

A model for the interaction of lexical and non-lexical knowledge in the determination of word meaning

Peter Gerstl

*IBM Germany, Scientific Center
Institute for Knowledge Based Systems
Schloßstr. 70
7000 Stuttgart 1
e-mail: gerstl @ ds0lilog.bitnet*

Abstract

The lexicon of a natural language understanding system that is not restricted to one single application but should be adaptable to a whole range of different tasks has to provide a flexible mechanism for the determination of word meaning. The reason for such a mechanism is the semantic *variability* of words, i.e. their potential to denote different things in different contexts. The goal of our project is a model that makes these phenomena explicit. We approach this goal by defining *word meaning* as a complex function resulting from the interaction of processes operating on *knowledge elements*. In the following we characterize the range of phenomena our model is intended to describe and give an outline of the way in which the interpretation process may determine the referential potential of words by the integration and evaluation of a variety of factors.

1 Introduction

A system with the capability of natural language understanding typically relies on knowledge about a restricted domain of application. For example, as a natural language component of an information system, it needs to be able to identify the relevant linguistic patterns. In case of an information system for flight scheduling words such as “plane”, “departure”, “late”, . . . will typically be more relevant than for example: “palm”, “this-tle”, “pine” which might be appropriate for a different domain. In any event, there will be a whole range of words that are commonly used in conversation and, thus, are independent from the choice of a specific domain. It is therefore desirable to have a multi-level architecture which can be adapted to different domains without being forced to redesign the whole system. A text understanding system based on this kind of architecture would provide the kernel functionality that allows it to couple principles and mechanisms not immediately dependent on the domain a specific implementation of the system will be used for. The main problem of such a modular architecture is how and where to draw the boundary between domain-independent and domain-specific knowledge. There are at least two more reasons which motivate a domain-oriented design strategy:

With regard to knowledge representation, the history in artificial intelligence research has lead from early enthusiastic plans of ‘general problem solving capabilities’ to more realistic applications of expert systems. One reason was the huge amount of data that would have to be represented together with a large set of regularities introducing a level of complexity which could not be handled in a realistic manner by the systems currently

available. Another problem is the inconsistency of data that would necessarily arise once a lot of different and sometimes conflicting information had to be integrated into a single knowledge base. Under this perspective task- and domain-orientation is a matter of rendering the knowledge base manageable and to allow reasoning processes to draw meaningful inferences on the basis of consistent data.

The second argument in favour of a domain-oriented design strategy comes from the area of lexical semantics. It is a well known fact that the meaning of a word depends on a multitude of contextual influences. In a very broad notion of context the task and the domain of a text understanding system may be considered a part of the context that licences an effective restriction of the 'semantic scope' of a single word. This again is mainly an argument of tractability which in this case helps to minimize the amount of lexical information needed. In our example it is a natural design decision to assume that the lexicon of a language understanding system as part of an information system about flight schedules does not have to account for the 'plant'-reading of "plane".

2 The variability of words

The way in which a word might contribute to the determination of the meaning of linguistic expressions is indeterminate in different ways. W. Labov calls the semantic potential of words enabling them to constitute various links between linguistic expressions and elements of the domain (semantic) *variability*. It is this potential which is responsible for the already mentioned context dependence of word meaning. An important goal of our model is to classify types of variability according to a set of more or less specific properties. The questions guiding this classification are: Which kind of representation does the variability affect; by which means are the variants related and, how can the referential potential be restricted in order to single out the intended meaning? In the linguistic tradition, these variability phenomena fall into one of four classes which we will sketch out below.

2.1 Morphological ambiguity

This class represents cases of identical *surface representations* of words extracted from a discourse. Depending on the kind of representation in which the natural language input is encoded (i.e. orthographic, phonetic, ... form) cases of *homography* or *homophony* may belong to this class or not. *Homonymy* as a specific instance of morphological ambiguity results from the identity of lexical base forms. In a lexicon using orthographic representations as is normally the case in dictionaries (singular nouns, infinitive forms of verbs, ...) there are for example two homonymous entries for "firm" (the adjective and the noun variant). In addition, morphological ambiguity captures the more general case of an inflected form that is identical to a base form or to another inflected form. An example is the occurrence of "saw" which depending on the context can be understood as noun, as base form of the verb "saw" or as inflected form derived from the base form "see".

