To appear in "Concepts and Fuzzy Logic"

Editors: Radim Belohlavek and George J. Klir

MIT PRESS

FINAL DRAFT

CHAPTER 9. Concepts and Natural Language

James A. Hampton

9.1 Concept, Word and Object

In this chapter, I am concerned with a particular kind of concept. The term "concept" can be used for a wide variety of mental and cultural entities. We can talk about the concept of art, the concept of a mathematical function like addition, the concept of democracy or the concept of a space-time continuum. In fact it can be argued that the history all of human intellectual endeavour, outside of the arts, is the history of the development of concepts. One particular domain in which concepts play a crucial role is that of enabling us to communicate our ideas via language, and it is this type of "lexical" concept that is the focus of this chapter. In the first section I describe the close association between language and concepts, and in the second section I present three ways in which the content of a concept can be determined or individuated. I then turn in Section 3 to the issue of vagueness in meaning and the relevance of fuzzy logic approaches.

As a starting assumption, we can think of lexical concepts as the meanings of single words (or idiomatic phrases) in a given language. Hence if in English we have the words "apple" and "sauce", we assume that each of these words names an underlying concept – respectively APPLE and SAUCE. When we wish to communicate to another person a fact, or a question or a comment, then we do so by composing a thought and then expressing it in our language. Our companion then hears the expression that we have uttered and uses it, together with the context in which it was uttered, to arrive at an understanding of what our intention was in making the utterance. Much of everyday

discourse is an elaborate ritual of asking after each other's health, joking around, commenting on news items that we believe the other would be interested in, raising questions about future events and so forth. It is very important to recognize that a great deal of language behaviour is not about the assertion of true beliefs. However an equally important function of language is that it *does* allow the expression and assertion of beliefs, and it is when language is used in this way that questions of logic and truth arise. Our present interest in the concepts that correspond to the meanings of nouns, verbs, adjectives etc. is thus going to derive from the question of how we can construct truth-bearing propositions from language. These propositions can then be asserted, denied, entertained, assumed, or generally allow us to be in a mental state that involves having a "propositional attitude" towards a given proposition.

It is fair to assume that language and thinking have evolved in close parallel in the evolution of human cognition. Our ability to think particular thoughts and to communicate them to others has led to the structure of propositional thought closely mirroring that of language. In other words the meaning of a sentence such as "John likes icecream" can be "translated" into a conceptual representation as something like

LIKES(John, icecream),

where LIKES is a two-place function that takes as arguments an agent (John) and a patient (icecream) and maps these onto truth values. For example, all the pairs (*x,y)* corresponding to someone (*x)* who likes something (*y)* will have a value of True for this function. Pairs in which someone does not like something would return a value of False.

This "language of thought" hypothesis was made explicit by Fodor (1975) in his book of that title. Fodor used the term "mentalese" for the hypothetical internal language in

which our thoughts are expressed. It can be argued that concepts are finer-grained than words. We have a single word “apple” that serves to name a concept that includes a whole range of individual varieties of apples, which an individual may wish to differentiate with different conceptual tokens. But generally speaking, as soon as a need arises to make conceptual distinctions, then words will appear in the language to allow those distinctions to be expressed and communicated. Others hearing those words will seek to discover their meanings, and in this way language and concepts develop and evolve in tandem.

9.2 Three Meanings of “Concept”

Concepts correspond to the building blocks or atoms from which we can construct thoughts with propositional contents. But where do these concepts come from, and how is their individual content determined? Three answers need to be distinguished, the objective, the inter-subjective and the subjective. Let us consider each in turn.

9.2.1 Objective Concepts and Externalist Theories of Meaning

If our use of language is to be more than “just talk”, then the words that we use need to label concepts that actually refer to real classes of things in the world. The first place to look for the determination of conceptual content is therefore the objective world. Philosophers such as Kant proposed that we must be born with certain conceptual capacities in order to be able to make any sense of the world at all (for a more recent development of this argument, see Keil 1981). It would in fact be very odd if in the evolution of the human brain there was not a survival advantage for individuals whose brains are able to conceptualize the experienced world in ways that correspond to the reality of the world out there. The ability to detect cause and effect, to monitor the

passage of time, to "parse" a visual scene into a set of individual objects standing in relation to each other in a three-dimensional space, would all be examples of this type of innate conceptual ability.

Locating conceptual contents in the external "real" world gives us a so-called "externalist" account of the content of concepts, a position commonly defended in philosophy (e.g., Fodor 1998; Rey 1983). Arguments for this position were offered independently by Putnam (1975) and Kripke (1980). To take one example, Putnam asked us to imagine that our planet has a twin in which the colorless, tasteless liquid that covers two-thirds of the planet and freezes to make ice and snow does not have the chemical formula H_2O, but instead XYZ. People on the two planets, he argued, would have identical mind and brain states, since their experience, their naming behaviour and everything else about them would be identical. But whereas "water" on earth *means* H_2O, on twin earth the equivalent word *means* XYZ. Hence the meaning of our everyday conceptual terms is determined not by what they lead us to think, but by what they *actually* refer to or denote in the outside world.

The externalist theory of meaning provides a solid foundation for the semantics of natural language, since the semanticist can focus on finding the logical relations that relate simple terms in the language, and their combinations, to classes and relations in the external world. A person can use language in order to express ideas, but it is possible that the person can be *wrong*, not just in the propositions that they express, but in their usage of the words and their meanings. Externalism thus also provides an account of the so-called "problem of error", or how it is possible for someone to have a false belief about the contents of a concept. As externalists point out, if we define the content of a concept

in terms of the content of a mental state (i.e. a psychological view of contents), then the concept means just whatever the person has in their head to represent that concept at the time. If you and I have different information representing the meaning of a word like "war" then the externalist would argue that we could never be said to disagree about any statement concerning war. Whenever you assert a statement about war that I would deny, then that could be just a sign that we have different concepts – we are talking at cross-purposes.

