

The Bermuda Triangle: Natural Language Semantics Between Linguistics, Knowledge Representation, and Knowledge Processing

Peter Bosch

Abstract

Linguistic parameters alone cannot determine the interpretation of natural language utterances. They can only constrain their interpretation and must leave the rest to other knowledge sources and other processes: language understanding is not just a matter of knowing the language, but also to a considerable degree a matter of logical inference and world knowledge. This is no news as far as the interpretation of referential expressions is concerned. Predicate expressions, however, tend to be treated as if they were functional or relational constants that are directly interpreted with respect to a model. In this paper an attempt is made to treat them too as referential. The only real difference is that the referents of predicate expressions are of a different type: concepts rather than first-order objects.

This generalized notion of reference gives us not only a natural way of understanding the interaction of knowledge representation and knowledge processing on the one hand and linguistic processing on the other, but also opens up a perspective for the modularization of Natural Language Understanding (NLU) systems that provides for a very high degree of independence of the more strictly linguistic component from the specific tasks and domains of a particular application. The overall result should be a less language-dependent knowledge representation and less domain-dependent linguistic components, i.e., overall improved portability of the modules.

Given that Knowledge Representation was principally driven by natural language concerns right up to the beginning of the decade, one would have expected substantial progress to have been made in the '80s on KR support for NL semantics. This seems not to have been the case.

Ronald J. Brachman (Proceedings of AAAI 1990:1088)

1 Introduction

It may seem that no substantial progress has been made in Knowledge Representation (KR) support for Natural Language (NL) semantics, and in view of the considerable amount of work that is actually carried out in this area, one may well get the impression that there is a kind of Bermuda Triangle where a good amount of this work simply disappears: the area between linguistics, knowledge representation, and knowledge processing.

A possible cause for this situation is that we have a serious problem with respect to the generality and portability of results. Even though logic and linguistics may interact smoothly in each individual research system that is built, there is probably too little attention for the question of how they ought to interact in general. Hence there are serious limitations to the portability of both the linguistics components and the knowledge representation modules, a "fresh start" is needed in just about every new project, researchers are overburdened with the essentially unnecessary repetition of work other have already done, and no time and resources are left to face the more fundamental questions and to make significant scientific progress. In other words: more work that mysteriously disappears in the Bermuda Triangle of natural

language understanding.

LILOG, I believe, is one of the few projects in which some progress with respect to general solutions and portability has been made, as is visible from the contributions in this volume. A very important part of the progress is found in the tools that have been developed and that incorporate many of the general insights that have been gained. Another respect in which LILOG has advanced further than many other projects is the modularity of the system. But there is still a long way to go until we have the kind of principled division of modules that enables the full generality and portability that we are after. What is needed, in particular, is a better understanding of the division of labour of the AI and NL components in NLU.

When knowledge representation does part of the job of linguistic semantics, language-specific information tends to seep into knowledge representation, with the result that the knowledge representation becomes unnecessarily dependent on features peculiar to natural language, or even to a particular natural language, and NL components become unnecessarily dependent on the task and domain inherent to a particular knowledge base. Conversely, when linguistic semantics delivers linguistically insufficiently analyzed structures to knowledge processing components, the latter are burdened with tasks for which they are ill-equipped to provide principled solutions.

2 Inferential Instability

Let us assume for the sake of the argument that the core problem of linguistic analysis is disambiguation: the translation of natural language input strings into unambiguous representations on which inferential processes can operate, i.e., translation into inferentially stable representations.

Part of the problem of how linguistic and logical modules interact then turns on the nature of the inferential instability that characterizes natural languages.

For our purpose it will suffice to review just a few of the factors that lead to inferential instability or ambiguity. One is structural equivocation, such as the ambiguity of quantifier scope in (1):

- (1) Five companies sold two hundred installations.

(cf. also Link and Schütze, this volume, Section 1.4). Another is ambiguity of syntactic analysis, as with respect to the attachment of the prepositional phrase in

- (2) Fred saw the woman with the binoculars.

(cf. also Bosch and Geurts 1990). Further, the occurrence of word forms that are morphologically ambiguous in the sense that they could belong to different syntactic categories may cause different readings. For instance Chomsky's infamous visiting in

- (3) Visiting relatives can be boring.

Finally, the ambiguity of sentences may be due to the occurrence of ambiguous words, like the notorious *bank* in

- (4) Pete went to the bank this morning.

Some of these sources of equivocation are better under control than others. But what they have in common and what makes them look less threatening than the cases to come is that in

each case there is, in principle, a finite disjunction of inferentially stable representations among which the linguistic analysis must make a choice.

This is no longer the case when we turn to indexical expressions, like the determiner *this* in the sentence

(5) This paper is ten pages long.

This sentence may be true or false, depending on what paper *this paper* refers to. But then it is usually said that this case is a matter of referential ambiguity and is not due to an equivocation in linguistic analysis.

Usually the difference between (4), which is taken to depend on the lexical ambiguity of *bank*, and the case of indexicals in (5) is regarded as reasonably clear - which, unfortunately, it is not (cf. Bosch 1990). Take for instance the colour adjective *red*, which shows clear equivocation in the following contexts:

(6) red tomato
red apple
red hair
red wine
red grapefruit
red traffic lights

We are here not concerned with an ambiguity of the kind we find with lexical items like *bank*. There is good reason for the suspicion that there are not just the six variants of the concept of redness that we find under (6) but that this list can be extended ad libitum, and hence the equivocation of *red* would rather resemble the referential ambiguity we find with indexicals.

