gefördert vom Bundesministerium für Forschung und Technologie

Verbmobil
Verbundvorhaben

# The VERBMOBIL Domain Model Version 1.0

J. Joachim Quantz[1]
Manfred Gehrke[2]
Uwe Küssner[1]
Birte Schmitz[1]

[1] Technische Universität Berlin
[2] Siemens AG München

**Report 29**
September 1994

September 1994

J. Joachim Quantz[1]
Manfred Gehrke[2]
Uwe Küssner[1]
Birte Schmitz[1]

Technische Universität Berlin
Fachbereich Informatik (13)
Projektgruppe KIT-VM11
Sekr. FR 5-12
Franklinstraße 28/29
D-10587 Berlin

Tel.: (030) 314 - 27778
e-mail: `kitvm11@cs.tu-berlin.de`

**Gehört zum Antragsabschnitt:** 11.10 Domainmodellierung

This work was funded/partially funded by the German Federal Ministry for Research and Technology (BMFT) in the framework of the Verbmobil Project under Grants 01 IV 101 Q 8 and 01 IV 102 A 0. The responsibility for the contents of this study lies with the authors.

## Abstract

This report describes the domain model used in the German Machine Translation project VERBMOBIL. In order make the design principles underlying the modeling explicit, we begin with a brief sketch of the VERBMOBIL demonstrator architecture from the perspective of the domain model. We then present some rather general considerations on the nature of domain modeling and its relationship to semantics.

We claim that the semantic information contained in the model mainly serves two tasks. For one thing, it provides the basis for a conceptual transfer from German to English; on the other hand, it provides information needed for disambiguation. We argue that these tasks pose different requirements, and that domain modeling in general is highly task-dependent. A brief overview of domain models or ontologies used in existing NLP systems confirms this position.

We finally describe the different parts of the domain model, explain our design decisions, and present examples of how the information contained in the model can be actually used in the VERBMOBIL demonstrator. In doing so, we also point out the main functionality of FLEX, the Description Logic system used for the modeling.

# Contents

# List of Figures

# 1 Introduction

In this report we describe Version 1.0 of the VERBMOBIL domain model. This version contains

1. a hierarchy of *speech-event types* relevant for the demonstrator scenario, i.e. appointment scheduling;
2. a *conceptual hierarchy* covering the word list used in the demonstrator;
3. *thematic relations* used for conceptually representing linguistic functor-argument structures;
4. a model of *calendar* information for representing temporal descriptions.

The main objective of this report is to describe both the current domain model and the methodology underlying the modeling. To understand this methodology it is useful to take a closer look at the various tasks supported by the domain model in VERBMOBIL. We therefore start with a sketch of the VERBMOBIL-demonstrator architecture from the perspective of the domain model (Section 2). In Section 3 we then present some rather general considerations on the nature of domain modeling and its relationship to semantics. Section 4 then describes the different tasks supported by the domain model. We will argue that these tasks pose different requirements, and that domain modeling in general is highly task-dependent. A brief overview of domain models or ontologies used in existing NLP systems (Section 5) confirms this position.

In Section 6 we describe the different parts of the domain model and explain our design decisions. Section 7 then gives examples of how the information contained in the domain model can be actually used in the VERBMOBIL demonstrator. We end the report with a brief overview over our future activities concerning domain modeling.

Some readers may wonder why we included a rather long "philosophical" section on the task of semantics. We included Section 3 mainly because of the heterogeneity of the research background of the various research groups involved in VERBMOBIL. The discussions at project meetings and workshops in the first year have shown that the different project partners start from rather different assumptions concerning objectives and methodology.[1]

[1] Section 3 is thus meant as an explanation of our own objectives and methodology. Those readers not interested in "philosophical" discussions might choose to skip this section.

Given the heterogeneity of research contexts, we think that two dichotomies are particularly obvious and important:

1. Logico-Linguistic (LL) semantics vs Artificial Intelligence (AI) semantics;
2. theory-oriented research vs application-oriented research.

Though there is a slight correspondence between these dichotomies, i.e. AI semantics tend to be application oriented, whereas LL semantics tend to be theory oriented, we think that separating both aspects is useful.

A distinction between AI semantics and LL semantics has been suggested in [Pinkal 88]. AI semantics are based on representation formats as *Semantic Nets* [Quillian 68], *Frames* [Minsky 75], or *Scripts*. LL semantics are based on Montague's initial work [Montague 74] and comprise *Montague Semantics* [Dowty et al. 81], *Discourse Representation Theory* [Kamp, Reyle 93] , and *Situation Semantics* [Barwise, Perry 83]. Pinkal sketches some of the main differences between both research areas concerning the subject of research, the ontological and methodological assumptions, etc. We think that the most important difference between the two research paradigms concerns the *status of semantic representations*. Not surprisingly, AI semantics is interested in the computational aspects of semantic representation, more precisely it addresses two questions:

1. Given an NL expression, how can the computer determine an appropriate semantic representation?
2. Given the semantic representation of an NL expression, how can the computer draw inferences from it?

LL semantics, on the other hand, is more interested in the logical aspects of semantic representations, e.g. in the problem of formally specifying the appropriate semantic representations of NL expressions. Appropriateness is here usually understood as correctly capturing the truth conditions of an expression.

In Section 3.1 we will argue for a distinction between the *interpretation task* and the *representation task*. Whereas AI semantics tend to focus on the interpretation task, LL semantics focus on the representation task. Since we think that the domain model in VERBMOBIL will be used mainly for the interpretation task, we opt for an approach more in the spirit of AI semantics. In contrast to many of the representation formats used in AI semantics, however, we use a Description Logic system (see Section 6.1) which combines the formal foundations of LL semantics with computational requirements.

Let us now turn our attention towards the dichotomy between theory-oriented and application-oriented research. Usually, theory-oriented research starts with defining the boundaries of its research—which aspects are to be accounted for and which aspects will not be addressed. Obviously, this definition of the object of research is itself influenced by the particular theory chosen (see [Kuhn 62]). Mostly, theory-oriented research analyzes a well-defined subject in depth, e.g. individual words of a language, specific constructions, etc.

Application-oriented research, on the other hand, cannot define the object of research itself, but rather has to deal with the phenomena relevant in the respective application. Thus when building NLP systems for some application it is more important to *somehow* cover *all* phenomena occurring in the application, then to cover few phenomena in detail.

Given these differences in methodology and objective, there is sometimes the tendency in both "fields" to disqualify research from the other field. Thus application-oriented researchers complain about the *irrelevance* of theory-oriented research for practical purposes, while theory-oriented researchers discard application-oriented research as *ad hoc* and criticize the *lack of theoretical foundation*.

It should be obvious, however, that both research strategies should complement each other. Applications can be a good opportunity to evaluate theories, and theories can provide sound foundations for applications. The problem is then to find the right "balance" between theory-oriented and application-oriented activities. We think that as far as the VERBMOBIL demonstrator is concerned, our research has to be *application oriented*. In Section 3 we will therefore develop some application-oriented requirements for domain modeling, taking into account the theoretical foundations in the field of formal semantics. Before doing so, however, we will briefly sketch the overall architecture of the VERBMOBIL demonstrator in so far as it is relevant for the task of domain modeling.

# 2 The Architecture of the VERBMOBIL Demonstrator

For the VERBMOBIL demonstrator a more or less sequential architecture is envisaged.[2] In the context of this paper the modules *Syntax*, *Lexicon*, *Semantic Construction*, *Semantic Evaluation*, and *Dialogue* are relevant. In the following

[2]Note that this section is meant only as a *description* of the chosen architecture. We neither intend to defend nor to criticize this choice in the following.

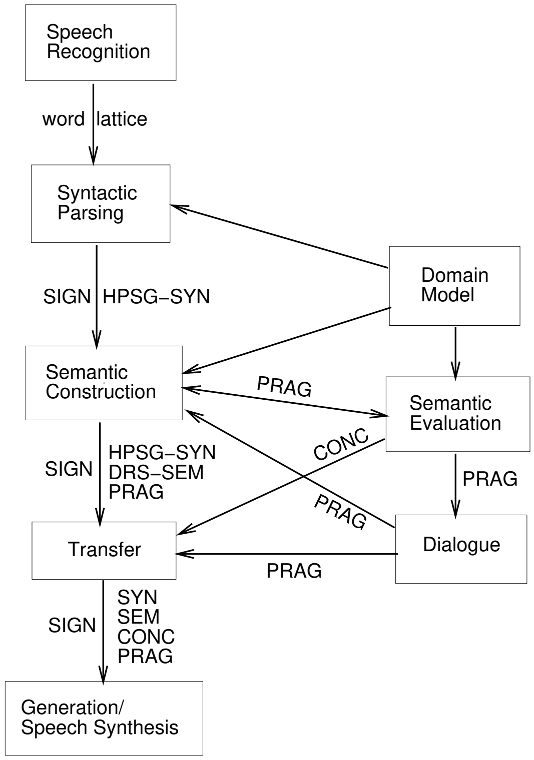


Figure 1: Simplified architecture of the VERBMOBIL demonstrator from the domain-modeling perspective.

we will briefly sketch the functionality of these modules and their relationship to the domain model (see also the architecture depicted in Figure 1).

Since VERBMOBIL is a speech-to-speech translation system, the input to the *syntactic parser* is not a sequence of words, as in text-based NLP systems, but rather a word lattice. Each path in this lattice is a sequence of words as analyzed by the speech recognizer. The syntactic parser then builds syntactic structures for some of these paths, eliminating other paths as being syntactically inconsistent.

A rather important aspect of this speech-based scenario is thus that the input to the parser is already highly ambiguous. To obtain an efficient performance of the overall system it is therefore crucial to eliminate "inconsistent" structures as early as possible. In addition to purely syntactic information the parser takes into account also semantic information in the form of selectional restrictions and semantic sorts.

Consequently, the *lexical entries* contain information about semantic sorts assigned to the lexemes, as well as syntactic and semantic selectional restrictions for the arguments they can take. To provide a *sort hierarchy*, i.e. a set of semantic sorts and functionality to test sort compatibility, subsumption, or disjointness is thus one of the tasks of the domain model in VERBMOBIL (see Section 4.1).

The output of Syntax are HPSG-style phrase structure trees, possibly alternative ones whenever alternative paths in the word lattice are syntactically consistent, or whenever one path is syntactically ambiguous. *Semantic Construction* takes these phrase structure trees and builds up semantic representations in the format of *Discourse Representation Structures* (DRSs). These representations contain discourse referents for the referring expressions and conditions capturing the semantic content of the utterance. Conditions are basically predicative propositions, where the predicate is a *semantic predicate* obtained from the NL expression and the arguments are discourse referents. Again, selectional restrictions of semantic predicates to their arguments and sort information about discourse referents are used to eliminate "inconsistent" structures.

The information contained in the DRS's is used as input for *Transfer*. If *Transfer* cannot determine an English translation on the basis of this information,[3] *Semantic Evaluation* further disambiguates the semantic representation. The basic idea of this architecture is to allow a variable depth of analysis, i.e. a transfer-guided disambiguation strategy. Thus ambiguities arising from the specific translation pair German/English do not influence Semantic Construction but are treated

[3] It is not yet clear whether the choice of particular English lexemes is performed by Transfer or by Generation.

by Semantic Evaluation on demand by Transfer. Furthermore, Semantic Evaluation can perform non-compositional and non-monotonic operations, e.g. type shifting or coercion (see below), whereas Semantic Construction proceeds mostly compositionally and monotonically.