It is a notorious problem in lexical semantics [Kooij 71] to justify the distinction between coincidental identity of forms (morphological ambiguity) and semantic variants of a single lexical item (chapter 2.2). In our approach it depends on the purpose of the system and, thus, is a design decision comparable to modelling conventions for the domain knowledge. Identical basic word forms which are in no relevant and transparent

way related to the representation of knowledge about the domain will be represented by different lexical items. The question whether two identical forms ‘collapse’ into a single entry is then directed by the choice of the domain and the task analogous to the way in which drawing a boundary between elements of the domain is motivated. Defining two word forms as being homonymous yields the consequence that once the occurrence in a discourse has been morphologically identified and thus mapped onto the corresponding lexical item it is no more possible to skip to a different homonymous variant. Words which shall not expose this behaviour should not be modelled as homonymous but as semantic variants of a single lexical item.

2.2 Polysemy and polyfunctionality

In the preceding chapter we outlined the situation where two lexical items realize the same word form. The potential of semantic variation encoded in a single lexical item is known as *polysemy*. It remains in effect once the appropriate lexical item has been identified by means of morphological processes. A special case of polysemy is what [Weber 74] calls *polyfunctionality* referring to the situation where two variants of a lexical item belong to different syntactic categories. Polyfunctionality occurs very frequently since many lexical items allow identical realizations which belong to different syntactic categories being related by means of conversion or other morphological processes which do not affect the word stem. Examples are the nominal and verbal reading of “point” and the variants of “clean” which are categorized as adjective, adverb and verb. The comparison with variants in a dictionary is not as effective as it was in chapter 2.1 because the entries in a dictionary tend to conflate phenomena we call polysemy and cases of variability outlined in chapter 2.3.

2.3 Metonymy and change of semantic type

In the previous chapter we mentioned that a polyfunctional expression can be analyzed as the realization of different categories in different contexts. The difference in semantic potential that arises from this variability sometimes not only involves the transition to a different semantic type in the sense of Montague grammar but it may be paralleled by a more or less extensive shift in conceptual interpretation (cf. the one-place versus two-place predicate reading of “drink”). Apart from ambiguities which are reflected by morphosyntactic properties of a lexical item (eg. its argument structure) there are instances of semantic variation allowing to change the interpretation of a class of linguistic items in a systematic manner. The crucial question is if this class should be characterized by lexical information or on the basis of regularities found in the domain. *Metonymy* is a specific instance of this type of variability where the different readings are related by elements of a set of fundamental relationships such as ‘part-whole’, ‘cause-effect’, etc. [Nunberg 78] investigates more general mechanisms that licence the use of a word in place of another in cases where both of them are related by means of a context-specific relation. The phenomena range from cases of systematic correspondence such as in the ‘newspaper-example’¹ to more idiosyncratic ones as the famous ‘ham-sandwich-case’². As Nunberg notes these are not cases of linguistic ambiguity because pointing to the sandwich would

¹The word “newspaper” can be used to refer either to the publisher or to the publication.

²A waiter might apply the expression “the ham sandwich is sitting at table 20” in order to identify a unique guest who ordered a ham sandwich.

serve the same purpose as the utterance of the complex phrase. Nevertheless, it is not clear if the relations are considered as instances of lexical or encyclopaedic knowledge.

An example similar to the 'newspaper-case' is the potential of words such as "school", "opera", ... to select one of the alternative meaning variants 'building', 'process', 'institution', etc. The approach outlined in [Bierwisch 82] derives this potential from systematic relationships between *concepts* representing entities of the domain. The conceptual knowledge about social institutions has to provide the background information that "the parliament is at the end of the street" is semantically well-formed though "?the government is at the end of the street" is not. Since the lexical specifications of "parliament" and "government" cannot account for the fact that it is naturally assumed that the former can be associated with a specific building whereas a similar assignment is not possible for the latter. A comparable argument may be found for the 'substance'-reading of words naming trees. The ill-formedness of "?this table is made of plane" in contrast to "this table is made of oak" results from the non-verifiability in the common-sense model of the domain. If an expert would affirm that it is quite common to use the wood from palm trees for the construction of tables we would probably change our model of the domain and licence the acceptability of the first sentence. This example is different from the 'school'/'newspaper'-cases since in addition to the conceptual shift it involves a modification of semantic properties of the underlying lexical items. This in turn is an argument in favour of a lexicalist position which would classify this type of variability as cases of polysemy. [Pustejovsky 90] shows that a lexicalist approach to event structure allows to systematically characterize a whole range of type-shifting phenomena together with their consequences with respect to well-formedness conditions.