The equivocation we observe under (6) clearly leads to a difference in truth-conditions. If we assume compositionality, and the adjective *red* were to make the same contribution to the truth-conditions in all these cases, the conditions for calling a tomato, an apple, hair, wine, etc. *red* should be the same, which they are not, and one would surely want to avoid the system's conclusion that tomato, hair, grapefruit, etc., are, at some level of granularity, all of the same colour. However, there is no level of granularity at which these colours are the same, nor are the differences we find in (6) in any interesting sense related to vagueness.

But a treatment that postulates lexical ambiguity of *red*, possibly with suitable subcategorization restrictions that distinguish tomatoes from grapefruits, i.e., the postulation of competing lexical meanings, paraphrased as "red for a tomato" versus "red for a grapefruit", does not look like a good solution. The case bears much more similarity to the case of indexicals than to lexical ambiguity in the sense that the variation in the adjective's contribution to the truth-conditions depends on reference to objects of experience, i.e., on our empirical knowledge of tomatoes, grapefruits, etc. And such knowledge should certainly not be included in the dictionary, nor in any other representation of linguistic knowledge. It is strictly the business of the representation of knowledge of the world.

Thus the lexical semantics of *red* should be the same in all these cases - as it is for an indexical determiner like *this*. The variation in the contribution the adjective makes towards the truth-conditions can instead be explained by assuming that *red* (more precisely, its semantic specification) is a function which, depending on contextual factors, refers to different concepts, i.e., yields different concepts as its value - quite analogous to the reference of indexicals, with the difference that *this* refers to objects, while *red* refers to concepts. This at least looks like a

reasonable approach at first blush (but compare Sections 6 and 7 below).

The point is not specific to adjectives or colour terms. Consider an example that may look quite different at the first glance. Rommetveit (1986) tells the story of a certain Smith who is mowing his lawn. His neighbour addresses him over the fence and asks "Are you not working today?". Smith's reply: "No, I've taken the day off". At the same time Mrs Smith is on the phone to a friend telling her "My husband is working in the garden". Obviously, if *working* is taken at face value, our NLU system will derive a contradiction or ascribe contradictory beliefs to Mr and Mrs Smith.

But this does not mean that the English verb *to work* is lexically ambiguous. The verb means the same, but it is used in either case to refer to a different concept, again due to the different contextual parameters, in particular the available similarities and contrasts in the subject domain. In the question of Smith's neighbour, for whom Smith's physical activity is plain to see, a concept of work as physical activity cannot be at issue; hence other concepts of work are more plausible candidates, e.g. work in the sense of doing one's job. In the telephone conversation we have the more explicit formulation "working in the garden", which, ordinarily, suggests physical activity and is taken to exclude the concept of work in the sense of doing one's job. - As in the case of the colour adjectives, these distinctions reflect distinctions that are not of a linguistic nature, but are the business of knowledge about our particular social world and hence should be accommodated in the concepts we have about that social world rather than in the dictionary.

A point that is absent, or at most marginally relevant, in the colour example enters the scene in the case of Smith's garden work: communicative intentions that are understood among the interlocutors. They too influence the selection of relevant contrasting concepts from which the intended concept must be distinguished. This becomes even clearer in the following case, which is really a matter of different granularity of the concepts involved. The example is taken from Winograd and Flores (1986). Depending on whether an utterance of

(7) There's water in the fridge.

is meant as a statement about where the addressee can find a cold drink or as a warning that moisture-sensitive chemicals should not be stored in the fridge, the truth of (7) may vary, while the fridge, its condition and contents remain the same.

If we want to get such cases of inferential instability under control we need to provide for a clear division of labour between intrinsically linguistic information and information from other sources. But before we progress in this direction, I want to make an attempt to weaken overdrawn expectations with respect to what linguistic meaning can do in processes of linguistic comprehension. I shall look at a simple example in the following section.

3 Dog biscuits and nut cake

No-one expects there to be a dog in a dog biscuit, but most people expect there to be a nut, and usually more than one, in a nut cake.

NLU systems may reasonably be expected to share this disposition of the majority of English speakers. Still, the disposition is not conditioned by any linguistic facts about English. It is a matter of our conceptual system, i.e., the knowledge representation. Concepts like those referred to by *dog biscuit* and *nut cake* play a role in the representation of those parts of the world where corresponding objects are relevant. They are means to represent beliefs like those expressed in English sentences like *Texans love dog biscuits* or *Dogs are usually fond of nut cake*, and the like.

But the English expressions *dog biscuit* and *nut cake* only hint at these concepts and do not determine them. If we do not already have those concepts, then there is no way of constructing them merely from the semantics of the English expressions - they would make your mind boggle and the result could be anything.

As already suggested, I want to regard the relation between English expressions and concepts as a relation of reference: literally in the same sense of a relation reference as between referentially used definite NPs and their objects of reference, e.g., between the English expression *The British Prime Minister* and, at the time of writing, a particular person of the name *John Major*. In neither case can the semantics of the English expression on its own make sure that the correct referent is identified. It can give some help in this respect by constraining the set of possible referents, but no more.

Whether or not the mechanisms of reference work to satisfaction is a matter of two kinds of factors: one is knowledge of English, the other is world knowledge (i.e., in particular knowledge of the domain from which the referents - concepts or first-order objects - are to be chosen, common-sense knowledge, and knowledge of the communication situation). The better we keep these factors apart, the more suitable our approach will be for unexpected cases and unintended applications, and hence the more generality and the more portability of both knowledge representation and linguistic components can be achieved.

