

Word senses are not bona fide objects: implications for cognitive science, formal semantics, NLP

Adam Kilgarriff*
ITRI
University of Brighton

August 14, 1996

1 Introduction

Those authors who have looked closely at the nature of lexical meaning and polysemy have invariably come to the conclusion that a model in which the senses of a word are discrete and unrelated objects is untenable. While many have examined models and implications of a word's senses being related (Wilks, 1975; Lakoff, 1987; Pustejovsky, 1991; Atkins & Fillmore, 1994; Copestake & Briscoe, 1995) the fact that word senses are not in general discrete has received far less attention.

To lexicographers, it is a commonplace that identifying a word's senses, or deciding where one sense ends and another begins, is more art than science. Recent developments in corpus-based lexicography support a model in which a word sense corresponds to a cluster of corpus instances for the word. In some cases, clusters will be distinct, and all data will fall into one cluster or another, so a 'discrete senses' model will be valid, but at other times, clusters will overlap, be ill-defined, and not cover all the data. The word-sense-as-corpus-cluster (WSCC) model stands in stark contrast to the near-universal assumption in cognitive science, NLP and formal semantics that the word sense is a basic object in the ontology.

In this paper, I first review the arguments for the WSCC position and then consider its implications for cognitive science and formal semantics. I close with some comments on its implications for NLP.

2 Word senses are not always discrete

Researchers investigating the compositionality of meaning, as in formal semantics, or relations between word meanings, as in Cruse's textbook (Cruse, 1986), require some basic objects for their theories to manipulate. At a first pass, these 'atoms of meaning' would be words – but some words are ambiguous, and the semantics of “all the banks

*This work is supported by the EPSRC, Grant K18931, SEAL. I would also like to thank Richard Power for his help in clarifying the views expressed.

collapsed” is clearly not the same in the cases where the bank is a river bank and a money bank, so, at a second pass, the atoms of meaning are word senses. Below I list a series of arguments which show that this second pass is still not satisfactory.

2.1 Argument from language learning

A child ordinarily learns a word through hearing it in situ. The child hears it used, in a number of sentences, in a number of situations (or, to be more exact, in a number of paired linguistic and extra-linguistic contexts), and (in a process well beyond our current abilities at cognitive modelling) abstracts away from the contexts to identify what contribution the new word is repeatedly making, and the resulting abstraction, paired with the sound (and/or look) of the word is then stored in the mental lexicon. I hope that much is uncontentious, the words ‘abstract’ and ‘abstraction’ being woolly enough to accommodate all theories.

A word may, of course, repeatedly make more than one contribution. *Moose/mousse* repeatedly makes light-dessert contributions, and repeatedly makes big-deer contributions, and the child confronts the still harder task of identifying two abstractions from a mixed bag of contexts.

The range and variety of contexts that language learners hear the word in determines the range of meaning the word has for those learners and hence, as those learners become the next generation of fully competent language users, for the lexicon of the language in general. Sometimes, the sentences and contexts all support a single, simple model of the word’s meaning. At other times, as in *mousse/moose*, they support multiple models.

In general, the learner confronts the following questions. Do all the contexts contribute to a single model? Are some of them distinct enough to support distinct models, or should the original model be extended or rendered more complex in some way to encompass a wider range of contexts? We do not know enough about the internal architecture of the mental lexicon to know whether these questions correspond to alternative semantic structures in the mind. If they do not, then word senses are not discrete to an individual language user.

Even if they do, given different individuals’ different linguistic and extralinguistic histories, it is not plausible that all the speakers of a language develop models which divide up the semantic territory of all words in the same way. Even if discreteness is a feature of an individual’s mental lexicon, it is certainly not a feature of the lexicon of the language.

2.2 Argument from (failure of) ambiguity tests

It would mean little or nothing to say that word senses were discrete, if there was no way of determining where there were two senses for a word and where there was only one which encompassed a measure of vagueness or indeterminacy in how it was used. Various tests have been proposed to distinguish the two cases, but all are found wanting (Zwicky & Sadock, 1975; Cruse, 1986; Kilgarriff, 1992; Geeraerts, 1993).

2.3 Argument from lexicographers

When writing a dictionary entry for any but the simplest of words, the lexicographer's second hardest task is to determine how its meaning is best divided into distinct senses.¹ As detailed accounts of the process by lexicographers show (Stock, 1983; Krishnamurthy, 1987; Atkins & Fillmore, 1994), the task is one of organising and structuring knowledge about the word and its use, never forgetting that the goal is to help the dictionary user understand how the word is used. If word senses were discrete, then, presumably, this task would simply be one of listing the discrete senses – but this is clearly far from lexicographers' experience, for very many words.

