept ‘tall man’ we have first to decide how we are to measure the height of a man: must he remove his shoes and hat? what about hair and prosthetic limbs? what about posture? Once we have resolved these conceptual issues we then still have to deal with sorites vagueness in setting the threshold for tallness.

Thus, according to my view, the elucidation of vague terms requires two distinct modes of analysis. First, conceptual vagueness should be tackled by identifying substantial points of controversy among possible definitions. Artificial concepts can then be introduced which correspond to particular choices with regard to these interpretational issues. If this analysis were carried out ex-

haustively (and I don't want to take a view on whether this is always possible) we would end up with a (possibly very large) set of alternative definitions each corresponding to a clarified sense of the original free from conceptual vagueness. However, these artificially defined concepts may still be affected by sorites vagueness. There is little point in attempting to eliminate this more essential form of vagueness by definition, since any such definitions — e.g. stipulating that one sense of 'tall woman' is equivalent to 'woman whose height exceeds 5'10" ' — would be more or less arbitrary and would not serve to explicate the relevant vagueness. However, for certain practical purposes it may be useful to make such stipulations. More significantly we can place quite strong constraints on consistent usage of sorites vague concepts by specifying exactly which objectively measurable property the threshold of applicability lies on. This will enable us to order senses of sorites concepts according to where they place the cut-off point.

Both types of vagueness also interact strongly with contextual phenomena of various kinds. Many concepts exhibit some form of *contextual variability*. For example 'large' in the context of 'large pond' has a different interpretation from in the context of 'large lake'. In cases such as this we see that a sorites concept may be affected by its context so that location of its albeit vague threshold is shifted. The range of possible interpretations for a conceptually vague concept can also be affected, not so much by their immediate syntactic context but by their more general context within a particular exchange of information.

Because it is largely independent of vagueness, issues of contextual vagueness will not be addressed in the current work. I shall assume that context can be eliminated or ignored. For instance we might suppose that some transformation can be carried out that replaces contextually variable concepts with non-contextual concepts and explicit constraints; or, more simply, we could just consider composite concepts such as 'tall man' and 'tall child' as if they were syntactically atomic (although perhaps related by certain meaning postulates, which could be formalised). Despite the fact that in many cases vagueness seems to be separable from context there may be cases where this distinction is blurred. A close connection between the phenomena is born out by the fact that formalisations of the logic of context (see e.g. (McCarthy 1993)) have much in common with the supervaluation approach to vagueness.

3 Formalising the Logic of Vagueness

In modelling vagueness in geographical information researchers have tended to adopt techniques which have been used within computer science. The most common approaches here are based on *multi-valued* logic (Łukasiewicz and Tarski 1930) and its generalisation *Fuzzy* logic (Zadeh 1965, Zadeh 1975, Goguen 1969, Dubois and Prade 1988). For example, (Foody 1992) and (Usery 1996) apply a fuzzy representation to geographical features can be found. An extensive discussions of the pros and cons of fuzzy logic can be found in (Elkan 1993) and a more philosophically oriented consideration of these issues can be found

in (Williamson 1994). 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. Moreover, in characterising the vagueness of geographical concepts, fuzzy logic is a palpably blunt tool. The variety and nuances of possible interpretations of a term such as ‘forest’ cannot be adequately characterised in terms of a probability that a particular piece of land should be counted as a forest.

The approach which I propound in this paper is to model vagueness in terms of *supervaluation semantics* (Fine 1975), which I believe can provide a much deeper insight into the nature of the vagueness intrinsic in many geographical concepts. The fundamental idea upon which this theory is based is that a vague language is one which can be made precise in many different and sometimes incompatible ways. A way of making a language precise is called a *precisification*. Each precisification p is identified with a precise and consistent interpretation, I_p , of the vocabulary of the language. In the simplest case this would be a classical propositional or 1st-order model. A supervaluation model then consists simply of a set of precisifications. Given a supervaluation model \mathcal{V} , a proposition which is true under every interpretation $I_p \in \mathcal{V}$ is called *super-true* or — in my own terminology — *unequivocally true*.

Supervaluation semantics by itself does not add anything interesting to logic. It is easy to see that those formulae that are unequivocally true in every model are just the classically valid formulae. However, the semantics does provide a framework within which we can define operators that articulate certain aspects of the logic of vagueness.

One possibility is to take a modal approach and represent vagueness in terms of propositional operators (Bennett 1998). $\mathbf{U}\phi$ means that ϕ is *unequivocally true* — i.e. true for all precisifications; $\mathbf{S}\phi$ can be read ‘ ϕ is *in some sense* true’ — i.e. true for some precisification. \mathbf{S} is the dual of \mathbf{U} and thus can be defined by $\mathbf{S}\phi \leftrightarrow \neg \mathbf{U} \neg \phi$. The inference rules justified by this interpretation are those of the simplest modal logic, $S5$, where \mathbf{U} takes the place of the usual modal \Box operator.