As a consequence of seeking the portability of domain specific knowledge we follow the lines of Bierwisch in distinguishing lexical and conceptual information. Yet we do not reject the lexicalist position since we consider semantic type-shifting effects as driven by regularities of the conceptual structure. In the following section, we generalize this position to a systematic distinction between *linguistic* and *non-linguistic knowledge*. This allows us to keep variants introduced by linguistic ambiguities systematically apart from those cases we classify as non-linguistic variations.

Following [Binnick 70] we define polysemous variants of a word as those cases of variability which are (at least in principle) distinguishable on the basis of linguistic properties of the corresponding lexical item. Polysemous variants of a lexical item thus differ in at least one morphosyntactic or semantic property. For non-linguistic variants introduced by means of metonymy or type-shifting linguistic properties of a lexical item do not help to identify the intended reading because the variants have exactly the same linguistic properties. This situation calls for the disambiguation potential of contextual information in order to reduce the 'semantic scope'³.

2.4 Contextual relativity

Vagueness and *indexicality* also belong to the class of variability phenomena. They are usually associated with specific groups of linguistic expressions (gradable predicates in case of vagueness and deictic expressions in case of indexicality). In contrast to the effects introduced so far, in these cases the class of potential referents cannot simply

³We consider the 'semantic scope' of a word as the possible range of interpretation implied by the literal use of a word. The more general notion of 'referential potential' additionally accounts for cases of conceptual shift as in the 'ham-sandwich' example.

be characterized by an enumeration of alternatives. As [Pinkal 80] points out it is an inherent property of vague predicates to provide a 'grey area' where the decision whether the predicate is applicable or not depends on the discourse context. It is even impossible to precisely delimit the area of positive or negative applicability. Following the lines of [Bosch 83] we do not consider vagueness as an isolated semantic property of a specific class of words but as an instance of the more general notion of *context-dependence*⁴. Since this is an aspect of the referential potential of words our model has to cover these phenomena as well.

Indexicality is the potential of deictic expressions to select their meaning by exploiting peculiarities of the discourse situation. It is similar to vagueness since the implied referential indeterminacy cannot be resolved independently from the specific discourse context. Yet, even deictic expressions are not immune to other types of variability. For example, the pronoun "I" may be used to refer to an entity which is somehow related to the speaker in a certain discourse situation. An example is the utterance of "I am over there" with the speaker pointing to a desk. The expression "I" in this case can be used to refer to the place where the speaker usually works. This is an example of a systematic shift in meaning motivated by a conceptual relation. It thus belongs to the phenomena described in the previous chapter.

Another instance of context relativity occurs in cases of *privative opposition*. This relativity results from the lack of semantic information for a specific word which could be provided by the use of a different word. According to [Zwicky/Sadock 75] "dog" is ambiguous between the readings 'male dog' and 'female dog' because it can be forced to provide both readings in sentences like "that is a dog, but it isn't a dog". In contrast to "that is a lion, but it isn't a lion" it seems that a meaningful interpretation can be found for the former (the one which forces the selection of different variants for both occurrences of "dog") whereas the latter leads to a contradiction. The choice of a variant could be forced by the use of "bitch" instead of dog which is not possible for "lion" since there is not regular lexical specification for something like "lioness". This illustrates the fact that lack of semantic information for a lexical item can under certain circumstances yield the same effect as a disjunction of alternative readings. This may occur whenever the semantic 'gap' can be filled by one of a small set of possible alternatives.

3 A classification of knowledge types

In order to have a precise representational basis for our model of word meaning, this chapter is intended to introduce the basic notions used to classify the relevant phenomena. The distinction between linguistic and non-linguistic knowledge mentioned in the preceding section constitutes the methodical basis of our model. Up to a certain degree, this distinction allows an independent examination of properties characteristic for only one type of knowledge. By assuming this distinction we do not claim that linguistic and non-linguistic knowledge are in a way fundamentally different. We use this distinction as a methodological tool that makes it possible to isolate certain aspects of word meaning not directly involving the whole range of both types of knowledge. In the course of stepwise extending the complexity of interrelations between linguistic and non-linguistic knowledge we will have to carefully analyze the tenability of this distinction.

⁴In general the notion of context dependence applies to referential expressions such as definite nominal phrases. We will restrict our attention to cases of context-dependence which apply to single words.