In addition to the sort hierarchy the domain model is thus supposed to provide a *conceptual hierarchy* reflecting the semantic difference between German and English. These concepts are related to German and to English expressions by $m$-to-$n$ mappings. Roughly speaking, we will, for example, introduce $n$ concepts for a German lexeme if this lexeme can be translated into $n$ English lexemes (see Section 4.2 for more details).

Note that the sort hierarchy is supposed to be a subset of the conceptual hierarchy, namely its *upper structure*. The idea behind this architecture is that only very general sort information is used in Syntax and Semantic Construction, whereas more fine-grained conceptual information is used by Semantic Evaluation and Transfer.

Since VERBMOBIL is supposed to translate utterances from dialogues, determining the communicative functions of expressions is crucial for their translation. The module *Dialogue* provides relevant information about the dialogue context and Semantic Evaluation computes the speech-event type of utterances (see Section 4.3).

# 3 Semantics in VERBMOBIL

In the preceding section we have already indicated some of the tasks for which the domain model is to be used. It should have become obvious that the information contained in the domain model is intimately related to *semantic information* in general. In particular the relationship between *lexical semantics* and the domain model has to be clarified.

In this section we will consider the place of semantics in the VERBMOBIL project and present our own perspective, which has major impact on our design of the domain model. We will basically argue for a distinction between an *interpretation* and a *representation* task, and claim that there is a tendency for AI semantics to focus on the interpretation task, whereas LL semantics tend to focus on the representation task. Given this distinction we show which interpretation-related and representation-related requirements to domain modeling arise in the VERBMOBIL context.

## 3.1 The Task of Semantics

In this section we will first address some fundamental semantic issues, namely the problem of "aboutness", the relationship between entailment and derivability, the distinction between representation and intepretation, and the role of semantics in Machine Translation in general and in VERBMOBIL in particular.

### Semantics and Aboutness

We start this section by briefly sketching the view of semantics usually taken in LL semantics, namely *truth-conditional semantics*. The basic idea of truth-conditional semantics is that the meaning of a sentence can be identified with its truth-conditions—to know what a sentence means, to understand a sentence is to know under what conditions the sentence is true (see [Davidson 67, p. 23f]).

Usually this approach to meaning is accompanied by a *correspondence theory* of truth à la Tarski: "Die Wahrheit einer Aussage besteht in ihrer Übereinstimmung (oder Korrespondenz) mit der Wirklichkeit" [Tarski 44, p. 143]. This leads proponents of truth-conditional semantics to statements like the following:

> We merely wish to emphasize that truth-conditional semantics, in contrast to the other approaches mentioned, is based squarely on the assumption that the proper business of semantics is to specify how language connects with the world – in other words, to explicate the inherent "aboutness" of language.
>
> (...) in truth-conditional semantics we answer a question of the form "What is the meaning of a sentence S?" by providing some sort of description of how things would have to be arranged in some relevant corner of the world in order for S to be true. [Dowty et al. 81, p. 5]

Note that the matter of aboutness is relevant to us for two reasons:

1. For one thing, AI semantics have been criticized for not establishing aboutness, e.g. Lewis' critique against *markerese* in [Lewis 72], Fodor's critique of *procedural semantics* [Fodor 78], or Searle's *Chinese Room Gedankenexperiment* [Searle 80].

2. Secondly, LL semanticians usually claim to establish aboutness by using model-theoretic semantics, as illustrated by the above quotation from [Dowty et al. 81]. Similar claims can be found in [Lewis 72], [Cresswell 78], or [Kamp, Reyle 93].

We think that both arguments are unjustified: neither does a model-theoretic semantics establish aboutness, nor is the failure to establish aboutness a shortcoming of semantic theories (neither of AI semantics nor of semantic theories in general). We will briefly sketch the arguments for this position in the following. A more detailed argument can be found in [Quantz 95, Chap. 2.4].

Usually, LL semantics "establish" the connection between language and world by a model-theoretic semantics. Such a semantics defines formal criteria to decide whether a given sentence is true or false in a given model. What seems to be overlooked sometimes is that:

1. the important point of model-theoretic semantics is the definition of an entailment relation, i.e. models are only a *means* to define entailment;

2. models are formal systems and *not* the world, i.e. they are as "far away" from the real world as the sentences of a language.

Let us briefly explain these two points and their consequences. Clinging to the aboutness metaphor, (1) means that it is not necessary for semantics to establish aboutness and (2) that model theory is not sufficient to establish aboutness.

We begin by considering the claim that model-theoretic semantics establish aboutness. Kamp and Reyle, for example, criticize semantic theories based on mental representations for their failure to really explain meaning:

> Indeed, explaining natural language solely by referring to mental representations would only shift the problem of meaning to another language. [Kamp, Reyle 93, p. 10]

They claim that model-theoretic semantic, on the other hand, relate natural language to the real world and therefore explain meaning. This argument relies on the assumption that the models underlying the model-theoretic semantics are somehow closer to reality than natural languages or mental representations, a position we find hard to accept. As Potts points out:

> Model theory is, rather, an exercise in *translation*. (...)
>
> Thus it is just a confusion to suppose that model theory can say anything about the relation of language to the world; it can, at best, only elucidate one language by reference to another. [Potts 75, p. 248]

Obviously, a model is just a representation, i.e. an element of a formal sign system and therefore as different from "reality" as Natural Language expressions or mental representations. In particular, the assumption that the world is ultimately a collection of predicates and constants standing in instanceship relations, reflects the philosophical state-of-the-art of *Logical Positivism*, a research paradigm more or less abandoned in recent philosophical development.[4] The philosophy of *internal realism*, e.g. [Goodman 68, Goodman 78, Putnam 87] stresses the fact that the structure of reality depends on the sign systems we use to perceive and describe it. Therefore aboutness can only mean mapping from one sign system into another and all our explanations will never leave the level of sign systems and reach an independently existing reality.[5]

We therefore reject arguments criticizing the failure of semantic theories to establish aboutness, such as the ones cited above or given in [Fodor 78] or [Searle 80]. Expecting a semantic theory to provide a connection between language and world is ultimately as self-destructive as the doctrine of skepticism. Note that both positions take an independently existing world for granted and then complain that our knowledge/semantics fails to reach this world. But we have to reject the inconsistent notion of a world completely independent from us, and acknowledge the fact that a world is always given to us via signs and is ultimately only experienced and understood through signs. Then the position that semantics does not reach out of the level of signs but rests inside, i.e., is from a certain point of view only syntax, is perfectly acceptable and not defective in any way.

The major contribution of model-theoretic semantics is therefore not to establish aboutness, but rather to provide a formal definition of an *entailment relation*. In fact, there are alternative ways of defining such entailment relations, e.g. [Leblanc 83] and though model theory is an especially fruitful framework for studying formal properties, it is in a way inessential. Note that similar arguments against the relevance of aboutness for semantics can already be found in [Morris 38, p. 58f] and [Wittgenstein 56, 293].

### Entailment and Derivability

Summarizing the above considerations, the main contribution of LL semantics is to provide a formal entailment relation on (sets of) semantic representations. We will now address two important aspects of such an entailment relation:

1. the information relevant for entailment;

[4] The world is, so to speak, not a Herbrand universe.

[5] Putnam shows, for example, by using the Löwenheim-Skolem theorem, that a model-theoretic semantics does not unambiguously determine the denotations of the expressions of a language [Putnam 80].

2. the relevance of the entailment relation for computational systems.

Let us first take a closer look at the nature of the entailment relation. Given a semantic representation 'm' of an NL sentence, what formulae are entailed by it? It is obvious that some formulae are entailed by 'm' alone, whereas others are entailed by 'm' together with additional information. A simple example might illustrate this:

(1) a. Bill has a brother and a sister.

b. $\exists x$Brother(bill,x) $\wedge$ $\exists y$Sister(bill,y)

We can derive from (b) alone that Bill has a brother, due to the semantics of the logical connective $\wedge$. In order to derive the fact that Bill has two siblings we need additional information, namely that brothers are male, that sisters are female, and that the sets of males and females are disjoint.[6]

There has been a slight tendency in LL semantics to concentrate on the first type of entailments, i.e. on *structural semantics*, and to neglect the second type of entailment, i.e. *lexical semantics*:

> *But we should not expect a semantic theory to furnish an account of how any two expressions belonging to the same syntactic category differ in meaning*. 'Walk' and 'run', for instance, and 'unicorn' and 'zebra' certainly do differ in meaning, and we require a dictionary of English to tell us how. But the making of a dictionary demands considerable knowledge of the world.. (...) These are matters of application, not of theory. [Thomason 74, p.48f]

Obviously, lexical semantics is crucial for any theory of meaning interested in modeling or explaining understanding, and this has been stressed by other researchers in LL semantics, e.g. [Partee 80, Pinkal 85]. Nevertheless, Thomason's remark is not entirely unjustified and one way of dealing with the underlying problem is to distinguish between *linguistic meaning* and *encyclopedic meaning*. The idea is to separate the information related to expressions into linguistic and non-linguistic information. In a sense this distinction corresponds to the philsophical distinction between *analytic* and *synthetic* knowledge and thus inherits the shortcomings of this distinction [Quine 51]. We will not address this issue any further, but the following two points seem fundamental to us:

[6] Note that via the Gricean Maximes we can also infer that Bill has not more than two siblings [Levinson 83].

1. Given a sentence (or its formal semantic representation), inferences can involve various, "arbitrary" pieces of information.

2. The semantic information used to draw inferences cannot be easily and uncontroversially divided into linguistic and non-linguistic information. If there is a borderline between linguistic and non-linguistic semantic information at all, it is fuzzy.

The first point implies that the more one knows, the more a sentence means to one (the more inferences one can draw). We will now sketch the consequences of this for computational semantics in general and for VERBMOBIL in particular.

In the preceding paragraphs we have used terms like 'to draw inferences'. We still have to explain the relationship between such a term and the notion of semantic entailment. In fact this is to explain the relevance of entailment relations for computational systems, a task we have mentioned above but not addressed so far.

Given an entailment relation $\phi \models \psi$, one can define a deduction system based on a purely syntactic relation of derivability ($\phi \vdash \psi$). These inference systems can then be evaluated with respect to the semantics: an inference system is sound and complete iff $\Gamma \models \gamma \Leftrightarrow \Gamma \vdash \gamma$ (see [Sundholm 83] for details). Thus from a system-oriented point of view, the entailment relation is just an "ideal", whereas the derivability relation is the "reality".[7] This is particularly important if systems provide only incomplete algorithms. Note that incomplete algorithms are necessary if computing entailment is undecidable or intractable [Garey, Johnson 79], and (efficient) termination has to be guaranteed, as is the case in real-world applications.[8]

Summarizing the above discussion about formal semantics, we can distinguish a theoretical point of view and a system-oriented point of view. Whereas a semantic entailment is a good basis for developing inference systems and is highly useful for theoretical investigations, it is on its own useless from a system-oriented point of view. Given a system, the meaning of a semantic representation 'm' *for the system* is characterized by the representations 'n' derivable from 'm'. What is important to note is thus that:

1. the meaning of a representation 'm' for a system is *not* determined by

   (a) 'm' alone;

   (b) a model-theoretic interpretation of 'm';

[7] The importance of deduction is also stressed in [Benthem 81] and [Galton 88].

[8] This is also one of the differences between simplified artificial scenarios and realistic applications pointed out in [Wahlster 94].