2.4 Argument against taking the dictionary too seriously

Dictionaries are social artifacts. The format of the dictionary has remained fairly stable since Dr. Johnson's day. The reasons for the format, and the reasons it has proved so resistant to change and innovation, are explored at length in (Nunberg, 1994). In short, the development of printed discourse, particularly the new periodicals, in England in the early part of the eighteenth century brought about a re-evaluation of the nature of meaning. No longer could it be assumed that a disagreement or confusion about a word's meaning could be settled face-to-face, and it seemed at the time that the new discourse would only be secure if there was some mutually acceptable authority on what words meant. The resolution to the crisis came in the form of Johnson's Dictionary. Thus, from its inception, the modern dictionary has had a crucial symbolic role: it represents a methodology for resolving questions of meaning. Hence "the dictionary", with its implications of unique reference and authority (cf. "the Bible").

So the dictionary has its own reasons for presenting meanings as discrete. If it presented meanings as vague and flexible, prototypes with barely circumscribable limits, it would be of little use for resolving disputes about meaning. Publishers expect to publish, bookshops expect to sell, and buyers expect to buy and use dictionaries which at least appear to fix meaning.

The well-ordered world of discrete meanings is an alluring one, one we have grown comfortable with since childhood. Sadly, this is no reason to believe it paints a true picture.

2.5 Argument from sense-tagging

If senses were discrete, one would expect to be able to assign most occurrences-in-context of a word to exactly one of the word's senses.² 'Most', rather than 'all', because puns would still be possible, and there may be other cases where two or more senses could apply without compromising the discreteness of the underlying senses.

¹According to a survey of lexicographers at Longman. The hardest task is choosing the wording for the definition.

²It does not immediately follow. If meanings are defined independently of the contexts the word occurs in, discreteness of meanings would not imply unambiguous sense-tagging. However, since Wittgenstein, and given the considerations regarding language learning discussed above, the potential for defining meanings independently of contexts appears very limited.

The little evidence there is indicates that people do not find the sense-tagging task easy. Jorgensen (1990) investigated whether subjects clustered a set of occurrences-in-context into the same clusters when they weren't, and (a week later) were shown dictionary definitions for the word's senses. On a scale from -1 (complete disagreement) to 1 (complete agreement) scores for different words ranged from -0.05 to 0.83. Ng and Lee (1996) found that the human taggers working in Princeton on the SEMCOR project assigned the same WordNet tag as their own team of taggers just 57% of the time. (See also Gale, Church, and Yarowsky (1992).)

The difficulty of the sense-tagging task suggests that the theory under which it would tend to be easy – that senses are discrete – is false.

2.6 Word sense as corpus cluster

All these arguments serve to undermine the assumption that word senses are discrete objects. The language-learning, lexicographers', and sense-tagging arguments point to an alternative conception, in which the occurrence-in-context is the basic object and a word sense is a construct built by the language learner or lexicographer, to make sense of a tolerably coherent subset of all the occurrences encountered. Occurrences-in-context are events of relatively unproblematic ontological status, in contrast to word senses, which have no clear ontological status.

The WSCC view is compatible with the discrete-senses view. It could be that, for almost all words, almost all of the time, the discrete senses view is adequate, since usages do in fact fall into distinct and disjoint clusters. If so, the implications of the WSCC view would be much reduced. While quantitative conclusions in this area are not easily drawn, the evidence from sense-tagging presented above does strongly suggest that this is not the case.

It seems likely that, for some words, corpus instances fall into distinct clusters and it is reasonable to treat word senses as discrete, whereas for others, the situation is less well-defined. *Image* is such a word, and I present the non-discrete meanings here (as presented in LDOCE (1987), see also Kilgarrieff (1993)) for illustrative purposes.

image *n* **1** [C(**of**)] a picture formed in the mind: *She had a clear image of how she would look in twenty years time.*

2 [C] a picture formed of an object in front of a mirror or LENS . . .

3 [C] the general opinion about a person, organisation, etc., that has been formed or intentionally created in people's minds: *The government will have to improve its image if it wants to win the next general election.* . . .

4 [(*the*)S(**of**)] a copy: *He's the (very) image of his father.* **5** [*the*+S+**of**] a phrase giving an idea of something in a poetical form, esp. a METAPHOR or SIMILE **6** *old use* likeness; form: *According to the bible, man was made in the image of God.*

3 Relationship to cognitive science

At the heart of cognitive science lie concepts. How they are learnt, how structured, how they relate to each other, and how they can be combined are all central concerns of cogni-

tive psychology. How they can be modelled formally, in computationally tractable ways, are basic research themes in Artificial Intelligence. How they relate to objective reality is, of course, fundamental to philosophy of mind. If we shift to a linguistic perspective, we are concerned with concepts seen as meanings of words. We re-name them as ‘entries in the mental lexicon (semantic and pragmatic aspects)’, as opposed to ‘general knowledge’ (or ‘objects in the language of thought’) and investigate how they are accessed (particularly in psycholinguistics) and how they combine to give sentence meanings and text meanings.

