



Do Word Meanings Exist?

PATRICK HANKS

Oxford English Dictionaries

1. Introduction

My contribution to this discussion is to attempt to spread a little radical doubt. Since I have spent over 30 years of my life writing and editing monolingual dictionary definitions, it may seem rather odd that I should be asking, do word meanings exist? The question is genuine, though: prompted by some puzzling facts about the data that is now available in the form of machine-readable corpora. I am not the only lexicographer to be asking this question after studying corpus evidence. Sue Atkins, for example, has said “I don’t believe in word meanings” (personal communication).

It is a question of fundamental importance to the enterprise of sense disambiguation. If senses don’t exist, then there is not much point in trying to ‘disambiguate’ them – or indeed do anything else with them. The very term *disambiguate* presupposes what Fillmore (1975) characterized as “checklist theories of meaning.” Here I shall reaffirm the argument, on the basis of recent work in corpus analysis, that checklist theories in their current form are at best superficial and at worst misleading. If word meanings do exist, they do not exist as a checklist. The numbered lists of definitions found in dictionaries have helped to create a false picture of what really happens when language is used.

Vagueness and redundancy – features which are not readily compatible with a checklist theory – are important design features of natural language, which must be taken into account when doing serious natural language processing. Words are so familiar to us, such an everyday feature of our existence, such an integral and prominent component of our psychological makeup, that it’s hard to see what mysterious, complex, vague-yet-precise entities meanings are.

2. Common Sense

The claim that word meaning is mysterious may seem counterintuitive. To take a time-worn example, it seems obvious that the noun *bank* has at least two senses: ‘slope of land alongside a river’ and ‘financial institution’. But this line of argument is a honeytrap. In the first place, these are not, in fact, two senses of a single word;

they are two different words that happen to be spelled the same. They have different etymologies, different uses, and the only thing that they have in common is their spelling. Obviously, computational procedures for distinguishing homographs are both desirable and possible. But in practice they don't get us very far along the road to text understanding.

Linguists used to engage in the practice of inventing sentences such as "I went to the bank" and then claiming that it is ambiguous because it invokes both meanings of **bank** equally plausibly. It is now well known that in actual usage ambiguities of this sort hardly ever arise. Contextual clues disambiguate, and can be computed to make choice possible, using procedures such as that described in Church and Hanks (1989). On the one hand we find expressions such as:

people without bank accounts; his bank balance; bank charges; gives written notice to the bank; in the event of a bank ceasing to conduct business; high levels of bank deposits; the bank's solvency; a bank's internal audit department; a bank loan; a bank manager; commercial banks; High-Street banks; European and Japanese banks; a granny who tried to rob a bank

and on the other hand:

the grassy river bank; the northern bank of the Glen water; olive groves and sponge gardens on either bank; generations of farmers built flood banks to create arable land; many people were stranded as the river burst its banks; she slipped down the bank to the water's edge; the high banks towered on either side of us, covered in wild flowers.

The two words **bank** are not confusable in ordinary usage.

So far, so good. In a random sample of 1000 occurrences of the noun **bank** in the British National Corpus (BNC), I found none where the 'riverside' sense and the 'financial institution' sense were both equally plausible. However, this merely masks the real problem, which is that in many uses NEITHER of the meanings of **bank** just mentioned is fully activated.

The obvious solution to this problem, you might think, would be to add more senses to the dictionary. And this indeed is often done. But it is not satisfactory, for a variety of reasons. For one, these doubtful cases (some examples are given below) do invoke one or other of the main senses to some extent, but only partially. Listing them as separate senses fails to capture the overlap and delicate interplay among them. It fails to capture the imprecision which is characteristic of words in use. And it fails to capture the dynamism of language in use. The problem is vagueness, not ambiguity.

