

Productivity, Polysemy, and Predicate Indexicality

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

Abstract This paper argues that at least some cases of productive language use, specifically cases of apparent variation in word sense, require a treatment at the conceptual level rather than a lexical semantic solution. It is argued that the lexical semantics should be left underspecified in these cases, and the observed variation in truth-conditions should be attributed to differences in conceptual representation that result from differences in the utterance context. This involves rather drastic changes in our conception of the semantics-pragmatics interface, which are discussed in this paper.

Introduction: The Semantics-Pragmatics Interface

Language comprehension is concerned in the first instance with linguistic utterances. But, as much recent work emphasizes (e.g. Trueswell & Tanenhaus 2005), language comprehension also makes essential use also of non-linguistic information as it may be available concurrently with the comprehension process. How exactly the processing of linguistic and non-linguistic information interact is still largely an open issue. In this paper I am exploring an architecture of the language processing system that is decidedly *modular* in that it clearly distinguishes different subsystems, but that is also *incremental* and *interactional* in the sense that it permits current processes to take advantage of information as soon as it becomes available to any one of the subsystems. Similar views were proposed in the language processing literature already in Marslen-Wilson & Tyler (1980) and Altmann & Steedman (1988), and are receiving additional support from recent eye-tracking studies (Hartmann 2005, Karabanov e.a. to app.).

The assumption of modularity and incrementality seems to lead to a better fit between theory and observation. But it makes it harder to maintain classic ideas about the semantics-pragmatics interface: It becomes more difficult, in particular, to understand the relation between compositional processes that construct the sentence-semantic content of an utterance, and subsequent pragmatic processes that would finally yield a truth-evaluable content adequate to the utterance context.

There may well be additional reasons, and not necessarily all of them empirical, to re-discuss the role of sentence-semantic contents or propositions (Bach 1999, Bosch 1982, Recanati 2002). I will skip these

arguments here and start directly by exploring an architecture that assumes that sentence comprehension uses information from the utterance context immediately and results in truth-evaluable utterance contents – without the intermediate construction of sentence contents. The idea, in brief, is that we dispense with semantic contents that are not context-relative.

Perhaps surprisingly it seems that already Gottlob Frege, in his later writings, held a similar view of sentence contents:

*Wenn mit dem Tempus Praesens eine Zeitangabe gemacht werden soll, muß man wissen, wann der Satz ausgesprochen worden ist, um den Gedanken richtig aufzufassen. Dann ist also die Zeit des Sprechens Teil des Gedankenausdrucks.*¹ (Frege 1918/1966:38f.; my emphasis, PB)

The currently probably most widely discussed, if not also the most widely accepted, theory that takes utterance context into account in the construction of semantic values is David Kaplan's theory of Demonstratives (1989). Kaplan distinguishes *characters* of linguistic expressions from their *contents*. If the linguistic expression is a declarative sentence, its *character* is a function that, much like one would expect of the 'meaning' of the sentence, yields for each context in which the sentence may be uttered the proposition it expresses relative to that context. This proposition then is the truth-evaluable *content* of the sentence for that context.

This way Kaplan can cater for many forms of explicit context dependence. In particular the reference of personal pronouns of the first and second person can be handled this way, but also many other indexical expressions, like *here*, *now*, *tomorrow*. While a majority of lexical expressions look as if they had constant semantic values, the exceptional class of indexical expressions is handled separately; they take specific contextual parameters as their values: *here* takes the place of the utterance as its value and *I* the speaker of the utterances, etc. It is not entirely clear though that Kaplan's account could be applied equally directly to the indexical interpretation of tense, as in Frege's example. And it seems likely that the account is not applicable to certain forms of implicit reference to the utterance context. John Perry (1998) and others have claimed that there are such implicit, or as Perry calls them "unarticulated" constituents that function indexically. He argues, for instance, that each utterance of the sentence

(1) It's raining,

contains an implicit place reference and hence cannot be assigned a semantic value without first making the place reference explicit: We cannot evaluate a statement made by (1) for its truth unless we know *where* it is supposed to be raining at the time of utterance.

¹ "If the present tense is used as an indication of time one needs to know when the sentence was uttered in order to grasp the thought correctly. The time of speaking thus is part of the expression of the thought".

2 More Context Dependence

Cases like in Perry's example form a serious problem for the architecture of the Semantics-Pragmatics interface. They seem to show that the utterance context cannot be satisfactorily modelled with the help of a set of parameters that could, as it were, be extracted from a list of expressions that occur explicitly in the linguistic utterance and that refer to features of the utterance context.

