Modes of Concept Definition and Varieties of Vagueness*†

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

March 13, 2005

Abstract

The paper considers the problem of defining concepts within formal ontologies. A number of distinct modes of definition are identified, which represent alternative viewpoints on the classification of objects. The nature of these modes is examined and interactions and correlations between them are considered. This analysis is used to characterise a number of key problems that confront the attempt to analyse the meanings of natural language concepts by means of a precise formal ontology. Some proposals are made as to how these problems might be overcome. It is argued that specification of precise formal semantics requires clear separation between aspects of meaning that are based on different modes of classification. But, to provide an adequate ontological framework to describe the semantics of natural language terms, additional formal apparatus is required to synthesise their meaning by combining the different modes. This meta-level ontological formalism must articulate the interactions between the different modes.

Keywords: Ontology, Concept Definition, Vagueness.

1 Introduction

The pervasive presence of vagueness and ambiguity in natural language is a perennial worry of ontologists. It presents a major challenge to the formalisation of intuitively understood terminology within the precise framework of logical representations.

Although ontological engineering is now a well-established and rapidly expanding area of information science, there are still many competing views of how ontologies should be developed, evaluated and utilised. A major schism exists between those who want to construct a single unifying ontology and those who argue that many different ontologies are required to cater for different terminological contexts and user groups. Philosophers and knowledge representation formalists from the field of symbolic AI (e.g. [9, 6]) have often sought to establish a single ontology of fundamental concepts that is expressive enough to describe all aspects of reality (at least at some coarse level of detail). By contrast, information system

*This work was partially supported by the Ordnance Survey.
†This document is a slightly modified version of [5].

designers and Semantic Web developers have tended to gear their ontologies to particular data-sets or user groups and have been sceptical about the possibility of achieving universal generality.

Coming from a tradition of formal logic and AI representations, it had puzzled me for some time, why information system designers should be quite so negative about the possibility of achieving precision and universality in ontologies. After all, ontologies are typically formulated using something like first-order logic, which is a paradigm of precision and pretty much immune to misinterpretation, except by the most hardened subversive philosopher.

Recently I have come to the view that the disagreement hinges on a difference in understanding of the word 'concept'. Those who are sceptical about the idea of precision and universality tend to regard a 'concept' as something rather close to a natural language term, as something which may be vague or ambiguous and may be disputed over by people with different world views. But for a (classical) logician, a concept is an abstract entity that is largely independent of the vagaries of natural language: only in idealised circumstances can a concept within a formal system be regarded as the referent of a natural term [12].

Having said this, the Semantic Web developers are still right in their own terms. If we want an ontology to help manage the content of the Web in a way that is accessible and meaningful to humans then we need an ontology that characterises the meanings of *words*, not only of formalised concepts. However, the problem remains that classical logic deals primarily with idealised concepts and only indirectly with natural language words. The phenomenon of vagueness and the persistent disagreements between different schools of philosophy and ontology design show us that the linkage between words and concepts is highly complex and involves a whole range of subtle pragmatic factors. The present paper attempts to make a very small step in clarifying the nature of these factors.

In the current work the focus is on conceptual terms used to classify objects of the physical word. In other words count nouns describing types of material object. In fact, I believe that much of the following analysis has rather general application to a broad range of both physical and non-physical vocabulary. However, it is likely that other important distinctions and semantic phenomena would come to light in an investigation of domains involving mentalistic and cultural concepts.

The remainder of the paper is organised as follows. First I shall distinguish a variety of different modes of characterisation which can be used to classify types of physical object. I shall then look at interactions between these modes and suggest that the meanings of natural language terms often involve a conflation of several modes of specification. While in most cases these modes are well-correlated in their application to the objects we usually refer to in ordinary discourse about our world, I suggest that in unusual cases this correlation will often fall apart. This is a major problem for the construction of precise ontological theories. This problem is largely obscured by the phenomenon of vagueness which blurs the assignment of precise meanings to natural language terminology. I proceed to analyse this vagueness into a number of distinct phenomena, which give rise to this blurring. Finally I suggest that if the enterprise of formal ontology construction is to overcome these difficulties it must develop a theoretical super-structure which can relate natural language terminology to the artificial concepts axiomatised within formalised theories.