We can now qualify assertions according to whether they hold in some or all precisifications. For example we might write $\mathbf{S}[\text{Wood}(\textit{‘Woodsley Clough’})]$ or $\mathbf{U}[\text{Wooded}(\textit{‘parcel1’})]$. We can also use these operators to specify dependencies between the meanings of vague concepts. Thus $\forall x[\text{Copse}(x) \rightarrow \mathbf{S}\text{Wood}(x)]$ means anything which is a copse (i.e. a small group of trees) is in some sense a wood. Similarly, $\forall x[\text{Wood}(x) \rightarrow \mathbf{S}\text{Forest}(x)]$ captures the intuition that any wood is arguably a (small) forest. If a copse is in some sense a wood and a wood is in some sense a forest, this does not mean that a copse is in some sense a forest; and indeed according to supervaluation semantics the formula $\mathbf{U} \neg \exists x[\text{Forest}(x) \wedge \text{Copse}(x)]$ is consistent with the previous two formulae. This illustrates the ability of the theory to model the blurring of concepts, while still

imposing strong constraints upon them.

In many geographical applications we will want to employ artificial concepts that are designed for empirical classification of land types and features. These will in many cases be precise sharpened versions of natural terms. The supervaluation operators enable us to relate these artificial concepts to their vague natural language counterparts. For instance, the following formula asserts that `Forest1` is a more precise version of the concept `Forest`:

$$\forall x[\text{Forest1}(x) \rightarrow (\mathbf{S} \text{Forest}(x))] \wedge \forall x[(\mathbf{U} \text{Forest}(x)) \rightarrow \text{Forest1}(x)]$$

By specifying such axioms, ‘soft’ constraints are placed on the meanings of natural concepts. Classifications in terms of artificial concepts can be combined with information containing natural concepts.

Supervaluation semantics allows one to specify a number of different entailment relations of varying strength. In (Bennett 1998) I defined seven different entailments each of which has a different force. The weakest of these is what I call ‘arguable’ entailment, which holds if there is any sense of the concepts in the formulae under which the implication corresponding to the entailment holds. This gives us entailments that hold under very flexible (perhaps even inconsistent) interpretations of the concepts involved. The strongest is ‘reliable’ entailment, which holds if: whatever senses the premisses are interpreted under, the conclusion holds in every sense. This can be used to derive entailments which must hold despite the presence of vagueness. For instance $\mathbf{S} \text{Desert}(x) \rightarrow \mathbf{U} \neg \text{Marsh}(x)$ might hold even where `Desert` and `Marsh` are very vague predicates. This ability to derive secure consequences involving vague concepts is perhaps the main advantage of supervaluation semantics over fuzzy logic, where fuzzy concepts cannot support completely reliable inferences.

Although modal operators allow many logical properties of vague concepts to be expressed, they do not provide any way of referring directly to individual precisifications or to classes of similar precisifications. But in order to carry out a detailed analysis of different senses of vague geographical concepts we shall often want to relate natural terms (such as ‘forest’, ‘desert’ etc.) to artificially sharpened concepts which give a more precise and objective characterisation of a geographical feature.

In a general *reified* supervaluation semantics we could associate arbitrary propositions with precisification variables and constants. Thus, $\text{InPrec}(p, \phi)$ would assert that ϕ is true according to precisification p . At the expense of some elegance we can achieve the same expressive power by simply supplementing each predicate and function of an ordinary 1st-order language by an additional argument place. For clarity I shall write this as a prefix to the predicate or function. For instance a `Swamp(x)` would be replaced by $p: \text{Swamp}(x)$ saying that, in precisification p , x is a swamp. If we use this approach we need not worry about axiomatising the logical predicate `InPrec`.

Since a precisification fixes the meanings of all the vague vocabulary of a language, a classification which makes precise only part of the vocabulary may be common to a class of precisifications. In a formalism with reified precisifications,

we can model this by introducing predicates of precisifications. For example, $\text{UNESCOF}(p) \leftrightarrow \Phi(p)$ might mean that the predicate UNESCOF applies to those precisifications satisfying some precise formal specification Φ of the UNESCO forestation classification given in Table 1 (towards the end of this paper). In this way any refined sense of a vague term may be identified with some subset of the space of all precisifications.

Given that the previous section made much of the distinction between between conceptual and sorites vagueness, I ought to explain how this is manifest in supervaluation semantics. In fact according to my analysis the difference lies not in the abstract semantics but only in the types of axioms that constrain the two modes of vagueness.

Whereas conceptual vagueness is analysed by sharpened definitions and axioms relating these definitions within the space of precisifications, sorites vagueness is analysed by specifying the observational properties relative to which a concept has a vague threshold. A straightforward and general way of specifying the relevant property is via the definition of comparative relations. For instance the relation taller could be defined in terms of height by:

$$(\forall pxy[p:\text{taller}(x, y) \leftrightarrow (p:\text{height}(x) > p:\text{height}(y))])$$

Clearly, 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:

$$\forall pxy[(p:\text{tall}(x) \wedge p:\text{taller}(y, x)) \rightarrow p:\text{tall}(y)]$$

This encodes a crucial consistency property on precisifications regarding their precise interpretation of the adjective ‘tall’.