Surveying the range of words and meanings that we use, we can discern certain domains in which externalist accounts of meaning work well. Words that relate to the natural world (e.g. the names of animal and plant species) tend to have different psychological properties from those that relate to the human-made world. These differences are particularly marked in regards to the focus of this book, namely partial truth and borderline category membership. For example, Estes (2004) investigated borderline cases of semantic categories, like tomatoes and olives as fruits. He measured the degree to which people see borderline cases of conceptual categories as having *partial* degrees of membership, as opposed to having *uncertain* membership. In the first case, a person might be quite certain that the tomato is only partly entitled to belong in the fruit category. So membership in the category can be a matter of degree (as in the fundamental axiom of fuzzy logic). On the other hand another person may be quite certain that a tomato has to be either a fruit or not a fruit, it is just that they are uncertain which it is. Estes (2004, Experiment 2) accordingly asked people first to judge whether an item was a clear member, partial member, or clear non-member of a category, and then to rate on a 10 point scale how confident they were in their decision. He found that

judgments of artifact kinds, like clothing, furniture, ships or weapons were more likely to use the partial response, and also more likely to be rated with greater confidence than were judgments of biological kinds. Within each domain there was however a negative correlation, with items with more partial responses also having lower confidence ratings. Other studies have also found very consistent differences between the natural kind or biological domain and the domain of artifacts (Diesendruck and Gelman 1999; Hampton 1998; Kalish 1995).

The externalist view of conceptual content may therefore work best for so-called natural kinds – types of thing that existed in the external world before any people came along.

9.2.2 Intersubjectivity and the Social World

While many biological kinds happily exist without the intervention of humans[1], there are many other lexical concepts that refer to kinds of things that are more or less entirely cultural artifacts. How the "correct" meaning of a term is then determined becomes a more complex matter. In many cases societies create "nominal kinds" in which the meaning of terms is stipulated by a public *dictat*. As parliaments or congresses pass laws, they create legal language which needs to be as clear and as precise as is possible. Hence the concepts that are central to such laws will usually be given some more or less explicit definition. To be found guilty of driving a car while under the influence of alcohol, the law doesn't concern itself with the vagueness and fuzziness of the concept of "under the influence of" and instead goes for a simple test of the concentration of alcohol within the blood. Fuzziness at the borderline can be handled by issuing warnings to those with up to

[1] Exceptions are the biological kinds that we have created through selective breeding, which include almost all our plant and animal foods, much of the contents of gardens and parks, and of course animals like cats and dogs and race horses.

10% excess, and prosecuting those where the concentration is sufficiently high to allow no doubt concerning the accuracy of the test result. In order for a modern society to function efficiently, a lot of common or everyday concepts have to be given stipulative definitions, be it in terms of the legal limits for late term abortions, the age at which it is safe for children to drink alcohol, or the age at which a person should be entitled to free bus travel. Vague intuitive concepts are replaced by rules that are a compromise between capturing the vaguely defined intuition intuition in a rather insensitive way, while also being easily understood and easily enforced.

It appears then that many of our concepts are defined for us as a result of explicit conventions adopted by people with the political power to make such decisions. The process is recursive and adaptive. As public attitudes to homosexuality have changed over the last 50 years, so there has been pressure to redefine other social concepts including family, parenthood and marriage.

Social forces on concepts also operate at a much less explicit level in the way that everyday word meanings evolve and adapt to changing social circumstances. Many lexical concepts, like FURNITURE or GAME or PHONE have neither an objectively determined content to be found in the world, nor a nominally stipulated content derived from explicit convention. Instead there is a dynamic equilibrium among different language users that keeps the meanings stable over the short term, but allows them to drift and adapt in the longer term. This equilibrium leads naturally to fuzziness in the application of a term, since at any particular point in time the usage will have a natural variability within the population. Furthermore in order to keep calibrated with others, individual's usage has to remain flexible and adaptable, consequently giving rise to

inconsistent and variable categorization or naming behaviour.

9.2.3 The Individual and Word Meanings

The third way in which one can try to determine the content of concepts is at the individual level. This approach puts the "owner" of the concepts first. It asks – what is it to have a particular concept?

How do you know what the word "game" means? Your knowledge of its meaning could be characterised in two ways. You have experience of other speakers referring to a certain class of activities as games (e.g. chess, tennis, poker), and you have noticed that whenever you have used the word to refer to members of that class, you have been easily understood and no one has corrected you. So there is a common understanding between you and other speakers on what is the *appropriate use* of the term. Note that this kind of conceptual knowledge is primarily extensional. It involves knowing the usual things that the term refers to. It need have relatively little intensional content, since I may just have learned the things that are called "game" as a list of items, without any further understanding of what they might have in common or what the basis is for calling anything at all a game.

The second way to characterize your knowledge of the word's meaning is in terms of the concept *intension*, the information that the concept carries about the likely properties of members of the class. You may not have actually come across many occasions on which the word is used to refer to other concepts, but you may have come across the word "game" in the context of discussions about the importance of different attributes of games – how playing games develops a sense of fair play and justice, how games are often competitive but allow aggression to be displayed in a harmless way without causing

personal affront to others. In other words you may know a lot about the role that games play in social settings and in people's lives, without actually knowing too much about which activities actually count as games. Having this intensional information is also the basis for using the word in creative and metaphorical extensions as in "God is playing a game with us", or "I was dealt a poor hand in the game of life".