In mainstream cognitive science models of language understanding, one step in understanding a sentence is lexical access. The meaning or meanings of each word are retrieved from the mental lexicon. If there is more than one meaning, the word will be disambiguated (or, the preferred meaning will be selected on a default basis), and an interpretation of the sentence will be built using the appropriate meaning. The meanings may be simple or complex: in sophisticated models, they may be structured entities that can be coerced by other elements of the sentence to take the form that is consistent in that context (see eg., Pustejovsky (1991), also contributions to the Workshop on “Breadth and Depth of Semantic Lexicons” at ACL ’96), so a single object retrieved from the lexicon may encompass more than one meaning. The paradigm requires that meanings are packaged into distinct units (possibly sharing features). For words such as *image*, where there is a substantial degree of arbitrariness in how any lexicographer might divide some areas of the word’s meaning into distinct packages, the paradigm is not satisfactory.

This critique has certain similarities with Lakoff’s attack on what he calls the ‘classical’ tradition (a label I adopt from here on), in which concepts are defined by necessary and sufficient conditions (Lakoff, 1987). There are two major differences between Lakoff’s critique and the point made here. Firstly, whereas Lakoff aims for a Kuhnian paradigm shift, rejecting the classical account in its entirety, I aim merely to demonstrate its limitations. The great progress that has been made on classical premises is not undermined by the non-discreteness of word senses. Rather, we establish that there will be some territory – particularly the mental lexicon, and NLP word sense disambiguation – where we must set them aside.

Secondly, Lakoff’s critique comes from a mentalistic perspective. He is primarily concerned with ‘language in the mind’, and the argument that mental concepts are not classical concepts. The WSCC position relates to ‘language in the world’, as exemplified in language corpora. One move to ‘save’ a classical account of meaning is to maintain that any individual’s lexicon had a determinate number of meanings for a word, even if there is no reason to believe that that individual’s meanings matches any other’s. From a language-in-the-mind perspective, that is a legitimate possibility that Lakoff must counter with psychological evidence that individual’s concepts are not like that. But from a language-in-the-world perspective, the defence misses the point, since the object of study is the lexicon of English which the native speaker learns, not the lexicon in any individual’s head.

The new paradigm promoted by Lakoff’s critique is prototype theory. Work in prototype theory includes issues of how prototypes combine, though clearly, prototype theory have a far less well-developed view of how sentence-meaning relates to word-meaning

than classical theory. It also includes an ongoing debate on what a prototype is, and it is by no means clear that the prototypes in prototype theory are any less discrete than concepts in classical theory. Prototype theory does not in any obvious way answer language-in-the-world based concerns about the discreteness of word meanings.

A related paradigm, to which similar considerations apply, is connectionism. Like prototype theory, connectionism draws its inspiration from mental phenomena. The ways in which concepts or word senses may be represented do not necessarily lend themselves to capturing the non-discreteness of word senses. Probably the best-known piece of connectionist work on word sense ambiguity, Cottrell (1989), presents no account of non-discreteness.

Within classical theory, facts about the lexical contexts a word occurs in play no part in defining what it means, and facts about its syntactic contexts play only a limited role. (The syntax of the word determines the form of the definition, but there is a limited number of syntactic categories.) For ambiguous words, lexical and finer-grained syntactic contexts will be useful for identifying which meaning any particular instance of the word has, but this is seen as a quite distinct from the semantic information that constitutes the word meaning. This position becomes untenable once the arguments above regarding the nature of word senses are acknowledged. If two senses for a word are distinguished on the grounds that the word occurs in two distinct lexico-syntactic patterns, and the core meaning takes on a different slant in the two cases, as in senses 1, 3, 4 and 5 of *image*, then the contexts play a central role in constituting the meanings. The contextual information that provides the evidence on which learners learn the meaning(s) of words, and which hearers use to determine what contribution the word is making to sentence-meaning, is part of the lexical representation for the word.

The implications for cognitive science are, then, that accounts of the mental lexicons should acknowledge the non-discreteness of word senses, and the role of lexical and syntactic contexts in constituting them.

4 Implications for Formal Model-Theoretic Semantics

Model-theoretic semantics operates with a model in which meanings are defined in terms of denotations. The challenge in formal semantics lies in finding the rules of composition. An example: the referent of a definite noun phrase such as “the cookie” is a member of the set denoted by the head noun (*cookie*). The denotation is a given, from which the analysis proceeds.

