

Predicate Indexicality and Context Dependence

Peter Bosch
Institute of Cognitive Science
University of Osnabrück, D-49069 Osnabrück, Germany
pbosch@uos.de

Abstract This paper argues that productive language use, specifically cases of apparent variation in word sense, or so-called polysemy, require a treatment at the denotational level rather than a purely lexical-semantic solution or a sentence-internal compositional solution. It is argued that the lexical semantics is much less specified than is often assumed and only contains structural constraints over the kind of conceptual entities that can be denoted by the lexical item, but does not contain the conceptual content. Among these constraints we are proposing a distinction between those that introduce index dependence and others that make the interpretation dependent on argument domains. Only the latter seem to be involved in polysemy.

1 The Problem

Modern mainstream linguistics may be characterized as an investigation not of languages, language, or language behaviour, but as an investigation of *linguistic knowledge*¹. This characterization of the subject matter of the discipline is one of the cornerstones of the Chomskyan research programme and has been responsible for much of the success of modern linguistic research.

The focus on linguistic knowledge distinguishes this line of work from preceding and also from other current research paradigms. It has led, under the influence of modern logic and the study of formal systems, for the first time in the history of linguistics, to rigorous formal characterizations of some of the structural properties of human language, particularly in the areas of syntax and semantics.

The central concept on which both syntax and semantics are built in the linguistic knowledge paradigm is the *sentence*, both as the basic notion of grammar and – in its guise as proposition or sentence meaning – as the basic notion of formal semantics. Sentences are generated according to syntactic rules from syntactic constituents; and propositions, alias sentence meanings or truth conditions, are built up compositionally from the meanings of their constituents, ultimately the lexical items, of the sentence. This is the core of the generative recipe – attributable in principle already to Gottlob Frege – that provides the principle for an explanation of the fact that arbitrarily many different natural language expressions, phrases and sentences, can be produced and understood by any human speaker: All you need is the primitive expressions, i.e., the lexical items, and the rules for their combination.

As the purpose of syntactic rules is to determine how constituents, and ultimately sentences, can be built from lexical items, a large part of this knowledge can straightforwardly be accommodated in the syntactic categories of lexical items themselves. In principle, if syntactic categories are made sufficiently complex, they may incorporate all of syntax in the form of syntactic category information and syntax can thus be "lexicalized", i.e., can be made part of the

¹ I am grateful for the opportunity to present an earlier version of this paper at the conference on Utterance Interpretation and Cognitive Models at the Université Libre de Bruxelles, 23-24. 06. 2006 and for comments and criticism from participants. In addition I want to thank Graham Katz for very helpful discussions, and Philippe De Brabanter and Mikhail Kissine for many very useful comments on an earlier draft.

lexical information that becomes available with each lexical item. This path has been taken in quite a number of current approaches to grammar at least to some degree.

One would expect that also *semantic information* could similarly be represented for each word in the lexicon. Putting aside for the moment the question of what such representations may look like and sticking to the basic idea that they must contain all the knowledge that the lexical item contributes to the meaning of the sentence, there is an obvious difficulty with lexical items that have more than one meaning. This is not a very interesting problem as long as we are concerned with a fairly small number of meanings per word, as in the case of ordinary lexical ambiguity, most typically homonymy. Lexical ambiguity is usually accommodated for by allowing for more lexical entries. This clearly makes sense when there are not only different meanings, but when there is also the intuition that we are dealing with different words that just happen to look or sound the same.

A clear indication would be additional morphological or syntactic differences: The expressions *to book a flight* and *a book of fairy tales* are reasonably assumed to contain two different lexical items, both pronounced alike and both spelt "book" but not otherwise related in any interesting sense, albeit in some etymological way. When we are concerned with the native speaker's linguistic knowledge, however, etymology is irrelevant anyhow.

Even when there is no difference in terms of morphological or syntactic properties, one would still occasionally assume two different lexical items. E.g. in the case of the two different meanings for the noun *caterpillar*, designating the animal or the vehicle, even though it may still be transparent to most speakers that one is derived metaphorically from the other.

The rule of thumb simply is that one would assume ambiguity and separate lexical entries, i.e., homonymy, whenever there is no obvious rule that would *productively derive one meaning from the other* – never mind actual historical origins. A speaker of English may learn and use either of the two homonyms without knowing the other, and no-one would expect the native speaker to use, or even understand, the form in one sense merely on the basis of already knowing the other.

A problem case is when we find uses of lexical items that don't differ much, or do not differ at all, in their syntax and somehow seem obviously related in their meaning, but not to the point of being semantically identical. The English verb *to open* is a stock-in-trade example²: In phrases like *opening a door*, *opening a letter*, *opening a bottle*, or *opening a book* the verb *open* seems to denote different types of actions each time. If in order to open a door you were to do the same thing to the door that you would do to a letter in order to open the letter, you would not only fail to open the door, but your action is also unlikely to be described as even an attempt to open the door.

Still, there is a clear intuition that we are looking at a productive mechanism in this case: Knowing one of these uses of *open* seems like a good start for understanding and using all the others, and there seems no obvious limit to the number of different types of object that a speaker of English could "open" while the kind of opening operation denoted would still be different each time³

This is where the problem starts: If, as it may appear, the same word may mean something different every time it combines with other words, how are we to systematically predict the meaning of the resulting phrase for unseen cases, and how are we to compute the meaning or truth conditions of sentences in which these words occur?

My argument in this paper is that polysemy has very little to do with lexical semantics and compositionality. Polysemy is neither in the lexicon, nor does it arise from an interaction of different words in a sentence. Instead, I shall argue that polysemy effects come about when

² Cf. earlier treatments in (Bosch 1995, 2007b) that already contain some of the thoughts also defended in this paper.

³ Perhaps one may still want to exclude some uses of *open* as non-productive or idiomatic, like *opening your arms to greet somebody* or *opening your heart to somebody*. I have no specific opinion on this, but even with a reasonable finite number of such idiosyncratic cases that may not be derived productively there is no doubt that there is still a very large class of uses that are productively related.

lexical information interacts with the denotation that an occurrence of a particular expression has in a particular context. Such interactions yield context-specific denotations for expressions, and it is these denotations that are the input for compositional processes, which eventually will give us something like the "meaning" of an utterance. This model is clearly different from the more generally accepted idea, according to which, first of all, abstract sentence meanings are constructed compositionally and are subsequently adjusted to the context in which the sentence occurs.

2 Accounting for polysemy

In Section 2.1 I want to briefly sketch one practically useful but theoretically unsatisfactory approach to polysemy that is commonly taken in lexicography, and in Section 2.2 I will give a summary of James Pustejovsky's Generative Lexicon theory, which is in principle a formal and compositional theory and overcomes essential shortcomings of many other approaches. As the discussion of Pustejovsky's account will show, I believe that it still falls short of satisfying certain architectural or modularity requirements that are important for a good understanding of polysemy in particular and natural language comprehension more generally.