2. but instead by ‘m’ together with
   (a) the information stored by the system;
   (b) the inference algorithms implemented in the system

In other words, a formal representation $\diamond \forall x P(x) \rightarrow Q(x)$ is for a system no more meaningful than ‘foo’ if it has no inference rules for dealing with $\diamond$, $\forall$, $P$ and $Q$.

**Semantics in Machine Translation**

So far we have sketched the task of semantics in NLP systems in general, without addressing the particular requirements in Machine Translation. Let us start by sketching a “naive” approach to MT in order to illustrate some relevant aspects. Such an approach can be characterized as follows:

1. Analyze the source-language expression and derive its meaning, i.e. a semantic representation.
2. Generate a target-language expression from this semantic representation, i.e. an expression expressing the same meaning as the source-language expression.

Note that this approach offers a lot of desirable advantages. The semantic representation functions as an interlingua, i.e. having an analyzer for German and a generator for English, one would “only” have to write an analyzer (or a generator) for Japanese in order to achieve Japanese/English (German/Japanese) translation.

Furthermore, this approach seems to capture the common-sense understanding that the meaning of an expression is the thing to be conserved in translation, i.e. is the *equivalent of translation*. Though theoretically attractive, this approach faces a number of problems, especially in connection with a truth-conditional approach to semantics:

1. there are still many unresolved issues in truth-conditional semantics, e.g.
   (a) a truth-conditional semantics of non-declarative sentences, which are crucial in the dialogue-oriented VERBMOBIL scenario, is still fragmentary;
   (b) as we will see below, truth conditions ultimately depend on lexical semantics and have to include encyclopedic knowledge;

(c) tautologies have identical truth-conditions, thus '2 + 2 = 4' would have the same meaning as, for example, '3 + 3 = 6';

2. in general, truth-conditions are not the only equivalent of translation; in the VERBMOBIL demonstrator scenario it seems much more appropriate to rely on the communicative function of an utterance as the most important equivalent of translation;

3. there are numerous occurrences of idiomatic expressions, for which a truth-conditional semantics is not appropriate;

4. the approach basically relies on the mapping relations $e_{SL} \rightarrow m$ and $m \rightarrow e_{TL}$ and does not "really" access the truth conditions.

Thus truth-conditional semantics seems to be neither sufficient nor necessary for MT.

Given the scenario for the VERBMOBIL demonstrator, it has been argued elsewhere [Schmitz 94] that the communicative function of an utterance, i.e. the speech-event type it realizes is in most cases the adequate equivalent of translation. Depending on the particular speech-event type, the semantic content is more or less important. For example, the date described in a proposal ('vorschlag') is crucial, whereas uttereances like

(2) Wie sieht es da bei Ihnen aus?

do not contribute much semantic content but are rather idiomatic realizations of certain communicative functions, e.g. 'aufforderung_stellungnahme' (see also Section 4.3).

The argumentation so far can be summarized as follows:

1. it is not clear to which degree a system has to understand an utterance in order to translate it. There are cases where "deep understanding" seems necessary, as well as cases where translation is possible without such a "deep understanding".

2. it is therefore not clear to which degree a semantic representation of an utterance has to be constructed, more precisely, how accurate a semantic representation has to reflect the full semantic content of an utterance.

In the following, we will argue that the degree of semantic representation should be largely determined by requirements arising from the *interpretation task*.

### Representation and Interpretation

The distinction between the representation task and the interpretation task can best be explained by considering a sequential architecture for NLP in which different formalisms are used to represent information on different levels (as is the case in the VERBMOBIL demonstrator). Given such a multi-level approach the *representation problem* is to represent the different interpretations possible at a given level and the *interpretation problem* is to produce the appropriate representations at each level from the input to that level [Allen 87, p. 9].[9]

It should be noted that first, the representation task cannot be adequately specified without taking into account the particular application for which the representation is to be used; and that second, the interpretation task depends on the chosen representation format—more precisely, whether a representation on a specific level is ambiguous depends on the representation format chosen for the subsequent level. Though both tasks are thus intimately related, it is useful to distinguish them. One might even conceive an NLP system in which the interpretation task is only guided by the system but has to be performed by a human user. An interactive MT system could ask the user to provide the information it needs to translate an expression into the target language; an interactive text archiving system could ask the user to provide the information it needs to appropriately classify a given NL text. For the design of such systems the representation task would be predominant, whereas the interpretation task could be neglected.[10]

Given this distinction, we think that in Machine Translation in general, and in the VERBMOBIL demonstrator in particular, semantics is mainly needed to support the interpretation task. Though semantic representation is used as an input for Transfer and Generation, there is no need to use it immediately for drawing inferences. Contrast this with Information Systems, in which the semantic content of a query has to be evaluated wrt the stored information; or with text understanding systems in which the information contained in a text has to be represented semantically in order to answer queries.[11]

[9] Note that a similar distinction can be made for NLP systems using a homogeneous representation of these levels, e.g. the sign structure in HPSG. Here the different levels are represented via different features in the sign structure.

[10] See, for example, Kay's ideas on the *Negotiator* architecture [Kay et al. 91].

[11] In Section 5 we will see that domain models or ontologies are used differently in different systems. Some are used only for the interpretation task or the representation task, some are used for both tasks.

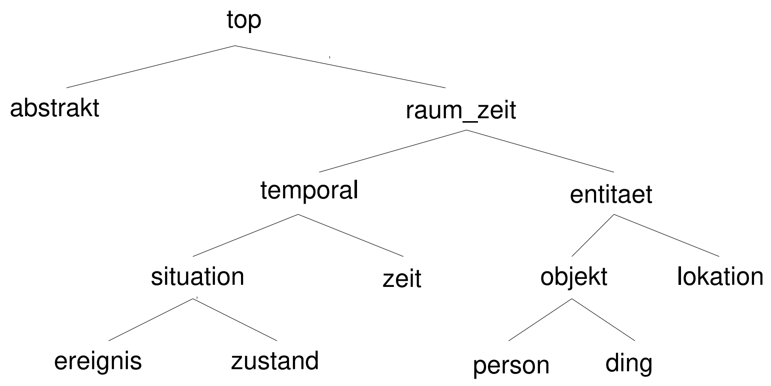


Figure 2: A simple ISA hierarchy.

### Semantics in VERBMOBIL

Thus we see the main task of semantics in the VERBMOBIL project not in providing an interlingua-like representation for translation, but rather in supporting the interpretation process, i.e. the disambiguation steps necessary in transforming German expressions into English expressions.

Note that such a view implies the following methodology:

1. identify the interpretation processes in VERBMOBIL, i.e. the transitions from one representation format to another which require disambiguation;

2. identify the semantic information needed to perform these disambiguations;

3. model this semantic information.

Instead of thus modeling semantic information in general, which would necessitate arbitrary decisions what to model and what to leave out, this approach provides us with a criterion to decide whether a piece of information should be included in the domain model or not. In Section 6 we will sketch four different tasks and their support by the domain model. Before doing so, however, we will address two general points, namely the function of *conceptual hierarchies* and the problem of *polysemy* and *type shifting*.

## 3.2 Conceptual Hierarchies

Conceptual hierarchies are the most obvious way to represent semantic information needed in the interpretation task. A simple ISA hierarchy, as shown in Figure 2 allows type checking and inheritance. We will briefly address six issues in the following:

1. the function of conceptual hierarchies;
2. the distinction between linguistic and non-linguistic semantic information;
3. the complexity of conceptual hierarchies;
4. the "distinction" between upper model and domain model;
5. the connection between conceptual hierarchies and LL semantics;
6. the status of conceptual hierarchies wrt NL expressions.

First, conceptual hierarchies are based on the intuition that some of our commonsense knowledge is hierarchically structured, a notion underlying representation formats as *Semantic Nets* [Quillian 68] or *Frames* [Minsky 75]. The main advantage of a hierarchical structure is that information can be inherited from general concepts to specific ones, and thus allows compact modeling and storage of information. For example, the fact that dogs bark can be inherited from the concept 'dog' to the more specific concept 'collie'.

It should be noted that conceptual hierarchies are used similarly within a new Machine Translation paradigm, namely *Example-Based Machine Translation*. Here conceptual hierarchies called *thesauri* are used to determine the semantic difference between the phrase to be translated and the stored examples translations, e.g. [Almuallim et al. 94].

Second, it should be noted that conceptual hierarchies can contain linguistic, as well as non-linguistic semantic information. As mentioned above in Section 3.1 we do not think that a clear-cut distinction between these types of information is possible (this will be illustrated when discussing literal meaning and polysemy below). Since we do not see any advantage of making such a distinction in the VERBMOBIL context anyway, we will not refer to it any further in the following.

Third, conceptual hierarchies can be more complex than simple ISA hierarchies. Semantic Nets, Frames, and Description Logics allow to specify "horizontal" links between concepts, e.g. expressing that events occur at a particular location and time. Using more complex hierarchies makes reasoning a lot more complex as well, and many NLP systems use only simple hierarchies in the interpretation process.

Fourth, a distinction is sometimes made between an upper model and the domain model. This distinction, which is rather popular in the literature on domain modeling, has been suggested in [Mann 85, Bateman 92] and realized in PENMAN, a text-generation project. It should be noted that the distinction is based on theoretical assumptions which might not be straightforwardly applicable to NLP in general. Basically, the PENMAN Upper Model reflects abstract semantic distinctions needed to categorize the propositional content of a text. In the framework underlying the PENMAN system, namely *systemic-functional linguistics*, this aspect of meaning is called *ideational meaning* [Halliday 70].

There are different views on the distinction between upper model and domain model. From a system-oriented point of view there is a rather pragmatic reason for the distinction—the upper model is provided by PENMAN, whereas the domain model has to be defined by the user, basically by refining the upper model. Thus the upper model contains all relevant aspects of actions and relations and inherits them to the domain model. In [Bateman 92] the upper model is motivated as an interface between linguistic and conceptual knowledge. For the VERBMOBIL demonstrator scenario we prefer a direct mapping from NL expressions into concepts and do not see the necessity of an additional interface level.

A comparison of the ontologies and the domain models used in existing NLP systems, as the ones sketched above, shows that they differ considerably wrt the basic categories they use (see Section 5). We think that these differences stem mainly from the fact that the respective models are used for rather different tasks (e.g. text understanding, translation, disambiguation, generation).

In the following we will not make a distinction between an upper model and a domain model. We will, however, distinguish a subset of the conceptual hierarchy. This subset, which we will refer to as the *sort hierarchy*, will be used by Syntax and Semantic Construction and contains rather general semantic categories (see Section 6.2). We do not use the term 'upper model' for it in order to avoid the connotations of this term stemming from the work in the PENMAN context.

Fifth, it should be noted that conceptual hierarchies in fact realize a sort of lexical semantics. Thus having an ISA relation between P and Q means logically $\forall x P(x) \rightarrow Q(x)$. A conceptual hierarchy can thus be seen as a list of *meaning postulates* [Dowty et al. 81] or *constraints* [Barwise 89]. Two things should be noted, however:

1. the ISA relations hold between concepts and not between lexemes or expressions of a natural language;

2. the ISA relations contain only part of the semantic potential of a concept, i.e. they do not contain *analytic definitions*. The amount of information contained in a conceptual hierarchy depends on the requirements of the task for which it is needed.

The first remark leads us to the relationship between the concepts in a conceptual hierarchy and the lexemes/expressions of a natural language. This is a rather important issue and is, in a way, the overture to the problem of polysemy and type shifting discussed in the next section.