2 Modes of Object Classification

Our intuitive classification of physical objects into kinds seems to depend on a wide spectrum of properties relating to very different aspects of an object's nature. In order to devise and organise an ontology of physical objects it is useful to identify the principle *modes* of description by which a classification could be made.

I propose the following breakdown of classification modes applicable to classification of objects in the physical world:

- Physical
  - Geometry
  - Material properties
- Historical
  - Originary — how the object came into being
  - Initiated — some event that the object has undergone
  - Periodic — types of event that the object repeatedly undergoes
- Functional
  - Actual functionality
  - Potential functionality
  - Intended functionality
- Legal/Conventional
  - Ownership
  - Associated rights and responsibilities
  - Traditional convention

Let us look at these different modes in more detail:

Physical Properties

The physical properties of an object can be specified in terms of the types of material of which the object is composed and the geometrical form of these portions of matter. The vocabulary of terms describing physical properties can be defined in terms of a geometrical theory of space and matter. Since relevant properties are often of a qualitative nature, I suggest that it may be easier to build a useful ontology on the basis of a qualitative theory of spatial and physical properties (e.g. [14] or [2]).

Historical Factors

For many classifying terms, the objects to which they apply depend (at least partly) upon the history of those objects. Fully comprehensive knowledge of an object would include its entire history from its origin to its current state. This knowledge would subsume all the modes of classification given above. However, in practise we rarely have this complete knowledge. Consequently, the criteria we use to classify objects generally depend either on the current state of the object or on fairly limited information about the object's history. Nevertheless this limited information is crucial to the applicability of certain classifying terms.

The way in which an object came into being, its *originary* history, is of particular importance. For instance, the term *artifact* may be taken as applying to objects which were created by intelligent agents. In other cases the applicability of a term may be *initiated* by some event occurring during an object's history: a king must be crowned, a 'listed building' must be 'listed'. Finally, an object's type may be determined by its undergoing some periodic event. This is quite common in geographic terms such as 'flood plain'. It is also relevant to the classification of buildings, whose type is often determined by activities which take place periodically within the building.

Functional Properties

The functional properties of an object consist of those properties that can be described in terms of actions involving the object as an instrument. Again these actions can be described by means of a logic of actions and events (for example the Situation Calculus [11], the Event Calculus [16] or the formalism given in [2]). In so far as types of action can be characterised by changes in physical properties, they can be defined directly in terms of the more basic theory of space, time and matter. However, in the case of artifactual objects fabricated for human use, action types that are considered significant in describing these objects are typically extremely complex when analysed as patterns of physical change. Consequently it may be expedient initially to formulate a theory of these actions at a high level of abstraction, rather than directly defining them from concrete physical properties. The dependence of these high-level actions on underlying physical reality could then be established at a later stage.

Functional properties of an object can be divided into three types:

- **Potential functions.** These are all the possible functions that *could* be achieved using the object as an instrument. The potential functionality of an object at a given time depends entirely on the physical properties of the object at that time and could in principle be defined in terms of these properties (though in practise this may not be feasible).

- **Actual functions.** These are the functions for which an object has been or is being used. To describe its actual functions, the history of an object must be taken into account but its precise physical nature need not be described since the relevant characteristics will be inherent in the actions which have been achieved using the object.

- **Intended functions.** These can be described in a similar way to actual functions except that we now refer not to functions that are (or have been) carried out but to those functions that the designer had in mind. It seems that having intended functions is arguably a defining characteristic of artifactual objects. However, within a human context natural objects may also sometimes acquire intended functions.