We account for possible similarities between both types of knowledge by using the same formalism for the representation of linguistic and non-linguistic knowledge. It is a variant of order-sorted predicate logic [Nebel/Smolka 89] which combines properties of the KL-ONE family of knowledge representation languages with properties of feature based unification grammars such as HPSG [Pollard/Sag 87]. In order to concentrate on the description of our model we will not go into the details of our formalism here⁵.

The central components of our model are the *lexicon* on the linguistic side and the *ontology* on the non-linguistic side. The lexicon and the ontology provide the 'basic building blocks' of linguistic and non-linguistic knowledge respectively. The elements of the lexicon are called *categories*; the elements of the ontology *concepts*. It is important to note that the lexicon in our model integrates specifications of syntactic categories (i.e. N, V, ..., N', ..., AP, ...) with lexical items.

The formal means for the description of categories and concepts are *sorts* which are related by means of *attributes* and *rules*. Attributes may be used to express characteristic properties or relationships motivating the choice of a specific distinction between sorts. Rules on the other hand are not considered as tools for the description of inherent and permanent properties but as representations of regularities which might arise under certain circumstances. Apart from that, the collection of attributive characterizations has to be consistent. That is not necessarily required for the system of rules as a whole. The fundamental organizational principle of subsumption relates categories and concepts are licensing the inheritance of attributes between sorts. The subsumption order does only apply between elements of one and the same type of knowledge. The notation used for the description of sorts is a feature-logic as in [Shieber 86] for categories and a simple relational notation for concepts. We represent the fact that *A* subsumes *B* as $A \sqsubseteq B$. Rules of grammar and rules of inference are represented by using simple predicate logic notation. Sorts, attributes and rules will be called *knowledge elements*. All of them put together constitute the *knowledge base* of our system.

According to our argument in favour of a design strategy specifically tailored to the domain and the task the system is intended for, the structure of the ontology must be covered by an appropriate theory about entities of the domain, their inherent properties and the diverse aspects in which they are related. We reiterate this methodological claim here since a similar argument can be applied to the organization of the lexicon. The typical task of a text understanding system is to facilitate the *analysis* and *production* of textual input. It depends on the capabilities required whether certain aspects of this task involve restrictions on the set of relevant linguistic phenomena⁶. The choice of lexical items depends on the domain at issue since the lexical inventory should cover at least the range of non-linguistic phenomena represented in the non-linguistic component of the system.

Categorial knowledge provided by the lexicon together with the *rules of grammar* constitutes the descriptonal apparatus for the classification of *expressions*. On the one hand expressions serve as input for linguistic processing and on the other hand they represent sequential patterns of written or spoken language. This intermediate status makes them

⁵Most of our assumptions about the representation of linguistic and non-linguistic knowledge are based on experiences gained from work in the LILOG-project at IBM Stuttgart. A description of formalisms and methods applied in this project can be found in [Geurts 90].

⁶One can for example reduce the computational complexity by limiting the relevant sentence-level constructions to simple propositional clauses if the system is not meant to deal with other types of modality. Even if this argument sounds quite trivial the determination of a set of requirements for the linguistic component are as important as they are for the representation of domain knowledge.

elements of *discourse knowledge*. Expressions belong to the type of knowledge which serves as a kind of record for the registration of linguistic interactions together with their spatio-temporal specifications. It directly corresponds to *episodic* knowledge on the non-linguistic side. Episodic knowledge has the same intermediate status as discourse knowledge since on the one hand it serves as a record of 'statements' and other 'experiences' with respect to the domain and on the other hand it is used to characterize entities from the domain as *individuals* on the basis of conceptual knowledge. Conceptual knowledge combines the structural information conveyed by the ontology with the additional information expressed by *rules of inference*. Individuals which result from the processing of certain linguistic expressions are called *referents* since they are open to further reference by linguistic means. The figure below shows the whole classification assumed as the basis of our model.