It is common practice in lexical semantics to define semantic relations on lexemes and NL expressions, e.g. [Cruse 86, Chap. 4]. It should be noted, however, that these definitions are ultimately based on the meanings of the expressions, and that the corresponding tests of truth conditions have to be defined rather carefully in order to avoid *contextual effects* (e.g. Cruse's definition of hyponomy [Cruse 86, p. 88f]).

Let us take a closer look at this problem by considering a concrete example. Suppose you are modeling a conceptual hierarchy and have to decide where to insert the concept expressed by the German lexeme 'Abendessen' (dinner). It becomes immediately obvious that there are a number of options, comprising, for example:

- 'Abendessen' as a specialization of FOOD;

- 'Abendessen' as a specialization of EVENT.

This "problem" has been discussed in different contexts in the literature, most often in connection with *polysemy* or *type shifting*. In the next section we will address this phenomenon in detail.

## 3.3 Polysemy and Type Shifting

In order to discuss the problem of polysemy and type shifting we need to briefly sketch a general framework of semantic interpretation and representation. We will distinguish two main tasks of semantic interpretation, namely the determination of

1. the conceptual content and

2. the referent

of an expression. This distinction, which roughly correponds to the distinction between *denotation* and *reference*, is sometimes obscured in existing semantic theories (see [Lyons 77, 7.2] and [Lyons 91, p. 10f]). Strictly speaking it is not an expression which refers, but rather the speaker by using that expression. As a consequence, expressions have reference only in so far as they are text-expressions, but not as system-expressions per se.[12] Thus it does not make any sense to speak of the reference of system-expressions like 'the president', 'documents', 'she', or 'office'.

Nonetheless it is perfectly sensible to speak of the meanings of such system-expressions. Given a denotational approach to semantics, we might say, for example, that 'president' denotes the class of all presidents, 'office' the classs of all offices, etc. Note that one could try to reduce denotation to reference by saying that the denotation of a lexeme 'l' contains all the objects to which a speaker can refer by using the expression 'this l'. According to this view a referring expression has a denotation, which is determined by its linguistic form, and it refers to a particular member of this denotation when used in a particular context (cf, for example, the distinction between meaning and interpretation in situation semantics [Barwise, Perry 83, p. 36f]).

In the following we will discuss tow main problems arising in connection with such an approach:

1. Given an expression, how can we determine its conceptual content (denotation)?

2. What exactly is the relationship between the conceptual content of an expression and its referent?

The first problem is related to the problem of *literal meaning and polysemy*, the second one to the problem of *type shifting or coercion*.

### Literal Meaning and Polysemy

In the following we will assume that the conceptual content of a referring expression is a concept and that its referent is an object (thus using DL terminology as sketched in Section 6.1). We will not address the question of what exactly a concept is, however. For the purpose of our argumentation it suffices to assume that a

[12] See [Lyons 77, p. 31] and [Brown, Yule 83, Sect. 1.3] for this distinction.

concept can be more specific than, equivalent to, or disjoint from another concept, and that objects are instances of concepts. This can be formalized in First-Order Predicate Logic or in Description Logics, but we do not think that this particular formalization is essential for our discussion.

Given the dominance of truth-conditional semantics in the last decades, it is not suprising that many approaches to lexical meaning rely on the truth-condition-testing method. Consider, for example, the definition of cognitive synonymy in [Cruse 86, p. 88]:

> X is a cognitive synonym of Y, if (i) X and Y are syntactically identical, and (ii) any grammatical declarative sentence S containing X has equivalent truth-conditions to another sentence S', which is identical to S except that X is replaced by Y.

Though we think that this is a useful test criterion for synonymy it is not a real definition since it reduces synonymy of lexemes to equivalence of truth conditions, i.e. to analyticity (see [Quine 51, p. 28ff]). In fact recourse to truth conditions does not help us much in modeling conceptual contents, because if we are uncertain about the exact meaning of an expression we will not be able to specify the exact truth conditions of a sentence containing it.

To illustrate the problem, consider some of the examples discussed in the literature. Bosch discusses the truth conditions of

(3) There will be coffee in the lounge at 10.30.

in [Bosch 85]. He claims that this sentence would be true if

1. there were a coffee break at 10.30;
2. there were a bag of coffee beans on one of the tables in the lounge;
3. there were coffee powder spilled on the carpet;
4. etc.

This example illustrates several points already mentioned in previous sections. For one thing, truth conditions, if understood as in Bosch's example, do not reflect our ability of understanding completely. "Our understanding of utterances in ordinary circumstances is almost always more specific than the truth conditions of the sentence uttered" [Bosch 85, p. 251].

Second, instead of assigning a single conceptual content to each expression, it seems more reasonable to say that an expression as a system-expression is ambiguous since it can express several different conceptual contents. Only in a particular context, i.e. given the expression as a text-expression, is a particular conceptual content activated. Thus we can list the concepts COFFEE_DRINK, COFFEE_BEANS, COFFEE_POWDER, COFFEE_BREAK, COFFEE_PLANT, etc. as possible conceptual contents of the lexeme 'coffee'.[13] Doing so has the advantage that each of these concepts can be classified at an appropriate place in the conceptual hierarchy, e.g. COFFEE_DRINK below DRINK, COFFEE_BREAK below EVENT, etc.[14]

But we immediately face another problem—how do we decide *how many* conceptual contents we have to assign to an expression? Again, consider an example taken from the literature, namely the meaning of the lexeme 'open' discussed in [Searle 83, p. 145]:[15]

(4) a. Tom opened the door.

b. Sally opened her eyes.

c. The carpenters opened the wall.

d. Sam opened his book on page 37.

e. The surgeon opened the wound.

Searle claims that

> ... the word "open" has the same literal meaning in all five of these occurrences. Anyone who denied this would soon be forced to hold the view that the word "open" is indefinitely or perhaps even infinitely ambiguous since we can continue the example; and indefinite ambiguity seems an absurd result. [Searle 83, p. 146]

He gives additional examples in which one might argue that the lexeme 'open' has a different meaning:

[13] Note that one problem in conceptual modeling is to find appropriate identifiers for the concepts. Though we think that this is a rather important issue, especially wrt useability of a modeling, we will not pursue it any further in this report.

[14] One consequence of this approach is that truth conditions are underdetermined by sentences. Truth conditions are only determined given a particular choice of the conceptual contents of the expressions occurring in the sentences (see, for example, [Pinkal 85, p. 12ff] or [Bosch 85, p. 251]).

[15] Bosch discusses the same problem and lists the 29 senses of 'open' in Collins Cobuild English Language Dictionary [Bosch 93].

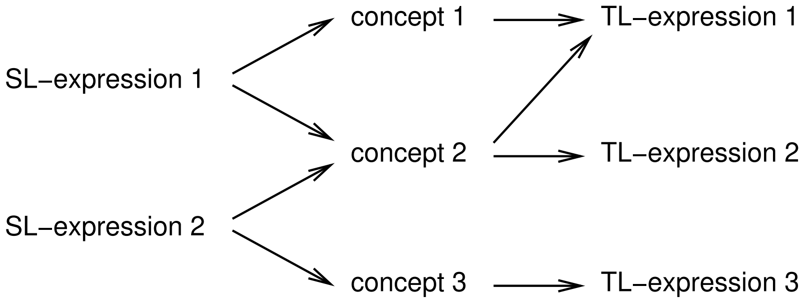


Figure 3: NL expressions and concepts are related by $m$-to-$n$ mappings.

(5) a. The chairman opened the meeting.

b. The artillery opened fire.

c. Bill opened a restaurant.

Searle uses these examples to show that understanding requires a *preintentional Background* and that the literal meaning of a sentence is thus not a context-free notion, but is relative to a set of preintentional Background assumptions and practices [Searle 83, p. 145f].

For our argumentation it is important to note that these examples show how linguistic meaning and encyclopedic information are intertwined. Only given our knowledge about the typical events described in the sentences do we understand the word ‘open’. How could we divide this knowledge into a linguistic-meaning part and a world-knowledge part?

Let us summarize the discussion:

1. the conceptual content of expressions is context dependent;

2. it is problematic to decide how many concepts to assign to an expression (truth-condition tests are helpful but do not resolve all problematic cases).

Since we are not primarily interested in truth conditions anyway, we propose to base the assignment of concepts to expressions on the requirements of our particular application, namely translation form German into English.

### Transfer-Oriented Conceptual Modeling

A straightforward criterion for deciding how many conceptual contents to assign to a source-language expression is to provide a separate concept for each possible target-language translation. Consider the examples (4) and (5)—the German translation for ‘open’ in (4) is ‘öffnen’, whereas in (5) it is ‘eröffnen’. This would be a reason to distinguish the conceptual content of ‘open’ expressed in (4) and (5). In general, we would thus have $m$-to-$n$ mappings between NL expressions and concepts, as illustrated in Figure 3.

Note that the structure of the conceptual hierarchy to be modeled thus depends on the particular source and target languages. Based on *contrastive analyses* it can be decided whether a concept has to be included into the domain model or not. Note further, that it is in principle possible to extend such a bilingually motivated hierarchy by adding new source and target languages, i.e. by adding mappings and integrating the additionally required conceptual distinctions [Hovy, Nirenburg 92, Henschel, Bateman 94].

Though the above criterion is a fundamental methodological principle underlying our domain modeling, it is not the only one. In adopting an $m$-to-$n$ mapping from expressions into concepts we have to perform *conceptual disambiguation* when processing utterances—when given a particular utterance, we have to select a particular conceptual content form the various possible ones.

As already indicated in the above discussion, the choice of a particular conceptual content is determined by the context in which an expression occurs. The examples have also shown that this context involves complex background knowledge. The major problem for conceptual disambiguation is

1. to make background knowledge available, i.e. to store it in the computer;

2. to provide fast access to backrground knowledge, i.e. to provide inference mechanisms drawing the *relevant* inferences.

In the absence of a satisfying solution to this task which comprises our entire common-sense knowledge one has to devise heuristics based on partial information and incomplete reasoning. The most popular way of doing conceptual disambiguation is by means of *selectional restrictions*. Roughly speaking, conceptual disambiguation is performed in this approach by considering the arguments of an expression, or vice versa the functor taking an expression as argument. Note that this is a rather limited way of taking context into account, but it is exactly this limitation which makes it computationally feasible.

For illustration consider again the examples (4) and (5). To distinguish between the two readings of ‘open’ we can check whether the argument of ‘open’ is a thing or an event. If it is a thing we map ‘open’ to the concept expressed by ‘öffnen’, if it is an event we map it to the concept expressed by ‘eröffnen’.

Having sketched the general approach we will now point out some of the problems arising in its application. Wrt design criteria ot should be kept in mind that we will have to insert concepts into the hierarchy needed for expressing the selectional restrictions. Thus we have two main criteria for deciding which conceptual distinctions to include in the domain model:

1. distinctions due to different lexicalizations in German and English;
2. distinctions relevant for selectional restrictions used for conceptual disambiguation.

In Section 4.2 we will illustrate this in more detail by considering concrete examples from the VERBMOBIL demonstrator corpus.

**Type Shifting**

Let us take a closer look at the strategy for conceptual disambiguation sketched above. The attentive reader may have noticed our rather informal description of selectional restrictions, i.e. the phrase ‘whether the argument of ‘open’ is a thing or an event’. Obviously, the argument of ‘open’ can be neither, but rather is a linguistic expression. Thus it would be more accurate to test whether the referent of the argument of ‘open’ is a thing or an event; or whether the conceptual content of the argument of ‘open’ is subsumed by THING or by EVENT.

This might seem a little pedantic but there are in effect two serious issues involved:

1. selectional restrictions should not be applied on the linguistic level, but rather on the conceptual level;
2. we still have to explain the exact realtionsship between the conceptual content of an expression and its referent.