I want to argue that the situation may actually be even more difficult than Perry's argument suggests. Not only are there *implicit* indexical constituents that make the semantic value of a sentence depend on the utterance situation, much in the way that Kaplan proposed for explicit indexicals, but also a large proportion of the *explicit* constituents that are not in Kaplan's class of indexicals depend on properties of the utterance context in the contribution that they make to the truth-evaluable content of the sentence. These constituents do not seem to have lexically fixed constant semantic values.

Take the verb *rain* in sentence (1). Is there a constant semantic value for it in all utterances of the sentence *It's raining*? Is there a semantic value for *rain* that is independent of utterance situation and speakers' intentions? Are the truth-conditions for (1) really the same when the sentence is uttered as a reply to questions such as those in (2)?

- (2)
- a. Is it still snowing?
 - b. Are you saying it's still drizzling?
 - c. Is it still pouring like this morning?
 - d. Can we go for a walk now?
 - e. Why did you bring the washing in?
 - f. Why did you call a taxi?

The kind of modification that such contexts bring about in the interpretation of a lexical item is *productive*: Modification that is made and understood automatically and with no effort. It remains unnoticed by the language user, and it can yield, in principle, infinitely many variants of arbitrarily fine granularity. The *rain* example already hints this direction. Also, and perhaps even more clearly, there is the familiar kind of variation that we find in the interpretation of common verbs like *run*. Different truth conditions result for the utterance, depending on the argument to which the verb is applied: No constant semantic value for *run* will cover the cases of a running sportsman, a running water tap, a watch, a lecture, a program (or programme), or a stocking that are "running".

3 Productive modification of lexical meaning

Productive modification of lexical meaning is not, as may be thought, appropriately modelled as lexical disambiguation, i.e., as the mapping of the occurrence of an expression onto one of a certain number of semantically different lexical entries. Lexical disambiguation is well suited particularly for cases of homonymy, for instance, an occurrence of *bass*

must be mapped onto either of the lexical items $bass_1$ or $bass_2$, depending on the intended meaning.

Lexical disambiguation implies a cognitive choice and is a task that inhibits comprehension processes. It should be distinguished from processes that lead to a differentiation of word senses. The former task is accomplished fairly reliably also without much contextual information while the latter is not (cf. Veronis 1998, 2001). It has also been shown that homonymous words, which require disambiguation, slow down lexical access, while polysemous words, which activate a multiplicity of word senses, speed up lexical access (Rodd e.a. 2002).

However, both the productive modification of semantic values and the straightforward choice between lexically different items have in common that they require additional non-lexical information.

Suppose we had two lexical entries for the English verb *work*, as it is used in the sentences *The telephone wasn't working this morning* and *The caretaker wasn't working this morning*.

(3) Lexical Representation

Lexical entry 1: $\llbracket work_1 \rrbracket = \lambda x WORK_1(x)$

Lexical entry 2: $\llbracket work_2 \rrbracket = \lambda x WORK_2(x)$

Each of the two senses is here given as a different concept, and in our knowledge representation or conceptual representation we would explicate the difference between the two concepts $\lambda x WORK_1(x)$ and $\lambda x WORK_2(x)$ by different inferences that they may license:

(4) Conceptual Representation

$\forall x (WORK_1(x) \rightarrow DEVICE(x) \dots)$

$\forall x (WORK_2(x) \rightarrow HUMAN(x) \dots)$

If we do not assume that *work* is lexically ambiguous in this way, we still have two different concepts for the two sentences at hand, and we still have the very same difference in the conceptual representation. Only the lexical entry would give us just a single denotation: A concept $\lambda x WORK(x)$ that has, as it were, less content and is a super-concept of $\lambda x WORK_1(x)$ and $\lambda x WORK_2(x)$.

Note that the conceptual or denotational differentiation of the occurrence of a lexical item in the case at hand, just as in the case of disambiguation, requires access to semantic values, i.e., to the intended reference situation, and not just to lexical entries. An utterance of (5) does not contain relevant information to differentiate the denotation of *work* between $\lambda x WORK_1(x)$ and $\lambda x WORK_2(x)$ as long we know nothing about the intended reference of the name *Charley*.

(5) Charley isn't working this morning.

If Charley is a computer, another interpretation would be appropriate than in a case where Charley is our caretaker. But the relevant information comes from the utterance situation, or the intended reference situation, not from the lexicon.

Although it is quite clear in this case that differentiation or disambiguation is brought about by the argument and not by the argument expression, many argument expressions are apt to mislead about this point, as the examples in (6) may demonstrate.