For the vast majority of words in use, including the two words spelled **bank**, one meaning shades into another, and indeed the word may be used in a perfectly natural but vague or even contradictory way. In any random corpus-based selection of citations, a number of delicate questions will arise that are quite difficult to resolve or indeed are unresolvable. For example:

How are we to regard expressions such as 'data bank', 'blood bank', 'seed bank', and 'sperm bank'? Are they to be treated as part of the 'financial institution'

sense? Even though no finance is involved, the notion of storing something for safe-keeping is central. Or are we to list these all as separate sense (or as separate lexical entries), depending on what is stored? Or are we to add a ‘catch-all’ definition of the kind so beloved of lexicographers: “any of various other institutions for storing and safeguarding any of various other things”? (But is that insufficiently constrained? What precisely is the scope of “any of various”? Is it just a lexicographer’s cop-out? Is a speaker entitled to invent any old expression – say, ‘a sausage bank’, or ‘a restaurant bank’, or ‘an ephemera bank’ – and expect to be understood? The answer may well be ‘Yes’, but either way, we need to know why.)

Another question: is a bank (financial institution) always an abstract entity? Then what about 1?

1. [He] assaulted them in a bank doorway.

Evidently the reference in 1 is to a building which houses a financial institution, not to the institution itself. Do we want to say that the institution and the building which houses it are separate senses? Or do we go along with Pustejovsky (1995: 91), who would say that they are all part of the same “lexical conceptual paradigm (lcp)”, even though the superordinates (INSTITUTION and BUILDING) are different?

The lcp provides a means of characterizing a lexical item as a meta-entry. This turns out to be very useful for capturing the systematic ambiguities which are so pervasive in language. . . . Nouns such as *newspaper* appear in many semantically distinct contexts, able to function sometimes as an organization, a physical object, or the information contained in the articles within the newspaper.

- a. The newspapers attacked the President for raising taxes.
- b. Mary spilled coffee on the newspaper.
- c. John got angry at the newspaper.

So it is with *bank*₁. Sometimes it is an institution; sometimes it is the building which houses the institution; sometimes it is the people within the institution who make the decisions and transact its business.

Our other *bank* word illustrates similar properties. Does the ‘riverside’ sense always entail sloping land? Then what about 2?

2. A canoe nudged a bank of reeds.

3. Ockham’s Razor

Is a bank always beside water? Does it have one slope or two? Is it always dry land? How shall we account for 3 and 4?

3. Philip ran down the bracken bank to the gate.
4. The eastern part of the spit is a long simple shingle bank.

Should 3 and 4 be treated as separate senses? Or should we apply Ockham’s razor, seeking to avoid a needless multiplicity of entities? How delicate do we want our sense distinctions to be? Are ‘river bank’, ‘sand bank’, and ‘grassy bank’ three different senses? Can a sand bank be equated with a shingle bank?

Then what about ‘a bank of lights and speakers’? Is it yet another separate sense, or just a further extension of the lcp? If we regard it as an extension of the lcp, we run into the problem that it has a different superordinate – FURNITURE, rather than LAND. Does this matter?

There is no single correct answer to such questions. The answer is determined rather by the user’s intended application, or is a matter of taste. Theoretical semanticists may be more troubled than language users by a desire for clear semantic hierarchies. For such reasons, lexicographers are sometimes classified into ‘lumpers’ and ‘splitters’: those who prefer – or rather, who are constrained by marketing considerations – to lump uses together in a single sense, and those who isolate fine distinctions.

We can of course multiply entities ad nauseam, and this is indeed the natural instinct of the lexicographer. As new citations are amassed, new definitions are added to the dictionary to account for those citations which do not fit the existing definitions. This creates a combinational explosion of problems for computational analysis, while still leaving many actual uses unaccounted for. Less commonly asked is the question, “Should we perhaps adjust the wording of an existing definition, to give a more generalized meaning?” But even if we ask this question, it is often not obvious how it is to be answered within the normal structure of a set of dictionary definitions.

Is there then no hope? Is natural language terminally intractable? Probably not. Human beings seem to manage all right. Language is certainly vague and variable, but it is vague and variable in principled ways, which are at present imperfectly understood. Let us take comfort, procedurally, from Anna Wierzbicka (1985):

An adequate definition of a vague concept must aim not at precision but at vagueness: it must aim at precisely that level of vagueness which characterizes the concept itself.