We begin with the second issue which will eventually lead us to the first issue.

In general it would seem that the conceptual content of an expression is a type or a concept, and that its referent is a token or a particular individual, which is an instance of the conceptual content. This conception underlies Nunberg’s excellent analysis of reference [Nunberg 78]. His first criticism concerns the fact that the referent of an expression is not necessarily an instance of its conceptual content (Nunberg uses the terms 'designatum' and 'nominatum'). He lists a number of examples, some of which have become quite famous, to show that complex functions may be involved in order to determine the referent of an expression given its conceptual content. He offers a rough disctinction between *normal*, *local*, and *metaphorical* word-uses (p. 55).

(6) We had *chicken* for dinner last night. (p. 91, (1))

(7) The *ham sandwich* is waiting for his check. (p. 91, (10))

(8) That *snake* (John) would say anything to get elected (p. 91, (15))

The use of 'chicken' to refer to chicken meat is normal, since it can be rationally used out of context. The use of 'ham sandwich' to refer to the person who ordered/ate a ham sandwich is rational only in that specific context and thus local. Finally, the use of 'snake' is metaphorical in (8). Note that Nunberg "introduced these categories only for the sake of convenienc. The boundaries between them are rough" (p. 91).

What Nunberg's examples show in any case is that words can be used to refer to a wide range of objects. Given our framework of conceptual content and reference, we are thus left with two options:

1. We can enlarge the range of conceptual contents associated with expressions. Thus one possible conceptual content of 'ham sandwich' would be 'person who ordered/ate a ham sandwich'. We could then continue to say that the referent of an expression is an instance of its conceptual content.

2. We can loosen the requirement that the referent of an expression is an instance of its conceptual content and allow additional relations between referent and conceptual content (instantiation being the normal or preferred relation).

Nunberg shows that similar patterns found with descriptive terms are mirrored in ostension (Sect. 1.5). We can, for example, point to a newspaper and utter

(9) Hearst bought that.

thereby using 'that' to refer to the newspaper company. We take this as evidence for a theory which keeps the number of conceptual contents associated with an expression small and allows referents to be related to conceptual contents by other means than instanceship.

Note that Nunberg's criticism goes even further. If we can use words to refer to such a wide range of objects, how do we determine the literal meaning of a word? In other words, how do we know what the conceptual content of a word is, if we cannot define it as the type of objects which can be referred to by using the word (Sect. 3.2). Nunberg believes that this question can only be answered by taking into account the "normal beliefs" of a speaker or a community (Sect. 4.1) and that hence a purely semantic theory of meaning is not possible. This again shows that a distinction between linguistic semantic knowledge and general world knowledge is problematic. We will not pursue this issue here but we have to keep in mind that we have no clear-cut method to decide how broad the range of conceptual contents assigned to an expression should be. We can only base our decision on our intuition concerning the distinction between normal and local uses.

Let us assume for the time being that in normal uses the referents of expressions are instance of the respective conceptual contents, whereas in local and metaphorical uses there will be complex functions between conceptual contents and referents. We will consider non-normal uses when discussing selectional restrictions below. In the case of normal uses we have to decide how to model the various conceptual contents of expressions adequately. Nunberg criticizes linguistic treatments of polysemy according to which the various conceptual contents of a word are explicitly listed in the lexicon, yielding multiple lexical entries for each word:

> All these syntactic treatments of polysemy are subject to two related criticisms: they complicate the grammar unnecessarily, and they are unrevealing, giving us descriptions in place of explanations. (...) There are several excellent reasons, however, for supposing that many of the multiple uses of polysemous words don't have to be listed, since they would be generated by pragmatic schemata in any event (p. 13).

This idea has been taken up by several authors. The fundamental principle is to provide a single conceptual content for an expression from which additional conceptual contents can be generated. Note that this generation is triggered by the context, for example if the "original" conceptual content does not satisfy the selectional restrictions imposed by a functor. The process of substituting the "original" conceptual content by a conceptual content consistent with the contextual requirements is called *type shifting* or *type coercion*. Formal meaning shifts motivated by the type-theoretic framework of Montague Grammar have been proposed in [Partee 87, Partee 92]. In [Dölling 92] this approach is taken up and applied to domain-specific shifting operations, e.g. from mass nouns to count nouns or from plurals to singulars. Pustejovsky proposes to use *qualia structures* to control type coercion [Pustejovsky 93a].

Having stated the problem in its general form, we will now analyze its relevance in the VERBMOBIL context. Let us address normal uses first. Though we agree that having a list of type shifting operations which generate conceptual contents is highly desirable from a theoretical point of view, it is not indispensable from a system-oriented point of view. In fact, considering performance requirements it might even be more efficient to explicitly store the alternative conceptual contents of a lexeme than to generate them during run-time.

In any case, from the perspective of Domain Modeling and Semantic Evaluation it does not make much difference whether the conceptual contents associated with a lexeme are explicitly stored in the lexicon or generated by additional rules.[16] In both cases the respective conceptual contents have to be represented in the conceptual hierarchy, and methods for chosing the appropriate conceptual content in a particular context have to be devised.

Note that the situation is different for more complex type shifts, especially as in local or metaphorical uses. It seems that these cases cannot be handled with a small list of type-shifting operations but require complex reasoning. We will briefly consider the relevance of this phenomenon in the VERBMOBIL context. Reconsider example (7)—at first glance, one might try to ignore this phenomeneon in the context of Machine Translation for two reasons. First, one might argue, the example is idosyncratic and presupposes a rather specific context. Second, the phrase 'the ham sandwich' has to be translated into a corresponding target-language expression, and not into an expression corresponding to 'the customer who ordered a ham sandwich', i.e. the type shift is irrelevant for translation.

Though both arguments contain important insights, they are unfortunately not valid. Whereas the ham-sandwich example is in effect a rather extreme case of type-shifting, less extreme cases occur more frequently than one might expect. In the dialogues collected so far, we can distinguish, for example, at least three major ways of using the noun 'Termin'

**event:** in the sense of an appointment, meeting:

[16] The situation might be compared to the use of meta rules as, for example, used for passivization. It should not make much difference to the syntactic parser whether passive verb forms are epxlicitly stored in the lexicon or generated by meta rules.

- einen Termin beim Zahnarzt haben
- den ganzen Tag hindurch Termine haben
- zu unserem Termin treffen

**temporal:** in the sense of a time point or a time interval:

- einen Termin frei haben
- den Termin freihalten

**abstract:** as a unit of information:

- den Termin notieren
- einen Termin vergessen

As regards the second argument, namely the irrelevance of type shifting for translation, we have already argued that translation does not always require deep understanding. But in this example the type shift has to be dealt with. As mentioned above, the type shift is triggered by the verb phrase 'is waiting for his check'. There are at least three pieces of information indicating the shift:[17]

1. the predicate 'is waiting' cannot be applied to sandwiches;
2. the anaphor 'his' cannot refer to a sandwich (it would have to be 'its' instead);
3. 'its check' with 'its' referring to sandwich is not meaningful.

The second point is, for example, crucial when translating the German sentence

(10) Das Schinkenbrot wartet auf seine Rechnung.

into English. In order to translate 'seine' as 'his' and not as 'its' one has to know that it refers to the customer not to the sandwich.

The information sketched in points (1) and (3) is usually encoded in form of *selectional restrictions* to which we will turn our attention now.

[17] A fourth piece of information is the definite description 'the ham sandwich', where 'the' either introduces a discourse object being so prominent in the discourse setting that just to move to that setting "instantiates" this object as e.g. in "There is an election to parliament next month. The candidates have already .. ." or that the referent of the definite description has been already established in a preceeding part of the discourse.

### Selectional Restrictions

The main problem of selectional restrictions is that they are often used on the linguistic level, whereas their justification stems from the conceptual level. Above we said that the predicate 'is waiting' cannot be applied to sandwiches and we regard this as a conceptual constraint. More precisely, we might say that the lexeme 'wait' can express a concept INTENTIONAL_WAIT, which requires intentional agents, i.e. humans, institutions, and perhaps animals.

Note that this does *not* imply that the lexeme 'wait' can only be combined with subjects referring to potential agents. One can, for example, say

(11) This work will have to wait until later.

A couple of remarks seem in order:

1. Should we say that the verb 'to wait' has the same conceptual content in (7) and (11)?

2. If they do have the same conceptual content, what is the status of the selectional restriction to intentional agents?

3. If they do not have the same conceptual content, what is the relationship between the two conceptual contents?

We think that there are arguments both for claiming that the conceptual content is the same and for saying that it is different. According to our understanding, concepts are used to classify "real-world entities" (e.g. objects or events), i.e. they are *general* and constitute *abstractions over particulars*. NL expressions are then used to convey these abstractions, and can even abstract further, i.e. the same expression can be used to express different conceptual abstractions. Note that this is an example of the *efficiency of language* as defined in [Barwise, Perry 83, p. 32ff].

Since we only have direct access to the linguistic level, and (more or less direct) to the level of particulars, the conceptual level is highly underdetermined. Thus, using the same word in (7) and (11) can be taken as evidence for the same underlying concept. But taking a closer look at the particular events described, one might also opt for two different concepts of 'waiting'. As explained above, we will mainly rely on contrastive analyses to decide this issue. If both the source and the target language use the same word to express "two concepts", we will identify these concepts and include only one concept in the domain model. This is the main reason for calling our approach *transfer-oriented domain modeling*.

There are still a number of open questions regarding the treatment of sentences like (7). In particular, example (11) has shown that we cannot have a *strict* selectional restriction requiring that the subjects of 'two wait' have to refer to intentional agents. The sentence

(12) The ham sandwich is waiting for a customer.

is perfectly meaningful. Thus the type shift cannot be triggered by the verb 'is waiting' alone. It seems to us that it is the verb phrase as a whole, which triggers the type shift, especially the anaphoric pronoun 'his' in 'waiting for his check'. Thus it is complex world knowledge about checks, customers, ham sandwiches, and friends which triggers the type shift in (7), blocks it in (12), and triggers it in:

(13) The ham sandwich is waiting for a friend.

Note that a type-shifting reading, in a sense analoguous to (13), is also possible in (12), and one might even cook up contexts which block type shifting in (7) and (13). Roughly speaking one has to test all possible combinations of conceptual contents for a functor andits argument to find out the consistent ones.

We draw the following conclusions from these considerations:

1. Selectional restrictions should be modeled on the conceptual level.

2. If selectional restrictions are used on the NL-expression level, they can only have the status of preference rules, i.e. they behave like defaults and not like strict constraints.

The main reason for using selectional restrictions on the NL-expression level are

1. to speed up the performance, i.e. to allow disambiguation without recourse to complex semantic analysis;

2. to cope with partial information, i.e. to allow disambiguation in the absence of complex common-sense knowledge.

In the next section we will show in detail how information from the domain model in general and selectional restrictions in particular are to be used in VERBMOBIL. It should be noted that the modeling of selectional restrictions, as the domain model in general, is largely determined by the specific requirements of the interpretation tasks and not by general semantic considerations.

# 4 Domain Modeling in VERBMOBIL

In this section we describe the different uses of the domain model within VERBMOBIL. So far we see four main tasks which are supported by the domain model:

1. sortal information and selectional restrictions are used in the parsing process in order to eliminate "incoherent" hypotheses from the word lattice;

2. sortal information and selectional restrictions are used by Semantic Construction and Semantic Evaluation to disambiguate polysemous expressions and structural/referential ambiguities; the resulting conceptual representation is the basis for Transfer and Generation;

3. conceptual information from the domain model is used in the determination of the speech-event type of an utterance; a hierarchy of speech-event types for the VERBMOBIL demonstrator scenario is also part of the domain model;

4. the information in the domain model is used to build up a context representation, i.e. a semantic representation of the utterances in a dialogue. This context representation is used in the interpretation process.