By comparing the extensions of sorites properties according to different precisifications we can define an ordering on the precisification space. This is similar to the metrical model for *margins of error* model of vague concepts suggested by Williamson (1994). Further discussion of this can be found in (Bennett 1998).

4 Spatial Vagueness

Supervaluation semantics is a very general approach to vagueness but it can only be useful for reasoning about a specific domain if the peculiar logic of that domain is adequately modelled. In the next section I shall turn to the detailed analysis of the concept of ‘forest’; but prior to that it will be useful to look at the more general question how vagueness affects the determination of spatial extensions. The application of supervaluation semantics to spatial concepts is relatively undeveloped, although it has recently been given some attention (Lewis 1993, McGee 1997, Kulik 2000, Varzi 2001)

The spatial properties that are easiest to understand semantically are those that can be defined in terms of properties of points — i.e. their extension consists of all points satisfying some given condition. Examples of such concepts are ‘the region of the Earth that is more than 1000m above sea-level.’ However, in general, a ‘region property’ will be associated with a property of a whole area or




Figure 1: Models of Vagueness and Extension

volume, which cannot be explicitly reduced to properties of individual points. For example a ‘lake’ is not simply made up of the set of points which are covered by water, it is rather a particular *maximal connected* set of water covered points. Indeed, maximal connectedness is one of the most important factors that enable us to individuate geographical features from attributed point data. However, only very basic types of feature can be regarded simply as maximal connected sets of points exhibiting a given property. Typically, whether a set of spatial points can be taken as the extension of a feature of a given type, is dependent on much more complex constraints (consider e.g. how we differentiate lakes from other hydrological features or how we might characterise a ‘building’). Some of these requirements may not even relate to physical properties (e.g. a ‘listed building’).

A further issue that complicates the identification of regions with sets of points is the status of boundary points. For many spatial concepts it is not clear whether boundary points should be counted as included in the region to which they refer. This may be regarded as an example of conceptual vagueness. However, it is also a general ontological issue which applies to spatial concepts that are not in other ways vague.

The ‘Egg-Yolk’ theory (Lehmann and Cohn 1994, Cohn and Gotts 1996a, Cohn and Gotts 1996b) directly models the notion of a vague (or uncertain) region in terms of its maximal and minimal possible extensions. The maximal extension is called the ‘egg’ and the minimal is the ‘yolk’, which is required to be a part of the egg (see Figure 1a). (The case where the yolk is equal to the egg is allowed, such cases corresponding to ‘crisp’ regions.) This analysis is simple and supports an account of some significant inferences involving relationship between vague regions. However, it cannot handle complex constraints on a region’s possible extensions between its maxima and minima. For instance, although a vague region such as an area of marshland might have maxima and minima as illustrated in Figure 1a, the area within the dotted line might not correspond to any reasonable precise interpretation of ‘marshland’.

Supervaluation semantics is much more general in that it has the potential to model arbitrary constraints on the distribution of possible extensions, as illustrated in Figure 1b. However, the possible extensions of natural vague concepts will not be completely chaotic since, according to supervaluation theory, they

correspond to a cluster of precise concepts with similar meanings. In the case of a purely sorites vague concept, where the vagueness is in the choice of a suitable threshold for some observable, the possible extensions will typically (though not necessarily) be contoured as shown in Figure 1c. Each contour corresponds to a more or less strict sense of a spatial concept. For instance different definitions of ‘marshland’ may require more or less water to be present. Where we have mixed vagueness we will have several sets of contours each corresponding to varying the threshold for some conceptually unambiguous but still sorites vague concept.

5 A Supervaluationistic ‘Forest’

We are now ready to employ the supervaluation theory to carry out a detailed analysis of geographical feature descriptions. Each type of geographical feature has many idiosyncracies in its particular meaning and in the ways that its meaning can be stretched or tightened to suit different purposes. Nevertheless, once an abstract logical analysis is given, forms of vagueness can be diagnosed that are present in a wide variety of geographical terms. Thus, the supervaluationistic dissection of ‘forest’ will serve as an example of how one might analyse similar concepts such as ‘desert’, ‘marsh’, ‘mountain’ or ‘lake’.

One of the most important aspects of the conceptual vagueness of the term ‘forest’ is the ambiguity between forests conceived of as natural features and forests as parcels of land upon which is legally or conventionally conferred the status of being a forest. Although it may be argued that forests are always originally identified with some natural feature, once forests are named (and thus probably also owned) additional conventional and legal mechanisms may be employed to individuate forests. Smith (1995, 2000) has investigated the ontology of conventional regions of this kind, which he calls *fiat* regions.