- (6) a. *cut*: hair, bread, lawn, cake, ...
 b. *open*: book, letter, door, bottle, buffet, ...

The fact that the information required for the differentiation comes from the intended reference and not from the argument expressions, and hence cannot be thought of as being recoverable from lexical knowledge, is brought home when we consider pronouns as argument expressions: Plainly you can't get a differentiated denotation for utterances of *cut it*, or *open it* until you know what the pronoun refers to.

In the kind of case we have been considering one could argue, as in the case of Kaplan's indexicals, that the relevant information for disambiguation is found in one specific feature of the utterance context: in our case the argument. But not all cases of productive modification are like this. Here is a type of case where the denotation of argument expressions won't help:

- (7) a. Where is Fred? He's working.
 (8) a. How can Fred afford these expensive holidays? He's working.

If *He's working* figures as an answer to *Where is Fred?* the utterance is interpreted as giving information about Fred's location, and if it figures as an answer to *How can Fred afford these expensive holidays?* it must carry information about Fred's financial situation.

- (7) b. $WORK_i(\text{fred}) \rightarrow \phi(\text{LOCATION_OF}(\text{fred}))$
 (8) b. $WORK_j(\text{fred}) \rightarrow \psi(\text{WEALTH_OF}(\text{fred}))$

The denotation of *work* can do this only if it is enriched by contextual knowledge, and in different ways for (7a) and (8a)². The difference between the semantic values of the two occurrences of *work*, i.e., the concepts $\lambda xWORK_i(x)$ and $\lambda xWORK_j(x)$, is truth-conditionally relevant and licences different inferences: Nothing follows, for instance, about Fred's location when *He is working* is an answer to *How can Fred afford these expensive holidays?*

The difference between $\lambda xWORK_i(x)$ and $\lambda xWORK_j(x)$ in (7) and (8) is stable within the utterance context, as is demonstrated by VP anaphora:

- (9) Fred is working and so is Pete.

An utterance of (9) cannot be interpreted as saying, e.g., that Fred is in his office and Pete can afford expensive holidays.

What then is the exact difference between the occurrences of *work* in (7) and (8)? It can't be a difference in lexical meanings (Kaplan's *characters*)

² It may be thought that the difference in interpretation between the two replies *He's working* should be a matter of Gricean implicature. I have no objection to this. Note, however, that this formulation is no more informative than what I'm saying.

- because the variation correlates with a change in the utterance context. So I conclude that it is a difference in the semantic values (*contents*) of the two occurrences of *work*. Following the footsteps of Frege, who argued that the semantic values of predicate expressions are concepts, i.e. truth functions, I call these semantic values *Contextual Concepts* (CCs) (Bosch 1991, 1997).

CCs are the contextual referents or semantic values of predicate expressions. They are truth functions that are completely defined for arguments in the intended context. And what's more, CCs are linguistically real: they define the required notion of identity in VP anaphora, VP ellipsis, coordination, question-answer coherence, and they define the units in counting (cf. Bosch 2006).

As for VP anaphora, VP ellipsis, and coordination we observe, as hinted above, that the property ascribed to Fred and Pete in (10)-(12) is identified with the interpretation of the first occurrence of *working* in these sentences.

- (10) Fred is working and so is Pete.
- (11) Fred is working, and Pete too.
- (12) Fred is working and Pete is working.

When these sentences occur as answers to a question like *Where is Fred?* then this context enforces a CC $\lambda x \text{WORK}_i(x)$ that is a sub-concept of a more general location concept, somewhat like "to be at one's workplace" – only as such can the CC figure in an answer to *where* questions. If, in a different context, the sentences (10)-(12) are used to answer a question like *Is Fred still unemployed?* the CC induced by that context, $\lambda x \text{WORK}_j(x)$ is a sub-concept of a concept of employment status that could perhaps be paraphrased as "to be employed".

Finally, also the identity of countable objects rests on CC identity. When we are talking of Fred, Pete, and a few others and I say:

- (13) I wonder how many of them are working.

then we are not allowed to re-interpret *be working* for each of the individuals we are counting. The question is how many of them have the one property in common that is identified by the contextual interpretation of *be working*. If the answer is *Two of them*, then this cannot mean that one is working in the sense of being at his place of work right now and the other in the sense of being able to afford expensive holidays or just in the sense of being employed. Whatever CC is chosen as the semantic value for *work*, it must be the same for all of them.

Note that CCs are not meanings or senses in any sense. This is clearly seen from (14).