Note that these tasks are partly independent from each other, partly related, and pose different requirements for the domain model. These different requirements have to be integrated into a single, consistent domain model.

## 4.1 Sort Hierarchy and Syntactic Parsing

Examining the word-hypotheses lattices produced by a recognizer besides perfect reproductions of the spoken utterances one finds

1. syntactically and semantically correct sentences, which do not match the spoken utterance; detecting such cases would require a precise pragmatic analysis, thus one just can hope that such cases some how fit into the ongoing discourse;

2. syntactically correct, but semantically deviant sentences, which can be detected by the application of selectional restrictions as shown below;

3. strings of words with no recognizable syntactic structure, which are discarded by the grammar.

In order to test the selectional restrictions most lexical entries are augmented by a "sort"-feature, whose value is associated with a concept of the domain model. The grammar has to provide for the percolation of the sortal information and to test it at appropriate grammar rules. Though these tests lead to wrong results in cases of type coercion or metonomies (cf Section 3.3), they are nevertheless necessary to detect recognizer results of the second type.

Consider the following examples where the first one is a recognizer mismatch, while the second is a case of type coercion (violations of sortal restrictions are emphasized):

(14) a. Dann muß ich *drei Uhr* auf den Freitag verschieben.
b. I have to postpone *three o'clock* to friday.

(15) a. Dann muß ich *Dreyer* auf den Freitag verschieben.
b. I have to postpone *Dreyer* to friday.
c. (lit.: I have to postpone the meeting with Dreyer to friday.)

If the process of semantic construction should not be burdened with the task to select the semantically sensible readings out of the parsing results, a solution to this dilemma, i.e. discarding meaningless utterances, but still have a chance to interpret cases of type coercion, is to apply the test of the selectional conditions as a soft constraint, which gives each selectional clash a penalty instead of immediatelly rejecting it. Thus the application of selectional restrictions results in an ordering, where the "sortally" best readings are presented first to semantic construction process.[18]

To provide tests for selectional restrictions the grammar rules have to be augmented by two kinds of features. At first the selectional conditions have to be percolated from the lexicon to the respective constituents. Secondly, whenever a daughter which serves as an argument or adjunct, is adjoined to a constituent, it has to be checked whether the selectional conditions are met.[19]

Additionally the lexical entries, with the exception of some classes of functional words, are augmented by their sortal information which are associated with

[18]Note that this ordering thus follows the strategy of *preferential disambiguation* as described in [Quantz, Schmitz 94].

[19]If this test should be applied to modifiers of nouns is still an open question. In the case of verb nominalizations, which still possess some kind of argument structure, such tests seem to be necessary, but what about the rest of the nouns? The extent to which the test of selectional restrictions has to be performed is additionally language-dependent. While German e.g. is quite free in this respect in conjoining NPs, in Japanese the nouns of a NP-conjunct must be subsumable under the same sort.

concepts of the domain model.[20] Lexical entries with an argument structure, such as verbs, adjectives, and prepositions also state conditions on the sorts of their arguments.

There are two types of tests to check selectional restrictions:

- A subsumption test between the sort of an argument position and the sort of the possible argument, where the first one has to subsume the second one. The same holds between the sort of a phrase and the sort of an adjunct.

- Checking whether a specific relation (a BACK-role) holds between two arguments, as it is the case for semantically empty verbs such as *to be* or *to have*. A distinctive feature for *to be* in the definitorial reading is that the sort of the subject has to be subsumed by the sort of predicate noun.

Since the sorts are BACK-concepts and the subsumption test is performed using the subsumption process of BACK, which differs from the notion of subsumption in unification, the test of the selectional restrictions is performed after the application of the parsing process.[21] The application of the selectional tests after the parser process has the drawback that the analysis is not directly cancelled, when a sortal mismatch is detected. But this strategy opens the opportunity to trigger further reasoning processes, as e.g. for interpreting type coercion on sortally illegal analyses if no correct analysis has been found. It should also be noted that the proper analysis of type coercion effects may depend on material analysed after the sortal mismatch.

As mentioned above the selectional conditions are soft constraints modeled as a bonus system. Each successfull application of such a condition increments the bonus counter by 1. Normalizing the total bonus figure with the number of all applications of selectional conditions of the utterance gives a measure for the selectional quality of an analysis. Thus we get the following "soft" quality criteria:

[20] At present, where the main emphasis is laid on the construction of the domain model, the association between sorts and concepts is simply a 1-1 mapping, which will be exchanged by a more sophisticated treatment of the lexeme-concept-mapping along the approaches described in [Bierwisch 83] and [Pustejovsky 89, Pustejovsky 93a].

[21] It might be argued that the typed feature structures as being used in HPSG or the mapping of ISA-Hierarchies into unifiable term structure (TFS) as proposed in [Mellish 88] allows a kind of subsumption inside the unification process. But these representation structures are less expressive than a description logic (DL) like BACK, so that the mapping from a DL-structure into a TFS to avoid the construction and maintenance of multiple models, which provide the domain knowledge, will throw away some necessary information.

$$\frac{N(successes)}{N(tests)} = \begin{cases} 1 & \text{literal meaning} \\ < 1 & \text{non-literal meaning or acoustic recognition error} \end{cases}$$

This bonus system needs one additional refinement in order to get a better treatment of pronouns. With the exception of 1st and 2nd person personal pronouns and some other pronouns, which refer explicitly to persons, the sort of a pronoun ought to be subsumable by any other sortal concept. The only concept that fulfills this condition is the concept NOTHING, which is very undesirable, because it also denotes a mismatch. Thus the remaining pronouns get a special signature onto which the subsumption test is not applied and their bonus is a figure slightly smaller than one. The reason for this smaller figure is twofold. On the one hand the sortal appropriateness of this analysis has still to be confirmed by the reference identification process of the DRT. On the other hand, if the grammar allows a pronoun being analyzed as a NP, as in

(16) a. an den (termin) habe ich nicht gedacht.
b. i forgot that date.

and when the definite article and the pronouns have the same form, then there are cases where it is by syntactic means nearly impossible not to analyse a "... Pron N ..."-sequence also as a sequence of two NPs. By the smaller figure for pronouns the interpretation of such a sequence as "... Art Noun ..." gets a better score, which is in most cases the appropriate reading.

Also prepositions need some special treatment for two reasons. Firstly, when a preposition serves just as a case marker for a PO-verb it lacks any sortal requirement on an NP. Therefore in this case no selectional test should be applied. Secondly, the sort of a preposition is usually its thematic role. Now the most frequent prepositions often denote a broad variety of thematic roles, where each one has distinguished requirements on the sort of its argument. If this subsumption test can not be performed on the fly, it results in numerous readings of the same syntactic structure. Therefore during the parsing only a very general sort is ascribed to a preposition, which is then refined, when due to the subsumption test the appropriate thematic role can be determined.

## 4.2 Conceptual Disambiguation and Representation

Information from the conceptual hierarchy is used both for conceptual disambiguation (interpretation task) and for conceptual representation (representation task).

Thus the conceptual information modeled in the conceptual hierarchy is motivated by requirements stemming from:

1. contrastive analyses for the transfer from German to English;
2. the determination of relevant information for resolving an ambiguous mapping from German expressions into concepts.

We will illustrate this by considering the translation of prepositions. First, a contrastive analysis yields for each German preposition a set of corresponding English prepositions. Having established the set of possible translations of a German preposition, the next step is to determine the information relevant for choosing the appropriate translation in a given utterance. In doing so, one has to decide whether to treat prepositions occurring in *(optional) complements* and in and *free adjuncts* homogeneously:

(17) ich muß sagen mir wär's dann lieber wenn wir die ganze Sache *auf Mai* verschieben.

(18) Wir treffen uns in der Eingangshalle des Czerczinsky *mit den Unterlagen*.

As Pollard and Sag point out, one difference between optional complements and free adjuncts is that the latter can occur with a wide range of heads and contribute a more or less uniform semantic content, whereas optional complements are more restricted and their semantic contribution depends idiosyncratically on the head [Pollard, Sag 87, p. 136]. Therefore, the translation of prepositions occurring in (optional) complements usually depends on the English verb chosen. For free adjuncts the translation of the preposition is more difficult, however, depending mainly on

1. the *internal argument*, i.e. the noun phrase constituting the prepositional phrase together with the preposition;
2. the *external argument*, i.e. the phrase modified by the prepositional phrase.

Both syntactic and semantic information about the internal and the external argument can be relevant for translation.

Having sketched this general methodology, it should be noted that there are two main problems concerning the details. For one thing, there are conceptual differences which are irrelevant for translation from German to English but seem too important to be ignored in the conceptual hierarchy. Thus the German preposition 'in' can be used with spatial and temporal noun phrases and is in both cases translated by the English preposition 'in':

(19) a. Wir treffen uns *in* der Eingangshalle.

b. We will meet *in* the hall.

(20) a. Wir treffen uns *in* zwei Wochen.

b. We will meet *in* two weeks.

This indicates that focussing solely on the linguistic level, i.e. transfer from German to English can produce "conceptually weird" hierarchies.

On the other hand, contrastive analyses might yield "overspecified" concepts. Consider the translation pair

(21) a. Mein Büro ist *im* zweiten Stock.

b. My office is *on* the second floor.

It could be argued that 'Stock' in German is conceptualized as a three-dimensional container, whereas 'floor' in English is conceptualized as a two-dimensional plane. Now we could either try to model this in the conceptual hierarchy by having conceptual relations specializing 'located', e.g. 'located_in_building_part'; or we could decide to use a general representation based on *located*. Note that the choice will have impact on the Transfer/Generation task. In particular, it is not yet clear whether the conceptual representation should unambiguously determine the English lexemes to be used, or whether the choice of lexemes is done by Generation on the basis of the conceptual representation.[22]

It should be noted that the conceptual disambiguation of an expression can thus be triggered by two different requirements:

1. in order to determine the corresponding target-language expression (representation task);
2. in order to determine the conceptual content of *another expression*, e.g. a functor or an argument of the expression (interpretation task).

## 4.3 Determination of Speech-Event Types

Given the demonstrator scenario, namely the appointment scheduling by business partners, the determination of the speech-event type of an utterance is crucial for an

[22]We have a slight preference for the latter approach, in which the choice of *on the second floor* instead of *in the second floor* is treated by Generation on the basis of English-specific, idiosyncratic, syntactic selectional restrictions.

adequate translation. First, it should be noted that this scenario is a highly conventionalized setting. This explains why most verbs occur frequently in stereotypical phrases, i.e. prefabricated utterances that express recurrent communicative functions [Coulmas 81]. It is a well know fact that the translation of theses phrases is—due to their intercultural variance—problematic.

The following examples will show the impact of the speech-event type of an utterance on its translation (see [Schmitz 94] for a detailed presentation). First, consider the following two utterances:[23]

(22) a. Ich muß sagen mir wär's dann lieber wenn wir die ganze Sache auf Mai verschieben. *Geht es da bei Ihnen auch?*

b. I must say that I'd rather postpone the whole matter until May then. *Would that suit you?*

(23) a. Wo sollen wir uns denn treffen? *Geht es bei Ihnen?*

b. Where shall we meet then? *Can we meet in your office?*

Though the second parts of these utterances are almost identical ('Geht es da bei Ihnen auch?' vs. 'Geht es da bei Ihnen?') their respective translations are completely different due to their different communicative function. In the first example 'Geht es da Ihnen auch?' is a request to comment on a proposed date, whereas in the second example 'Geht es bei Ihnen?' is a query whether a meeting can take place in the office of the dialogue partner. We thus have to know the speech-event type of these utterances in order to chose the adequate translation.