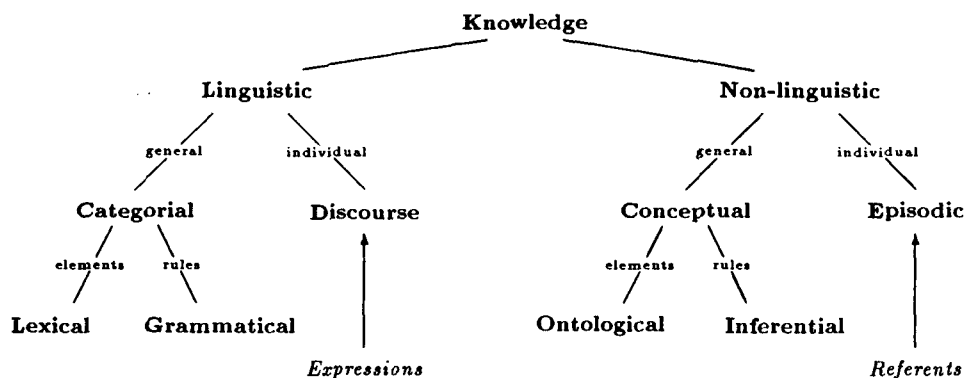


Figure 1: The classification of knowledge types.

4 Word meaning

The task of our model is an approach to the various aspects of word meaning which are responsible for the variability effects described in section 2. In order to reduce the complexity of the linguistic domain we restrict the relevant linguistic expressions to those representing simple word forms which cannot be further decomposed by morphological processes other than inflection. As a consequence, the *granularity* for the representation of linguistic knowledge treats basic morphemes as minimal elements. Another consequence is the width of the temporal grid which specifies the minimal 'temporal distance' between elements of discourse knowledge. We introduce temporal *indices*, allowing to subdivide the linguistic input into a chain of word-level segments. Each index uniquely identifies a gap between two words and directly corresponds to a set of intermediate results in the course of processing the input. These results are what we call the *context*. Formally speaking, a context is a set of *factors* marked by a temporal index specific to a certain stage of processing at which the system is observed. Factors are functions between knowledge elements or between instances thereof. They are elements of specific contexts and therefore differ from attributes because their existence is strictly tied to a certain stage of processing associated with a temporal index. As a result of iterated forwarding a factor may remain applicable during a sequence of processing steps. Factors may be classified according

to their origin. Factors which are directly derived from knowledge elements are called *primitive factors*. Depending on the type of knowledge elements involved we distinguish two modes of origin. Primitive factors can be ...

- *selected* as instances of attributes or
- *established* by the application of rules.

Complex factors are derived from primitive ones by one of the following operations:

- the *restriction* of the domain and/or range of a primitive factor
- the application of set-theoretical operations on primitive factors
- the *functional composition* of primitive factors

We present this classification because our analysis of word meaning crucially depends on the notion of contextual factors. It is the main goal of our project to reconstruct word meaning as the result of the interaction of processes which cope with an effective integration of various linguistic and non-linguistic factors primitive and complex in nature. Since we investigate word meaning under the aspect of the potential of words to refer to representations of entities of the domain, word meaning in our terminology is a complex factor which links elements from discourse knowledge (expressions representing words) to elements from episodic knowledge (referents). The temporary status of factors is responsible for the fact that for the identification of the referential meaning of a word the whole context has to be taken into account. The linguistic notion of 'word meaning' therefore derives from the analysis of subsets of factors that result from the intersection of contexts present in a sufficiently large group of different uses of the same word.

4.1 The constituents of reference

In order to characterize the interrelation between factors introducing variability effects (*productive factors*) and those limiting the 'semantic search space' (*restrictive factors*) we need to examine the way in which word meaning can be decomposed into a small number of factors⁷. A segmentation of the interpretation process according to our classification of knowledge types leads to three *components* which by application of functional composition constitute word meaning. Since components are derived by functional decomposition of a complex factor (word meaning) they are factors as well. Components may be further analyzed as the results of set-theoretic operations on basic factors some of which limit and some of which extend the 'semantic scope' of a word form⁸. Factors extending the scope of interpretation are directly responsible for the variability effects described in section 2. Factors constraining the scope of interpretation are the topic of chapter 4.2. The following list introduces the three components of meaning together with examples of the relevant productive factors. An interesting criterion for the classification of productive factors is whether they are established by the application of rules or selected from attributes between knowledge elements.

⁷We do not assume contexts to be finite but our approach relies on the fact that a finite subset of the context suffices to describe word meaning precisely enough to demonstrate the requirements a system with reasonable disambiguation capabilities has to fulfil.

⁸In fact the same function may in one stage of the interpretation process serve as productive factor and in another as a restrictive factor. Thus, the property of productivity or restrictivity cannot definitively assigned to specific factors. More precisely speaking, it is property of factors dependent on the current stage of processing represented by the temporal index.

item by means of the semantic specification.