- (14) Fred is working for her, and so is Pete.

The relevant observation here is that the reference of *her* must be the same for both conjuncts: Both Fred and Pete must be working for the same

woman³. The reference of *her* is part of the specification of the CC that is asserted of both Pete and Fred. – If this observation is correct then it would follow that the identity of CCs depends on identity of reference. This would exclude the option of regarding CCs in any sense as meanings – at least under most current accounts of "meaning".

4 How are Contextual Concepts computed?

The computability of the content of indexicals along the lines of the Kaplan approach rests on the assumption that we are concerned with a process of variable saturation: One variable at a time that is evaluated by the context (speaker, listener, place, time, etc.).

This assumption won't work for the context dependence of predicate expressions we have been discussing, i.e., for *predicate indexicality*, as one might call it. The reason is that there are arbitrarily many parameters with respect to which CCs can differ from each other, and there can be no finite parameterisation of contexts (hence no context-independent identity of contexts, cf. Bosch 1982). – So how can we compute CCs, given whatever little information about contexts we have?

First of all there is an important difference to note: The computation of CCs differs for different information structure status of the constituent we are concerned with (Bosch 1999). *Anaphoric constituents* are interpreted by reference to objects that are already given in the established discourse representation. There is no modification and no lexical semantics involved. *Focal constituents*, on the other hand, are interpreted via their lexical entries, plus disambiguation and modification. The "disambiguation" and modification here are however not linguistic, but conceptual processes: We are concerned more with reasoning than with linguistic semantics. I will return to this point below.

But let me first explain the relevance of the distinction between *anaphoric* and *focal constituents*. Let's take as an example a question like (15) that is being asked about our friend Fred. *Working* would here have a *focal* occurrence: Its interpretation cannot be derived from any preceding discourse context, but depends largely on the lexical semantics of the expression plus, possibly, other factors that are not our concern at the moment. If now (15) is answered by (15a), then the interpretation of *working* in (15a) must be exactly the same as in (15) – whatever the interpretation in (15) may have been – on pains of (15a) not being an answer to the question. This is what it means to say that the occurrence of *working* in (15a) is *anaphoric*. – Now suppose the answer is not (15a) but (15b), which would have exactly the same effect in the dialogue and be truth-conditionally fully equivalent; perhaps it would sound even more

³ Readers who have any doubt about this observation, as one of the reviewers had, are reminded that sloppy identity readings of VP anaphora are irrelevant here. The argument only requires that there is at least one reading of (14) in which the referential identity of the pronoun is decisive. One situation in which such a reading is unavoidable is one in which the pronoun is accompanied by a pointing gesture.

natural as a reply to (15). But the word *working* does not occur in (15b). So, clearly, if both (a) and (b) mean the same in this context, then the lexical semantics of *working* in (a) can't be making a semantic contribution to the interpretation of (a).

- (15) Is Fred working?
a. No, he isn't working. He's sick.
b. No, he isn't. He's sick.

Coming back to our question about the computation of focal and anaphoric constituents, we can now say that anaphoric constituents select their denotation from a small and finite domain: The current discourse representation. The information used in their interpretation is discourse-structural information, as in the interpretation of anaphoric pronouns, such as salience status.

Focal constituents on the other hand are mapped onto their semantic values in two steps:

- (i) The expression selects a lexical entry, which contains a pointer to a lexical concept, and
- (ii) conceptual processes and information from non-linguistic sources complete the construction of semantic values. – This requires that the lexical concept that is identified in (i) is part of a conceptual representation and is linked to other concepts via its internal structure, via subsumption, and possibly also via various axioms. It may still be underspecified with regard to what is the eventual semantic value.

The properly speaking semantic part of this interpretation ends after step (i) and the more complex and clearly more interesting part is in step (ii), which is the business of conceptual processing or reasoning – and with no specifically linguistic ingredients.

So far I have given a rough sketch of how semantic values are computed in the case of anaphorically occurring predicate expressions and I have assigned the computation of semantic values for focal constituents to conceptual processes. This means that what I referred to earlier as *productivity of interpretation* ends up in the conceptual system. – Well, not entirely: The semantic values of sub-sentential constituents that are computed via conceptual processes may (and ordinarily should) still combine with each other compositionally in the construction of a semantic value for the entire utterance – and here we are of course back in the semantics. The crucial point is rather that, in order to compute the semantic values also for sub-sentential constituents, we need to step outside of the semantics and get into conceptual processing before we can return to compositional semantic processing. If you like, you may call this "compositional pragmatics".