In axiomatising the vague term ‘forest’ it is clear that the natural and fiat interpretations will obey rather different axioms. Hence, any adequate analysis should split this concept into two specialisations. The following axioms, which employ the reified precisification notation, ensure that in any precisification `Fiat_Forest` and `Natural_Forest` are sub-concepts of `Forest` and that all forests are of one of these two types (they do not rule out the possibility that something may be both):

- $\forall px[p:\text{Fiat_Forest}(x) \rightarrow p:\text{Forest}(x)]$
- $\forall px[p:\text{Natural_Forest}(x) \rightarrow p:\text{Forest}(x)]$
- $\forall px[p:\text{Forest}(x) \rightarrow (p:\text{Fiat_Forest}(x) \vee p:\text{Natural_Forest}(x))]$

Though free from a certain ambiguity, the predicates `Fiat_Forest` and `Natural_Forest` are still extremely vague, each will correspond to a wide range of possible senses and further subdivisions and axioms will be required to explicate these. In the rest of the analysis I shall deal only with ‘natural’ forests, since these seem to be vague in a greater variety of ways; however, the semantics of fiat forests is no doubt also very complex. To avoid cumbersome terminology,

the word ‘forest’ shall henceforth be used to mean ‘natural forest’ and the formal predicate **Forest** shall be used in place of **Natural_Forest**.

In clarifying the concept of ‘(natural) forest’ we immediately encounter a second fundamental ambiguity that affects this and many similar geographical concepts. When used with an article (‘a forest’ or ‘the forest’) the term typically refers to a particular integral feature whose boundary (albeit vague) is determined by the meaning of the concept. However, it can also be used in an adjectival sense to describe an arbitrary region as ‘forest’. These two uses are not really due to vagueness but rest on a logical distinction that ought to be explicit in any ontology of geographical descriptions.

Though ontologically distinct, features and corresponding land-type concepts have strong logical interdependence which must be formally specified (see (Eschenbach 2000)). Let us use the predicate **Forest** as a vague feature type and **Forested** as the corresponding vague land-type classifier and see what axioms one would expect to link the two concepts. The situation is complicated by the fact that **Forest** refers to a feature which is a three-dimensional material object consisting of trees distributed in space, whereas **Forested** is a predicate of regions of land. To relate these two types of entity we will need to characterise the mapping between a feature and its *terrestrial extension*, which I assume to be a two-dimensional region.

We might be inclined to say that a region is ‘forested’ just in case it is part of the extension of some forest. However, this definition suffers from a problem of granularity, since the extension of a forest may include pockets which are not at all forested or are too small to even be legitimate candidates for such a description. One might hope to avoid this problem by taking forested as the more basic property and then defining a forest as a feature whose terrestrial extension is a maximal self-connected forested region. However, since ‘forested’ is a land type, the objects that fall under this predicate are terrestrial regions not physical objects like forests. Clearly a type of physical object cannot be defined from a predicate whose arguments are purely spatial: in order to carry through a definition we shall have to add some physical ingredient.

Given any region we can determine, without excessive controversy, the vegetative material that is present in that region. It then seems reasonable to assume that if this area is the terrestrial extension of a forest (*qua* natural feature) then all properties of that forest supervene on properties of that quantity of vegetative material.² Thus the properties of forests (including their identity criteria) are determined by the vegetation they comprise, which in turn is determined by their terrestrial extension.

It must be noted that the terrestrial extension of a forest is not determined directly by its vegetation. That is, if we were to simply shrink-wrap the vegetation and project this volume vertically onto the earth’s surface we would define extension at too fine a grain size to correspond with our intuitions of the extent of a forest. Rather the extension of a forest is normally understood as

²Here I am ignoring the temporal aspects of forests but I do not see any obvious impediment to reconstructing this analysis, albeit in a rather more complex form, taking into account temporal persistence and evolution.

including areas which are surrounded by trees but are not actually beneath any branch or above any root. Thus, terrestrial extension is a subtle function of vegetation distribution which will vary according to what precise interpretation is placed upon the vague ‘forest’ concept. However, if we can define a predicate **ForestExtent**, which holds of all those regions that are in some artificially precise sense the terrestrial extension of a forest, we will thereby also determine the physical constituents of each forest and an identity criterion for forests.

In terms of the parthood relation **P** and a predicate **CON** meaning that a region is self-connected, I define a **ForestExtent** as a maximal connected forested region of sufficiently large area:

$$p:\text{ForestExtent}(r) \equiv_{def} p:\text{Forested}(r) \wedge \text{CON}(r) \wedge \text{area}(r) \geq p:\text{minfa} \wedge \neg\exists r' [p:\text{Forested}(r') \wedge \text{CON}(r') \wedge \text{P}(r, r')]$$

The scope of the precisification variable p ensures that under any given precisification the meaning of **ForestExtent** is logically determined by the meaning of **Forested** under that same precisification. This consistency requirement within definitions supports various patterns of reliable inference that hold whatever reasonable sense we give the concepts.