4 Lexical Meaning

In order to sort out the contribution of linguistic factors to referent identification, i.e., the constraints imposed by the linguistic properties of an utterance, we require a full linguistic analysis of the relevant expressions, i.e., the exhibition and exploitation of all linguistic regularities (cf. also Lang, this volume, Section 2.2).

For instance we should not take two English expressions of the form

(8a) x has been opened

(8b) x is open

at face value and attempt to map them onto two concepts like HAS_BEEN_OPENED(x) and IS_OPEN(x) in the knowledge representation, even though it is straightforward enough to write suitable axioms that state the logical relationship between these concepts and model the usage of the English expressions.

Such an approach would ignore the fact that there is a linguistically regular semantic relation between the two expressions, which is a matter of English rather than of any particular conceptual representation. It is, in fact, not only a regular relation between the two expressions in (8), but between classes of English expressions. Relevant parameters are the aspectual classification (cf. Vendler 1967) of the verb *to open* as an accomplishment verb (or perhaps, more specifically, a causative verb), the relation between active and passive, a semantics of English tenses, a semantics for morphological derivation of causative verbs from adjectives (*to open* from *open*, *to clean* from *clean*, *to clear* from *clear*, or *to tighten* from *tight*, *to brighten* from *bright*, etc. - cf. Jespersen 1956, Vol. 6, Ch. 20.5), and perhaps others.

The latter regularity has been captured by Dowty (1979:206f) in a Montague-based approach by the following rules:

S23. If $\alpha \in P_{\text{ADJ}}$, then $F_{23}(\alpha) \in P_{\text{IV}}$, where $F_{23}(\alpha) = \alpha$ **en** if α ends in a non-nasal obstruent, $F_{23}(\alpha) = \alpha$ otherwise.

T23. F_{23} translates into: $\lambda x[\text{BECOME } \alpha'(x)]$

Whether the regularity should actually be stated in this form is another question. We neither have the required generalization concerning the aspectual properties of the English perfect tense nor the generalization concerning the passive voice¹. More importantly, however, the formulation still allows for the extraction of further, more general, regularities concerning the relationship between causative verbs, accomplishment verbs, and transitivity (cf. Bierwisch 1982:71-75 for proposals of the intended kind).

Such additions and amendments are not easily made in the format Dowty uses. The device of lexical templates, as used in the LILOG lexicon, would facilitate such generalizations. The point, to be sure, is not just economy of representation, but rather the extraction of linguistically relevant generalizations about notions like causativity, accomplishment, and transitivity – i.e., purely linguistic generalizations that not only define semantic relations among lexical items of English, but also constrain the set of concepts that can serve as the referents of the relevant classes of expressions.

5 Contextual Concepts

Thus the task of the linguistic analysis is to exhibit as much linguistic regularity as possible. And the role of linguistic regularities is to provide constraints that limit the class of concepts to which a linguistic expression can refer. This is what lexical semantic representation should do.

If this notion about the relation between linguistic expressions and concepts is to work, i.e., if lexical semantics is to constrain the class of potential referents of an expression, the following assumptions about concepts are required:

- The language in which the semantic analysis is formulated must form a subset of the language in which concepts are represented. More specifically, the semantic language must be capable of describing classes of concepts. This need not mean that each of the constants of the semantic language refers to a concept. Concepts may be more complex than semantic constants and they may be more specific: they may contain attributes and attribute values that are not available in the semantic language.
- Concepts must have structured representations so that it can be verified whether a concept satisfies constraints given in terms of semantic constants. The structure may be internal or external (in the sense that different concepts are logically related to each other)².
- There must be concepts available from the knowledge base among which the linguistic expression can choose its referent, and there must be operations on concepts available that can build or modify concepts ad hoc to satisfy the constraints given in the semantic analysis of a linguistic expression.

When concepts from a more or less permanent knowledge base are used in a particular discourse, they are usually further constrained and modified by world knowledge of all varieties as it becomes available in the course of the text or discourse to yield the locally relevant inferentially stable representations. The result of such processes of modification are

¹ This is not to be taken as a criticism of Dowty, who was not discussing the examples in (8) but only the relation between the adjective *open* and the verb *to open*.

² Note that we are here concerned with conceptual representations and not with semantic representations. In the latter case there is a long-standing controversy about the internal or external representation of the relevant structure (cf. Bierwisch 1982:63; Fodor et al. 1980), and I have to leave the question undecided here to what extent arguments carry over to the issue of conceptual representation.

discourse-specific concepts or **contextual concepts**, for short CCs³. We have seen examples above in Sec. 2: the specific concepts related to *red*, *work*, and *water* that we discussed are CCs rather than unmodified abstract concepts that could be formed purely from the constraints that lexical semantic specifications impose.