5 Some applications

In the following I want to look at the consequences of our proposals from the preceding sections for the treatment of some relevant phenomena,

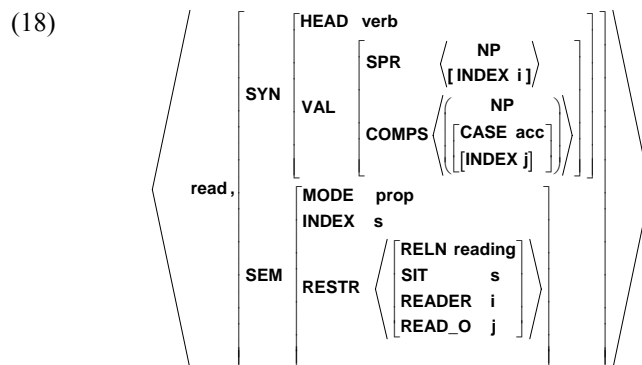
here in particular the transitive-intransitive alternation and the interpretation of argument expressions, as exemplified in (16) and (17) respectively.

- (16) a. Fred is reading a novel.
 b. Fred is reading.
- (17) a. Fred began reading the novel.
 b. Fred began the novel.
 c. Fred began.

How many lexical entries and how many lexical meanings for the verbs *read* and *begin* do we need?

5.1 Transitive-intransitive alternation

Let me first discuss the difference between (16a) and (16b). At least intuitively, one would like to have just one lexical entry for the verb *read*. Perhaps like the one represented in (18). Here the direct object is syntactically optional, but is co-indexed with the semantic object. Although we don't mention explicitly the object that is being read every time that we talk about reading, the relevant concept of reading still requires something that is being read as much as it requires a reader.



The denotation of the verb *read* is a lexical concept, which may be partially specified in the conceptual representation as in (19).

$$(19) \quad \llbracket read \rrbracket = \lambda y \lambda x \exists s (\text{READING}(s) \wedge \text{READER}(x,s) \wedge \text{READ_O}(y,s))$$

What happens in the comprehension process if the linguistic utterance mentions no object that is being read (no READ_O), as in (16b)? Nothing very serious. Note that (19) forms part of the conceptual representation. This is the level of representation where information from all sources available, not just from the linguistic utterance, gets integrated into to the comprehension process. If the utterance does not say what it is that is being read, the listener will still assume that such an object exists in the READING situation, because it is conceptually required. This is clear from a "bridging inference", as in (20).

(20) Fred is reading. But he's bored by the book.

The definiteness of *the book* demonstrates that the referent is regarded as identifiable and familiar. The assumption would be reasonable then that some object read by Fred was either already assumed in the comprehension of the first sentence in (20) or is at least easily accommodated. This follows directly from our concept of reading situations.

(21) $\forall s(\text{READING}(s) \vdash (\exists x\text{READER}(x,s) \wedge \exists y\text{READ_O}(y,s)))$

The verb *walk*, for comparison, does not denote a concept that would introduce a reading object or would allow easily for the accommodation of one, cf. (22).

(22) [#]Fred is walking. But he 's bored by the book.

Similarly, if the *READER* is not explicitly mentioned, as in (23), we still may assume that there is a reader, albeit a generic reader:

(23) This book reads beautifully. The reader never gets bored.

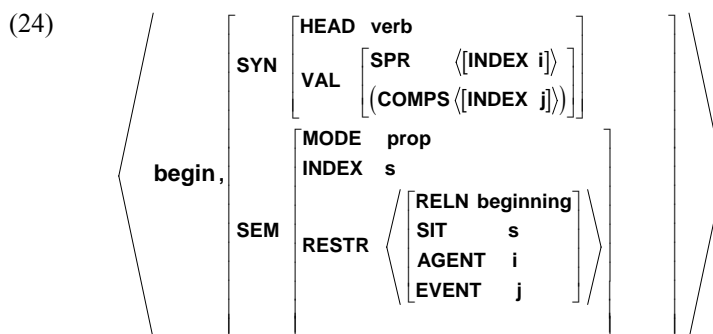
Since, as would seem plausible in the absence of other evidence, the intransitive, the transitive, and the ergative construction of *read* all denote the same concept, we assume no lexical semantic difference. The difference is exclusively syntactic⁴.

5.2 Argument (re-)interpretation

The second set of cases I want to look at are those in (17), repeated here.

- (17) a. Fred began reading the novel.
 b. Fred began the novel.
 c. Fred began.