2.1 World knowledge and the listing of senses

How do lexicographers handle polysemy? Looking up the verb *to open* in the *Collins Cobuild English Dictionary* (1990 edition) we find, next to – literally – thirty-seven other "senses" for this verb, the entries in (1) and (2), which are fairly representative of the rest.

- (1) If you open a letter or parcel, you cut or tear the envelope, or remove the wrapping.
- (2) If you open a book, you move its covers apart to read or write on the pages inside.

Somehow this reads like instructions for how to open letters, parcels, or books, rather than descriptions of word meanings. But even though this way of explaining word meanings is not the line a theoretical linguist would embrace, it is probably on the right track for practical lexicography, because such lexical entries are clearly useful for the non-native dictionary user: They attempt a classification of different uses of a word, exemplifying each, and thus provide useful albeit indirect hints at the word's meaning.

But why is the Collins approach not attractive for the theoretical linguist? I can see two reasons. One is that no distinction is made between the description of *what a word means* and *the objects or events the word is applied to*, i.o.w., between linguistic knowledge on the one hand and world knowledge or conceptual knowledge on the other. The Collins word senses of *open* just seem to represent a taxonomy of those events that are all described in English as *openings*. – Now this point may not be seen as critical by all linguists. Some linguists, like many psychologists, believe that word meanings just *are* concepts.

The other reason for being less than happy with the Collins approach is that it just seems to produce lists of different word senses and thus misses generalizations. By merely listing different uses of a word one does not even get round to asking the question of how these different uses arise and if some or perhaps all of them may be the result of some productive or generative mechanism.

Following up the second line of criticism one may like to view the apparent differences in the sense of *open* as the result of something like a compositional process: the "sense" that the verb takes on in a particular phrase would differ depending on the type of argument that the verb takes in that phrase, while there is in fact only one lexical sense and just one lexical entry for the verb *to open* from which these various compositional senses arise.

This is, roughly, an approach that James Pustejovsky has been proposing. It has the advantage that it gets away from a mere listing of word senses and can model polysemy phenomena as the result of a productive process. On the other hand, as we shall see, Pustejovsky includes a good deal of conceptual as well as contingent world knowledge into his lexical semantics and would probably not endorse the first of our two critical remarks about the lexicographer's approach.

2.2 Polysemy resulting from compositional processes

2.2.1 The Generative Lexicon approach

Pustejovsky (1995:221) discusses the three sentences in (3) to illustrate the apparent polysemy of the English verb *to open*.

- (3) a. Mary opened the letter for her mother.
 b. The rangers have opened the trail for the summer.
 c. John opened the door for the guests.

The lexical entry, or rather just the semantic part of it, that Pustejovsky (1995:221) suggests for the English verb *to open* is given by the feature structure in Diagram 1.

Here we find three attributes: the Event Structure (EVENTSTR), the Argument Structure (ARGSTR), and the Qualia Structure (QUALIA) of the verb. EVENTSTR tells us that the verb semantically relates two events, one of which is of the semantic type **process**, the other of the type **state**, and that the **process** must precede the **state**. ARGSTR gives the semantic arguments of the verb, coindexing the first argument with an agent argument, and the second with the result argument in the QUALIA. This second argument is furthermore marked as being of the semantic type **physobj** and its QUALIA are given, in abbreviated form, by the feature FORMAL, whose value is the semantic type **entity**. The QUALIA features FORMAL and AGENTIVE are respectively used to distinguish the verb's meaning in a larger domain or for information on the origin or bringing about of an opening. This QUALIA representation is not specific to the verb *to open*, but is a semantic type. Pustejovsky (1995:76) calls it the "direct causative lexical conceptual paradigm" (**dc-lcp**) that represents the QUALIA structure of all direct causative verbs. I could not honestly say I have fully understood these QUALIA, but in the case at hand the information they are supposed to give is that a relation of the type **open_result** relates ARG2 to the **state** e_2 and a relation of the type **open_act** relates the process e_1 to both ARG1 and ARG2. In plain English this comes to something like this: Opening something results in a particular state of the thing that is being opened, and the opening is a process that involves an agent and a physical object.

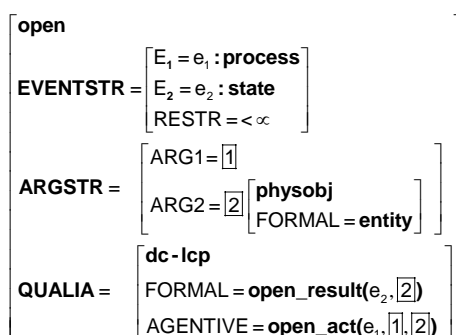


Diagram 1

The semantics of the verb thus represented is obviously quite neutral with respect to the three senses in (3). ARG1 would respectively take on Mary, the rangers, or John as its value, and ARG2 would be the letter, the trail, or the door; the process, e_1 , would presumably in all cases be a **process** of opening; only the **state**, e_2 , would differ: the way these sentences are typically read – in any case the way that Pustejovsky wants us to read them – the letter in (a) would not only be open after Mary opened it, but would be available for reading, the trail in (b) would not be open for being read, but would be available to walkers to walk on, and the door in (c) would be open for the guests to walk through.

This comes to about the same level of detail as the explanations in the Collins dictionary and one may wonder if this much detail should be part of a semantic interpretation, and if these variants in the interpretation of *open* would indeed enter the truth conditions of all utterances of the sentences in (3 a.-c.) We will return to this point below.

Pustejovsky in any case believes that these differences between the typical or likely interpretations of the sentences in (3) ought to be accounted for in the compositional semantics, and attempts to derive them from an interaction between the lexical semantics of *open* in Diagram 1 and, for instance, the lexical semantics of the noun *letter* in Diagram 2 (Pustejovsky 1995:222).




Diagram 2

The noun *letter* thus behave semantically as what Pustejovsky calls a "dot object": something that belongs simultaneously to more than one semantic type, in the case at hand the semantic type **physobj** and the semantic type **info**. The dot object is again a lexical conceptual paradigm, and its TELIC role says that the thing is for reading.

Now, in order to see the relation between *opening* and *letter*, as it is presumably intended in (3 a.), we have to look more closely into the arguments of the **open_result** relation in the lexical entry of *open*. Diagram 3 gives a more detailed representation of the corresponding parts of the feature structure in Diagram 1.