Functional properties are intimately associated with states and events which occur in time; and hence functional classification is often entangled with the historical modes of classification. We can distinguish two different ways in which a functionality is manifest over time:

- **Static.** The functionality is provided by a constant property of the object which is present throughout some time period (e.g. supporting, sheltering, enclosing).
- **Dynamic.** The functionality arises through agency, instrumentality or facilitation of some type of action or activity (e.g. a canal lock has a dynamic function since it must open and close as part of a complex event that allows passage of a boat).

Legal and Conventional Properties

In dividing the natural world into objects, we are not always guided by boundaries which are themselves present as natural phenomena. Thus many kinds of geographic object are of a *fiat* nature [18, 17]. More generally, there are a host of other legal and conventional attributes that are relevant to the individuation of geographic objects and their classification into kinds. For instance we often divide a single built structure (say a row of terraced houses or a block of flats) into sub-structures (e.g. individual houses or flats) on the basis of ownership.

Often divisions according to ownership correlate with more concrete divisions (e.g. houses within a terrace are not connected by interior doors). However, this need not be the case. From a physical point of view, the boundaries between provinces, counties and even gardens may be arbitrary.

The ascription of legal and conventional boundaries and attributes may be regarded as a historical act, and in this sense they could be thought of as originary or initiated historical properties. However, the event that creates a fiat object is not one in which that object takes an active part, so conventional properties cannot be determined just by looking at the history of the object itself. I believe that this difference justifies the idea that the legal/conventional aspects of objects should be treated as a distinct mode of classification.

3 Relationships Between Modes

Given that there are many ways to classify objects, an ontology designer must ask what are the relationships between these modes: is one more primary than the others; are there dependencies and/or redundancies among the different modes of description.

3.1 Inter-Definability

For a concept to be given a precise formal definition according to a definite semantics, one must characterise the conditions for it to be truthfully applied in all possible situations (or *possible worlds* in the terminology of modal logic). This idea has been developed in the works of Padoa [13] and Tarski [19] to give various criteria for the definability of one concept in terms of a set of other concepts (see also [4]).

In [3, 4] I used the method of Padoa to argue that within the idealised confines of a formal system, any set of primitives that can characterise the spatio-temporal and material structure of all possible spatio-temporal histories is in fact theoretically adequate to characterise the meanings of any (precise) concept. This is because it is not possible to have two models that are identical in respect of all physical, temporal, material and possibilistic properties, and yet

differ in the extension of some other property; otherwise the property would be beyond all empirical observation of verifiability even from a god's eye perspective on the totality of all physical possibilities.

Because of this, it seems that all classifications of an object could be derived from spatio-temporal, material and possibilistic properties. In other words from a combination of the *physical* and *historical* modes of classification. This would involve reduction of any functional properties and even fiat enactments to purely materialistic terms. But, of course this assumes that functional and social/legal/conventional terminology is (or can be made) completely precise. This may not be the case. Since complete and accurate knowledge of the material and historical condition of objects is seldom available, natural modes of classification focus on more limited and readily apparent properties of an object.

In natural terminology the functionality of objects is described not by exhaustive description of the physical properties of events involving the object, but rather by high-level vocabulary that is tailored to those distinctions that are most informative in terms of what humans want to say about object functionality. Hence, we have a large vocabulary for describing events, activities and process, which is highly abstracted from the details of the underlying physical changes that are involved.

Having recognised the different modes of classification, we can appreciate an important obstacle to achieving a comprehensive and consistent ontology. That is the problem of combining the very different modes into a single body of defining criteria. While reduction of the various modes to a unified ontology may be possible at a deep theoretical level, this does not mean that a formal ontology should be framed exclusively in terms of a single mode. Rather, each of the main modes should be supported within the ontology as coexisting descriptive frameworks.

3.2 Exemplar Clustering and Merging of Modes

In ordinary situations, objects that exhibit one property, will very often also exhibit another property and vice versa. Moreover this may happen even though there is no necessary connection between the properties. The fact that two properties almost always occur in conjunction may be because of patterns and regularities in our limited experience of reality that are essentially contingent. Thus properties may appear to be closely connected simply because we encounter certain combinations of properties much more frequently than others.