The question here is how we can lexically specify the verb *begin* so that (17a) – (17c) are all licensed and interpretable in the intended sense. A suitable lexical entry is provided by the feature structure in (24) and an appropriate partial conceptual representation follows in (25).



⁴ The lexical entry in (18) licenses only the transitive and intransitive, but not the ergative construction. I take it that there is a semantically empty lexical rule that derives the ergative structure which takes care of the argument linking.

$$(25) \quad \llbracket begin \rrbracket = \lambda y \lambda x \exists s (\text{BEGINNING}(s) \wedge \text{AGENT}(x,s) \wedge \text{EVENT}(y,s) \\ \wedge \neg \exists z (\text{PART_OF}(z,y) \wedge \text{EARLIER_THAN}(z,s)))$$

Both (24) and (25) say that the concept $\llbracket begin \rrbracket$ requires an event argument. Accordingly the comprehension process would attempt to recover information about the event argument from utterances of sentences like those in (17) or, alternatively, from the utterance situation.

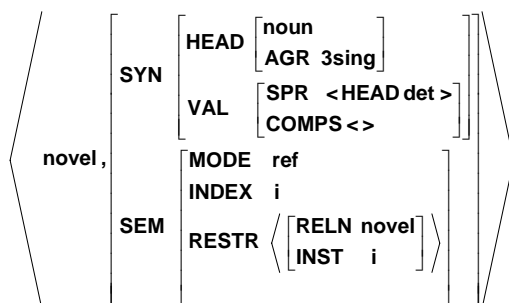
This is straightforward for (17a): Reading is an event (as should be represented in the conceptual representation by a subsumption relation for the concept $\llbracket read \rrbracket$), and so is reading a novel, which should be represented as a subconcept of $\llbracket read \rrbracket$. (17c) is treated along the same lines as already (16b): There is no reason to specify every relevant conceptual ingredient linguistically if it can also be retrieved from other sources. Hence *Fred began* is perfectly normal in contexts where the required event can be recovered from contextual knowledge, as must be the case for an utterance of (26):

(26) Everybody was waiting. But not Fred. Fred began.

When we consider (17b), however, we are encountering a new problem: What do we do with the surface object argument of *begin*, the NP *a novel*? Can we interpret *a novel* as referring to an event?

Following the general line suggested in this paper, we assume underspecified lexical representations, as in (27) for *novel*, and leave the conceptual content to the conceptual representation, as in (28).

(27)



$$(28) \quad \llbracket novel \rrbracket = \lambda x \text{NOVEL}(x)$$

Admittedly, (28) is not any more informative than the semantic restrictions in (27). But we must take care, because the noun *novel* is highly polysemous. We can refer to a physical object as "a novel" (when the weight of the book is at issue), we can refer to a set of ideas as "a novel" (when we are saying that an author is working on one, while there is not a single printed copy of it yet, nor even a finished typescript), we can refer to a publication as "a novel" (when we are talking about an author's most recent novel), etc. If we want to maintain just one lexical entry for *novel* it must remain underspecified in many respects (though eventually we will probably need to add some such attribute as perhaps RELATED_

TO_A_BOOK-LENGTH_PROSE_NARRATIVE.)⁵

How are these different senses of *novel* represented? The simple proposal I have to make (Bosch 1997) is that the various senses each represent a different *view* of our underspecified concept $\llbracket novel \rrbracket$. We may view a novel as a physical object, as a text, an object of the publishing industry, etc. Each of these views is represented in the knowledge representation by different subsumption relations. The concept $\llbracket novel \rrbracket$ may be subsumed by concepts like PHYSICAL_OBJECT, TEXT, PUBLICATION, etc. In each such view the resulting CCs for *novel* (in all cases sub-concepts of the lexical concept $\llbracket novel \rrbracket$) inherit different attributes from the different superconcepts under which they are being subsumed. Allowing, as we do, multiple inheritance relations, nothing prevents the formation of CCs that are subconcepts of NOVEL and are also subordinate to TEXT and PHYSICAL_OBJECT. This latter concept would be required, e.g., in the comprehension of (29).

(29) Peter was reading the novel that he had found at the bus stop.

Which *views* of a concept are actually possible is a further question that we cannot start investigating here (but cf. Bierwisch 1982:93 for discussion).

Returning to our current problem about *novel* as the grammatical object of *read*, we can now say that in one class of CCs a novel is viewed as an event, i.e., it is subsumed by a higher concept EVENT. These event concepts of *novel* are still different from each other: Reading events, writing events, production events, etc. These *novel*-CCs are partially specified by their subsumption relations as in (30) via which they inherit some of their attributes.