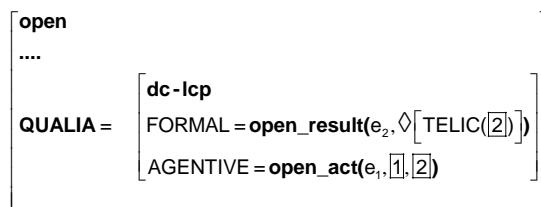


Diagram 3

The TELIC quale from the representation of the **state** that is the result of the opening act is now included under a possibility operator as the second argument of **open_result** and thus tells us what the purpose of this state is: to make possible the purpose given in the TELIC quale of the second argument of *open*, i.e., the object opened. This would then, by unification of the feature structures (Pustejovsky speaks in this case of "co-composition"), lead to a semantic representation for the phrase "open the letter" as in Diagram 4, which tells us that the point of the letter-opening operation is to make the reading of the letter possible (Pustejovsky 1995:223).

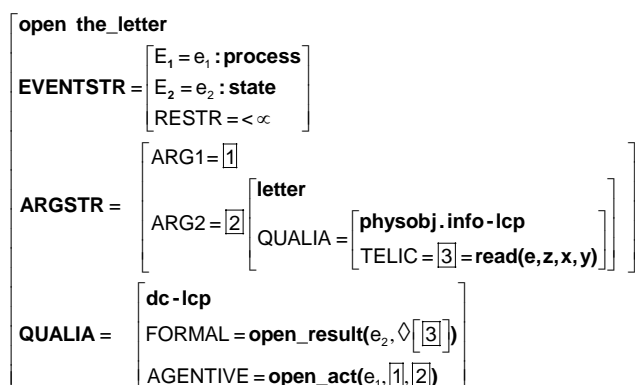


Diagram 4

Assuming a lexical representation for the semantics of *door* as in Diagram 5 we can then, quite analogously, derive the semantic representation in Diagram 6 for sentence (3 c.).

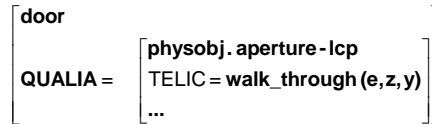


Diagram 5

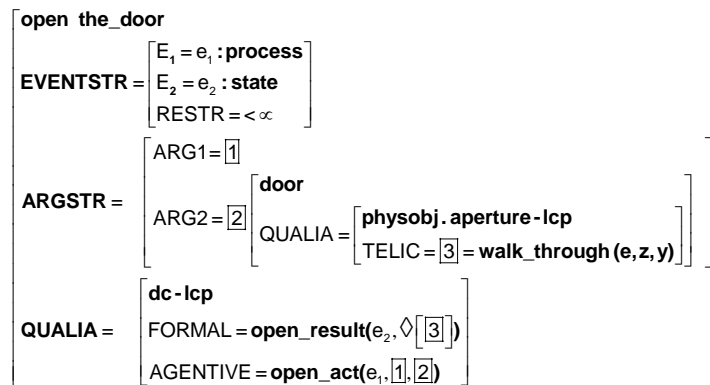


Diagram 6

There is a lot of detail and background missing in this rendering of Pustejovsky's proposal, but for current purposes at least the outline of the account should be clear enough.

2.2.2 Problems with the Generative Lexicon

Pustejovsky does not list different word senses but accounts for polysemy by compositional effects that come about in the interaction of the verb *open* with its noun arguments. What is the nature of these compositional effects?

First of all it may be surprising, at least from a linguistic point of view, how much world knowledge or, if you like, conceptual knowledge is found in Pustejovsky's lexical representations. Does it really matter to the understanding of the phrase "open a letter" that letters should have something like an intrinsic purpose of being read?

Suppose that it does. What then, if the discourse context gives us explicitly different information? For the postman the purpose of a letter is rather to be delivered, and for the filing clerk perhaps letters must, in the first instance, be sorted and filed. Similarly, not nearly all doors are for walking through; think of cupboard doors, wardrobe doors, etc. But they are still doors. And a carpenter may sometimes have to "open a door", but not as an object of the type **physobj.aperture-lcp**, but rather as a simple **physobj**, because the lock inside the door needs repair. Walking through the door is not at issue then either. – In sum, one may wonder about how much of all this knowledge should be part of lexical knowledge and how much of it is just contingent common sense knowledge that, first of all, is not really linguistic in nature and, second, may change from one discourse context to the other.

For instance, one may ask if the plausible inference from (3 a.), that the purpose of opening the letter is to make it possible that it be read, is in any sense part of the *semantics of the sentence*, i.e., part of the truth conditions that any declarative utterance of the sentence would have to respect. I find it hard to believe that this should be so.

If this inference were part of every proposition that the sentence denotes the sequences in (3 a'.) should yield a contradiction, which they patently do not.

- (3) a'. Mary opened the letter for her mother. The point was just to show her that it was empty.
- a''. Mary opened the letter for her mother. The point was to take out the cheque and then throw it away⁴.

The knowledge about purposes that is at issue here may well be part of the understanding of *particular occurrences* of the sentence *in particular contexts*. But Pustejovsky models it as part of *the semantics of the phrase or sentence* and as derived by compositional processes. If we want to include this kind of knowledge into the model of the comprehension process – which would certainly make sense – then it should more correctly be treated as defeasible pragmatic knowledge, like Gricean implicatures, rather than as a proper part of the semantics.

It should be uncontroversial that world knowledge and conceptual knowledge of all kinds plays an essential role in natural language understanding, and certainly also for polysemy. The issue is not *whether*, but *how* to take this knowledge into account. Modelling it as part of linguistic knowledge strikes me as simply wrong.

Pustejovsky's compositional processes are strictly sentence-internal and do not interact with discourse or utterance context. Influence from discourse or utterance context can only be taken into account *after* the system has delivered a sentence meaning or a semantic value for the sentence. However, by including part of the knowledge that may in fact vary from one discourse or utterance context to the other in the lexical entries, the Generative Lexicon in fact models language comprehension by implicit reference to a *default context*, and any adjustment to contexts where the defaults do not hold would have to take place after the system has completed operations.

This type of architecture for language understanding, which assumes, first, the construction of sentence meanings or propositions, and then, as a further step, an adjustment to contexts, is of course a fairly common position, even though it is rarely made explicit (but cf. e.g. Cutler & Clifton 1999) and tends to be more of a common background assumption. Also the classic division of labour between semantics and Gricean pragmatics (challenged though in more recent Neo-Gricean thinking) naturally fits in with this architecture, where conversational implicatures can only be derived from propositions; so the propositional result of compositional processes must be completed first, before the pragmatic component can start operating.

3 A denotational approach to comprehension

3.1 Utterances and sentences

The problems about lexical meaning and polysemy so far discussed have arisen and have been formulated in the context of the compositional semantics of sentences. Sentences are objects of linguistic knowledge, certainly in the sense that native speakers of a language can by and large recognize sentences of their language as such. And possibly, although this is a point where we had better be careful, speakers can also decide – within certain limits – what particular sentences could mean, how to paraphrase them, or what other sentences would be semantically compatible.