In the following I want to demonstrate that CCs are of immediate relevance not only for knowledge representation but also already in the course of linguistic analysis. That regular (i.e. first order) objects of reference, as represented in discourse representations, play a role in the resolution of pronominal and full-NP anaphora is uncontroversial. But we can show analogously that concepts, CCs, are required for similar purposes, e.g., for VP anaphora, which is in English often expressed by the form *and so* + AUX + NP (e.g., *and so {does/has/is/will/...} Fred*). Consider the following cases:

- (9a) John loves his mother and so does Fred. [...and Fred loves his mother]
- (9b) John loves his mother and so does Mary. [...and Mary loves her mother]
- (9c) John loves her and so does Fred. [...and Fred loves the same woman as John]

Early linguistic accounts used to assume that VP anaphora can be explained on the basis of a substitution of the VP anaphor for an underlying repeated VP. This could possibly be defended for cases like (9a). But already for (9b) the account must allow for variation in the form of the underlying expression (here the possessive pronoun, which in the antecedent is masculine and in the anaphoric repetition feminine). Cases like (9c) finally demonstrate that any account based on linguistic form is wrong. Here the interpretation of the anaphor contains a specific reference to a particular person and this person must be the same as the one referred to in the antecedent. If Fred loves another woman than John, the sentence is false. Hence the actual reference of *her* in the sentence token is relevant for the interpretation of the VP anaphor.

What the *and-so-do* expressions pick up anaphorically thus can be neither the linguistic expression nor the meaning of the linguistic expression of the antecedent (which could not contain specific references), but must be the interpretation of the antecedent expression with respect to a particular context, i.e., a CC. In (9a) and (9b) this may be the concept of loving ones own mother⁴, which one may argue could be either compositionally derived or culturally entrenched. The concept required in (9c), however, can be neither; it is clearly formed for a purely ad hoc purpose and contains the reference to a specific woman, so that the same CC has obviously no value beyond the context immediately given.

A widespread misuse of the construction exemplified in (9) above is its use as an ambiguity test, the idea being that sentences like those in (9) are unmarked only in case the "meaning" of the antecedent VP is identical to the "meaning" of the anaphor. Hence if there is an ambiguous expression in the antecedent, then the reading that is intended in the antecedent must also be intended in the anaphor. If another reading is intended in the anaphor, we get a marked sentence, often a pun. Thus if we try to read (10) in a way that the first *port* is meant to be a harbour and the second fortified wine, we would get a marked interpretation and thereby proof of the fact that *port* is indeed ambiguous.

³ In Bosch (1985) and other places I have used the term "contextual notions" in order to make a distinction between permanent "concepts" and the more fluctuating "notions". The fact that the structure and content of both is the same is a good reason, however, not to keep up an unnecessary terminological distinction and to speak simply of "concepts" in both cases. The adjective "contextual" can still be added occasionally, when it is important to emphasize that there is a process that leads from the permanent concept to the actual CC.

⁴ Or the concept of loving John's mother - but let us ignore this complication for the moment, since it contributes nothing to the current argument.

Similarly, if we wonder if *hit* is ambiguous between a specification as intentional and non-intentional action (an issue discussed in Lakoff 1970 and Catlin and Catlin 1973), the impossibility of reading sentence (11) with one interpretation in the antecedent and another in the anaphor should prove that this is indeed an ambiguity. Compare

(10) The ship was steering for the port and so was the Captain.

(11) Fred hit the wall and so did Pete.

We have seen already that meaning is not what is at issue in VP anaphora, and hence ambiguity is not at issue either. We are concerned with CCs, and since ambiguous lexical items would also ordinarily refer to different CCs, the assumption of CCs readily explains that ambiguous lexical items show the same effect.

Interestingly, in (11) we can observe that the acceptability of the sentence is subject to assumptions that are made about the context. And this is of course where the actually relevant CCs come from. If we assume that Fred trips over the carpet, hitting his head against the wall (unintentionally), and Pete furiously hits the wall with his fist (intentionally) then indeed the sentence is not a suitable description of the situation. The crucial difference, however, is not intentionality, but rather contextual standards for what may count as the same, i.e., what may in that context be subsumed under the CC which in that context is the referent of *hit*. Suppose Fred trips over the carpet as above and Pete imitates Fred's clumsiness, hence goes through the same moves, but quite intentionally⁵. In that situation the sentence is a true and unmarked description of the two events, despite the difference in intentionality. The difference is just that in this latter context the two events may count as tokens of the same type and hence fall under the same CC. - Thus if we want to have an intuitive test for the identity of CCs, VP anaphora will serve us fine⁶.

6 Operations on Concepts

6.1 Bierwisch's proposal

A pioneering effort in the study of the role of concepts for natural language understanding and their interaction with linguistic semantic specifications was made in Bierwisch (1982). In particular, Bierwisch proposed an account for a set of metonymy phenomena to illustrate his more general points.

First, observe the variation in the interpretation of (12), depending on whether *Faulkner* is intended to refer to the man's pronunciation, his actions, or his literary work.

(12) Faulkner is hard to understand.

Bierwisch says that we are concerned with a **conceptual shift** in the interpretation of the proper name; the interpretation is shifted to different conceptual domains: that of spoken utterances, of actions, or of literary works. Correspondingly, we get a **differentiation** in the respective concepts of understanding: auditory comprehension is another differentiation of the concept of understanding than the understanding of human action or the understanding of works of literature.

⁵ The example is from Catlin and Catlin (1973); cf. also Lakoff (1970), Bosch (1979).

⁶ There are a number of similar test constructions available (cf. Zwicky and Sadock 1975), which support the overall validity of the point (cf. also Bosch 1985).

Neither shifting nor differentiation are a matter of linguistic semantics. In Bierwisch's formalization (cf. below) the lexical semantic representation in both cases only contains variables that are bound by one or the other operator (abstractor or existential quantifier respectively). The application of the abstractor to an expression of a suitable category or the instantiation of the variable constitute the actual processes of shifting and differentiation and yield what I have called above the contextual concept, i.e., the inferentially relevant representation.