So the two aspects, the extension and the intension, are both represented to a greater or lesser degree in individuals who possess a particular concept. Having a full representation of both aspects means that not only could you generate a list of examples of games, but you could also tell me what are the common characteristics, purposes and social functions of games. These two aspects of concepts have been captured in the laboratory through models of concept learning that focus on either the learning of exemplars (Nosofsky 1988) or of intensional prototypes or schemas (Ahn et al. 2000; Hampton 2006). Both models have means for classifying novel cases in the category. Exemplar models rely on a notion of similarity. A new activity would be classed as a game if it had more similarity to other games than it did to members of other contrasting categories. Intensional models would determine how closely the new activity matched the expected attributes of a game, and the relational links between the attributes, before deciding if the degree of match was sufficient. In both cases, then, membership in a category has the potential to be partial or borderline.

In the determination of word meaning, both extensional and intensional learning will involve a similar interchange with other language users and other sources of cultural information. To remain a successfully functioning member of a linguistic community, speaking a common dialect of a language, an individual has to calibrate her use of words

to the usage of the other members of the community (in just the same way as different regional accents arise through mutual calibration within interacting speech communities.)

9.2.4 Extensional Knowledge May Not Depend on Intensions

In the preceding section, I suggested that people may have both extensional and intensional understanding of a term such as game. As described earlier, there are good arguments that for at least some domains of concept the intension (as defined here) may not be relied on to determine the extension. These domains are those in which we have some notion that there are real classes in the external world that exist independently of human cognition. Biological species are an example, as are chemical elements and compounds. Following Putnam and Kripke's arguments about external determination of conceptual content, there have been psychological studies that confirm that in these cases people consider their own concepts to be possibly incorrect. If asked whether a particular object is made of gold, most people would have to admit that they don't actually know what *makes* something gold rather than something else. There just is this stuff called gold, and chemists have figured out what it is, what properties it has, why it has the properties that it does, and how you can test something to see if it is made of gold, and if so what purity etc. So as a fully signed-up well educated member of your language community, you are ready to admit that there are domains of common words in your vocabulary that you don't actually know the meaning of, at least in the sense of providing any of the above information with any confidence. This prospect has been captured in the proposal of a theory of Psychological Essentialism (Medin and Ortony 1989). The theory makes two proposals. First, that we believe that some conceptual kinds have an "essence" that makes them the kind of thing they are. Gold is a stuff with some kind of atomic

essence, and tigers are creatures with some kind of genetic essence. The second proposal is that most often we lack information about what the essence may be. So our conceptual representation of the concept contains an empty placeholder – a promissory note that says that there is someone you could go to who can tell you the answer (an idea that Putnam referred to as the "linguistic division of labor"). A corollary of the proposal is that people should be willing to defer. They should be prepared to change their notion of a concept in the face of properly qualified expert opinion to the contrary.

A set of studies by Braisby (2001, 2004) examined the idea that people should defer to experts. In (Braisby 2001, Study 2) people were told that a number of food-stuffs (apple, potato, salmon and chicken) had been genetically modified, introducing either 0, or 50% of genetic material from other species either in the same superordinate category (e.g. fish for salmon) or from a different superordinate (e.g. a mammal for salmon). They were also told that according to biologists the food was (or was not) still (say) salmon. Participants tended to defer to the expert opinion around 75% of the time, suggesting that there is a majority of people who are willing to accept expert opinion. Willingness to defer however was reduced when the genetic modification was more extreme (using material from an unrelated animal or plant), with only 63% now agreeing with expert opinion that the stuff was still of the same kind. (The design was such that 50% agreement would indicate no effect of deference to the expert).

In a control study (Study 3), Braisby changed the source of the "expert" opinion from biologists to other shoppers. In this case there was still significant deference, but the shift dropped to about 60% average compliance to the opinion of other shoppers.

Another study that tested for beliefs in essences was run by Kalish (1995). In one of

his reported studies, Kalish posed the following kind of question. Suppose there are two people who are visiting the zoo and they see a creature in a cage, but there is no information available about what kind of creature it is. They discuss what kind of creature they are looking at, and disagree. Is there a matter of fact that will decide who (if either) is correct, or is it possible for both of them to be right? Kalish used a variety of examples from both biological and artifact categories, and also included clear cases where it should be obvious that the answer is a matter of fact (e.g. mathematical cases) and others where it should be clearly a matter of opinion (e.g. whether B+ constitutes a "good grade"). Kalish found that people were more inclined to consider a biological kind to be a matter of fact, and an artifact kind to be a matter of opinion. Like Braisby (2001), however, the results were not overwhelmingly strongly in favor of essentialism. There was still a substantial minority of people who considered that even questions of biological categorization could be a matter of subjective opinion.

9.2.5 Relating Extensions and Intensions

I have suggested that full possession of a given lexical concept is likely to involve knowledge of both the extensional class to which the term refers, and the intensional information of what belonging in the class entails in terms of characteristic attributes and causal structures. An intriguing result by Malt and others (1999) has shown that in certain domains the link between extension and intension may be surprisingly weak. A common assumption, for example Gardenfors (2000), is that the referential use of different language terms in a given domain should map onto the intensional similarity structure of the domain. That is to say that if a set of objects is laid out in terms of the similarity of one to another so as to form a "similarity space" with proximity between items

proportionate to their similarity, then the set of terms with the same name should form a convex region of that space. Concept terms should label sets of things that are similar to one another, and different from things with other labels. Malt and others (1999) took photographs of a set of 60 different containers, including cups, mugs, cans, pots, juice boxes, lunch boxes etc. Participants came from three different language groups, speaking English, Spanish and Chinese. When they rated the pairwise similarity of the containers, by different criteria of similarity, there was good consensus across languages in terms of which pairs were similar. However when it came to giving a name to each container, the three languages divided up the 60 containers in very different ways. While the correlation of overall similarity judgments across languages was above 0.90, the correlation in naming categories was between 0.35 and 0.55. Malt and Sloman (2007) and Sloman and Malt (2003) argue that naming behavior is subject to other historical/linguistic factors so that the way in which a thing is named does not give a direct path to its conceptual content. Support for this idea can also be found in folk-taxonomy for biological kinds where early mis-classifications have been resistant to correction. Thus, according to Wikipedia, poison oak is not an oak, and poison ivy is not an ivy.