(2b) The non-linguistic constituent of lexical meaning

conc₇₅	SITUATION	\sqsubseteq	...
	time	:	conc₅ spatio-temporal properties
	location	:	conc₃
conc₂₃	PERCEPTION	\sqsubseteq	SITUATION
	actor	:	conc₈ The filler of the actor role visually perceives
	theme	:	conc₁₃ the filler of the theme role
	instrument	:	conc₃₅ by using the filler of the instrument role
conc₃₄	REALIZATION	\sqsubseteq	SITUATION
	actor	:	conc₈ The filler of the actor role
	proposition	:	conc₁₉ realizes the filler of the proposition role
conc₃₉	VISTING_SITUATION	\sqsubseteq	SITUATION
	visitor	:	conc₂ The filler of the visitor role
	visited	:	conc₄ ...

(3) Individuation

This last factor in the chain of meaning components becomes established by the application of rules of inference. It maps concepts onto referents. The productive factor of individuation is *referentiality* extending the range of possible referents a concept can be individuated to. Referentiality here serves as cover term for the phenomena described in chapter 2.4 together with cases of conceptual variation which qualify as ‘ad-hoc-anaphora’ because they succeed to identify a unique referent in a specific context but cannot be characterized as instances of general principles guiding a shift in conceptual interpretation¹⁰.

Additional parameters of the discourse situation (the time and location of the utterance as well as a proposition r_1) allow to *establish* an individuation which maps **conc₃₄** onto a referent r_2 with the following properties:

$$\begin{aligned}
 \text{actor}(r_2) &= r_{\text{speaker}} \wedge \\
 \text{proposition}(r_2) &= r_1 \wedge \\
 \text{location}(r_2) &= r_3 \wedge r_3 \subseteq s_{\text{discourse}} \wedge \\
 \text{time}(r_2) &= r_4 \wedge r_4 < t_{\text{discourse}}
 \end{aligned}$$

The three components of word meaning can be considered intermediate steps of the interpretation process. They may be analyzed and described in isolation since their interaction results from the way in which the range of the preceding component fits to the domain of the following. The task of the interpretation process on this background is to find a ‘path’ leading from an expression to an individual which under consideration of all the available contextual factors qualifies as plausible candidate for the referential meaning of the expression.

4.2 How the semantic scope can be restricted

The crucial question now is how the diverse components interact in order to reduce the range of word meaning by the exclusion of implausible variants. Here we pick out three

¹⁰Nunberg’s ham-sandwich is an example instance of this kind of context specific ad-hoc-anaphora.

example groups of factors which in a typical situation may support the reductive factors of meaning components and such help to reduce the referential potential of a word.

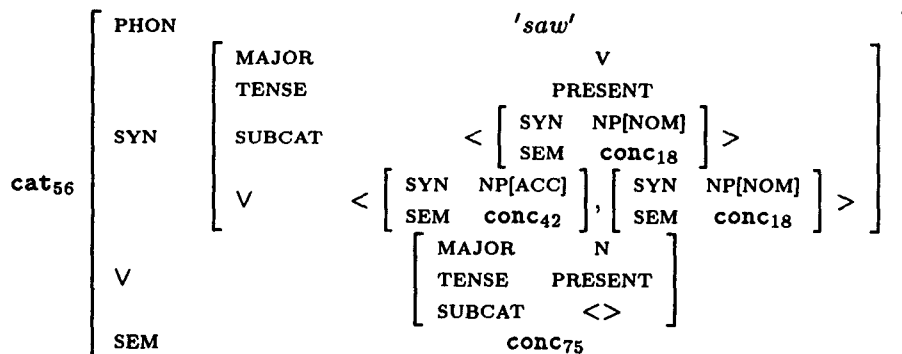
(1) Word-specific factors

The first group are factors which result from structural relationships expressed by morphosyntactic attributes and rules of grammar. As we mentioned in chapter 2.3 these factors only affect variabilities which are introduced by the linguistic part of our knowledge base. Factors of this group thus may help to resolve cases of morphological ambiguity, polysemy or polyfunctionality but they have no effect on variants that result from metonymy, change of semantic type or other instances of contextual relativity.

Consider the following part of discourse:

“I tried to find a possibility to escape. Then I saw a hole in the fence.”