Another case of conceptual shift Bierwisch discusses is the interpretation variants that the word *school* exhibits in the following contexts⁷:

- (13)
- a. The school made a major donation.
 - b. The school has a flat roof.
 - c. He enjoys school very much.
 - d. School is one of the pillars of our civilization.

While in (a) we are concerned with an institution, in (b) it's a building, in (c) a certain totality of events, and in (d) the relevant concept is, as it were, a generic variant of the institution concept from (a), the "institution as a principle", as Bierwisch puts it.

The lexical semantic representation (SEM) of the word *school* is the same for all these cases; in Bierwisch's formulation:

- (14) $SEM("school") = \lambda X [PURPOSE\ X\ W]$
 where $W=PROCESSES_OF_LEARNING_AND_TEACHING$

The conceptual shift is brought about by the application of certain functions that map SEM into a contextual concept (or, as Bierwisch calls it, an utterance meaning, *m*); in the case at hand these are the following functions:

- (15)
- a. $\lambda X [INSTITUTION\ X\ \&\ SEM\ X]$
 - b. $\lambda X [BUILDING\ X\ \&\ SEM\ X]$
 - c. $\lambda X [PROCESS\ X\ \&\ SEM\ X]$
 - d. $\lambda X [ENTITY\ X\ \&\ SEM\ X]$

Applying (15a) to (14), e.g., yields the contextual concept

- (16) $\lambda X [INSTITUTION\ X\ \&\ PURPOSE\ X\ W]$
 where $W=PROCESSES_OF_LEARNING_AND_TEACHING$

This proposal shows a way of keeping the linguistic meaning constant and avoiding the postulation of any ambiguity, and still account for the differences in inferential potential between the different occurrences of the expressions in question.

⁷ Actually, Bierwisch discusses the German word *Schule*, not its English equivalent *school*, and the following sentences are my translations of his German examples. This must be mentioned because in German, unlike in English, all four sentences have *Schule* with a definite article. I cannot, however, discuss the implications of this difference in this paper.

6.2 Some elaborations

An aspect of the proposal I want to elaborate in this section is the role of the functions under (15). Bierwisch mentions the fact that words like *parliament*, *opera*, *university*, and many others allow for a very similar, if not identical, set of conceptual shifts. But it would be awkward if we had to explicitly state these functions in the representation of each of these concepts. A more plausible, and quite conventional way of representing these interpretation options in the knowledge representation is in terms of relational and functional attributes, i.e., roles and features, of the concept INSTITUTION (if this is what we call the concept under which all the relevant concepts fall - note that this is, of course, not the concept of "institution" that Bierwisch uses in (15a) and (16)). In other words: we have a general concept of institution and say that each of its sub-concepts or instances has an associated LOCATION, an associated SET_OF_EVENTS, and an idea or PRINCIPLE on which it rests. These three attributes are thus inherited by all concepts that fall under the INSTITUTION concept and allow for the predication that, at least by default, schools as well as parliaments or operas have these attributes in common.

In addition we must state that conceptual shifts from a concept to one of its attributes are regularly possible; very much in analogy to the rhetorical figure of *pars pro toto* we here get another metonymical shift: *concept for attribute*. The fact that there are differences between the shifts that *school* and *opera* allow, is again accounted for quite regularly by the device of multiple inheritance: the concept OPERA is not only subsumed by INSTITUTION, but also by WORK_OF_PERFORMING_ART, and hence has attributes that the school concept has not. This explains the additional shifts that *opera* shares with *symphony*, *drama*, *comedy*, etc. and which are not available for *school*.

Whether indeed the shifting-principle "concept for attribute" is correct is hard to say at the current stage of investigation, both for empirical and for conceptual reasons. We neither know enough about attributes in knowledge representation languages, nor have we looked at sufficiently many cases.

But let us consider some consequences. It may seem at first glance that there is an important difference between the *concept for attribute* and the *pars pro toto* shift: given a part, the corresponding whole of which it forms a part seems uniquely determined; but given a concept it is usually unclear which of its attributes the shift should go for. But the actual difference is smaller than it may look: also a part is not inherently a part of a particular whole and hence does not by itself determine a corresponding whole it is a part of. There is not significantly more determinacy in one case than in the other. If someone is called a "big mouth", the whole which the mouth is a part of could theoretically (but already limited to a 'natural' set of choices) be the face, the head, the body, or, as usually intended, the whole person.

Here as well as in the case of the school example we need additional information to solve what Bierwisch (1982:76, 92f) calls the "selection problem": How do we determine which of the available concepts is to be selected as the actual referent for a linguistic expression in a given context? In the simplest case, we have already interpreted another linguistic expression whose selection restrictions solve the problem. If we read that "The new school is being built in Parks Road", the selection restrictions introduced by the concept of building would determine a concept of school in the sense of school building. But already here one may worry that compositionality of interpretation is endangered: we cannot first interpret each composite expression independently and then compute the interpretation of the whole sentence, but have to take the interpretation of some components into account while we interpret others. But this is nothing out of the ordinary. We may uphold compositionality by simply increasing ambiguity at the lower levels of analysis, i.e., first list all possible interpretations for each component and then see which can be eliminated in the course of their combination. Not a sensible strategy,

and certainly not in the line of psychological realism, but it can preserve compositionality when this is important. The result of the combination may of course not always eliminate ambiguity entirely, and this is where things get more interesting. Because here the extra-sentential context comes into play and we require control strategies for the use of the various knowledge sources involved. In Sec. 7 below, we shall look at one part of this problem.