Second, consider the following parts of an interpreted dialogue [Bade et al. 94, App. II, No. 25] (again pauses and noises have been deleted from the transcription):

(24) **PIL:** Dann würde ich doch sagen, am Freitag, … dem sechsten Mai würde das gut passen. Vielleicht so gegen neunzehn Uhr?

**GRA:** (Then) on Friday, May sixth around seven?

In the VERBMOBIL data utterances like 'ich schlage vor', 'ich würde vorschlagen', 'ich hätte vorzuschlagen', 'dann könnte ich Ihnen aber noch anbieten' are frequently used to introduce a proposal of a date. These verbal phrases are generally not compositionally interpreted into the target language.

[23]Note that we have deleted information concerning pauses and noises from the transcriptions.

These examples, and the VERBMOBIL data in general, indicate that a compositional translation of the propositional content of utterances is not adequate for certain speech-event types. Instead in these cases it is rather the speech-event type itself, together with additional information, e.g. a temporal description of the date, which is the equivalent of translation. It is therefore important to determine the speech-event type of an utterance, in order to

1. determine the equivalent of translation;
2. decide whether certain phrases are used stereotypically or compositionally;
3. chose adequate lexicalizations for the respective speech-event type in the target language.

The domain model contains a hierarchy of speech-event types, such as 'Vorschlag' (proposal), 'Annahme' (acceptance) or 'Ablehnung' (declination), which will be presented in detail in Section 6.5.

Speech-event types as we define them here comprise more information then the mere illocutionary act in Austin's terms [Austin 62]). Each speech event type is defined in such a way that it combines a certain proposition with a particular illocution. These speech-event types resemble Wittgestein's *Sprachspiele* [Wittgenstein 56]. It is obvious that the set of language games is unlimited, therefore we only defined a set of speech event types that are characteristic for the scenario, i.e. appointment scheduling dialogues (for a definition of the speech-event types used for the VERBMOBIL-demonstrator cf. [Schmitz, Jekat-Rommel 94]).

In order to automatically assign speech-event types to utterances, we take the following knowledge into account:

**micro-structural information:** Here we consider syntactic information, semantic information concerning both words and sentences as well as local pragmatic information of *a single utterance* (without its context).

**macro-structural information:** Here we consider global pragmatic information resulting from the overall discourse structure and the background knowledge about calendar structures.

Some of this information, namely the conceptual content of (parts of)the utterances and calendar information is provided by the domain model, whereas other information is provided by Syntax or Semantic Construction. Our approach for the automatic assignment of speech event types to utterances is described in [Schmitz 94].

Though the speech-event hierarchy itself is domain specific, determining the speech-event type of an utterance is crucial for cooperative dialogues in general. It should be noted that the general approach used for the determination of speech-event types, namely the combination of micro-structural and macro-structural information modeled as preference rules can be applied to other scenarios as well.

## 4.4 Context Representation

Finally, one task of the domain model is to provide the concepts used to build up a context representation. The basic idea underlying this context representation is to provide macro-structural information useful for processing subsequent utterances, e.g. for determining their speech-event type of for resolving anaphora or ellipses. Though the general purpose of a context representation is thus clear, there are at least two open questions concerning the details.

First, in the demonstrator scenario the dialogues contain mostly English utterances and only few German utterances. In particular, only one speaker produces German utterances, whereas the other one speaks English. As long as there is no analysis of the English utterances, a context representation is neither feasible nor useful, since the German utterances are only represented in isolation. It has been agreed to use dialogues conducted entirely in German to test the contextual representation until an English analysis is available.

Second, it is not yet clear what exactly has to be represented in the context representation. A minimal representation comprising the speech-event type of an utterance and the temporal referents has been suggested at various workshops and project meetings (e.g. by Birte Schmitz and Christa Hauenschild):

(25) a. Dann lassen Sie uns doch noch einen Termin ausmachen.
o1 :: sprecher:s1 and init_terminabsprache and
no(vorgaenger_sprechhandlung)

b. Wann wär's Ihnen denn recht?
o2 :: sprecher:s1 and aufforderung_vorschlag and
vorgaenger_sprechhandlung:o1

c. Also ich dachte noch in der nächsten Woche, auf jeden Fall noch im April
o3 :: sprecher:s2 and positiver_vorschlag and
vorgaenger_sprechhandlung:o2 and
zeit_referenz:t1 and zeit_referenz:t2

t1 :: in_monat:april
t2 :: in_woche:14

d. Ja, am Dienstag den sechsten April hätt' ich noch einen Termin frei allerdings nur nachmittags
o4 :: sprecher:s1 and positiver_vorschlag and
vorgaenger_sprechhandlung:o3 and zeit_referenz:t2
t2 :: am_wochentag:dienstag and am_datum:6 and
in_monat:april and am_tageszeit:nachmittag

e. Geht es da bei Ihnen auch?
o5 :: sprecher:s1 and aufforderung_stellungnahme and
vorgaenger_sprechhandlung:o4

Such a representation can be used as a basis both for Transfer/Generation and for determining the speech-event type of the subsequent utterance. By combining the temporal descriptions in the various utterances it can be detected whether a proposed date is refined (the temporal description can be consistently added) or whether a new date is proposed (the temporal description is inconsistent with the temporal description of the previously proposed date).

Note that such a representation should be seen as the minimal basis, which might be augmented if more information is needed. For example, it has to be seen whether the exact determination of temporal referents is necessary. Consider the following example:

(26) a. Donnerstag in der ersten Juni-Woche ist bei mir schon belegt , wie wär's denn dann am Montag drauf am siebten Juni ?

b. am siebten Juni habe ich morgens leider eine Konferenz , wenn Sie also nachmittags kommen könnten das wäre okay

c. leider hab' ich erst Montag Abend Zeit und ich denke daß das dann zu spät wird , deshalb schlage ich vor daß wir das auf Donnerstag verschieben

Obviously 'Donnerstag' in (c) refers to Thursday, June 10. While it is easy to detect that (c) contains a new proposal and not a refinement ('Montag' and 'Donnerstag' cannot be consistently combined), it takes more complex reasoning to determine the part of the previously proposed date which is, so to speak, kept constant (namely the week in this example). Note that it is not crucial to determine this information in order to translate the utterance—'Donnerstag' is rendered by 'Thursday' regardless to which Thursday it refers.

Interestingly, there are examples in the dialogues in which misunderstandings arise wrt the exact reference of a proposed date:

(27) ja am Dienstag den sechsten April hätt' ich noch einen Termin frei allerdings nur nachmittags. geht es da bei Ihnen auch?

(28) oh das is' schlecht , da habe ich um vierzehn Uhr dreissig einen Termin beim Zahnarzt. aber Donnerstag Vormittag so um neun wär' mir recht

(29) ist das der achte April?

(30) Oh ich dachte eigentlich an den fünfzehnten April. eine Woche später

In principle, default rules could be used to copy all the parts of a previously focused date which are not explicitly specified for a newly proposed date. Or more accurately, one might copy the "more general information", e.g. week, month, and year if a new day is proposed; also the day if a new time is proposed, etc. We will have to investigate this problem more closely in future work. Here we just want to emphasize that it is still an open question to which degree such a determination is needed for the VERBMOBIL project. From a representational point of view, i.e. considering the representational basis for Transfer/Generation it seems to be dispensable (whereas it would be indispensable for automatic appointment-scheduling systems). Determination of reference might be needed for the interpretation of subsequent utterances, however.

A major issue concerns the resolution of anaphora and ellipses. Consider the following examples:

(31)
a. der Termin den wir neulich abgesprochen haben
b. am zehnten an dem Samstag *da kann* ich doch nich';
c. wir sollten *einen anderen* ausmachen.

First consider the verb 'kann' in (b).[24] Usually, the verb 'können' subcategorizes for an infinitive construction, e.g. 'ich kann teilnehmen'. To parse (b)

1. we either have to treat it as an ellipsis and have to decide whether to "add"a missing infinitive such as 'teilnehmen';

[24]This phenomenon has been analyzed in detail by Jürgen Kunze at the TP11 Workshop in Berlin (6.–8. April 94).

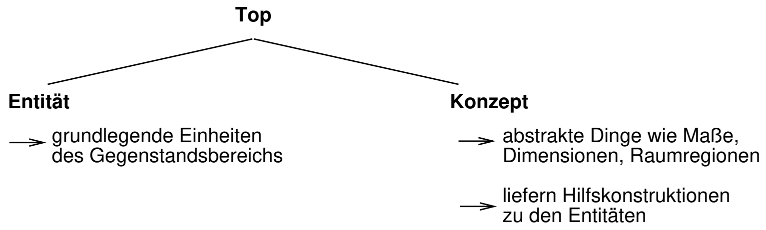


Figure 4: Topl-level categories in LEU/2 [Klose et al 92].

2. or we have to add a lexical entry allowing 'können' to act as an intransitive verb.

The second solution is more adequate in the VERBMOBIL demonstrator scenario since it avoids complex reasoning necessary for resolving the "ellipsis" but irrelevant for translating the utterance.

Similarly, it is not clear whether the anaphor 'da' in (b) has to be resolved. It is important to recognize that 'da' is used anaphorically and that it refers to a temporal entity (this information is, for example, used to determine the speech-event type). However, the information that 'da' refers to Saturday, 10 seems to be irrelevant for translation. Again, one might implement default rules stating that anaphoric temporal expressions always refer to the date currently focussed on (this seems to be a quite accurate heuristics wrt the collected dialogues).

We think that these examples show that the relevance of such phenomena within VERBMOBIL has to be carefully analyzed. Whereas in some cases a deep analysis and a resolution of ambiguities is necessary for translation, there will also be cases where translation can be achieved without such a deep analysis. It might be more adequate to provide underspecified representations for these cases in order to avoid redundant time-consuming inference processes. In other words, one has to decide whether a particular disambiguation is required by the application, i.e. translation, or by the formalism itself, i.e. HPSG or DRT. In the latter case it is more appropriate to adopt the formalism to the translation task instead of doing irrelevant disambiguations.

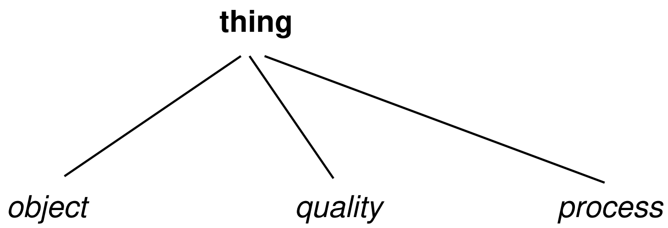


Figure 5: Top-level categories in PENMAN [Penman 89].

# 5 Designing Ontologies

There are a number of NLP systems which use knowledge bases containing conceptual hierarchies or ontologies. A comparison of the respective basic categories used in

- LEU/2 [Klose et al 92] (Figure 4),
- PENMAN [Penman 89] (Figure 5),
- KBMT-89 [Goodman and Nirenburg 91] (Figure 6),
- an CYC [Lenat, Guha 90] (Figure 7)

reveals various possibilities of modeling these top-level categories. We think that the main reason for the differences between these ontologies stems from the fact that that they are used for rather different tasks in the respective systems.