The precisification-relative area constant **minfa** gives the minimal size of the extension of a forest. Typically a forest is taken to be a rather large feature covering many square kilometers of land. Where the area is smaller the feature is likely to be called a ‘wood’. If it is smaller still, say less than a hectare, one would perhaps qualify it as a ‘small wood’ or use another term such as ‘spinney’. We could define ‘wood’ for example by using a ‘minimal wood area’ constant **minwa** and perhaps also a maximal wood area constant. The constraint $\forall p [p:\text{minfa} \geq p:\text{minwa}]$ means that in every precisification forests must be at least as large as woods. Nevertheless, there need be no such consistency between different precisifications, what is a wood in one particular precisification could be counted as a forest in another. One might also want to place constraints on the shape of legitimate extensions of woods and forest, since a predominantly linear distribution of trees is not normally considered a forest however large an area it covers. Finding reasonable shape constraints is surprisingly difficult and is beyond the scope of the present work.

Having defined ‘forest’ in terms of ‘forested’ we need to consider how observable measurements of the physical world determine which regions fall under the concept ‘forested’; or rather, given our supervalueation methodology, we need to elucidate how these observables relate to different precise interpretations of ‘forested’. It will be useful at this stage to introduce a vague definition which seems to take into account all of the most salient requirements of the concept. I shall say that a region is forested if it is *densely covered by trees*. Although this definition does not make the concept any more precise, it does focus our attention on the key sources of vagueness in any characterisation of forested: what is a tree? how should we measure the density of trees? and, when can a terrain type which is intrinsically vague and granular be said to ‘cover’ a given area?

Apart from the case of metaphorical usage of ‘forested’, which shall not be considered here (though it could be argued that metaphor sheds considerable

light on the meanings of vague terms), it is uncontroversial that the distribution of trees is the primary factor in determining whether a region is forested. Other land properties and vegetation may be relevant to properties of a forest but are not essentially relevant to its extension. However, the term ‘tree’ is itself to a certain extent vague.

My dictionary (the Concise Oxford 1999) defines ‘tree’ as “A woody perennial plant, typically with a single stem or trunk, growing to a considerable height and bearing lateral branches.” This definition exhibits both conceptual vagueness, as to which plant forms or species count as trees, and also sorites vagueness, in that ‘considerable height’ prescribes a vague threshold relative to an objective physical property. Artificially precise definitions of ‘tree’ may be given either genetically, in terms of a set of tree species, or by stipulating conditions of physical stature. In the literature on vegetation mapping these modes of classification are referred to respectively as *floristic* and *physiognomic* (see e.g. (USGS 1994b) for a discussion of these classifications). Perhaps the most intuitively reasonable classifications can be obtained by some combination of floristic and physiognomic constraints. The supervaluation approach is well suited formalising logical relationships among natural and artificially defined tree concepts. For instance the formula $\forall px[p: \text{Tree57}(x) \rightarrow p: \text{Tree}(x)]$ asserts that *Tree57* is a sharpened version of the natural, unrefined concept of tree.

Once we know what a tree is we can try to formulate possible precise versions of the notion of a dense coverage of trees. This turns out to be a surprisingly difficult problem. There are various possible ways one can quantify trees within an area. Practical forest mensuration techniques employ at least the following: the number of individual tree specimens, the total volume of vegetation, the cross-sectional area occupied by tree trunks (at some stipulated height from the ground). In each case the measure may be applied to just the dominant tree species, to all vegetation or to some restricted sub-class. Another consideration is that these measures only work well for regions of a certain minimal size. A very small region will, most likely, lie outside every tree or within or on the edge of an individual tree; and, in each case density cannot sensibly be measured. In fact it seems to be generally true that to determine whether a point belongs to the region of a particular land-type we often have to look not only at what is present at that point but at what is present in some significantly extended region including the point.

Given that one can specify measurement schemes which quantify the trees in a region (provided the region is sufficiently large) it might at first sight seem a simple matter to divide this quantity by the area of the region to arrive at a measure of tree density. Indeed we could define a family of possible measures and articulate the interdependencies between them in a supervaluation-based logic. However, if we calculate, by whatever means, the tree density of an arbitrary region we will get a value that is an average over the whole area. But the region might include one or more parts which are very densely forested while other parts might be completely treeless. In order to identify forests we must have some way of separating regions of high and low tree density. But since our method of calculating tree density requires one to start with a predefined region

we come up against a chicken and egg problem.

I have identified several ways of tackling this problem each of which has a rather different flavour. One is to impose on the terrain a tessellation of appropriate granularity. Density is then only computed directly for the cells of this grid — larger regions of dense tree structure are constructed as sums of these units. Another approach is to partition the land by reference to a more easily measurable secondary indicator. Both these methods are pragmatic solutions which are widely adopted in actual geographical field work. I will comment further on them in the next section.