We must be careful here, because sentences are constructs of linguistic theory and these constructs abstract from many factors that influence meaning. Thus it is only in a derived way that we can say that a *sentence* is true or false with respect to a particular state of affairs. More correctly one would have to say, rather, that a particular *statement made by uttering a sentence on a particular occasion*, or, if we want to ignore speech-act aspects, a particular *occurrence of a sentence*, is true or false with respect to a particular state of affairs. It is also only in a derived way of speaking that we could say that sentences are paraphrases of other sentences. It is not

⁴ There is an obvious additional problem here with the polysemy of "letter", which I am trying to ignore because it does not affect our argument.

sentences or expressions, but utterances or occurrences of sentences and expressions, that figure in linguistic behaviour and thus carry meaning, and only *utterances or occurrences of expressions* can refer to other things in the world. Linguistic expressions as such cannot. If we are interested in how linguistic knowledge relates to linguistic behaviour or how linguistic knowledge relates to other parts of cognition, or simply in how somebody in particular understands an expression that is uttered in a particular situation, we are not concerned with expressions but with *utterances or occurrences of linguistic expressions*.

How is all of this relevant to questions of apparent variation in lexical denotation, alias polysemy? Polysemy obviously does not show up as long as you look at an isolated word, but it shows up when you compare different occurrences of the word in different contexts. In computational linguistics, or at least in many of its applications, the context of a word is operationalized as a window of a certain number of words that occur to the left and to the right of the target word. And for many purposes of *Word Sense Disambiguation* this simple notion of context is quite sufficient. A linguist of course would complain, correctly, that even if this approach to disambiguation works fine, there is no good linguistic reason why it should. There is no linguistic reason why a word should take on a different sense just because of its neighbourhood – *pace* Firth, who had probably quite different things in mind when he produced his *bon mot* "You shall know a word by the company it keeps" (Firth 1957:11).

This complaint loses its force when context is no longer viewed flatly as a bag of words in the target word's neighbourhood but takes syntactic structure into account. If, instead of mere proximity relations, occurrence within the same phrase or sentence, and structural relations like argument or modifier status enter the picture, then linguistic selection mechanisms come into play and perhaps one may think of explanations of polysemy in terms of compositional semantic mechanisms. Approaches to polysemy in theoretical computational linguistics, like Pustejovsky's (1995) or Briscoe and Copestake's (1999), can be understood along these lines.

In these accounts problems of polysemy are explained as resulting from compositional semantic processes. Examples like (4) do not seem to require any information from a context beyond the sentence or the phrase to illustrate what they are intended to illustrate, i.e., that the adjective *fast* makes a different contribution to the semantics of its phrase, depending on the head of the phrase.

- (4) a. Fred had a fast run.
b. Fred had a fast car.
c. Fred had a fast typist.

This makes it look as if in the course of the compositional construction of the meanings of phrases or sentences, word meanings were interacting with the meanings of other words, ultimately yielding the "meaning of the sentence" or proposition.

3.2 Denotation matters

But this picture is misleading. It may give a roughly correct account for sentences that approximate the ideal of sentences like "Snow is white" or "Ice floats on water", where no specific reference is involved and discourse dependence is negligible. But sentences that figure in human communication are *ordinarily context-dependent* in the sense that they make only partially explicit what is being said. Truth conditions for occurrences of sentences in contexts will thus ordinarily have to take discourse as well as non-linguistic contextual information into account.

The information that a *sentence* makes available, i.e., word meanings and syntactic structure, only *constrains*, but does not determine, truth conditions. And similarly the lexical semantics only *constrains* the denotation of words. There is indeed good reason to suppose that the contribution that the words in a sentence make to the truth conditions of occurrences of the sentence is not a direct contribution of their lexical semantics, but rather a contribution of the denotation that the words have in the actual utterance context.

This notion is supported, for instance, by cases where an adjective combines not with a lexical noun, but with a pronoun or a proper name. Pronouns or proper names do not have lexical entries that could establish a relevant semantic difference. Consider the sequences in (5).

- (5) a. What did you think about Fred's last run? Yeah, that was a fast one.
- b. What did you think about Fred's last car? Yeah, that was a fast one.
- c. What did you think about Fred's last typist? Yeah, that was a fast one.

Here we get different truth conditions for each occurrence of the identical second sentence in (a)-(c), but this difference cannot be explained by any interaction between the adjective *fast* and its head. There just is no difference. Instead we must assume that the effect takes place at the level of denotations: *one* has a different *denotation* in each of these cases, brought about via its anaphoric link to the interpretation of the preceding sentence.

Similarly, the predication of *fast* in (6) remains vastly underspecified as long as we lack information about the denotation of *Lizzy*. As soon as we know that Lizzy is, e.g., a car, or a horse, or a woman, the truth conditions for an utterance of this sentence are tightened.

- (6) Lizzy is really fast.

The relevant information here is not denotation in the sense of the *identification of a referent*. For suppose that Lizzy refers to a particular person, your friend Lizzy. Then we still don't know whether (6) is intended to mean that Lizzy is a fast runner, a fast typist, or a fast driver, etc., because neither the lexical representation of the name *Lizzy* nor the identification of the intended referent or any general knowledge about the referent will help with this decision. Even though we know that Lizzy works as a typist, drives a Beetle, and runs Marathons we *still* don't know how to read *fast*. When (6) occurs in the context of a bicycle race in which Lizzy is taking part, it may rather be intended to say something about her speed as a cyclist.

What seems to be playing a role, rather than mere reference, is the conceptual representation of the intended referent for the current occurrence of the expression: *Lizzy as a cyclist, as a typist, etc.* However, such representations typically do not result from sentences, but require information from discourse as well as preceding and concurrent non-linguistic information.

So if the examples in (4) give the impression that the head nouns *run, car, and typist* are decisive for the interpretation of *fast*, then this impression is only superficially correct: These nouns and their lexical semantics may give a (usually pretty good) *hint* at the relevant denotation or conceptual representation of the referent, but they do not determine it.

3.3 Indexicality

Before we proceed we must clarify some terminology and introduce a little more theory about context. A common and conventional view in formal semantics is that context is a set of parameters that, together with other parameters, and independent of sentence meanings, influence the truth of a sentence. This conception, which is close to ideas in Montague (1968/1974) or Lewis (1970), is represented in Diagram 7.

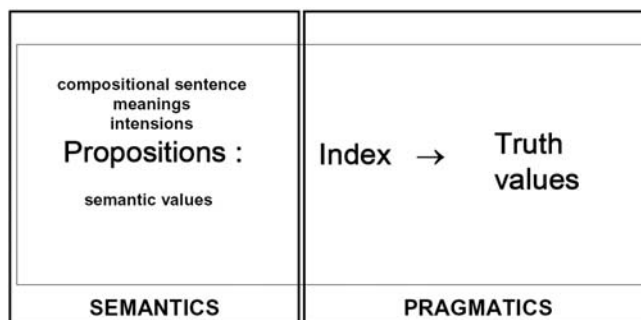


Diagram 7

Here the semantics is concerned with the fixing of semantic values, perhaps propositions, which are extension-determining functions, i.e., intensions in the sense of Carnap (1947), and which are built compositionally from the grammatical constituents of the sentence. The pragmatics would then determine truth values for propositions, given an index. The index is any packet of various parameters on which the truth of the sentence depends and may include contextual features, but will also include what Lewis called "contingency": the dependence on the world at which the sentence is evaluated.