The objects that are relevant to our life are very highly clustered compared to a continuous distribution of random objects. Thus, buildings tend to have similarities of structure as well as similarities in their functionality and history. Complex artifacts such as buildings have a particular physical structure as a consequence of the way they were created. In fact it is almost inconceivable that an object conforming to the typical physical structure of, say, a house could have come into being in any way other than having been 'built' by intelligent agents. This clustering of exemplars seems to be a major factor in the merging and entanglement that occurs between the various modes of classification.

We have no ordinary vocabulary that can distinguish between a 'house' that was 'built' and a similar structure that arose by random physical processes. Nevertheless, from a purely logical point of view we must accept that it is *possible* for objects to exist which have the same physical structure as some type of artifact and yet arose through some freak of nature. While this is a somewhat subtle philosophical point, its acknowledgement is significant for the systematic construction of ontologies. Its recognition will enable us to separate general

semantic properties of concepts from accidental correlations arising from particular contingencies of the actual world. To emphasise this point and also to give us a concept that will be useful in formalising an ontology of objects, I coin the term *artifable* to mean an object that has the physical structure of some kind of artifact and yet was created independently of any intelligent activity.

The fact that objects considered in ordinary discourse typically conform to contingent correlations of properties will be called *exemplar clustering*. This arises from the fact that the development and usage of natural language concepts takes place in the very specific context of the particular physical world in which we find ourself.

In terms of the analysis of different modes of classification, the conflation of logically distinct properties means that our intuitive understanding of a natural language term will often span more than one mode. Descriptions according to several different modes may coincide for the vast majority of exemplar objects, and yet in atypical cases the different modes would lead to contradictory classifications.

4 Ambiguity and Vagueness

As explained in the introduction, one of the main aims of the current work is to shed some light on the underlying factors that make up the phenomenon of conceptual vagueness. To this end, the following distinct varieties of ambiguity and vagueness can be identified:

- Simple Ambiguity
- Vagueness
  - Threshold (or *sorites*) Vagueness
  - Partiality
  - Deep Ambiguity
    - Mode Ambiguity (caused by exemplar clustering)
- Prototype Implication (caused by exemplar clustering and the etiquette of effective communication)

Simple Ambiguity

The simple form of ambiguity occurs where a term has more than one distinct meaning. For example 'bank' could refer to an inclined land form or a monetary institution. This is the most straightforward example of a case where there is not a direct correspondence between lexical terms (i.e. words) and the conceptual referents (i.e. meanings) of those terms. This complication has already been addressed in many ontologies, which make a distinction between lexical and conceptual entities (e.g. CYC).

Vagueness

In the classification given above, 'vagueness' appears as a general term covering three types of phenomenon that I believe contribute to what we ordinarily consider to be vagueness. Given that in essence vagueness is contrary to precision, it is not surprising that it is difficult to pin down exactly what semantic or pragmatic considerations operate beneath the veil of vagueness. However, I believe that at least the following three factors play a part.

Threshold (*Sorites*) Vagueness

This is the type of vagueness that has received most attention from philosophers and logicians. It arises when qualitative concepts are employed to distinguish among objects which in fact exhibit continuous (or fine grained) variation in the observables relevant to the ascription of the concept. This is the phenomenon underlying the famous paradox concerning whether, if we remove one grain of sand from a heap, we shall always be left with a heap (the *sorites* paradox).

Although the formalisation of sorites vagueness is still controversial, substantive proposals have been made. For instance, supervaluation semantics can be applied (e.g. [7, 1]).

Partiality

In ascribing many natural language terms, it seems that we are selective in what aspects or what part of an object we consider. For instance we may describe a frog as green, even though its mouth and tongue are clearly pink and its eyes are black. The objects of the world do not divide neatly into examples and non-examples of sharply defined properties.[1] In the case of artifacts, the requirement that an object be 'constructed or modified by an intelligent agent' often applies only to part of the object. This phenomenon appears to be similar in its mode of operation to the more studied *sorites* vagueness, so it is possible that a similar formal framework could give an account of partiality.