The selection problem, however, also crops up when several competing concepts are relevant to one interpretation. Formally, we may regard each entity as an instance of one or the other concept or of several concepts simultaneously. In the latter case we get multiple subsumption and multiple inheritance. This can be illustrated quite conveniently with Bierwisch's Faulkner example. The individual Faulkner may be classified in many, perhaps arbitrarily many ways: as an author, the speaker of a language, a rational agent (and, just to please Plato, also as a featherless biped). Each of these concepts have attributes that their instance, Faulkner, may inherit. In particular the AUTHOR concept should provide a role WORKS_OF, the SPEAKER concept should provide an analogous role UTTERANCES_OF and a feature PRONUNCIATION_OF, and the rational AGENT should have associated ACTIONS.

If Faulkner as an instance of the concepts mentioned inherits these attributes, we have all the objects of understanding that are required for the suggested interpretations of sentence (12) above. Whether and to what extent it is possible to mix these variants of the interpretation of Faulkner is a matter also of the tolerance of the available concepts of understanding. A sentence like

(17) His pronunciation as well as his stories are hard to understand.

sounds like a bad pun; but

(18) He doesn't seem to care about others and whether or not they understand him. Nobody understands his stories, nor even his pronunciation.

already seems a good deal more acceptable.

This, I would like to suggest, is one of the points where context plays the major role in the business of the formation and modification of concepts in discourse: in determining what may and what may not count as the same from the point of view introduced by a particular discourse context. Specifically: is there a discourse perspective that allows for the construction of a CC for *understanding* that accepts both stories and pronunciation as arguments, i.e. a discourse perspective for which pronunciation and stories are both subsumed under one particular concept which characterizes the locally relevant objects of understanding? See also Sec. 5 above and Bosch (1990). But it is well beyond the present paper to push these issues any further.

6.3 An extension

Let us now try and see how far we get with Bierwisch's proposal, slightly reformulated in terms of the theory of knowledge representation as above, when we apply it to the notorious case of the colour adjective *red* from Sec. 2.

Within Bierwisch's framework, we are, in (6), clearly not concerned with a conceptual shift in the interpretation of *red*, but rather with a differentiation, i.e., with the modification of a concept of redness that depends on the variation of the concepts of the various objects of which redness is predicated. But is this the correct approach? Do we really want different concepts of red? Are we not rather concerned with the same concept of redness, except that it applies to the various objects in a different manner?

The point is, I believe, that concepts fundamentally do two things:

- they specify certain properties of the objects they comprise, either as inherent properties or as default properties
- they specify certain options for further specifications of these objects.

If we compare the concepts of a tomato and a grapefruit in this respect, we find that both have a default specification for the colour property: red and yellow respectively. Both concepts are supplemented by an axiom that the colour of the unripe fruit is green, otherwise as in the default specification (i.e. red or yellow). Apart from the actual values of the default colour parameter, this is not information one would explicitly represent for the specific cases of tomatoes or grapefruits, but rather higher up in the sort lattice, with the concept of fruit, from where the sorts tomato and grapefruit inherit this information.

The predication of redness of a tomato thus is only an explicit confirmation that the colour default holds, i.e., that the tomato is ripe. But what does the predication of redness do for the grapefruit? It can either be interpreted in a conceptually incoherent manner that leads to the further question of how the grapefruit was turned red. Alternatively, there may be a richer concept of grapefruits that has a second colour attribute, for the colour of the pulp, with the disjunctive specification: either red or white, so that the predication of redness is related to the slot for the colour of the pulp.

If this is the structure according to which we arrive at the different interpretations for red tomato and red grapefruit, then the cause of the difference is exactly where it intuitively ought to be: in different structures of the concepts of tomato and grapefruit. Different empirical knowledge about the two kinds of fruit provides for different answers to the question: In what respect can this fruit be red and what follows from it? Corresponding conceptual representations tell us that for hair, *red* articulates a choice between blonde, brown, black, and red, hence that the attribution of red excludes these alternatives; for wine the set of alternatives is rather red, white, and rosé.

In all these cases *red* articulates a particular choice from a set that is already given with the concept that is to be modified. If the concept is too poorly specified and leaves open what kind of modification is conceptually anticipated, then only a superficial interpretation is possible and inferences like those indicated cannot be drawn.

The type of interaction we have just observed between conceptual representation and explicit linguistic assertion is not restricted to the fairly permanent concepts in the example but can also be found with conceptual representations that are built up and vanish in the course of a discourse. The most general formulation of the phenomenon probably is that in the course of a discourse questions (and often also sets of potential answers) are built up to which subsequent portions of the discourse provide answers, i.e., the classical rhetorical notion of a *quaestio*. Properly worked out, this view leads to a significant reduction of the tasks of linguistic semantics. It makes clear that what is understood from an utterance (i.e., the inferentially stable representation that results from it) is not in the first instance determined by the semantics of the sentences uttered, but is rather the result of a modification the semantics of these sentences brings about with respect to concepts and assumptions that were already available in advance.

7 Semantic Composition and Reference to Concepts

Even if we take a declarative approach to lexical semantics and knowledge representation we still have to ask what information is used when, i.e., ask about the relevant control structures. This is particularly important in view of the fact that much of the relevant information needed in NLU is in principle available from more than one source, though often in a different form

and with different consequences for the result. The point is that a robust system, as the human NLU system clearly is, cannot rely on just one source of information for one purpose nor can it rely on just one strategy of understanding. To put it bluntly: when logic fails or gets too complicated, experience can often help out, and conversely. But logic and experience do not always yield the same solutions for identical problems.