David Kaplan (1978) proposed a somewhat different and currently probably more popular division of labour (cf. Diagram 8) that leaves none of indexicality or context dependence to pragmatics, but only matters that are concerned with speech acts, implicature, or appropriateness. Kaplan's intention was to get some facts about indexical expressions, like *I*, *here*, *now*, right. But there is also the nice side effect that his account yields more intuitive semantic values for sentences. For instance, the sentence "I'm smoking." intuitively has the same content, i.e., *says the same thing*, in my current situation, as the sentence "Peter Bosch is smoking at 10:29 p.m. on Saturday, 12 January, 2008." This common *content* is modelled by Kaplan as the proposition that both sentences denote relative to my current utterance context. In the Montague-Lewis story, however, a sentence containing indexicals, would not determine a proposition in the intuitive sense of "what is said". Here the content of the sentence is constructed purely compositionally, before parameters of the index are taken into account. Our two sentences would thus always have different semantic values, where the value of the first is more variable than the value of the second. – The difference between the two accounts must not be over-emphasized, since any form of indexicality that can be modelled by one of them can also be modelled by the other. But the Kaplan version gives us, as also Lewis (1981) points out, a more intuitive notion of semantic values.

In order to identify the semantic value of a declarative sentence with a proposition or content (in the sense of *what is said* by the sentence in the utterance situation), semantic values must depend on the relevant features of the utterance context. The content then determines truth values for different circumstances of evaluation, in particular depending on what things are like in the world at which the utterance is to be evaluated. – This view is represented in Diagram 8, and here is a short extract from Kaplan(1978:84) that summarises it:

"Just as contents (or intensions) can be represented by functions from possible worlds to extensions, so characters can be represented by functions from contexts to contents. The character of 'I' would then be represented by *the function (or rule, if you prefer) which assigns to each context that content which is represented by the constant function from possible worlds to the agent of the context*. The latter function has been called an 'individual concept'. Note that the character of 'I' is represented by a function from contexts to individual concepts, not from contexts to individuals" [emphasis by Kaplan].

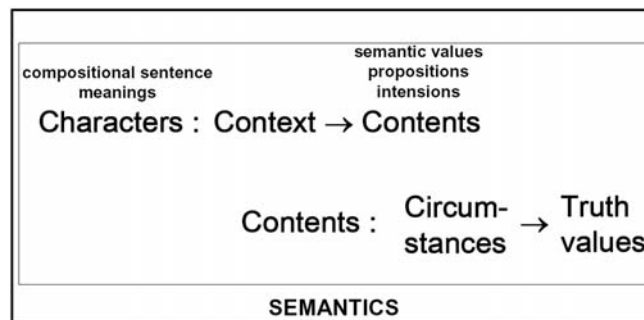


Diagram 8

The context-dependent expressions that Kaplan focuses on are those that he originally called "demonstratives". He mentioned as examples *I*, *you*, *here*, *now*, *that*, and the adjectives *actual* and *present* (Kaplan 1978:81). Later he explained the notion more generally and changed the terminology from "demonstrative" to Peirce's term "indexical": "What is common to the words or usages in which I am interested is that the referent is dependent on the context of use and that

the meaning of the word provides a rule which determines the referent in terms of certain aspects of the context. The term I now favor for these words is 'indexical'." (Kaplan 1989:490).

Non-indexical expressions, on the other hand, have a constant character, while "expressions containing demonstratives will, in general, express different concepts in different contexts. We call the concept expressed in a given context, the *Content* of the expression in that context." (Kaplan 1978:91) An important point to note is that *contents* are not referents or extensions – despite the misleading expression by Kaplan himself – but are concepts, functions, or rules that determine referents or extensions in particular circumstances.

3.4 Predicate indexicality

We can now return to our suggestion from Section 3.2, i.e., that the variation in the adjective interpretation may not be due to the lexical semantics of its nominal head but rather to the *denotation* of the nominal, which is a concept or property.

Having reviewed Kaplan's notion of context dependence, I now want to propose that the class of indexical expressions is actually considerably more comprehensive than Kaplan suggested. Not only may one think of including such additional expressions as Kaplan (1989:489) mentioned himself, namely *my, you, he, his, she, it, this, that, tomorrow, yesterday*, and other plausible additions of the same kind, like *neighbour, enemy, friend, nearby, right, left, on the other side*, etc., but perhaps it makes sense to ask if not *nearly all lexical expressions are in principle open to indexical behaviour*, in the sense that their lexical meaning (character) determines their denotation (content) *only relative to the context of their occurrence*. The contents of saturated expressions (names) would be individual concepts and also the contents of unsaturated expressions (predicates such as verbs, nouns, adjectives, adverbs, but perhaps also generalized quantifiers) would be concepts. In the simplest cases, like complete VPs, they are truth functions (not sets, but characteristic functions), in other cases they are more complicated functions of functions, etc.

Such a broadening of the notion of indexicality may be seen as posing a problem in the sense that it requires a much broader notion of index. Kaplan could make do with an index of just four parameters: *agent, time, position, world*, because it is only these context parameters on which the content of his indexicals depend. If, however, every lexical expression is viewed as potentially context-dependent, then the required index would become potentially very large. But possibly this threat is only apparent. As Montague (1974:98) pointed out (for *his* notion of an index), "It is not necessary to consider [contexts of use] in their full complexity; we may instead confine our attention to those among their features which are relevant for the discourse in question. Thus it will suffice to specify the set of all complexes of relevant aspects of intended possible contexts of use. We may call such complexes *indices*".

Computing the character function we are thus in each case only concerned with a small number of specific parameters. These, we may assume, are lexically given as part of the character of each lexical item, or of classes of lexical items. In order to compute the contextual denotation of the indexical *I*, we need the value of the speaker or *agent* parameter in the current context, and for a gradable relative adjective – to take an example that is probably uncontroversial – the lexical entry would specify parameters for the *standard* or *comparison class* (or possibly both), for which we have to find a value in the context.

Thus there is no danger that the suggested broadening of the notion of index would force us into assuming contexts as proper entities in the theory and into having to provide identity criteria for them. There is no need to assume that contexts are completely parameterized, but we may stick to the more convenient view that any two contexts count as the same, in the computation of a character function, as long as they do not differ with respect to the currently relevant index parameters⁵.

Quite pragmatically, this means that in the computation of the character function of an expression the current context representation is inspected only for those index parameters that are the lexically specified parameters of the expression that is being processed.