Deep Ambiguity

What I call *deep ambiguity* is the phenomenon whereby a term seems to refer to a cluster of overlapping but distinct concepts. In such a case we find that a large number of different properties are relevant to the ascription of a term, and yet none of them can be regarded as necessary. This type of concept is known in the philosophical literature as *open textured* [20].

Mode Ambiguity

The possibility of identifying and classifying objects according to different modes represents a kind of deep ambiguity. The merging of classification modes within natural language terms presents a serious obstacle to the formulation of adequate identity conditions within a formal ontology. The problem is what I call *identity condition schism*. It occurs when conditions for identity and continuity are incompatible between two (or more) different modes by which an object type could be classified.

Consider the term 'road'. A road such as the M1 (UK's first motorway) has both a physical existence, as a certain strip of tarmac, and also a functional and to some extent a legal/conventional existence as the major road connection between Leeds and London. Moreover, the physical, functional and conventional modes of road identification can conflict in certain circumstances. For instance when the motorway system around Leeds was extended, the M1 was partly re-routed, so the course of the M1, considered as a conventional/legal entity, was altered. However, the physical road itself remained unchanged, except for alterations of road signs.

[1] In fact the world does not divide neatly into objects at all. But the problem of individuating objects likes beyond the scope of the current paper, which focuses on the terminology used to classify kinds of object, and naively assumes that the objects are already there to be classified.

Prototype Implication

Authors in the philosophical and logical literature on vagueness, have often attempted give an account of vague concepts that is compatible with the existence of borderline cases. However, in ordinary communication, concepts are very rarely applied to borderline cases.

An unqualified ascription of a property to an object carries with it an assumption of prototypicality. Because we normally do not wish to mislead, we phrase our descriptions in such a way that the reality of the situation is close to what we would consider to be a typical instance of our description. Thus, limits of our concepts are not often tested in ordinary contexts of communication and cognition. Only in politics, marketing, writing references and similar rhetorical contexts do we readily apply concepts to borderline cases.

Formalisms for modelling prototype implication have been developed in AI under the heading of Default Reasoning [15, 10]. However, little application of such techniques has been made to the interpretation of natural language terms with respect to ontologies.

5 Formalising the Word-Concept Interface

The primary aim of this paper was to diagnose a variety of problems facing the definition of concepts, not to solve these problems. However, I will make some general suggestions as to how one might construct an ontology that can handle vagueness-related phenomena.

It is my view that the apparatus of classical logic, within which ontologies have traditionally been defined, cannot by itself account for the meanings of the conceptual terms of natural language. This would be expecting too much from a rigid logicist framework. Although there may be technical results which tell us that the classical formalisms do have sufficient expressive power in a technical sense, direct formalisation of natural language terminology in these systems is not practical. Handling the problems associated with vagueness and ambiguity requires ordinary logical theories to be organised within a principled theoretical superstructure. This will provide mechanisms that mediate between natural language concepts and the logically defined concepts of formal ontologies.

In order to formulate a precise ontology where concepts have an unambiguous semantics, we must separate different modes of classification. Within the limited perspective of single mode it is much easier to achieve a consistent axiomatisation which covers a range of different concepts. In order to do this effectively, it will be expedient to start by formalising the key primitives relevant to each particular mode. Once this fundamental mode ontology has been established we can give coherent formulations of a range of different concepts from the perspective of this mode. We will then use an auxiliary superstructure to represent the conflation of different modes in the intuitive sense of natural language terms.

I suggest that some form of supervaluationist semantics may be used to account for various types of vagueness [1]. The nature of the semantic apparatus required will need to account for differences between the nature of 'sorites vagueness' as opposed to 'deep vagueness', and the special case of mode ambiguity will require particular attention. Perhaps some kind of statistical apparatus is also needed to model the distrubution of possible interpretations of natural langauge vocabulary. 'Prototypicality assumptions' can be formulated within some kind of non-monotonic reasoning framework, but existing approaches would need to be modified to be used in the context of an ontology [8].