With respect to the problem of the interaction of linguistic knowledge and world knowledge we have been discussing, the question that matters is this: at what point in the course of processing an utterance is world knowledge and conceptual knowledge used? On the assumption of strict compositionality, as briefly hinted above, we should first carry out the linguistic analysis in full and stand by and watch an explosion of spurious ambiguities. And only when the disaster is complete, we should use conceptual and world knowledge and start eliminating the irrelevant ambiguities. The opposite extreme is to first work out, on the basis of conceptual knowledge, domain knowledge, and knowledge of the communication situation, including discourse strategies, etc., what would be a likely thing to be said and then carry out some linguistic analysis for confirmation. Clearly, neither strategy is very attractive as it stands. In the following I want to look into the processing of some types of adjectives in order to show that the question of control strategies may well merit detailed investigation and that there may well be specialized standard control strategies for particular constructions.

Compare the interpretation options for the adjective *red* in the (a) and (b) variants of the following sentences:

(19a) This is a red shirt.

(19b) This shirt is red.

(20a) This is a red tomato.

(20b) This tomato is red.

(21a) This is red hair.

(21b) This hair is red.

(22a) This is red wine.

(22b) This wine is red.

The (a)-sentences straightforwardly subsume a particular object under a particular composite concept, referred to by the phrase *red* *_*. When such a concept is available in the knowledge base it would usually differ from the mere unification of the semantic specifications of the adjective and noun. The latter are only constraints on the referential process that selects the concept from the knowledge base and are not by themselves equivalent to a concept. This difference becomes obvious when we compare (19a) to the other (a)-sentences: red shirts don't seem to be anything over and above what one could predict from their redness and their shirtiness. *Red hair*, *red wine*, and to a possibly lesser degree, *red tomato*, seem to refer to concepts that are much more strongly influenced by our experience of their instances.

The (b)-sentences use the two concepts, red and *_*, independently. *This* *_* is a referential NP and refers to a particular object that is selected from the knowledge base by means of the semantic specifications of the noun *_*. In a zero context such an object is simply an instance of a concept of *_* taken from the knowledge base⁸. The representation of this referent is

⁸ This is probably the most plausible place where stereotypes could play a role: if no contextual specifications are available then such an object is represented as an instance of the stereotype shirt, tomato, etc.

consequently modified by the predication of a concept of red, which, again, is selected independently from the knowledge base. The effect can be seen in the (b)-sentences and in how they differ from the (a)-sentences: in cases like (19) there is no clear difference, because there is no relevant composite concept of red shirts. In (20) - (22) however, the difference between the composite concepts and the independently interpreted expressions with their separate concepts is clearly visible.

When I point, say, to a clay tomato with red glazing, an utterance of (20b) is much more plausible than an utterance of (20a). And an utterance of (21a) in view of hair coloured in bright red is distinctly odd, while an utterance of (21b) would be fine.

Consider further

(23) Red wine is not always red. Often its colour is much closer to brown.

(24) Red hair is not usually red. It is rather very light brown with a touch of orange.

The first sentences in (23) and (24) are obviously not self-contradictory, which they would have to be if they were interpreted fully compositionally, i.e., without the control strategy just explained.

The effects of this control strategy become even clearer in (non-intersective) relative adjectives that are used predicatively. For example,

(25a) This is a good dentist.

(25b) This dentist is good.

The second may be used to say that a particular dentist, one among several present, is a good dart player. The former cannot be used for that purpose, at least not in unmarked intonation⁹. The relative adjective *good* in (25b) can be interpreted relative to whatever you please in the context, but the same adjective attributively used, as in (25a), must be relative to the professional qualities of dentists. In other words: in (25a) we are concerned with the composite concept of a good dentist, which is not the case in (25b).

This hypothesis is further confirmed by non-subjective adjectives like *former*, *alleged*, *future*, *apparent*, etc., that is, adjectives for which it does not hold that

$$a(n) \propto_{\alpha \in \text{ADJ}} _ \in N \text{ is a } _$$

For these adjectives there is no equivalence between "This $_$ is α " and "This is an $\alpha _$ " for the plain reason that an $\alpha _$ is not necessarily a $_$.

Non-subjective adjectives cannot be used in predicative position. Compare

(26a) * This police man is former.

(26b) This is a former police man.

The reason is clear: if, according to our proposed control strategy, predicatively used adjectives are not directly (on the level of compositional semantic processes) applied to the relevant noun, but first refer to a concept, then non-subjective adjectives must fail in this position. They do not have concepts as their interpretation, but only modify concepts. They do not refer to functions from things into the truth-values, but to functions that take concepts as

⁹ The reading becomes possible when we assume that *good* is accented and there are several dart-playing dentists around, not all of them good dart players.

arguments and yield concepts as values. But such second order concepts cannot be the referents of predicate expressions, because they cannot be applied to the referents of subject terms. Non-subjective adjectives can however be used attributively, because in that position they can apply to concepts.

We may conclude from these observations that there is something like a **standard control structure** in the understanding of simple sentences, which provides for a referential interpretation, i.e., an anchoring in the discourse representation or knowledge representation, of the two major constituents, in the simplest case of the main functor and the main argument of the sentence, before any composition of semantic specifications between the two takes place, i.e., before the function is actually applied to the argument.