(30) $\lambda x \text{NOVEL}_i(x) \subset \lambda x \text{TEXT}(x)$
 $\lambda x \text{NOVEL}_j(x) \subset \lambda x \text{PUBLICATION}(x)$
 $\lambda x \text{NOVEL}_k(x) \subset \lambda x \text{PHYS_OBJ}(x)$

And indeed all attributes of these higher concepts are inherited (by default) by the concepts they subsume. Accordingly, whatever attributes we can generally attach to physical objects, we can attach to novels under the appropriate physical object view of novels, and analogously for other superconcepts of $\llbracket novel \rrbracket$ on the respective views.

For our current problem the question is if novels may indeed be viewed as events. In general, anything may presumably be viewed as anything we please, provided the resulting concept does not become inconsistent or remains without sub-concepts or without instances for other reasons. With sufficiently general concept specifications, however, this happens less frequently than one may expect. As for our case at hand, i.e., for the subsumption of the lexical concept NOVEL under EVENT, it follows first of all, that under this view we can attribute any properties that we can generally attribute to events also to novels, and secondly that, under this view, a novel may figure as an event argument, as it does in (17b).

⁵ This adopts an idea of Manfred Bierwisch (1982), who used "has as its goal processes of education" for the polysemous *school*.

The approach to productive polysemy I am here proposing receives additional support when we look at examples like (31), that was discussed by several authors (Godard & Jayez 1993, Copestake 2001).

(31) Fred began the tunnel.

It was claimed that this could mean (if anything) that Fred began building the tunnel, but not, for instance, that he began walking through the tunnel. Even though, admittedly, this would certainly not be the most plausible interpretation when (31) is considered out of context, it seems like an entirely natural interpretation if the appropriate context is provided. All you need to know is that there is this somewhat weird group of people who call themselves "The Royal Tunnel Walkers Society" who spend their weekends travelling to the mountains and walking through tunnels, regarding this as some kind of an exciting sport. Having introduced these people into our context, the desired interpretation of (32) should come about fairly naturally.

(32) The first group began the tunnel at 5:15.

What we need in our conceptual representation is a view of a tunnel as an event. The role of the context here is simply in suggesting event types that may be related to tunnels, like building tunnels or, in our case, walking through tunnels, and of which our group could be the agent.

The proposal about polysemy I have been making differs somewhat from the 'generative lexicon' approach (Pustejovsky 1995). The central difference is, I believe, that I make a clear distinction between the lexical representation on the one hand, and conceptual or factual information on the other. In Pustejovsky's lexical entries we find a considerable amount of conceptual and contingent information, which is not linguistic information and which really belongs into the conceptual system and the representation of factual knowledge. My proposal in this paper follows rather the line of a 'disquotational' view of lexical semantics (Fodor & Lepore 1998), in an attempt to keep linguistic and conceptual information apart and allocate them to different modules of the cognitive system.

6 The semantics-pragmatics boundary - from the perspective of semantic minimalism

According to the approach taken in this paper, language comprehension is fed by information from many sources, only some of which are properly speaking linguistic. A considerable part of the semantic productivity that we observe in language is attributed to the workings of the conceptual system. This becomes apparent when we consider how linguistic utterances are mapped onto truth-evaluable contents that cannot be derived from the current linguistic input alone, but contain conceptual material from other sources: preceding discourse, background knowledge, the intended reference situation, and the current utterance situation. I want to suggest that it is worth continuing this line investigation, if only for the reason that the classic semantics-pragmatics interface – first constructing a

semantic representation and then using pragmatic knowledge – holds little promise for an account of productive language use.

Acknowledgements

Various earlier versions of this paper were presented at a workshop on "The (In-) Determinacy of Meaning: Issues in Formal Pragmatics" organised by Regine Eckardt and Markus Egg at the 27. Jahrestagung of the Deutsche Gesellschaft für Sprachwissenschaft in February 2005 in Cologne, at the Semantikzirkel at the Zentrum für Allgemeine Sprachwissenschaft, Berlin in a talk on 3rd June, 2005, at the Sixth International Symposium on Language, Logic and Information, 12-16 September 2005, in Batumi, Georgia, and in a talk at the DIP Colloquium at the Institute for Logic, Language and Computation in Amsterdam on 18th Nov, 2005. I am very grateful for the opportunity to present these ideas on these occasions and all the comments and criticism that I received at these meetings, in particular to Michiel van Lambalgen who had prepared a critical reply to my talk at the DIP Colloquium and to two anonymous reviewers who read the paper for his volume.