6 Conclusion

I have suggested that many of the difficulties encountered by ontology designers stem from the fact that one is trying to use formal tools suited for handling precise concepts to characterise the meanings of natural language terms that are far from precise. This leads to disputes over the correctness of ontological specifications and can also lead to logical inconsistency.

In order to overcome these difficulties, more attention needs to be paid to the subtle and complex relationship between natural language vocabulary and the symbols of a logical representation. If ontologies are to provide a mechanism for automated interpretation and processing of conceptual information that is compatible with our ordinary understanding of natural vocabulary, they must incorporate an adequate formalisation of the *word-concept interface*.

References

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

[2] B. Bennett. Space, time, matter and things. In C. Welty and B. Smith, editors, *Proceedings of the 2nd international conference on Formal Ontology in Information Systems (FOIS'01)*, pages 105–116, Ogunquit, 2001. ACM.

[3] B. Bennett. The role of definitions in construction and analysis of formal ontologies. In *Sixth Symposium on Logical Formalizations of Commonsense Reasoning*, AAAI Spring Symposium, Palo Alto, 2003. `http://www.comp.leeds.ac.uk/qsr/pub/BennettAAAsymp03a.ps`.

[4] Brandon Bennett. Relative definability in formal ontologies. In A. Varzi and L. Vieu, editors, *Proceedings of the 3rd International Conference on Formal Ontology in Information Systems (FOIS-04)*. IOS Press, 2004.

[5] B. Bennett. Modes of concept definition and varieties of vagueness. *Applied Ontology* 1(1), 17–26, 2005.

[6] E Davis. *Representations of Commonsense Knowledge*. Morgan Kaufmann, San Mateo, 1990.

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

[8] P. Gärdenfors. *Conceptual Spaces*. MIT Press, Cambridge MA, 2000.

[9] J. R. Hobbs and R C Moore, editors. *Formal Theories of the Commonsense World*. Ablex, 1985.

[10] J McCarthy. Circumscription — a form of non-monotonic reasoning. *Artificial Intelligence*, 13(1):27–39, 1980.

[11] J McCarthy and P J Hayes. Some philosophical problems from the standpoint of artificial intelligence. In B Meltser and D Mitchie, editors, *Machine Intelligence 4*, pages 463–502. Edinburgh University Press, 1969.

[12] Drew McDermott. Artificial intelligence meets natural stupidity. *SIGART Bull.*, (57):4–9, 1976.

[13] Alessandro Padoa. Un nouveau système irréductible de postulats pour l'algèbre. In *Proceedings of the Deuxième Congrès International des Mathématiciens*, pages 249–256, Paris, 1902.

[14] D.A. Randell, Z. Cui, and A.G. Cohn. A spatial logic based on regions and connection. In *Proc. 3rd Int. Conf. on Knowledge Representation and Reasoning*, pages 165–176, San Mateo, 1992. Morgan Kaufmann.

[15] R Reiter. A logic for default reasoning. *Artificial Intelligence*, 13(1,2):81–132, 1980.

[16] M. Shanahan. The event calculus explained. In M.J.Wooldridge and M. Veloso, editors, *Artificial Intelligence Today*, volume 1600 of *LNCS*, pages 409–430. Springer-Verlag, 1999.

[17] B. Smith. Fiat objects. *Topoi*, 20(2):131–148, 2001.

[18] Barry Smith. On drawing lines on a map. In Andrew Frank and Werner Kuhn, editors, *Spatial Information Theory — Proceedings of COSIT'95*, number 988 in Lecture Notes in Computer Science, pages 475–484, Berlin, 1995. Springer.

[19] A. Tarski. Some methodological investigations on the definability of concepts. In *Logic, Semantics, Metamathematics*. Oxford Clarendon Press, 1956. trans. J.H. Woodger.

[20] F. Waismann. Verifiability. In A. Flew, editor, *Logic and Language*, pages 122–151. Doubleday, New York, 1965.