We'll give a sketch of an analysis of the meaning of “saw” in this example on the basis of the knowledge elements introduced in the previous section. The rules of grammar suppress the nominal reading of “saw” since the principles of \bar{X} -syntax require the constituent “a hole in the wall” to be 'absorbed'. Morphological rules do not support the disambiguation process. On the contrary, their productive potential causes the introduction of the variant derived from the base form “see”. The variant cat_{56} is ruled out because of incompatibilities between its subcategorization frame and the syntactic environment of “saw” in our example.



The intransitive polysemous variant fails because of the same reasons as the nominal homonymous variant. The transitive reading of cat_{56} would force an optional prepositional argument to be headed by “into”¹¹. Thus, the polysemous variant $conc_{42}$ can be singled out purely on the basis of word-specific factors if the rules of grammar do not account for the adjunction of a locative PP with the head “in”. In case the grammar licences the existence of a prepositional adjunct, our model of the domain would have to contribute the restrictive factor that the concept associated with “the fence” does not fit with conditions on ‘sawing’-events. Since the reductive influence of word-specific factors ends with the selection of polysemous variants we cannot expect a further restriction of the ‘semantic scope’ without additionally considering other types of factors.

¹¹In order to simplify the example this alternative does not occur in the feature structure of cat_{56} .

(2) Selectional restrictions

A different group of factors belongs to the semantic level¹² of our model. Factors in this group neither are clear instances of linguistic regularities nor of non-linguistic ones. They are partially linguistic and partially non-linguistic in nature and therefore considered as complex factors derived by the integration of elements from both types of knowledge. They may help to reduce variabilities affecting categorization or lexical meaning. Yet, like word-specific factors they do not constrain contextual relativity.

Consider the following part of discourse:

"I saw { a piece of wood
a cup of coffee } in the bathroom."

The semantic specification for the internal argument of "saw" leads to a concept **conc₄₂** representing a class of entities which qualify as fillers of the corresponding role in the conceptual representation of the 'sawing'-event.

conc₇₅	SAWING	\sqsubseteq	SITUATION	
	actor	:	conc₁₈	for example: a human being
	object	:	conc₄₂	for example: a concrete object
	instrument	:	conc₂₉	for example: a set of tools

The compatibility between the semantic specification of the internal argument of the two polysemous readings of "saw" and the type specification of a role belonging to the corresponding SEM-value account for the existence of *selectional restrictions*. The conceptual representation of "a piece of wood" must be compatible with **conc₄₂** in order to establish the lexical meaning leading to the concept **SAWING_EVENT**. In the case of "this hand-saw saws well" the external argument would be because of requirements on fillers of conceptual roles have to be mapped on the instrument role of **conc₇₅**. The situation is more tricky if we compare the instances of the external argument of **cat₂₃** in the following example:

- (1) The policeman saw an accident.
- (2) *The ball saw an accident.
- (3) The automatic traffic control camera saw an accident.
- (4) ?The morning saw an accident.

An interesting aspect of this phenomena is that selectional restrictions may be cancelled by contextual factors or by means of rhetoric devices. The example sentences show how difficult it might be to identify an obligatory set of selectional restrictions. Comparing (1) with (2) suggests being an instance of the concept **PERSON** as a reasonable choice. Example (3) however shows that the critical property is something like 'having an optical sensing mechanism capable of detecting objects'. Sentence (4) might imply a metaphorical interpretation in spite of its apparent semantic illformedness. This again yields an argument in favour of a domain-driven design strategy for the semantic level linking between categories and concepts.

¹²The semantic level is the 'interface' between linguistic and non-linguistic knowledge represented by the SEM values in lexical items together with a set of rules which attune semantic specifications of argument structure to attributes of the corresponding conceptual definitions.

(3) The set of possible referents

The last group of factors exemplified here are the only means available to reduce the semantic scope resulting from contextual relativity. As we saw in chapter 2.4 contextual relativity is a fundamental property all referential expressions have in common. Since the 'semantic scope' introduced by variabilities of this type cannot be subdivided into a set of alternative readings neither lexical nor conceptual information does help to restrict the range of individuation.

The only way out of this dilemma is to derive a set of *possible referents* from knowledge about the domain and from information occurring in the preceding part of discourse. The latter calls for an investigation of discourse properties on the basis of pragmatic devices such as the Gricean conversation principles. *Bridging phenomena*¹³ as generalizations of anaphoric binding are promising candidates for an approach to the determination of possible referents. C. Sidner emphasizes: "anaphor interpretation can be studied as a computational process that uses the already existing specification of a noun phrase to find the specification of an anaphor" [Sidner 83, p.269]. The actual limits of a set of possible referents thus very much depend on the inferential capabilities of our system to reconstruct the conceptual relationships underlying text coherence. The notion of *focus* presented in the work of Sidner certainly plays a crucial role in the reduction of the computational complexity a computation of all possible bindings would involve if realistic discourse situations were to be considered.