References

- Altmann, G.T.M. & Steedman, M.J. (1988) Interaction with context during human sentence processing. *Cognition* 30:191-238.
- Bach, K. (1999) The semantics–pragmatics distinction: What it is and why it matters. In K. Turner (ed.) *The Semantics–Pragmatics Interface from Different Points of View*. Elsevier. Oxford. Pp. 65–84.
- Bierwisch, M. (1982) Semantische und konzeptuelle Repräsentation lexikalischer Einheiten. In R. Ružička and W. Motsch (eds.): *Untersuchungen zur Semantik*. Akademie-Verlag, Berlin. Pp. 61-99
- 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. 1982. Pp 282-285.
- Bosch, P. (1991) The Bermuda Triangle: Natural Language Semantics Between Linguistics, Knowledge Representation, and Knowledge Processing. In O. Herzog & C.R. Rollinger (eds.): *Integrating Linguistics and Logic. Final Report on the IBM Germany LILOG Project*. Springer. Berlin. Pp. 243-258.
- Bosch, P. (1997) Dynamik kontextueller Konzepte. In C. Umbach, M. Grabski & R. Hoernig (eds.): *Perspektive in Sprache und Raum*. Deutscher Universitätsverlag. Wiesbaden. Pp. 195-209
- Bosch, P. (1999) Presupposition, Focus, and Lexical Semantics. In *Semantica lessicale. Studi Italiani di Linguistica Teoretica e Applicata*, XXVIII,2: 249-264

- Bosch, P. (2006) *VP Anaphora at the Semantics-Pragmatics Interface*. Talk at the 2006 Milan Meeting, June 15-17, Palazzo Feltrinelli, Gargnano (slides at <http://www.cogsci.uos.de/~pbosch/>)
- Copetake, A. (2001) The semi-generative lexicon: limits on lexical productivity. In *Proceedings of the 1st International Workshop on Generative Approaches to the Lexicon*. Geneva.
- Fodor, J.A. & Lepore, E. (1998) The emptiness of the lexicon. *Linguistic Inquiry* 29: 269–288
- Frege, G. (1918) Der Gedanke. repr. in: Günther Patzig (ed.): *Gottlob Frege: Logische Untersuchungen*. Vandenhoeck & Ruprecht, Göttingen. 1966. Pp. 30-53.
- Godard, D. and J. Jayez (1993) Towards a proper treatment of coercion phenomena. *Proceedings of the Sixth Conference of the European Chapter of the Association for Computational Linguistics (EACL-93)*. Utrecht. The Netherlands. pp. 168–177.
- Hartmann, N. (2005) *Processing Grammatical Gender in German*. Bachelor Thesis, unpubl. Univ. of Osnabrück, Cognitive Science, http://www.cogsci.uni-osnabrueck.de/~CL/download/Hartmann_GramGender.pdf.
- Kaplan, D. (1989) Demonstratives. In J. Almog, H. Wettstein and J. Perry (eds.) *Themes from Kaplan*. New York. Oxford University Press. Pp. 481-563.
- Karabanov, A., Bosch, P. & König, P. (to app.) Eye Tracking as a Tool to Investigate the Comprehension of Referential Expressions. In S. Featherston & W. Sternefeld (eds.): [*Proceedings from the Workshop on Linguistic Evidence*. Tübingen 2006]
- Marslen-Wilson, W. D., & Tyler, L. K. (1980) The temporal structure of spoken language understanding. *Cognition*, 8, 1-71
- Perry, John (1998) Indexicals, Contexts, and Unarticulated Constituents. *Proceedings of the 1995 CSLI-Armsterdam Logic, Language and Computation Conference*. Stanford. CSLI Publications.
- Pustejovsky, J. (1995) *The Generative Lexicon*. The MIT Press. Cambridge, MA.
- Rodd, J., Gaskell, G. & Marslen-Wilson, W. (2002) Making sense of semantic ambiguity: Semantic competition in lexical access. *J. of Memory and Learning* 46:245-266.
- Recanati, F. (2002) Unarticulated Constituents. *Linguistics and Philosophy* 25: 299-345
- Trueswell, J. C. & Tanenhaus, M.K. (eds.) (2005) *Approaches to Studying World-Situated Language Use*. MIT Press. 2005
- Veronis, J. (1998) A study of polysemy judgements and inter-annotator agreement. In *Programme and Advanced Papers of Senseval Workshop*, Herstmonceux Castle, UK.
- Veronis, J. (2001) Sense tagging: does it make sense? In *Corpus Linguistics 2001 Conference*.