A third approach is to identify a region of uniformly high tree coverage by ensuring that all significantly large sub-regions maintain the required density. This is not easy to specify, but the following definition seems to capture the idea reasonably well:

$$p:\text{Forested}(r) \equiv_{\text{def}} \forall d[(\text{Disc}(d) \wedge \text{area}(d) = p:\text{fgran} \wedge \text{area}(d \cap r) \geq p:\text{fbgran}) \rightarrow \text{tree-density}(d \cap r) \geq p:\text{fdthresh}]$$

The specification is in terms of discs certain size (i.e. having an area of $p:\text{fgran}$), which is deemed appropriate (according to precisification p) to the granularity with which forest density can sensibly be measured. For all such discs overlapping the region by more than a certain amount (given by the ‘forest border granularity’ parameter, $p:\text{fbgran}$), the density in the area of overlap must be higher than a certain threshold, $p:\text{fdthresh}$. I restrict attention to discs since we will probably want to admit the possibility larger sub-regions which have a lower tree density but are fragmented or very linear in form. The formula is complicated by the requirements of modelling the edges of a forested region to avoid unwanted shrinkage of legitimate extensions. This requires special attention and makes the definition far less intuitive than one would like. Despite being rather artificial, the definition does succeed in providing a conceptually unambiguous characterisation of a forested region. The variability of reasonable values for the density and granularity parameters model purely sorites vagueness.

A further way that one could address the issue of dense coverage is to identify high density areas in terms of the local distribution of vegetation. That is we construct forested regions by a consideration of the locations of individual trees and their spatial relationship to other neighbouring trees. Finding sensible ways of grouping trees is far from straightforward but may lead to a fruitful analysis. There are certainly a large number of incompatible ways in which this could be done, and the lack of any reason to chose a particular approach is perhaps the main reason why this kind of classification has not been widely studied. However, using the framework of supervaluation theory one can explore many possible definitions without committing to any one.

A simple example is the following definition, in which a forested region is defined as one such that all of its points are within a certain distance of a tree:

$$p:\text{Forested}(x) \equiv_{\text{def}} \forall(\pi \in x)[\exists t[p:\text{Tree}(t) \wedge \text{dist}(\pi, t) \leq p:d]]$$

This gives a family of possible precise concepts determined by the parameter d which is a function of the precisification p . There is a strong logical constraint on this family in that the extension of this concept with a small d is always




Figure 2: Possible forest demarcations for a given tree distribution

a subset of the extension determined by any larger d . Hence we would get a contoured distribution of possible extensions similar to that shown in Figure 2. An obvious problem with this definition is that each isolated tree constitutes a small forest. One way of countering this would be to require that each point of a forested region is close to some moderately large number of trees. Another similar approach would be to construct the region as the sum of the convex-hulls of sets of ‘sufficiently close’ trees.

The *add hoc* nature of my characterisations of ‘forested’ may indicate that the concept is in need of deeper ontological analysis. On the other hand it might be that there is an essential conceptual vagueness in the concept, so that no complete finite specification of its potential interpretations is possible. Nevertheless the artificial concepts do seem to capture much of what is intended by the naive description of ‘forested’ and give a reasonable account of the parameters of its intrinsic sorites vagueness. Moreover, I suggest that intuitive classifications of areas into forested and non-forested will not greatly deviate from classifications obtainable by plugging reasonable parameters the artificial concepts.

A further facet of the interpretation of the word ‘forest’ which is worth a brief mention is its connotation relative to other terms for similar geographical features. I am thinking here of the contrast between, for example, ‘forest’ and ‘jungle’. Jungle almost always refers to a tropical or sub-tropical vegetation cover, whereas forest is more general but with a suggestion of a temperate climate. The logic of connotations is likely to be extremely tangled; but in so far as any clear differences in sense can be identified they can be straightforwardly encoded within the supervaluation framework.

6 Comparison with the Geographer’s Forest

Let us round off our examination of forests by considering how the supervaluationist approach can be related to a particular definition of forest that has been widely used in geographical applications. Table 1 shows a *physiognomic* classification, of levels of forestation that was proposed in (UNESCO 1973) and later adopted in (USGS 1994b). The range of different terms employed in the

Plant-form/Height	Percent Canopy Cover of Vascular Vegetation			
	100%–60% (interlocking)	60%–25% (touching)	25%–10% (spaced)	10%–1%
Trees >5m	Forest	Woodland	Sparse Woodland	Sparsely Vegetated
Shrubs/Trees 0.5–5m	Shrub-land		Sparse Shrub-land	
Shrubs <0.5m	Dwarf Shrub-land		Sparse Dwarf Shrub-land	
Herbs	Herbaceous			

Table 1: A physiognomic classification of vegetation types (UNESCO 1973)

table illustrates the way that a precisification (or class of precisifications) is not merely associated with a collection of senses of individual terms but with complex system of logical constrains concerning the meanings of multiple inter-related concepts.

This classification carries with it a lot of implicit conceptual baggage which may not be compatible with other ways of defining forests. For instance, any precisification satisfying it must enforce the constraint that woodland and shrub-land are necessarily disjoint. There is also some lack of specificity in the classification. It is not clear whether a population of fairly widely spaced tall trees growing among a dense cover of small shrubs should be counted as ‘sparse woodland’ or ‘shrub-land’. So some precisifications satisfying UNESCOF might require height to take precedence over density while others could require the converse.

One facet of the problem of characterising forests that is not tackled at all by the UNESCO classification is the question of how forest boundaries should be demarcated. The scheme seems to assume that a candidate area has already been identified which can then be measured in terms of average vegetation height (presumably the height of a species which is in some sense dominant in the area) and the percentage canopy coverage over that area.