5 Conclusion

Up to now we merely picked a collection of phenomena with an impact on the referential potential of words ranging from lexical properties to 'genuine' discourse phenomena. We presented a moderately general framework designed in a way that each of these phenomena has a place to fit in. The example analysis of "saw" illustrated the interaction of elements of our model and thereby pointed to problematic aspects that yet have to be resolved. A crucial problem is to distinguish between linguistic aspects of lexical meaning which the lexicon has to account for from non-linguistic aspects which derive from relations in conceptual structure. The next step in the evaluation of our model requires the determination of criteria which help to keep linguistic and non-linguistic aspects of the semantic level apart. The ultimate goal of this distinction is to narrow down the flow of information that passes through the semantic 'interface' between linguistic and non-linguistic knowledge. If this strategy succeeds it should be possible to adapt the conceptual knowledge to different domains of application not affecting linguistic knowledge as the basis for capabilities of discourse understanding.

References

- [Bierwisch 82] M. Bierwisch, "Semantische und konzeptuelle Repräsentation lexikalischer Einheiten", *R. Ružička / W. Motsch (eds.), Untersuchungen zur Semantik, Studia Grammatica XXII*, Berlin, 1982.

¹³In the terminology of [Clark/Haviland 77].

- [Binnick 70] R. I. Binnick, "Ambiguity and vagueness", *Papers from the Sixth Regional Meeting of the Chicago Linguistic Society*, Chicago, 1970, pp. 147-153.
- [Bosch 83] P. Bosch, "Vagueness is Context-Dependance: A solution to the sorites paradox", *T. T. Ballmer / M. Pinkal (eds.), Approaching Vagueness*, Elsevier Science Publishers B.v. (North-Holland), 1983, pp. 189-210.
- [Clark/Haviland 77] H. H. Clark and S. E. Haviland, "Comprehension and the given-new contract", *R. O. Freedle (ed.), Discourse Production and Comprehension*, Norwood, New Jersey, 1977.
- [Geurts 90] Bart Geurts (ed.), "Natural Language Understanding in LILOG: An Intermediate Overview", *IWBS Report 137*, IBM Germany GmbH, Scientific Center, 1990.
- [Kooij 71] J. G. Kooij. "Ambiguity in Natural Language. An investigation of certain problems in its linguistic description", Amsterdam, London, 1971.
- [Nebel/Smolka 89] B. Nebel / G. Smolka, "Representation and Reasoning with Attributive Descriptions", *IWBS Report 81*, IBM Germany GmbH, Scientific Center, 1989.
- [Nunberg 78] G. D. Nunberg. "The Pragmatics of Reference", *PhD Thesis*, Reproduced by the Indiana University Linguistics Club, 1978.
- [Pinkal 80] M. Pinkal, "Semantische Vagheit: Phänomen und Theorien: Teil I" *Linguistische Berichte*, 1980.
- [Pollard/Sag 87] C. Pollard / E. A. Sag, "Information-based Syntax and Semantics: Volume I", *CSLI Lecture Notes Number 19* Stanford, 1987.
- [Pustejovsky 90] J. Pustejovsky, "Semantic Function and Lexical Decomposition", *Schmitz / Schuetz / Kunz (eds.), Linguistic Approaches to Artificial Intelligence*, Lang, 1990, pp. 243-303.
- [Shieber 86] S. M. Shieber. "An Introduction to Unification-Based Approaches to Grammar", *CLSI Lecture Notes Number 4*, Stanford, 1986.
- [Sidner 83] C. L. Sidner. "Focusing in the comprehension of definite anaphora", *M. Brady / Robert C. Berwick, (eds.), Computational Models of Discourse*, Cambridge, Mass. and London, 1983.
- [Weber 74] H. J. Weber, "Mehrdeutige Wortformen im heutigen Deutsch: Studien zu ihrer grammatischen Beschreibung und lexikographischen Erfassung", Tübingen, 1974.
- [Zwicky/Sadock 75] A. M. Zwicky and J. M. Sadock, "Ambiguity tests and how to fail them", *J. P. Kimball (ed.), Syntax and Semantics*, pp. 1-36, San Diego, 1975.