9.3 Vagueness and Meaning

Having laid out a brief account of how language and concepts inter-relate, I turn in this section to issues of vagueness and truth in natural language concepts. There are many theories of vagueness in the philosophical literature. The interested reader is referred to the collection of papers with an excellent introductory overview by Keefe and Smith (1997), and the highly readable book by van Deemter (2010). Broadly speaking vagueness arises when the applicability of a concept to the world is not precisely

determined. Much of the philosophical discussion has focused around the issue of the logic of sorites series. The sorites paradox (from the Greek for heap) concerns the applicability of the word “heap” to a quantity of sand. If one starts with a heap of sand, can the removal of a single grain change the heap into a “non-heap”? If the answer is “no”, then you remove the grain, and then ask the question again. Paradoxically it seems that removing a single grain cannot change the categorization of the pile, whereas removing many grains can. Discussion of the best logical treatment of this problem is still ongoing, and continuous-valued logics are one of the currently favored types of model (Smith 2008), although others find them problematic (Keefe 2000). I do not attempt to review this large literature. Instead I try to show how the different means through which the content of concepts is determined can lead to different problems of vagueness in language use.

9.3.1. Vagueness in Real-World Categories

It might be argued (indeed it has been argued) that the world itself cannot be vague. If you believe in the reality of the world (as most of us do), then you can believe that there are atoms, composed of elementary particles, and that these atoms move and cohere together to create the world. There is nothing else but this. Apart from the niceties of Heisenberg’s Uncertainty principle we don’t need to worry about the world of things and stuff being imprecise or vague. But now consider the idea of properties. Are properties real in the same way that stuff is real? Or are properties a creation of our minds that we use for describing and explaining the behavior of atomic stuff? A deep question (Mellor and Oliver 1997)! The question is important because there are many important terms even in a scientific description of the world that lack precise definitions, and so are open

to problems of vagueness. The concept of *species* in biology has been central in the development of the theory of evolution, in the development of taxonomic classification and an understanding of the relation of one type of organism to another. Yet, the concept has no precise definition (Mayr 1982). Neighboring species of trees, for example, can form hybrids of differing degrees of viability and fertility, so that the question of whether one particular organism is of a given species may become a matter of degree.

More generally however, concepts that take their contents from the external world are much less likely to be vague than those that derive from social norms or from individual mental contents. Science is our theory of the external world, and scientific concepts are developed in order to improve the precision of scientific descriptions and enable the construction of better theories. When we are uncertain about the truth of statements that are expressed in precise language, then we are in a position of *epistemological* uncertainty. When the weather forecasters say that there is a 50% chance of rain tomorrow, they mean something quite precise, in fact something like the following. The forecasters have taken the current dynamic and static properties of the weather as it is now, and linked it to a set of situations in the past that have similar properties in all relevant respects. Within that set, 50% of the situations went on to a situation in which it rained the next day.

Vagueness for concepts with external content is therefore most likely to arise from epistemological uncertainty rather than from vagueness in the concepts themselves.

9.3.2 Vagueness in Nominal Categories

Vagueness for our second kind of concept determination is also likely to be restricted. When society needs to reduce vagueness, then explicitly stipulated meanings are

provided, which will be designed to show as little vagueness as possible. A vaguely expressed law will inevitably result in lengthy and costly litigation.

9.3.3 Vagueness in Everyday Language

Probably the most common source of vagueness, and hence of partial truth, comes from the way in which language relates to the world. As we saw in the previous section, most of our vocabulary is the result of a process of mutual coordination among the users of the language. It shifts with the changes in culture and across different generations, and for the most part there is no particular reason why anyone should need to stipulate a precise meaning. According to Dunbar (1996), language evolved as an advanced form of social bonding, similar to the grooming behavior of other primate species. Word meanings only need to be as precise as is necessary to achieve successful communication, and if the primary function of conversation is in any case not to communicate facts but to share a warm social experience, then it is not surprising that there has been little pressure on word meanings to be more precise.

There are also good reasons (or good adaptive pressures) for retaining vagueness in language terms. The world is a very complex and constantly changing place. Word meanings have to be imprecise so that we can still achieve reference in the face of things for which we have no word that fits well.

Vagueness in meaning is closely associated in psychology with the prototype theory of concepts (see Chapter 4). According to Hampton (2006) there are four aspects to concepts that lead to them being considered as prototypes, two based on extensions and two based on intensions.

On the extensional side, there are the phenomena of borderline cases and of

typicality. The things to which a term refers (the extension) form a fuzzy set with some cases that are not clearly in or out of the category. This fuzziness can be measured empirically by asking a group of people to judge whether the term applies (a measure of vagueness suggested by Black 1937). McCloskey and Glucksberg (1978) provided additional evidence by asking the same group to re-make the judgments a few weeks later. Borderline cases, where people disagreed, were also likely to be the same cases where people gave responses on the second occasion that were inconsistent with their original judgments. Being a borderline case is therefore not just a matter of lack of consensus, but also of individual uncertainty. We will call this aspect of categorization "graded membership" to indicate that belonging in the extension is not all-or-none but graded. Typicality is a related phenomenon, first identified by Rosch and Mervis (1975). Even if we just consider the set of items for which everyone agrees that a term applies, there may still be differences in how *typical* or *representative* the item is of that category. Penguins and robins are both given 100% endorsement as birds, but at the same time most people agree that robins are more typical as birds than penguins.