⁵ Cf. Bosch (1982) for a discussion of related problems.

3.4.1 An example

The denotation of an occurrence of an adjective, like *fast*, would be a concept. Or, more appropriately, a function from concepts to concepts⁶. But since *fast* is a gradable relative adjective, its denotation should make a different contribution to the truth conditions of the occurrence of a sentence, depending on the speed we are comparing with. This relativity is a lexical semantic property of the adjective and hence a property of its character. The comparative speed parameter, having as its value probably a class of entities we are comparing with, must occur in the lexical entry as an implicit variable.

For our example *fast* the character is represented in the lexical entry, as for instance in Diagram 9 (cf. Pollard & Sag 1994:55, Müller 2007:77ff). The information given in that feature structure⁷ is that *fast* is a relation between the semantic content of its head (the value of MOD N') and the comparison class, which is the value of the COMP feature. The latter is identified with the value of the COMP feature of the context index C-IND (which would, for other indexical expressions, contain such features as SPEAKER, ADDRESSEE, LOCATION – cf. Pollard & Sag 1994:332) and which must be recovered from the index.

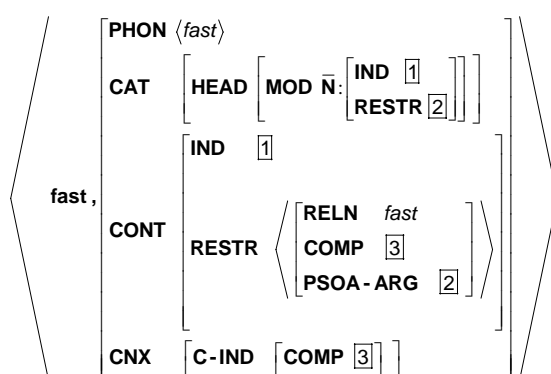


Diagram 9

Thus for an occurrence of sentence (7) in a discourse context concerned with a national competition, *fast* would denote a different relation than for an occurrence in the context of a recreational running event, the difference being due to a difference in the value of the COMP attribute. Similarly for (8) *fast* would typically denote a different relation in 1981, the time of the first Personal Computer, than in 2008; and also here it is due to the difference in the value of the COMP attribute that is identified with the comparison class that the index, C-IND, makes available.

(7) This was a fast run.

(8) This is a fast computer.

These denotational differences in (7) and (8) are, however, not a matter of any polysemy of the adjective, but – as we have modelled them here – a matter of *indexicality*.

In addition to the variation of the adjective interpretation in each of these sentences, there is also a difference between the adjective denotation in (7) in comparison to (8), which is due to the difference in the head that is modified by the adjective. This is the difference we already discussed in Section 3.2. And here we *are* concerned with polysemous variation in the denotation of the adjective: fastness of runners and fastness of computers are not only

⁶ Cf. the discussions in Partee (1995), Kamp (1975) and Siegel (1976). Note, however, that I am here only using the dependence of the interpretation of relative adjectives on context as an uncontroversial example of indexicality. I am trying to ignore all other questions that are controversial in the analysis of relative adjectives.

⁷ For a predicative use of adjectives, the lexical entry would obviously be different. But this is not our current concern.

established with respect to different comparison classes, but are different concepts⁸. The difference is a difference in their argument domain.

This is expressed in the lexical feature structure in Diagram 9 by including the semantic restrictions of the head noun, via coindexing, as an argument of the *fast* relation. Note, again, that these are restrictions on the *denotation* of the head and not restrictions on lexical expressions. The difference between (7) and (8) would be just the same if the head of the NP were in both cases the indefinite pronoun *one*, as long as the concepts denoted by the pronoun are still a run (in a national athletics competition or a recreational context) and a computer (in an occurrence related to 1981 or 2008).

The *indexicality* of *fast*, i.e., the dependence of its denotation on a comparison class, is obviously not independent of its polysemy or *context dependence*. We cannot, for instance, interpret an utterance of (8) in the sense that the computer in question is fast compared to a run in a recreational running competition. The comparison class cannot be selected independent of the kind of thing the adjective is applied to. This is not surprising since parameters of the index, such as the comparison class for gradable adjectives, are parameters of the context. The difference to other context parameters is that *index parameters are lexically specified* and are specified differently for different classes of lexical items. Context dependence comes in, however, via the arguments; for attributive adjectives, via the denotation of the nominal head.

3.5 Division of labour in language comprehension

We argued that the comparison class for gradable relative adjectives is a lexically specified index parameter. Evaluating this parameter may still be a tricky task as it involves knowledge about particular contexts. But this is no news. Language comprehension involves many cognitive resources of which linguistic knowledge is only one, and it is probably good policy to not only assume a modular conception of linguistic knowledge itself but also to adhere to a modular conception of language comprehension. So what we should have in a lexical entry, which is strictly part of linguistic knowledge, is only phonological, morphological, and syntactic information, plus, as far the semantics is concerned, the character function. The latter provides restrictions on denotations: their semantic type and possibly constraints on their arguments.

Let us now return to Pustejovsky's example that we already discussed. – What, should a lexical entry for the English verb *to open* look like? A rough proposal in terms of a feature structure is given in Diagram 10 for transitive *open*.

The semantics is mainly found in the CONT(ent) feature, which specifies a relational concept *opening* with an event argument (SIT), which remains implicit, and two syntactically realized arguments, AGENT and THEME. This is all the semantics that is required in the lexicon. The remaining tasks of comprehension are the job of the conceptual system.

Accordingly we would find in the conceptual representation (or the more abstract regions of the knowledge representation; the distinction does not matter here) a representation for one or more concepts that satisfy the RESTR imposed by the lexicon and that integrates our concept into the overall knowledge representation by specifying inference relations. In the case at hand the relevant kind of inference relation is a straightforward entailment, sometimes misleadingly called "meaning postulate": If an agent opens a theme, then we are entitled to assume an action involving this agent, and this action causes an opening of that theme. This entailment can be represented as in (9).

⁸ Concepts in this sense, which vary contextually, formed the basis of an account for vagueness I proposed in Bosch (1983) and were discussed further as "contextual concepts" in later papers (Bosch 1995, 2007a).