(USGS 1994a) surveys a number of methods which field workers use to elicit stand boundaries from other relevant and more directly measurable factors, such as climate, and topography. Similarly, in relation to soil-type boundaries, Mark and Csillag (1989) note that boundaries are often introduced on the basis of surface features that are correlated with but not essential to soil classification. Influences of inessential features on forest stand demarcation from aerial photographs are hinted at by the results of Johnston and Lowell (2000). Although indirect methods may be effective for many purposes they do not elucidate how to partition vegetation-types in terms of properties of the vegetation itself and hence, from an ontological point of view they are suspect because they define something in terms of factors that are only contingently related to the phenomenon in question.

Another way in which geographical information systems and other land surveys often avoid the difficulty of boundary identification is to employ some form of atomic area (usually the cells of a grid) as a minimal unit for which a land-

type is determined. Using this ontology the difficulty of assigning a boundary to an intricate natural object is largely avoided. Instead, measurements are applied to whole cells (or random samples from cells) and a land-type inherits its boundary from the already given boundaries of a group of cells. This is fine as long as we always take a coarse view of the world, where we can treat the cells as atomic units. However, if we are in the business of accounting for the different senses of a term like ‘forest’ we will also want to account for perspectives that go right down to the tree level of individual trees. For example we might have data that tells us that a garden is within a dense forest but we cannot infer that the garden contains trees unless the atomic units that have been classified as forest are smaller than the garden.

Given the slipperiness of the concept ‘forest’ one might assume that the problem of boundary demarcation would have been tackled exhaustively within the subject of forestry. However, this does not seem to be a major concern in the literature of that field (the standard text book (Husch, Miller and Beers 1963) considers only technical problems of surveying a boundary not and does not mention conceptual problems in defining boundaries). In fact, this is not surprising when we consider the nature of forestry and the kinds of information it requires. For most purposes a forester can assume that his forest consists of a collection of *stands* whose boundaries are well-defined. The properties of each stand can then be determined by random sampling techniques; and from these measurements, economically important quantities such as ‘forest volume’ can then be derived by simple computations or by the use of empirically verified tables. The problem of determining boundaries is not of great importance because the statistical approach to measurements works with any reasonable bounding of the forest area, and, in all but exceptional cases, mitigates the effect of any uncertainty in this boundary.

Thus, while the identification of boundaries is a major concern in the ontology of geographical features, for certain purposes they can be taken for granted. However, in situations where we need to identify land-types for some high-level evaluation or planning problem, the problem of demarcation is crucial, and ontological analysis serves a useful function in providing a basis upon which consistent reasoning can be carried out.

In the analysis of the last section, I avoided those aspects of the vagueness of ‘forest’ which are associated with persistence through time. In what we might call the ordinary usage of ‘forest’, it is applied to an area that is densely covered by trees at the time the description is made. However, in the context of forest management and ecological classification it is quite common to regard a land-type in terms of an cyclical or progressive process that takes place in some area of land over a considerable period of time. From this perspective it is perfectly natural to consider an area as ‘forested’ even when it all its trees have been felled and carted off to the log mill. In this sense ‘forested’ would include areas of land at all stages in the arboricultural process. Thus for many geographical applications one will need to carefully differentiate between senses of forest which employ conflicting notions of its temporal status.

7 Conclusion

The geographical sciences face increasing pressures to assimilate huge quantities of information and exploit this consistently for a host of diverse applications in industrial, environmental and social management. The need to provide a firm foundation for the interpretation and manipulation of this data has led to recent interest of geographers in ontological questions (e.g. (Frank 1997)). But, while there is a high level of awareness of issues uncertain and imprecise information the difficulty of taking account of the intrinsic vagueness of natural concepts does not seem to have been fully appreciated. As I have shown, an adequate analysis of a single basic geographical feature type may involve enormous complexity.

Vagueness is often regarded as a phenomenon which defies detailed theoretical explication; but, in this paper, I have attempted to show that a concerted analysis can reveal logical constraints underlying the apparently nebulous meanings of vague concepts. I have suggested supervaluation semantics as a powerful theoretical tool by which one may make inroads into the tangled semantics of ill-defined concepts by articulating what is fixed and what is variable among a space of possible precise senses. A key aspect of my analysis is the division of the phenomenon of vagueness into conceptual and ‘sorites’ modes, which allows one to separate the concerns of analytical ontology from problems of formalising the logic of pure sorites vagueness.

I must confess that in presenting various formal definitions purporting to characterise senses of the word ‘forest’, I have often met with some scepticism. Critics have argued that since the term is so obviously vague in such a numerous variety of ways, any attempt to pin down its meaning formally is completely hopeless. It is for the reader to decide whether the analysis given in this paper cuts anywhere near the heart of the meaning of ‘forest’ or merely wanders among its many branches. It is true that vagueness, especially in its pure sorites form, is as yet very poorly understood, and there is no consensus on how inference should be conducted in the presence of vague concepts. However, to equate what is not understood with what is unintelligible is to deny the value of philosophy. Many geographical concepts may be extremely vague but I do not believe they are completely unprincipled.