This takes us back to Wittgenstein’s account of the meaning of *game*. This has been influential, and versions of it are applied quite widely, with semantic components identified as possible rather than necessary contributors to the meaning of texts. Wittgenstein, it will be remembered, wrote (*Philosophical Investigations* 66, 1953):

Consider for example the proceedings that we call ‘games’. I mean board games, card games, ball games, Olympic games, and so on. What is common to them all? Don’t say, “There must be something common, or they would not be called ‘games’ ” – but *look and see* whether there is anything common to all. For if you look at them you will not see something common to all, but similarities, relationships, and a whole series of them at that. To repeat: don’t think, but look! Look for example at board games, with their multifarious relationships. Now pass to card games; here you find many correspondences with the first group, but many common features drop out, and others appear. When we pass next to ball games, much that is common is retained, but much is lost. Are they all ‘amusing’? Compare chess with noughts and crosses. Or

is there always winning and losing, or competition between players? Think of patience. In ball games there is winning and losing; but when a child throws his ball at the wall and catches it again, this feature has disappeared. Look at the parts played by skill and luck; and at the difference between skill in chess and skill in tennis. Think now of games like ring-a-ring-a-roses; here is the element of amusement, but how many other characteristic features have disappeared! And we can go through the many, many other groups of games in the same way; can see how similarities crop up and disappear. And the result of this examination is: we see a complicated network of similarities overlapping and criss-crossing; sometimes overall similarities, sometimes similarities of detail.

It seems, then, that there are no necessary conditions for being a bank, any more than there are for being a game. Taking this Wittgensteinian approach, a lexicon for machine use would start by identifying the semantic components of **bank** as separate, combinable, exploitable entities. This turns out to reduce the number of separate dictionary senses dramatically. The meaning of **bank**₁ might then be expressed as:

- IS AN INSTITUTION
- IS A LARGE BUILDING
- FOR STORAGE
- FOR SAFEKEEPING
- OF FINANCE/MONEY
- CARRIES OUT TRANSACTIONS
- CONSISTS OF A STAFF OF PEOPLE

And **bank**₂ as:

- IS LAND
- IS SLOPING
- IS LONG
- IS ELEVATED
- SITUATED BESIDE WATER

On any occasion when the word 'bank' is used by a speaker or writer, he or she invokes at least one of these components, usually a combination of them, but no one of them is a necessary condition for something being a 'bank' in either or any of its senses. Are any of the components of **bank**₂ necessary?

"IS LAND"? But think of a bank of snow.

"IS SLOPING"? But think of a reed bed forming a river bank.

"IS LONG"? But think of the bank around a pond or small lake.

"IS ELEVATED"? But think of the banks of rivers in East Anglia, where the difference between the water level and the land may be almost imperceptible.

"SITUATED BESIDE WATER"? But think of a grassy bank beside a road or above a hill farm.

4. Peaceful Coexistence

These components, then, are probabilistic and prototypical. The word “typically” should be understood before each of them. They do not have to be mutually compatible. The notion of something being at one and the same time an “(ABSTRACT) INSTITUTION and (PHYSICAL) LARGE BUILDING”, for example, may be incoherent, but that only means that there two components are not activated simultaneously. They can still coexist peacefully as part of the word’s meaning potential. By taking different combinations of components and showing how they combine, we can account economically and satisfactorily for the meaning in a remarkably large number of natural, ordinary uses.

This probabilistic componential approach also allows for vagueness.

5. Adam sat on the bank among the bulrushes.

Is the component “IS SLOPING” present or absent in 5? The question is irrelevant: the component is potentially present, but not active. But it is possible to imagine continuations in which it suddenly becomes very active and highly relevant, for example if Adam slips *down* the bank and into the water.

If our analytic pump is primed with a set of probabilistic components of this kind, other procedures can be invoked. For example, semantic inheritances can be drawn from superordinates (“IS A BUILDING” implies “HAS A DOORWAY” (cf.1); “IS AN INSTITUTION” implies “IS COGNITIVE”(cf.6)).