Typicality and graded membership look very similar, and in early work on the problem of prototypes they were often confounded, leading to criticisms of the whole approach (Osherson and Smith 1981, 1982, 1997). It makes more sense to keep the ideas separate since the two measures have different ranges (Hampton 2007). According to prototype theory, both typicality and graded membership reflect how close an item is to the central tendency or prototype of the fuzzy class. However while typicality is a continuously increasing function of similarity to the prototype, graded membership remains at zero until the borderline region is reached, then rises continuously to one,

where it then remains constant as typicality continues to increase. The problems of how to use fuzzy logic to reflect both typicality and graded membership are discussed in Chapter 8 of this volume. Strictly, since logic is concerned with truth, the truth value of a statement should be mapped to its graded membership. However certain phenomena (such as compensation within concept conjunctions) require that typicality can also influence truth values (see Chapter 8).

Turning to the intensions of lexical concepts the standard way to measure or access these is to interview participants about the attributes or features that they consider to be part of the meaning of the term. In other words, individuals are invited to reflect on the meanings of their words. Their answers are then listed out as set of descriptive properties. For birds, this list might include "has wings", "flies", "has feathers" and "has two legs". The properties that are listed by more than a few participants are then assembled into a list which (at least for Hampton 1995b, 2006) constitutes the prototype of the category. Note that I take the notion of a prototype to be intensional – a set of correlated features or attributes – and not extensional (such as the best or maximally typical exemplar). The reason here is that some concepts may not match anything in the world very well. They can be used to refer, but they have no typical examples. (Jokey examples would be Gourmet English cooking recipes, or cool university professors. The best real-life example falls short of the most typical example that could be imagined in a different possible world. Hence the need for an intensional basis for prototypes).

What about intensions leads to the prediction of vagueness? First, it can often happen that the set of properties generated in this way does not provide a set of defining features that can be used to identify the reference or extension of the concept. As Wittgenstein

(1953) famously noted, the members of a class identified with a natural language term may share "family resemblances" but have no common element definition. There is no set of features *F* such that their conjunction provides a necessary and sufficient definition of the class. As a corollary of this result, we find that people generate many properties in the concept intension that are just "characteristic" of the set. They describe what a typical member of the class would be like. Hence the properties of birds that they fly and sing are not true of all birds. Flying is not necessary for a bird, and nor is it sufficient (many insects and some mammals also fly). To explain the lack of conjunctive definitions, and the failure of people to differentiate features that are necessary for membership (birds are the only two legged creatures with feathers) from features that are just generic to the class (birds generally fly), prototype theory proposes that we represent the meaning of a concept of this kind by representing what is common or typical about the class to which the word refers.

9.4 Some Cases of Vagueness in Language

In this section, I explore a number of cases of vagueness in the use and meaning of language terms. In each case the fact that the relation of language to the world is vague invites the logician to use some variety of fuzzy logic to describe the truth of statements involving these terms.

9.4.1 Adjectives

The most commonly debated form of vagueness in the philosophical literature is that which relates to simple scalar adjectives such as tall or bald. As a man Joe's height increases from, say 5 ft (152 cm) to 6 ft (183 cm), so the "truth" of the statement "Joe is tall" increases from zero (It is definitely false that Joe is tall) to one (It is definitely true

that Joe is tall). Since the underlying reality that is being described is a continuous scale of height, with interval and ratio scaling properties, the problem of vagueness arises in deciding how this continuous variable should be converted to a binary truth value of Tall or Not Tall.

Briefly, there are several options that have been proposed. There are three main varieties.

(a) The *epistemological approach* (Williamson 1994) stipulates that there is in fact a precise height at which it is correct to say that someone is tall. This height could depend on numerous contextual factors, and so it is beyond our cognitive capacities to discover it. Vagueness in whether it is correct to call Joe tall is therefore of the same kind as uncertainty about whether it is correct to say that the population of London at this precise moment in time is greater than 7.5 million. There is arguably a fact of the matter about the population of London, but so many individual cases would need to be resolved (people in the process of being born and dying, people in temporary residences but planning to remain, people who have moved out but still have their addresses, etc.) that it is practically impossible to know the answer with more than a certain level of precision.

(b) The *supervaluation approach* proposes that there is a range of heights for which "*x* is tall" is neither true nor false (in the related *subvaluation approach*, it is both true and false). On any given occasion on which the term is used however, the term can be made precise by choosing a particular cut-off point – for example over 5 foot 10 inches is tall, and less than or equal to 5 foot 10 inches is not tall. Statements employing vague terms can then be definitely true if they are true on every possible way of making the term precise (this is to be "super-true"). Hence a statement like

"Either Joe is tall or he is not tall."

is necessarily true, and a statement like

"Joe is both tall and not tall."

is necessarily false, in spite of the vagueness of "tall".

This approach handles this kind of sentence well, at least if one shares the common intuition about the necessary truth and falsity of the two statements. There are however empirical studies that suggest that people may sometimes entertain contradictory statements as being true. In experiments performed by Alxatib and Pelletier (2010), participants had to make judgments about a line-up of men set against height markers on a wall. For each of the five men, they had to judge whether the following four statements were True, False, or Can't Tell:

a. #1 is tall

b. #1 is not tall

c. #1 is tall and not tall

d. #1 is neither tall nor not tall

For cases of borderline tallness, participants were quite willing to judge that (c) was correct. Some 54% saw it as true, and of these a substantial number also thought that (d) was true for the same individual. The intuition that (c) must always be false was not shared by the students who took part in the study. . (The results were however consistent with a simple fuzzy logic account using the Minimum function for conjunction. Someone who is tall 0.5, will be "not tall" also at 0.5, and the conjunction of these will also be true to degree 0.5.)

(c) The third option is to use a continuous valued truth function, so that as the

likelihood of someone *x* being called Tall increases (as her height increases), so does the truth of the statement "*x* is tall". As outlined in earlier chapters, fuzzy logic has a wide and powerful array of tools for handling this type of case. I just list a couple of the common objections taken from (Keefe 2000).