References

- Alker, S., Joy, V., Roberts, P. and Smith, N.: 2000, The definition of brownfield, *Journal of Environmental Planning and Management* **43**(1), 49–69.
- Bennett, B.: 1998, Modal semantics for knowledge bases dealing with vague concepts, in A. G. Cohn, L. Schubert and S. Shapiro (eds), *Principles of Knowledge Representation and Reasoning: Proceedings of the 6th International Conference (KR-98)*, Morgan Kaufman, pp. 234–244.
- Cohn, A. G. and Gotts, N. M.: 1996a, The ‘egg-yolk’ representation of regions with indeterminate boundaries, in P. Burrough and A. M. Frank

- (eds), *Proceedings, GISDATA Specialist Meeting on Geographical Objects with Undetermined Boundaries*, Francis Taylor, pp. 171–187.
- Cohn, A. G. and Gotts, N. M.: 1996b, Representing spatial vagueness: a mereological approach, in J. D. L C Aiello and S. Shapiro (eds), *Proceedings of the 5th conference on principles of knowledge representation and reasoning (KR-96)*, 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.
- Elkan, C.: 1993, The paradoxical success of fuzzy logic, *Proceedings of the National Conference on Artificial Intelligence (AAAI-93)*, pp. 698–703.
- Eschenbach, C.: 2000, Ontology, predicates and identity, in S. Winter (ed.), *Geographical Domain and Geographical Information Systems*, Vol. 19 of *Geoinfo*, Institute for Geoinformation, Vienna University of Technology, pp. 33–36.
- Fine, K.: 1975, Vagueness, truth and logic, *Synthèse* **30**, 263–300.
- Foody, G.: 1992, A fuzzy sets approach to the representation of vegetation continua from remotely sensed data: an example from lowland heath, *Photogrammetric Engineering and Remote Sensing* **58**(2), 221–225.
- Frank, A.: 1997, Spatial ontology: a geographical information point of view, in O. Stock (ed.), *Spatial and Temporal Reasoning*, Kluwer, Dordrecht, chapter 5, pp. 321–352.
- Goguen, J.: 1969, The logic of inexact concepts, *Synthese* **19**, 325–373.
- Husch, B., Miller, C. and Beers, T.: 1963, *Forest Mensuration*, Wiley. 3rd edition 1982.
- Johnston, D. and Lowell, K.: 2000, Forest volume relative to cartographic boundaries and sample spacing, unit size and type, *Geographical and Environmental Modelling* **4**(1), 105–120.
- Kulik, L.: 2000, Vague spatial reasoning based on supervaluation, in S. Winter (ed.), *Geographical Domain and Geographical Information Systems*, Vol. 19 of *Geoinfo*, Institute for Geoinformation, Vienna University of Technology, pp. 73–80.
- 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.
- Lewis, D.: 1993, Many, but almost one, in J. Bacon, K. Campbell and L. Reinhardt (eds), *Ontology, Causality, and Mind*, Cambridge University Press, pp. 23–38.

- 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.
- Mark, D. and Csillag, F.: 1989, The nature of boundaries on ‘area-class’ maps, *Cartographica* **26**, 65–78.
- McCarthy, J.: 1993, Notes on formalizing context, *Proceedings of the thirteenth international joint conference on artificial intelligence (IJCAI-93)*.
- McGee, V.: 1997, Kilimanjaro, *Canadian Journal of Philosophy* **23**, 141–195. supplementary volume.
- Ordnance Survey: 2000, *Land-Line User Guide v3.0*. http://www.o-s.co.uk/downloads/vector/landline/Lline_w.pdf.
- Smith, B.: 1995, On drawing lines on a map, in A. Frank and W. Kuhn (eds), *Spatial Information Theory - Proceedings of COSIT'95*, number 988 in *Lecture Notes in Computer Science*, Springer, Berlin, pp. 475–484.
- Smith, B.: 2000, Fiat objects, *Topoi*. In this volume.
- UNESCO: 1973, International classification and mapping of vegetation, A report of the United Nations Educational, Scientific and Cultural Organization (UNESCO), Series 6, Ecology and Conservation.
- Usery, E.: 1996, A conceptual framework and fuzzy set implementation for geographic features, in P. Burrough and A. Frank (eds), *Geographical Objects with Undetermined Boundaries*, GISDATA, Taylor and Francis, chapter 5, pp. 71–85.
- USGCRP: n.d., GCDIS: global change data and information system, United States Global Change Research Program web site. www.globalchange.gov.
- USGS: 1994a, *Field Methods for Vegetation Mapping*, United States Geological Survey, NPS Vegetation Mapping Program. <http://biology.usgs.gov/npsveg/fieldmethods/index.html>.
- USGS: 1994b, *Standardized National Vegetation Classification System*, United States Geological Survey, NPS Vegetation Mapping Program. <http://biology.usgs.gov/npsveg/classification/index.html>.
- Varzi, A.: 2001, Vagueness in geography, *Philosophy and Geography*.
- 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.