6. The bank defended the terms of the agreement.

What’s the downside? Well, it’s not always clear which components are activated by which contexts. Against this: if it’s not clear to a human being, then it can’t be clear to a computer. Whereas if it’s clear to a human being, then it is probably worth trying to state the criteria explicitly and compute over them. A new kind of phraseological dictionary is called for, showing how different aspects of word meaning are activated in different contexts, and what those contexts are, taking account of vagueness and variability in a precise way. See Hanks (1994) for suggestions about the form that such a phraseological dictionary might take.

A corpus-analytic procedure for counting how many times each feature is activated in a collection of texts has considerable predictive power. After examining even quite a modest number of corpus lines, we naturally begin to form hypotheses about the relative importance of the various semantic components to the normal uses of the word, and how they normally combine. In this way, a default interpretation can be calculated for each word, along with a range of possible variations.

5. Events and Traces

What, then, is a word meaning?

In the everyday use of language, meanings are events, not entities. Do meanings also exist outside the transactional contexts in which they are used? It is a convenient shorthand to talk about “the meanings of words in a dictionary”, but

strictly speaking these are not meanings at all. Rather, they are ‘meaning potentials’ – potential contributions to the meanings of texts and conversations in which the words are used, and activated by the speaker who uses them.

We cannot study word meanings directly through a corpus any more satisfactorily than we can study them through a dictionary. Both are tools, which may have a lot to contribute, but they get us only so far. Corpora consist of texts, which consist of traces of linguistic behaviour. What a corpus gives us is the opportunity to study traces and patterns of linguistic behaviour. There is no direct route from the corpus to the meaning. Corpus linguists sometimes speak as if interpretations spring fully fledged, untouched by human hand, from the corpus. They don’t. The corpus contains traces of meaning events; the dictionary contains lists of meaning potentials. Mapping the one onto the other is a complex task, for which adequate tools and procedures remain to be devised.

The fact that the analytic task is complex, however, does not necessarily imply that the results need to be complex. We may well find that the components of meaning themselves are very simple, and that the complexity lies in establishing just how the different components combine.

6. More Complex Potentials: Verbs

Let us now turn to verbs. Verbs and nouns perform quite different clause roles. There is no reason to assume that the same kind of template is appropriate to both. The difference can be likened to that between male and female components of structures in mechanical engineering. On the one hand, the verbs assign semantic roles to the noun phrases in their environment. On the other hand, nouns (those eager suitors of verbs) have meaning potentials, activated when they fit (more or less well) into the verb frames. Together, they make human interaction possible. One of their functions, though not the only one, is to form propositions.

Propositions, not words, have entailments. But words can be used as convenient storage locations for conventional phraseology and for the entailments or implications that are associated with those bits of phraseology. (Implications are like entailments, but weaker, and they can be probabilistic. An implicature is an act in which a speaker makes or relies on an implication.)

Consider the different implications of these three fragments:

7. the two men who first climbed Mt Everest.
8. He climbed a sycamore tree to get a better view.
9. He climbed a gate into a field.

7 implies that the two men got to the top of Everest. 8 implies, less strongly, that the climber stopped part-way up the sycamore tree. 9 implies that he not only got to the top of the gate, but climbed down the other side. We would be hard put to it to answer the question, “Which particular word contributes this particular implicature?” Text meanings arise from combinations, not from any one word

individually. Moreover, these are default interpretations, not necessary conditions. So although 7' may sound slightly strange, it is not an out-and-out contradiction.

7' *They climbed Mount Everest but did not get to the top.

Meaning potentials are not only fuzzy, they are also hierarchically arranged, in a series of defaults. Each default interpretation is associated with a hierarchy of phraseological norms. Thus, the default interpretation of *climb* is composed of two components: CLAMBER and UP (see Fillmore 1982) – but in 10, 11 and 12 the syntax favours one component over the other. Use of *climb* with an adverbial of direction activates the CLAMBER component, but not the UP component.

10. I climbed into the back seat.

11. Officers climbed in through an open window.

12. A teacher came after me but I climbed through a hedge and sat tight for an hour or so.

This leads to a rather interesting twist: 13 takes a semantic component, UP, out of the meaning potential of *climb* and activates it explicitly. This is not mere redundancy: the word 'up' is overtly stated precisely because the UP component of *climb* is not normally activated in this syntactic context.