First, when combining the truth of two propositions, the *relation* between the propositions has to be included, they cannot be treated as always independent. Thus the truth of the statement

"Ken is tall and Joe is not bald."

could be captured by a minimum rule, since the two statements are independent.

However if Joe is known to be 1 inch taller than Ken, then the statement

"Ken is tall and Joe is not tall."

should (arguably) always be false, since it would be irrational to say that the taller man is not tall, but that the shorter man is.

Thus even though the truth of the components may be equivalent (Joe is not bald = Joe is not tall), the outcome is affected by the lack of independence.

Second, fuzzy logic assigns to each proposition a precise degree of truth. So "Joe is tall" may be true at 0.7. But this use of the full power of the real numbers is inappropriate for capturing the true vagueness of language. Not only is it vague as to whether Joe is tall, but it is also vague as to just how vague it is! Fuzzy logicians can respond by assigning a fuzzy interval to be the truth value of a statement, but the interval itself would need to have fuzzy limits if the vagueness of vagueness is to be fully captured.

Talk of the vagueness of vagueness is obviously quite mind-bending, but the issue of so-called "higher-order" vagueness is a critical problem for logical treatments of

vagueness. There is no evidence that statements can be neatly divided into those that are definitely true, those that are definitely false, and the rest. In fact there is evidence to the contrary. Although the following study used noun categories rather than adjectives, it is likely that the results will generalize to any semantic decisions about concepts.

Hampton (2008) describes a study that examined whether people know when things are vague and when they are not. Two groups of participants were given a set of borderline cases to categorize, such as whether an olive is a fruit or whether sociology is a science. One group had to make a clear categorization True or False, so if they were uncertain they had to make a guess. The second group were allowed to use three responses. They were instructed that if they were 100% certain about the answer they should pick Definitely True, or Definitely False as their response. In all other cases they should pick Unsure. They were encouraged to use the latter category whenever they were not completely clear in their minds about the answer. The list included clear cases to enable them to use all three responses freely. The critical test came when the same individuals returned two weeks later and performed the test again, under exactly the same conditions. Now the performance of the first group with just two options should reflect the vagueness in the categorization task. If the truth of a statement like “Olives are Fruit” is vague, then there is a good chance that a person would be inconsistent in the two responses that they gave. This, after all, was the result that McCloskey and Glucksberg (1978) reported, as described previously. The second group however had the advantage of being able to place any inconveniently tricky statements into the Unsure response category. They should therefore be able to respond much more consistently to the categorization task, using their three response categories in the same way on each

occasion.

In fact, across several experiments using this methodology, I found no evidence at all that the group with three response options would be more consistent in their responding than the group with only two. The three-response group's data were analyzed by averaging the likelihood of changing one's mind about a Definitely True and the likelihood of changing one's mind about a Definitely False response. This average came out as exactly the same as the likelihood that someone in the two-response group would simply change a True to a False or vice versa.

It would appear then that logical treatments of vagueness have to be careful about not introducing sharp or precise values at any stage.

9.4.2 Vagueness in Noun Categories

A great deal of the psychological work on vagueness/fuzziness in language has focused on noun categories (Hampton 1998; McCloskey and Glucksberg 1978; Rosch and Mervis 1975). As detailed above, categories such as fruit, fish, furniture or sport provide an excellent set of materials with which to test the nature of vagueness in concepts. The advantage is that each category contains somewhere between 30 and 100 different category "members". People can be interrogated about what things they would place in the category, they can judge how typical those things are, they can provide ratings of whether each of the putative members possesses each of the intensional properties of the category, and so forth. In a large scale data collection exercise by De Deyne and others (2008) a large dataset has been collected including a wide range of psychological measures for a reasonably large sample of different semantic categories. This dataset, which can be freely downloaded from their website at Leuven University has started to

provide valuable tests of different quantitative models of how variability in categorization (and hence fuzziness of semantic sets) arises.

Why should noun categories be vague? We can identify the following possible sources of vagueness in a semantic category such as FRUIT, as seen in the levels of disagreement, within-rater inconsistency, and/or expressed lack of confidence in a categorization decision about an item such as TOMATO.

a) Different beliefs exist within the population. Perhaps 60% of people firmly believe that a tomato is a fruit, and 40% firmly deny it. Fuzziness then only exists at the group level, and we would expect disagreement, but no inconsistency and no lack of confidence. This difference in belief could either reflect a difference in what people think a fruit needs to be, or a difference in what people believe about the properties of tomatoes, or both.

b) Some people may feel *ignorant* of the facts about either tomatoes or fruits, or both. In that case they may hazard a guess, and we should find all three indicators of vagueness.

c) Categorization may be sensitive to the context or the perspective in which the judgment is made. What counts as a sport for TV sports channels may differ from what counts as a sport for an elementary school sports event. If the categorization task provides no context for the decision, then different contexts may be imagined by different people, and we should find both disagreement and inconsistency.

d) Categorization may be inherently probabilistic. When judging whether a tomato is a fruit, people generate a prototype for fruit, based on memories, knowledge and whatever contextual influences may be in play. They then generate a similar

representation for tomatoes. Finally then compare the one with the other, and determine an asymmetric similarity function of the tomato concept to the fruit concept. (For a suggestion of how the asymmetric similarity could work, see Hampton 1995a). This degree of similarity will contain random noise owing to the attributes that have been retrieved and the weights that have been attached to them. The similarity is then compared to a threshold value (also subject to random variation) and a positive decision results if the similarity is greater than the threshold.