Recent work on vagueness and ambiguity in formal semantics (Pinkal, 1995) would appear to link formal semantics to the nature of word senses. However this work is within the denotational paradigm, and, for a word to be vague is for it to have different denotations in different contexts, in such a way that further ‘precisification’ of the context picks out the appropriate denotation. The challenge is to formally represent meanings of vague and ambiguous words in such a way that the appropriate truth conditions for the sentences they occur in can be built from the meanings of the words. The vagueness of terms like *tall* and *fast* provide the paradigm examples: the indefiniteness of *image* is, for formal semantics, too ill-defined to be of interest. Model-theoretic semantics addresses

the relation between language and the extralinguistic world, and focuses on linguistic data with clear relationships with the extralinguistic world. By contrast, word senses are usually hard to tease apart just where there is no clear relation between uses of the word and the external world.

Formal semantics defines polysemy and vagueness as problems you can ‘cash out’ as issues of there being (contextually defined) disjoint sets of objects in the world. So it has a way of dealing with them, which can be brought into action once we know which words are polysemous and which vague, how, under what circumstances, and so forth. But that leaves substantive questions about the nature of word senses untouched.

5 Implications for NLP

The WSCC position sheds new light on what is seen by many as the key NLP task of word sense disambiguation (WSD).

NLP applications should choose the word senses between which they want to distinguish with care. The word senses a dictionary lists are not objective truths about the language. Rather, they are written for the dictionary’s intended audience, in the context of lexicographic traditions. It is a common error to assume that the set of word senses given in a (monolingual) dictionary is relevant for any NLP application. For most if not all NLP applications, to disambiguate between all the word senses in such a set would be a great misuse of effort, and doomed to failure.

For people and for NLP systems, word senses are not permanent fixtures of the language-understanding situation. Rather, they are abstractions from particular sets of usages, often in specific domains, distinguished only when there is a human or computational reason for doing so.³

Reference

- Atkins, B. T. S., & Fillmore, C. J. (1994). Starting where the dictionaries stop: the challenge of corpus lexicography. In Atkins, B. T. S., & Zampolli, A. (Eds.), *Computational approaches to the lexicon*, pp. 350–393. Oxford University Press, Oxford.
- Copetake, A. A., & Briscoe, E. J. (1995). Semi-productive polysemy and sense extension. *Journal of Semantics*, forthcoming.
- Cottrell, G. W. (1989). *A Connectionist Approach to Word Sense Disambiguation*. Pitman, London.
- Cruse, D. A. (1986). *Lexical Semantics*. CUP, Cambridge, England.
- Gale, W., Church, K., & Yarowsky, D. (1992). Estimating upper and lower bounds on the performance of word-sense disambiguation programs. In *Proceedings, 30th ACL*, pp. 249–156.

³For an account of a survey on the relation of word-sense ambiguity to NLP applications, and a fuller discussion of the implications of WSCC, see Kilgarriff (1996).

- Geeraerts, D. (1993). Vagueness's puzzles, polysemy's vagueness. *Cognitive Linguistics*, 4(3), 223–272.
- Jorgensen, J. C. (1990). The psychological reality of word senses. *Journal of Psycholinguistic Research*, 19(3), 167–190.
- Kilgarriff, A. (1992). *Polysemy*. Ph.D. thesis, University of Sussex, CSRP 261, School of Cognitive and Computing Sciences.
- Kilgarriff, A. (1993). Dictionary word sense distinctions: An enquiry into their nature. *Computers and the Humanities*, 26(1–2), 365–387.
- Kilgarriff, A. (1996). 'I don't believe in word senses'. *Computational Linguistics*, submitted.
- Krishnamurthy, R. (1987). The process of compilation. In Sinclair, J. M. (Ed.), *Looking Up: An Account of the COBUILD Project in Lexical Computing*, chap. 3. Collins, London.
- Lakoff, G. (1987). *Women, Fire and Dangerous Things*. University of Chicago Press.
- LDOCE (1987). *Longman Dictionary of Contemporary English, New Edition*. Edited by Della Summers. Harlow.
- Ng, H. T., & Lee, H. B. (1996). Integrating multiple knowledge sources to disambiguate word sense: An exemplar-based approach. In *ACL Proceedings*.
- Nunberg, G. (1994). The once and future dictionary. Presentation at *The Future of the Dictionary Workshop*, Uriage-les-Bains, France.
- Pinkal, M. (1995). *Logic and Lexicon*. Kluwer, Dordrecht.
- Pustejovsky, J. (1991). The generative lexicon. *Computational Linguistics*, 17(4), 409–441.
- Stock, P. F. (1983). Polysemy. In *Proc. Exeter Lexicography Conference*, pp. 131–140.
- Wilks, Y. A. (1975). A preferential pattern-seeking semantics for natural language inference. *Artificial Intelligence*, 6, 53–74.
- Zwicky, A. M., & Sadock, J. M. (1975). Ambiguity tests and how to fail them. *Syntax and Semantics*, 4, 1–36.