13. After breakfast we climbed up through a steep canyon.

7. Semantic Indeterminacy and Remote Clues

Let us now look at some examples where the meaning cannot be determined from the phraseology of the immediate context. These must be distinguished from errors and other unclassifiables. The examples are taken from a corpus-based study of *check*.

Check is a word of considerable syntactic complexity. Disregarding (for current purposes) an adjectival homograph denoting a kind of pattern (*a check shirt*), and turning off many other noises, we can zero in on the transitive verb *check*. This has two main sense components: INSPECT and CAUSE TO PAUSE/SLOW DOWN.

Surely, as a transitive verb, *check* cannot mean both 'inspect' and 'cause to pause or slow down' at the same time? 14 and 15 are obviously quite different meanings.

14. It is not possible to check the accuracy of the figures.

15. The DPK said that Kurdish guerrillas had checked the advance of government troops north of Sulaimaniya.

But then we come to sentences such as 16–18.

16. Then the boat began to slow down. She saw that the man who owned it was hanging on to the side and checking it each time it swung.

Was the man inspecting it or was he stopping it? What is 'it'? The boat or something else? The difficulty is only resolved by looking back through the story leading up to this sentence – looking back in fact, to the first mention of 'boat' (16').

16' "Work it out for yourself," she said, and then turned and ran. She heard him call after her and got into one of the swing boats with a pale, freckled little boy...

Not it is clear that the boat in this story has nothing to do with vessels on water; it is a swinging ride at a fairground. The man, it turns out, is trying to cause it to slow down ('checking' it) because of a frightened child. This is a case where the relevant contextual clues are not in the immediate context. If we pay proper attention to textual cohesion, we are less likely to perceive ambiguity where there is none.

17. The Parliamentary Assembly and the Economic and Social Committee were primarily or wholly advisory in nature, with very little checking power.

In 17, the meaning is perfectly clear: the bodies mentioned had very little power to INSPECT and CAUSE TO PAUSE. Perhaps an expert on European bureaucracy might be able to say whether one component or the other of **check** was more activated, but the ordinary reader cannot be expected to make this choice, and the wider context is no help. The two senses of **check**, apparently in competition, here coexist in a single use, as indeed they do in the cliché *checks and balances*. By relying too heavily on examples such as 14 and 15, dictionaries have set up a false dichotomy.

18. Corporals checked kitbags and wooden crates and boxes...

What were the corporals doing? It sounds as if they were inspecting something. But as we read on, the picture changes.

18'. Sergeants rapped out indecipherable commands, corporals checked kitbags and wooden crates and boxes into the luggage vans.

The word *into* activates a preference for a different component of the meaning potential of **check**, identifiable loosely as CONSIGN, and associated with the cognitive prototype outlined in 19.

19. PERSON check BAGGAGE into TRANSPORT

No doubt INSPECT is present too, but the full sentence activates an image of corporals with checklists. Which is more or less where we came in.

8. Where Computational Analysis Runs Out

Finally, consider the following citation:

20. He soon returned to the Western Desert, where, between May and September, he was involved in desperate rearguard actions – the battle of Gazala, followed by Alamein in July, when Auchinleck checked Rommel, who was then within striking distance of Alexandria.

Without encyclopedic world knowledge, the fragment ...*Alamein in July, when Auchinleck checked Rommel* is profoundly ambiguous. I tried it on some English teenagers, and they were baffled. How do we know that Auchinleck was not checking Rommel for fleas or for contraband goods? Common sense may tell us

that this is unlikely, but what textual clues are there to support common sense? Where does the assignment of meaning come from?