The fact that people are inconsistent in categorization argues that differences in belief are not the only source of vagueness. A study by Hampton, Dubois and Yeh (1996) suggests that the context may not be that critical either. In this study, participants categorized lists of items in 8 different categories. Four different groups were used (three only were reported in the published article). One had no context – they simply decided if the items (like tomato) belonged in the categories (like fruit). A second group were given "linguistic" instructions. They had to imagine that we were interested in how speakers of US English used these words, so did they think that they would spontaneously refer to the item (tomato) using the category name (fruit). A third group had the task set in a more technical legalistic setting. They were told to imagine that they were advising a governmental agency in setting up the basis for tax and import regulations concerning different classes of products. In this case, should the tomato count as a fruit or not? Finally a fourth group were set a pragmatic version of the task. They were asked to place things in the fruit category in the context of a search engine. Thus if they thought people would expect to find tomato listed under the search term fruit, they should include it in the category.

In spite of our attempt to produce widely different categorization contexts, the results were almost entirely negative. Context had only a few, category-specific effects. For example in one context people were invited to place items in the furniture category for the purposes of an inventory database for a department store. This context generated some shifts from the standard no-context control, which reflected people's knowledge of how items are displayed in department stores (for example refrigerators are found with electrical goods and not with furniture). Overall however, providing a context had zero effect on either reducing disagreement, reducing inconsistency, or reducing the high correlation between probability of categorization and context-free ratings of the typicality of items in the categories.

It is most likely then, that the major source of vagueness in semantic categorization comes from the final source listed above. Deciding on category membership is an inherently probabilistic process, subject to random processes in the collection and assessment of the evidence involved.

9.4.3 A Model of Vagueness in Semantic Categorization

Verheyen, Hampton and Storms (2010) provide a formalization of a quantitative probabilistic model for capturing the fuzzy nature of categorization. It assumes that categorization is a probabilistic process in which each individual assesses the similarity of any item relative to their own individual threshold for making a categorization response. Formally, the model uses the Rasch model developed in the domain of psychometrics (Rasch, 1960). The probability of a positive categorization of item i in Category A by participant p, $P(A,i,p)$, is given by

$$P(A,i,p) = \frac{e^{\alpha(\beta_i - \theta_p)}}{1 + e^{\alpha(\beta_i - \theta_p)}}. \quad (9.1)$$

The function defined by Eq.(9.1) increases from 0 to 1 as the difference between β_i and θ_p becomes more positive. The β_i parameter represents the position of item i on a scale that measures similarity to the category prototype. The parameter θ_p is the threshold for participant p. Thus as an item's similarity exceeds the person's threshold to a greater extent, so the probability of making a positive response also increases.

The model was fit to data from 250 participants who each made yes/no categorization judgments on lists of 24 items for each of 8 categories, some 48,000 data points. The fit of the model was good, and it was shown to correctly estimate the degree of non-modal responding across the scale (the amount of disagreement among participants), as well as the degree of inconsistency observed in previous studies. Most impressively, the estimates of the β_i parameters for particular items correlated at around .95 with an independent measure of the typicality of those items for the category. A further result was to show that an individual's threshold parameter θ_p showed consistency across different categories. Correlations were fairly low, but were significantly greater than zero, indicating that some individuals tend towards a liberal view and some a conservative view when categorizing across a range of different categories (Gardner 1953).

We also examined whether the parameter α in (6) needed to be allowed to vary from item to item, or whether a single value constant value could be used. This parameter is a measure of the degree of vagueness or fuzziness in an item's categorization. As it increases, so the slope of the S shaped curve relating probability of categorization to similarity becomes steeper. Hence high values of α indicate a sharply defined item-

category decision, and low values a vaguely defined (more probabilistic) item and category. We found that different values were required for different categories, but that overall the best model used a constant α for all items *within* the same category.

The value of this model for studies of fuzziness in categorization are that it can incorporate differences in individuals, differences among items and differences between categories or category domains. It therefore provides a general methodology for testing important hypotheses about the source and nature of variability and fuzziness in human categorization.

9.5 Final Remarks

The meanings of the substantive words in human languages are almost always vague. I have reviewed how within a language there are different kinds of words that owe their contents respectively to the external world, to experts in a society, or to individuals who form part of a group of language users. The kind of vagueness that is found is accordingly different in its source. I have reviewed the difference between knowledge of a concept's extension and knowledge of its intension, and the existence of cases where the two appear to be only weakly coordinated. Finally, I have argued that the major source of vagueness in language is the fact that people's conceptual contents are determined through a dynamic process of keeping calibrated with other language users. This process generates a certain optimal level of vagueness which allows successful communication while providing the flexibility for meanings to adapt as the social and physical world around us changes.

References

Ahn, W. K., N. S. Kim, M. E. Lassaline, and M. J. Dennis. 2000. Causal status as a determinant of feature centrality. *Cognitive Psychology* 41: 361-416.

Alxatib, S. and F. J. Pelletier. 2010. The psychology of vagueness: Borderline cases and contradictions. Unpublished MS, Cambridge MA: MIT.

Black, M. 1937. Vagueness: an exercise in logical analysis. *Philosophy of Science* 4: 427-455.

Braisby, N. R. 2004. Deference and essentialism in the categorization of chemical kinds. In *Proc. of the 25th Annual Conf. of the Cognitive Science Society*, ed. R. Alterman and D. Kirsch, 174-179. Malwah, NJ: Lawrence Erlbaum.

Braisby, N. R. 2001. Deference in categorisation: Evidence for essentialism? In *Proc. of the 23rd Annual Conf. of the Cognitive Science Society*, ed. J. D. Moore and K. Stenning, 150-155. Malwah, NJ: Lawrence Erlbaum.

De Deyne, S., S. Verheyen, E. Ameel, W. Vanpaemel, M. Dry, W. Voorspoels, and G. Storms. 2008. Exemplar by feature applicability matrices and other Dutch normative data for semantic concepts. *Behavior Research Methods* 40: 1030-1048.