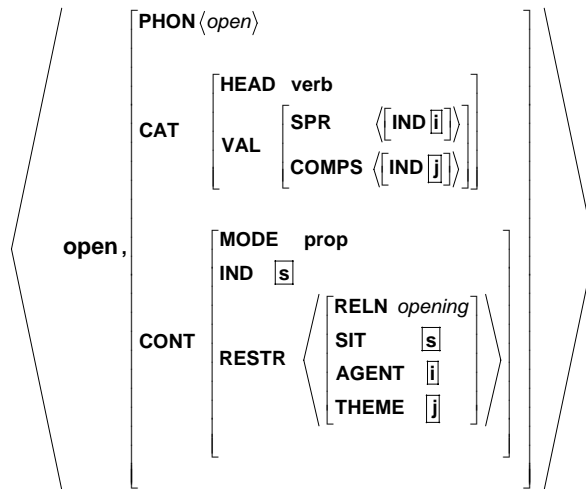


Diagram 10

- (9) $\lambda x \lambda y [\exists e (\text{opening}(e) \ \& \ \text{theme}(e,x) \ \& \ \text{agent}(e,y))$
 $\vdash (\exists e_1 \ \exists e_2 (e_1=e \ \& \ e_2=x \ \& \ \text{cause}(e_2,e_1) \ \& \ \text{action}(e_1) \ \& \ \text{agent}(e_1,y)$
 $\quad \quad \quad \& \ \text{opening}(e_2) \ \& \ \text{theme}(e_2,x)))]$

Elsewhere the conceptual representation will also tell us something about things opening, roughly as in (10), which incidentally would also qualify as the concept denoted to by the intransitive use of the verb *to open*.

- (10) $\lambda x [\exists e ((\text{opening}(e) \ \& \ \text{theme}(e,x))$
 $\vdash (\text{transition}(e) \ \& \ \exists s_1 (\text{preceding_state}(e,s_1) \ \& \ \text{- open}(s_1) \ \& \ \text{theme}(s_1,x))$
 $\quad \quad \quad \& \ \exists s_2 (\text{resulting_state}(e,s_2) \ \& \ \text{open}(s_2) \ \& \ \text{theme}(s_2,x)))]$

Applying this to the examples in (3), repeated below, we are entitled to infer that Mary's action causes the letter to be open, the rangers' action causes the trail to be open, etc. What it "means" to know that a letter, a trail, a door is *open*, or for that matter a shop or a bottle or an exhibition, i.e., what inferences we are entitled to is still another piece of knowledge which surely must be represented somewhere in our conceptual or knowledge representation, but this knowledge is unrelated to linguistic knowledge. And since we are here concerned with semantics, this may well be the point at which to stop.

- (3) a. Mary opened the letter for her mother.
 b. The rangers have opened the trail for the summer.
 c. John opened the door for the guests.

There is one important issue that we have not yet touched upon: in what sense are we concerned in the comprehension of (3) with *context dependence*, as I claimed above?

So far we have shown that the variation in inferences permitted by (3 a.-c.) is apparently due to our knowledge of what it means when a letter, a door, or a trail is open. We saw that the lexical items *letter*, *door*, *trail*, etc. play no other role than introducing the denotations that they happen to have in a particular occurrence as arguments of the opening relation. Any other linguistic expression that does the same job in a particular utterance context, like pronouns or proper names, will produce the same result for comprehension. Here the current proposal differs from Pustejovsky's, where the variation in understanding (3 a.-c.) was derived from compositional relations between lexical items in a sentence and where contextual denotations play no role at all.

One plausible complaint against the current proposal may be that the denotation for the transitive verb *to open* that I suggested, i.e., $\lambda x \lambda y [\exists e (\text{opening}(e) \ \& \ \text{theme}(e,x) \ \& \ \text{agent}(e,y))$, is vastly underspecified and thus could not account, for instance, for classic effects in gapping, such as the apparent incoherence of (11).

(11) # He opened a conference yesterday and she a boutique.

Does this not show that the characters or lexical meanings are far more specific than I claimed? No it does not. The relevant difference is one in the denotations, the *contents*, which are indeed much richer than the character of *to open*. They are the values of interpreting a verb occurrence in a particular context and take any or all denotational knowledge into account that context and knowledge representation make available, such as knowledge about boutiques and conferences and about opening boutiques or conferences.

Once the reader has reached the end of "He opened a conference" in (11), the conceptual representation for the VP would look somewhat like (12). But the concept of *conference opening* that is here predicated of the agent cannot be the content of the missing verb in the continuation "and she a boutique". What is needed to make (11) coherent would be a super-concept for the theme of *open* that subsumes both the relevant *conference* concept and the relevant *boutique* concept. If, as would be likely in a default context, the *conference* concept is represented as subsumed by concepts like *social event*⁹ or eventually *event* and the *boutique* concept is subsumed by a super-concept like *retail shop* and eventually by either *business* or *location*, then no unification would be possible and no suitable restriction for the theme of the verb *open* would be available. This is what makes (11) look incoherent.

(12) $\lambda x \lambda y. (\exists e (\text{opening}(e) \ \& \ \text{theme}(e,x) \ \& \ \text{conference}(x) \ \& \ \text{agent}(e,y)))$

The discourse context may, however, suggest different conceptualizations. Suppose the context is as in (13). This contextualization would facilitate a conceptualization of both openings as *opening events of the kind that the agency manages*, and hence would give us a suitable denotation for *open* with a vastly underspecified theme¹⁰.

(13) Mary's and John's event management agency is really busy these days. He opened a conference yesterday and she a boutique.

As the construction of conceptual representations is strictly a matter of conceptual representation and not of linguistic knowledge I don't want to go into it any deeper here. I only want to suggest that this is where systematic relations of conceptual subsumption have their place, as I have just shown in connection with (11) and (13). I believe that this is also the place where some of the conceptual relations belong that Pustejovsky, to my mind mistakenly, puts into his lexical entries.

3.6 Indexicality vs context dependence

In the case of the adjective *fast* in Section 3.4 we saw that the computation of the context-dependent interpretation was a matter of both indexicality and context dependence, where the index dependence goes via the comparison class and context dependence is computed via the denotation of the nominal head. The latter variation in content I claimed is a case of polysemy.

Things are different for the verb *to open*: there is no contextual index attribute in the lexical entry for this verb, and the link to the context is provided only via the argument expressions of the verb and their denotations, or, as we may say alternatively, via the argument domain for which the contextually relevant *opening* concept is defined.

⁹ I'm here using italics for the names of concepts, as I did in the feature structures. Of course this does not identify concepts with linguistic expressions, which in the text are also represented by italics. The text will always make clear what we are referring to, concepts or linguistic expressions.

¹⁰ In previous work (Bosch 1995, 2006) I discussed such contextual denotations for VPs as "contextual concepts".

What is common to both cases is that the context dependence of either *fast* or *open* is established via their argument domains. The difference is that for *fast*, but not for *open*, there is, in addition, an implicit lexically fixed parameter that accesses the context index.

So what we would like to say is that *fast* behaves in important respects like Kaplan's indexicals. Also for them the relevant parameter that is interpreted with respect to the context, is lexically fixed: *I* looks for the speaker, *now* for the current time, *nearby* looks for the current place etc. and for *fast* the indexical parameter looks for a comparison class. - This is different for *open*, which does not lexically provide for index parameters.