- From internal text evidence, in particular the collocates? Relevant are the rather distant collocates *battle*, *rearguard actions*, and perhaps *striking distance*. These hardly seem close enough to be conclusive, and it is easy enough to construct a counterexample in the context of the same collocates (e.g. **before the battle, Auchinleck checked the deployment of his infantry*).
- From the domain? If this citation were from a military history textbook, that might be a helpful clue. Unfortunately, the extract actually comes from an obituary in the *Daily Telegraph*, which the BNC very sensibly does not attempt to subclassify. But anyway, domain is only a weak clue. Lesk [1986] observed that the sort of texts which talk about pine cones rarely also talk about ice-cream cones, but in this case domain classification is unlikely to produce the desired result, since military texts do talk about both *checking equipment* and *checking the enemy's advance*.
- From real-world knowledge? Auchinleck and Rommel were generals on opposing sides; the name of a general may be used metonymically for the army that he commands, and real-world knowledge tells us that armies *check* each other in the sense of halting an advance. This is probably close to psychological reality, but if it is all we have to go on, the difficulties of computing real-world knowledge satisfactorily start to seem insuperable.
- By assigning *Auchinleck* and *Rommel* to the lexical set [GENERAL]? This is similarly promising, but it relies on the existence of a metonymic exploitation rule of the following form:

[GENERAL_i] checked [GENERAL_j] = [GENERAL_i]'s army checked
(= halted the advance of) [GENERAL_j]'s army.

We are left with the uncomfortable conclusion that what seems perfectly obvious to a human being is deeply ambiguous to the more literal-minded computer, and that there is no easy way of resolving the ambiguity.

9. Conclusion

Do word meanings exist? The answer proposed in this discussion is “Yes, but . . .”

Yes, word meanings do exist, but traditional descriptions are misleading. Outside the context of a meaning event, in which there is participation of utterer and audience, words have meaning potentials, rather than just meaning. The meaning potential of each word is made up of a number of components, which may be activated cognitively by other words in the context in which it is used. These cognitive components are linked in a network which provides the whole semantic base of the language, with enormous dynamic potential for saying new things and relating the unknown to the known.

The target of ‘disambiguation’ presupposes competition among different components or sets of components. And sometimes this is true. But we also find

that the different components coexist in a single use, and that different uses activate a kaleidoscope of different combinations of components. So rather than asking questions about disambiguation and sense discrimination (“Which sense does this word have in this text?”), a better sort of question would be “What is the unique contribution of this word to the meaning of this text?”

A word’s unique contribution is some combination of the components that make up its meaning potential, activated by contextual triggers. Components that are not triggered do not even enter the lists in the hypothetical disambiguation tournament. They do not even get started, because the context has already set a semantic frame into which only certain components will fit.

A major future task for computational lexicography will be to identify meaning components, the ways in which they combine, relations with the meaning components of semantically related words, and the phraseological circumstances in which they are activated.

The difficulty of identifying meaning components, plotting their hierarchies and relationships, and identifying the conditions under which they are activated should not blind us to the possibility that they may at heart be quite simple structures: much simpler, in fact, than anything found in a standard dictionary. But different.

References

- Church, K.W. and P. Hanks. “Word Association Norms, Mutual Information, and Lexicography”, in *Computational Linguistics* 16:1, 1990.
- Fillmore, C.J. “An alternative to checklist theories of meaning” in *Papers from the First Annual Meeting of the Berkely Linguistics Society*, 1975, pp. 123–132.
- Fillmore, C.J. “Towards a Descriptive Framework for Spatial Deixis” in *Speech, Place, and Action*, R.J. Jarvella and W. Klein Eds. New York: John Wiley and Sons, 1982.
- Hanks, P. “Linguistic Norms and Pragmatic Exploitations, Or Why Lexicographers need Prototype Theory, and Vice Versa” in *Papers in Computational Lexicography: Complex '94*, Eds. F. Kiefer, G. Kiss, and J. Pajzs, Budapest: Research Institute for Linguistics, 1994.
- Pustejovsky, J. *The Generative Lexicon*. Cambridge MS: MIT Press, 1995.
- Wierzbicka, A. *Lexicography and Conceptual Analysis*. Ann Arbor MI: Karoma, 1985.
- Wierzbicka, A. *English Speech Act Verbs: A Semantic Dictionary*. Sydney: Academic Press, 1987.
- Wittgenstein, L. *Philosophical Investigations*. Oxford: Basil Blackwell, 1953.