Below the sentential level however, i.e., in our examples within the NP, semantic composition takes place first and only then is there a referential interpretation, i.e., reference to a concept or object referent¹⁰.

For the mapping of lexical semantic specifications onto concepts these control facts have as their immediate consequence that lexical items that figure as independently referential, such as predicative adjectives, directly refer to concepts as they are available from the discourse representation or the more permanent knowledge representation. However, since non-subjective adjectives have no concepts as their interpretation, the interpretation of, say, (19a) must fail. For the same reason of independence of interpretation, the concepts of red in (19b), (21b), (22b), (23), (24), and (25b) are not influenced by their respective arguments, while the attributive adjectives have no independent interpretation but select a composite concept which they determine together with their noun. Hence there are no contradictions in (23) and (24), and there are different interpretation options for the (a)- and (b)-sentences in (19)-(22) and (25).

8 Conclusion

The central point of this paper is to try to give a direction to work on the interaction of linguistic analysis and knowledge representation in knowledge-based NL Systems. I have tried to argue and to demonstrate that without a full linguistic analysis there is little hope that we shall ever have reasonably general and portable language modules in NL systems. It has also become clear, I hope, that this is not a trivial task but requires a decent amount of empirical research for many years to come. But the linguistic research required is not isolated research in pure linguistics, but close cooperation with work on knowledge representation and – although this is a point I have not argued for – psychological work on conceptual systems, is imperative.

The most difficult problem to overcome, I believe, is that the most generally held belief in the scientific community with respect to our problem is that the distinction between linguistic and conceptual facts is arbitrary and hence not a proper research question, but a matter of pragmatic decisions. It is this belief more than anything else that inhibits further progress of the kind Brachman found lacking.

¹⁰ Such control strategies are closely connected with syntax, in particular word order, but also with intonation. The data in the text about the interpretation of adjectives are further strengthened by data from French, which, by and large only allows postposed attributive adjectives, which are related to predicative adjectives pretty much the same way as in English. In cases however, where a separate interpretation of adjective and noun must be prevented, because the intended composite concept cannot be regularly composed by a modification of the noun concept by the adjective concept, the attributive adjective is preposed. Often in these constructions the adjectives assume another meaning than they usually have. Compare, e.g., *ancien roi* vs *roi ancien*, *belle femme* vs *femme belle*, *brave homme* vs *homme brave*, etc. (cf. Chevalier et al., 1964). Also in accordance with our above observation there is no possibility of accenting preposed adjectives in French.

References

- Bierwisch, M. (1982): Semantische und konzeptuelle Repräsentation lexikalischer Einheiten. In: R. Ruzicka and W. Motsch (eds.): Untersuchungen zur Semantik. Akademie-Verlag, Berlin, pp. 61-99
- Bosch, P. (1979): Vagueness, ambiguity, and all the rest. In: M. van der Velde, W. Vandeweghe (eds.): Sprachstruktur. (= Akten des 13. Linguistischen Kolloquiums, Gent 1978). Niemeyer, Tübingen. pp. 9-19
- Bosch, P. (1983): Agreement and Anaphora. Academic Press. London, New York
- Bosch, P. (1985): Lexical meaning contextualized. In: G. Hoppenbrouwers, A. Weijters, P. Seuren (eds.): Meaning and the Lexicon. Foris, Dordrecht. pp. 251-258
- Bosch, P. (1990): Indexicality and representation. In: R. Studer (ed.): Natural Language and Logic. Springer-Verlag, Berlin, Heidelberg, New York. pp. 50-61
- Bosch, P. and Geurts, B. (1990): Processing definite NPs. *Rivista di Linguistica* 2:177-199
- Brachman, R.J. (1990): The future of knowledge representation. In: Proceeding of the AAAI 1990. MIT Press, Cambridge, MA. pp. 1082-1092
- Catlin, J.-C. and Catlin, J. (1973): Intentionality: a source of ambiguity in English? *Linguistic Inquiry* 3: 504-508
- Chevalier, J.-C., et al. (1964): *Grammaire Larousse du Français Contemporain*, Librairie Larousse, Paris
- Dowty, D.R. (1979): Word Meaning in Montague Grammar. Reidel, Dordrecht
- Fodor, J.A., Garrett, M.F., Walker, E.C.T., Parkes, C.H. (1980): Against definitions. *Cognition* 8: 263-367
- Jespersen, O. (1956): *A Modern English Grammar on Historical Principles*. Allen and Unwin, London/Munksgaard, Copenhagen
- Lakoff, G. (1970): A note on vagueness and ambiguity. *Linguistic Inquiry* 1: 357-359
- Lang, E. (1990): Primary perceptual space and inherent proportion schema. *Journal of Semantics* 7:121-141
- Rommetveit, R. (1986): Meaning, Context, and Control. Discussion Paper at the ESF Workshop in Zurich, Sept. 1986
- Vendler, Z. (1967): Verbs and times. In: Z. Vendler: *Linguistics in Philosophy*. Cornell University Press, Ithaca, NY. pp. 97-121
- Winograd, T. and Fernando, F. (1986): *Understanding Computers and Cognition*. Ablex, Norwood, NJ
- Zwicky, A.M. and Sadock, J.M. (1975): Ambiguity tests and how to fail them. In: J.P. Kimball (ed.): *Syntax and Semantics*, Vol. 4. Academic Press, New York. pp. 1-36