Diesendruck, G. and S. A. Gelman.1999. Domain differences in absolute judgments of category membership: Evidence for an essentialist account of categorization. *Psychonomic Bulletin and Review* 6: 338-346.

Dunbar, R. 1996. *Grooming, gossip, and the evolution of language*. London: Faber & Faber.

Estes, Z. 2004. Confidence and gradedness in semantic categorization: Definitely somewhat artifactual, maybe absolutely natural. *Psychonomic Bulletin and Review* 11: 1041-1047.

Fodor, J. A. 1975. *The language of thought*. New York: Crowell.

Fodor, J. A. 1998. *Concepts: Where cognitive science went wrong*. Oxford: Clarendon Press.

Gärdenfors, P. 2000. *Conceptual spaces*. Cambridge MA: MIT Press.

Gardner, R. W. 1953. Cognitive styles in categorizing behavior. *J. of Personality* 22: 214-233.

Hampton, J.A. 1995a. Similarity-based categorization: The development of prototype theory. *Psychologica Belgica* 35: 103-125

Hampton, J. A. 1995b. Testing Prototype Theory of Concepts. *J. of Memory and Language* 34: 686-708.

Hampton, J. A. 1998. Similarity-based categorization and fuzziness of natural categories. *Cognition* 65: 137-165.

Hampton, J. A. 2006. Concepts as Prototypes. In *The Psychology of Learning and Motivation: Advances in Research and Theory,* Vol. 46*,* ed. B. H. Ross, 79-113. Amsterdam: Elsevier.

Hampton, J. A. 2007. Typicality, graded membership and vagueness. *Cognitive Science* 31: 355-383.

Hampton. J. A. 2008. The Rumsfeld Effect: the unknown unknown. Paper presented at the Annual Convention of the Psychonomic Society.

Hampton, J. A., D. Dubois, and W. Yeh 2006. The effects of pragmatic context on classification in natural categories. *Memory & Cognition* 34: 1431-1443.

Kalish, C. W. 1995. Essentialism and graded membership in animal and artifact categories. *Memory & Cognition* 23: 335-353.

Keefe, R. 2000. *Theories of vagueness*. Cambridge: Cambridge University Press.

Keefe, R. and P. Smith. 1997. Theories of vagueness. In *Vagueness: a reader,* ed. R. Keefe & P. Smith, 1-57. Cambridge: MIT Press.

Keil, F. C. 1981. Constraints on knowledge and cognitive development. *Psychological Review* 88: 197-227.

Kripke, S. 1980. *Naming and necessity*. Cambridge, MA: Harvard University Press.

Malt, B. C. and S. A. Sloman. 2007. Artifact categorization: the good, the bad, and the ugly. In *Creations of the mind: Theories of artifacts and their representation,* ed. E. Margolis and S. Laurence, 85-123. Oxford: Oxford University Press.

Malt, B. C., S. A. Sloman, S. Gennari,, M. Shi, and Y. Wang. 1999. Knowing versus naming: Similarity and the linguistic categorization of artifacts. *J. of Memory and Language* 40: 230-262.

Mayr, E. 1982. *The growth of biological thought: Diversity, evolution, and inheritance*. Cambridge, MA: Harvard University Press.

McCloskey, M. and S. Glucksberg. 1978. Natural categories: Well-defined or fuzzy sets? *Memory & Cognition* 6: 462-472.

Medin, D. L. and A. Ortony. 1989. Psychological essentialism. In *Similarity and analogical reasoning,* ed. S. Vosniadou and A. Ortony, 179-195. Cambridge: Cambridge University Press.

Mellor, D. H. and A. Oliver, eds. 1997. *Properties*. Oxford: Oxford University Press.

Nosofsky, R. M. 1988. Exemplar-based accounts of relations between classification, recognition, and typicality. *J. of Experimental Psychology: Learning, Memory, and Cognition* 14: 700-708.

Osherson, D. N. and E.E. Smith. 1981. On the adequacy of prototype theory as a theory of concepts. *Cognition* 11: 35-58.

Osherson, D. N. and E.E. Smith. 1982. Gradedness and conceptual conjunction. *Cognition* 12: 299-318.

Osherson, D. N. and E.E. Smith. 1997. On typicality and vagueness. *Cognition* 64: 189-206.

Putnam, H. 1975. The meaning of 'meaning'. In *Mind, language, and reality: Philosophical papers,* Volume 2*,* ed. H. Putnam, 215-271. Cambridge: Cambridge University Press.

Rey, G. 1983. Concepts and stereotypes. *Cognition* 15: 237-262.

Rasch, G. 1960. *Probabilistic models for some intelligence and attainment tests*. Copenhagen: Danish Institute for Educational Research.

Rosch, E. R. and C. B. Mervis. 1975. Family resemblances: studies in the internal structure of categories. *Cognitive Psychology* 7: 573-605.

Sloman, S. A. and B. C. Malt. 2003. Artifacts are not ascribed essences, nor are they treated as belonging to kinds. In *Conceptual representation*, ed. H. E. Moss and J. A. Hampton, 563-582. Hove, UK: Psychology Press.

Smith, N. J. J. 2008. *Vagueness and degrees of truth*. Oxford: Oxford University Press.

van Deemter, K. 2010. *Not exactly: In Praise of vagueness*. Oxford: Oxford University Press.

Verheyen, S., J. A. Hampton, and G. Storms. 2010. A probabilistic threshold model: Analyzing semantic categorization data with the Rasch model. *Acta Psychologica* (in press).

Wikipedia: http://en.wikipedia.org/wiki/Toxicodendron, accessed 24.06.2010.

Williamson, T. 1994. *Vagueness*. London: Routledge.

Wittgenstein, L. 1953. *Philosophical investigations*. New York: Macmillan.