Finally, let me point out that the difference that we found between *fast* and *open* is corroborated by their behaviour in VP anaphora. Cf. (14) vs. (15)

- (14) a. Fred just loves speed. His car is really fast and so is his bike.
b. Fred and John went back to their homes. Later in the evening Fred went for a drink nearby, and so did John.
- (15) # Fred opened the door to let the carpenter in and so did the carpenter to replace the lock.

(14a) shows that the denotation of *fast* remains constant as it must under VP anaphora, still containing a variable for the comparison class though that is in each case evaluated with respect to the local argument: cars and bikes. The sentence does not ordinarily mean that the bike is anywhere near as fast as the car. What the bike and the car have in common is that each is fast with respect to its own comparison class.

This would be the same for *nearby* in (14b)¹¹ adjusting its indexical location parameter in the two clauses for Fred's and John's location: Fred went to drink near *his* home and John had a drink near *John's* home. The denotation of "went for a drink nearby" is still the same in application to John and to Fred, with a difference merely in the value of the location parameter of the index.

The same does not work out for the context dependence of *open*. Here we are not dealing with indexicality. The lexical entry of *open* has no contextual index attribute and there is no index dependence in the contextual opening concepts. The context-dependent denotation of *open* comes about purely on the basis of differences in the domain of the *opening* concept. Once the argument expression has been evaluated as a door (in the sense in which Fred opens a door to let somebody in, i.e., a door as *some kind of a boundary*), the contextual concept that stays identical in VP anaphora, has been fixed. Since the more plausible reading for the second clause of (15), however, is that the carpenter opened the door in order to get at the inside of the door – and here the door must be conceptualized differently, as *some kind of bounded object* rather than as a *boundary* – the sentence becomes incoherent. The variation between the two clauses of (15) is not a variation in the evaluation of an index parameter, but it is due to a difference between the conceptual representations of the arguments.

4 Conclusion

We saw in our discussion of the Generative Lexicon approach that the different contributions that different occurrences of the verb *to open* make to the truth conditions or to comprehension can probably not be accounted for in terms of compositional relations between constituents. I argued that the apparent variation in word sense is in fact the result of processes at the denotational level, i.e., conceptual rather than linguistic processes. Still, if we want to know what the contribution of linguistic knowledge to these processes is, we must find out how linguistic and conceptual processes interface.

Part of the solution I am proposing is a broadening of David Kaplan's notion of indexicality by allowing, in principle, for all lexical expressions to encode part of their meaning as a dependence on an index parameter. Indexicality is a lexical property of expressions and thus a

¹¹ The example is due to Graham Katz.

genuine part of linguistic knowledge. Indexical expressions thus carry their index parameter with them in all their denotations. *Fast* means in all contexts: fast with respect to some comparison class. Different values for the index parameter between two different occurrences do not lead to anything like different "senses" or polysemy.

We saw, however, that indexicality can only do part of the job of relating expressions to contexts. The other part is a form of context dependence that hinges upon the argument domains of the concepts denoted in particular occurrences and on the conceptual representation of arguments. This part, despite its consequences for genuine linguistic processes and despite its interactions with indexicality, seems to be the business of conceptual representation. And this is also where the roots of polysemy are: In context dependence in the narrow sense that is not indexicality and that is not reflected in any form in lexical entries.

5 References

- Bosch, P. (1982) The role of propositions in natural language semantics. In: W. Leinfellner, E. Kraemer, J. Schank (eds.): Proceedings of the 6th International Wittgenstein Symposium. Hoelder, Pichler Tempsky. Wien. pp 282-285.
- Bosch, P. (1983) Vagueness is context dependence. A solution to the Sorites Paradox. In: T.T Ballmer & M. Pinkal (eds.): Approaching Vagueness. North Holland. Amsterdam. pp 189-210.
- Bosch, P. (1995) Meanings and contextual concepts. In: M. Bierwisch & P. Bosch (eds.) Semantic and Conceptual Knowledge. Papers from a joint workshop of the Max Planck Arbeitsgruppe 'Strukturelle Grammatik' and the IBM Institute for Logic and Linguistics, 21.-23.4.1994 Berlin. Arbeitspapiere des Sonderforschungsbereichs 340, Bd. 71. Universität Tübingen. pp. 79-99
- Bosch, P. (2007a) Lexical & contextual concepts (Talk at the 8th Szklarska Poreba Workshop on Lexical Concepts, Szklarska Poreba, Poland, 25.02.2007)
- Bosch, P. (2007b) Productivity, polysemy, and predicate identity. in B. ten Cate & H. Zeevat & (eds.) Language, Logic and Computation. Proceedings of the Sixth International Tbilisi Symposium on, 2005. pp. 58-71
- Briscoe, T. & A. Copestake (1999) Lexical rules in constraint-based grammar. Computational Linguistics 25:4, 487-526
- Carnap, R. (1947) Meaning and Necessity. Chicago UP.
- Cutler, A. & C.C. Clifton, Jr. (1999) Comprehending spoken language: A blueprint of the listener. In C.M. Brown & P. Hagoort (eds.): the Neurocognition of Language. Oxford UP. Oxford. pp. 123-166
- Firth, J. R. (1957). A synopsis of linguistic theory 1930–55. In Studies in Linguistic Analysis: 1-32. The Philological Society. Oxford.
- Kamp, J.A.W. (1975) Two theories about adjectives. In E.L. Keenan (ed): Formal Semantics for Natural languages. Cambridge UP. pp. 123-155
- Kaplan, D. (1978). On the logic of demonstratives. J.Phil. Log. 8: 81-98
- Kaplan, D. (1977/1989) Demonstratives. In J. Almog, H. Wettstein and J. Perry (eds.) Themes from Kaplan. New York. Oxford University Press. 1989. pp. 481-563 (Ms. 1977)
- Lewis, D. (1970) General semantics. Synthese 22 (1970) 18-67.
- Lewis, D. (1981): Index, context, and content. In: Kanger, S. & Öhman, S. (eds.) Philosophy and Grammar. Reidel. Dordrecht.
- Montague, R. (1968/1974) Pragmatics. in R.H. Thomason (ed.) Formal Philosophy. Selected papers of Richard Montague. Yale UP. New Haven. 1974. pp. 95-118
- Müller, S. (2007) Head-Driver Phrase Structure Grammar. Stauffenburg. Tübingen.
- Partee, B.H. (1995) Lexical semantics and compositionality. In: An Invitation to Cognitive Science - 2nd Edition: Vol. 1. L. R. Gleitman & M. Liberman (eds). Cambridge, MA. MIT Press.
- Pollard, C. & I. Sag (1994) Head-Driven Phrase Structure Grammar. University of Chicago Press. Chicago.
- Pustejovsky, J. (1995). The Generative Lexicon. Cambridge, MA. MIT Press.
- Siegel, M. (1976) Capturing the Russian adjective. In B.H. Partee (ed.) Montague Grammar. Academic Press. NY. pp. 292-309