LEU/2 is a system modeling text understanding and its knowledge base is used to support different types of inferencing, namely

- to analyze linguistic surface structure,
- to compose an internal conceptual representation of the contents of the text,
- to support reasonable question answering, and
- to provide access to implicit text knowledge. [Klose 93, p. 63]

PENMAN is a system for Text Generation system and its knowledge base is used mainly to support generation. KBMT is a system for Machine Translation, in

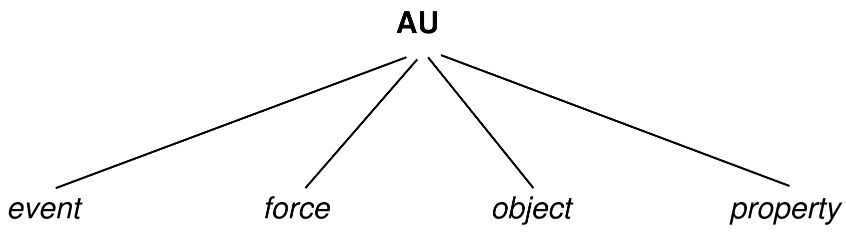


Figure 6: Top-level categories in KBMT [Goodman and Nirenburg 91].

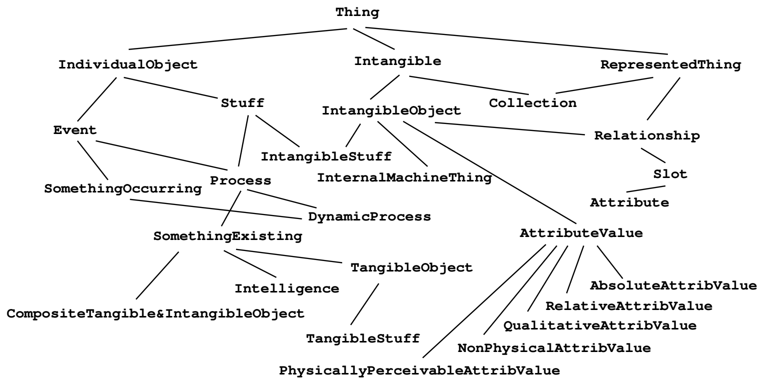


Figure 7: Top-level categories in CYC [Lenat, Guha 90].

which the ontology acts like an interlingua. The CYC project aims at developing a huge, task-independent knowledge base, which contains common-sense knowledge required for all types of "intelligent systems".

Investigations in a quite different area, namely *cognitive development* , may give additional insights how to represent conceptual knowledge and to relate it to language. In [Keil 79, Keil 89] the development of ontological knowledge is investigated, where natural kinds are the subject of the first monograph while the second one concentrates more on nominal kinds. Phenomena, where ontological knowledge shows up, are:

- sentences, which are neither true nor false, but anomaluous such as *"The table is hungry."*,

- naturalness and similarity of classes of concepts, where e.g. a class constituted by (**Human, Animal, Plant**) is easier to grasp than (**Human, Minerals, Number**), while in a class, e.g. consisting (**Human, Animal**) the concepts are more similar to each other than in (**Human, Plant**),
- copredication, i.e. not all pairs of predicates are sensible, as e.g.
  **+** *X is big and red*,
  **-** *X is ungrammtical and waterproof.*

All these different phenomena address the problem of predicability, namely which predicates can be combined with which natural language terms so that the predicate-term-combination is sensible and a truth value can be assigned to it [Sommers 65]. As a consequence predicability imposes quite strict criteria on the organisation of ontological knowledge in a highly structured manner namely a non-tangling tree.

> "In Sommers' theory, there is a more structured relation between predicates and terms. Two terms are of the same type if and only if all predicates that span term $t_1$ also span $t_2$. Similarily, two predicates $P_1$ and $P_2$ are of the same type if they span exactly the same sets of terms. If term $t_1$ is spanned by a subset of the predicates that span $t_2$, then $t_1$ is a member of a category *C* while $t_2$ is a member of a category subordinate to *C*. An essential part of this theory is that terms and predicates always sort themselves out in the same manner; that is, either two terms share exactly the same predicates, or one shares a subset of the other's predicates, or they share no predicates at all. There can never be a case where two terms have intersecting sets of predicates."
> [Keil 79, p.14]

Applying these criteria[25] to a tree of predicates gives a tree that has no diamonds in it[26], i.e. no subordinate concept has two or more mother concepts immediately preceding it. It should be kept in mind that these criteria apply to concepts, not to words, so that a word like *bat* will express two concepts, one below **Artifact**

[25] This so called "M-constraint" states that no tree should include an "M" or "W"-shaped substructure.

[26] I.e. the occurence of such "diamonds" in the real world, makes it hard to learn and understand phaenomena related to it. This may be one reason, why quantum physics, e.g. the wave-particle dichotomy, is not easily understandable.

and one below **Animal**. Such a separation of a word into several concepts is not only necessary for polysemous words, but also for heterotypical words where the concepts they express are related by more functional aspects [Sommers 71, Bierwisch 83].

In numerous developmental studies Keil has shown that though the tree structure becomes more elaborated and partly restructured during a child's growing up, the tree structure remains surprisingly stable not only between children and adults from the same social environment, but also between different cultural surroundings [Keil 79]. The most upper part of an ontology according to these studies is shown in Figure 8.

The developmental studies reported in [Keil 89] are centered around nominal kinds, such as moral acts, artefacts etc.. The setting in these studies was to give the children definitions of such terms in the form of little stories and then to ask them, whether the action or thing can be an X. In general the results indicate that there is a shift from mere perceptual similarity over characteristic features to definitorial conditions (necessary and/or sufficient ones). But this shift does not occur as a general and simultaneous advance in all conceptual domains, but is on the contrary quite domain specific. Keil concludes from these findings, that the construction of a causal model or theory for a domain is the driving force of this shift. And if there is no such causal model for a domain, then also adults use perceptual similarity or characteristic feature as a fallback strategy.

# 6 VERBMOBIL Domain Model Version 1.0

## 6.1 The Description-Logic Systems BACK and FLEX

We use the Description-Logic (DL) systems BACK and FLEX to model the VERBMOBIL domain model. In this section we will briefly sketch the main ideas of DL-systems. BACK V5 is described in detail in [Hoppe et al. 93]. FLEX is an extension of the BACK system, which is currently developed at TU Berlin to meet the specific requirements arising in the VERBMOBIL application [Quantz et al. 94].

Description Logics (DL) can be seen as a formal elaboration of the ideas underlying *Semantic Networks* [Quillian 68] and *Frames* [Minsky 75]. Following the debate on the lacking formal foundation of these representation formats in the mid-70's, Brachman proposed the representation language KL-ONE, which is described in an overview, which was circulated in the beginning of the 1980's and was finally published in 1985 [Brachman, Schmolze 85]. In the last decade, sub-

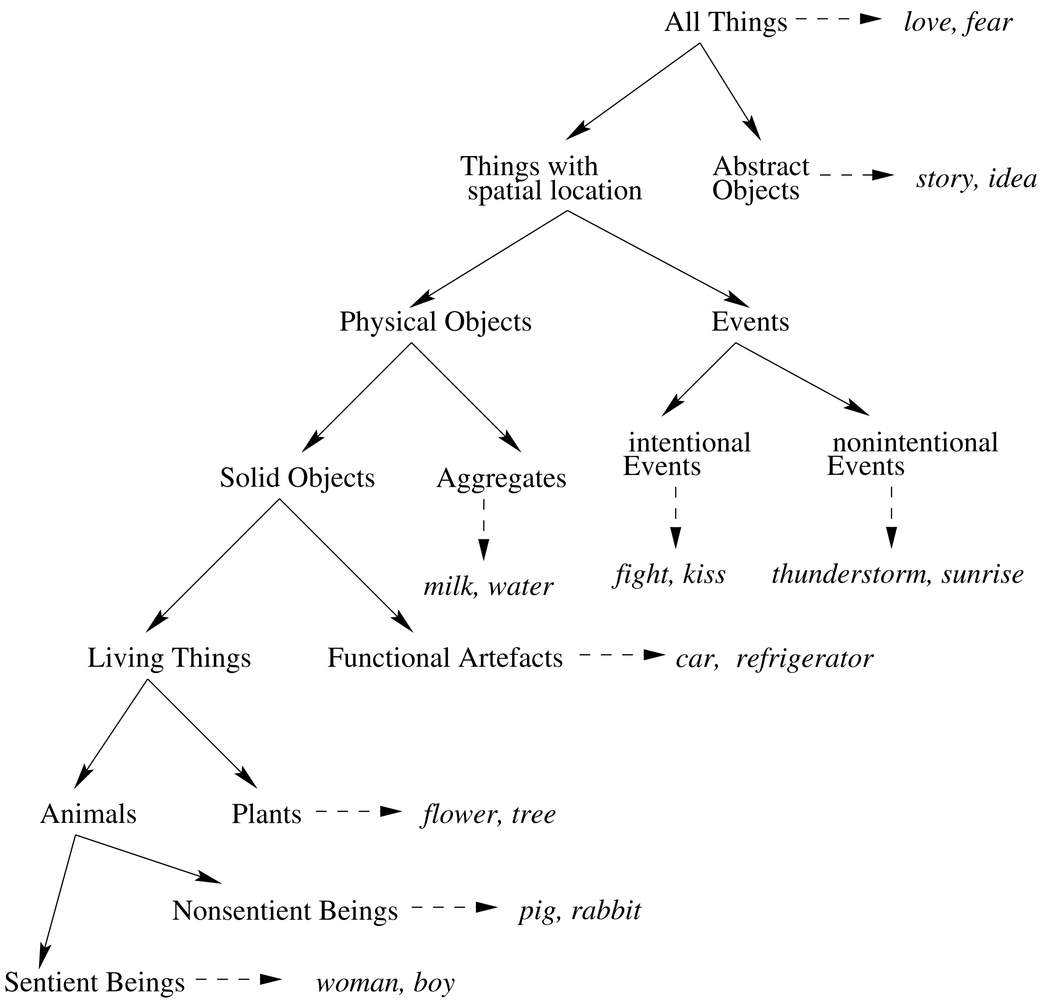


Figure 8: The ontology in [Keil 79, p.16].

| | | |
|---|---|---|
| product | :< | anything |
| chemical product | :< | product |
| biological product | :< | product & not(chemical product) |
| company | :< | some(produces,product) |
| produces | :< | domain(company) |
| chemical company | := | company & all(produces,chemical product) |
| some(produces,chemical product) | => | high risk company |
| toxipharm | :: | chemical product |
| biograin | :: | biological product |
| chemoplant | :: | chemical company |
| toxiplant | :: | atmost(1,produces) & produces:toxipharm |

Figure 9: A sample DL modeling.

stantial theoretical research concerning complexity and decidability of different dialects of Description Logics has been conducted (e.g. [Donini et al. 91]). In parallel, several DL systems have been implemented and used in various applications (e.g. [Brachman et al. 91, Hoppe et al. 93, Quantz, Schmitz 94]).

In DL one typically distinguishes between *terms* and *objects* as basic language entities from which three kinds of formulae can be formed: *definitions*, *descriptions*, and *rules* (see the sample modeling on page 50 below). A definition has the form $t_n \doteq t$ and expresses the fact that the name $t_n$ is used as an abbreviation for the term t. A list of such definitions is often called *terminology* (hence also the name Terminological Logics). All DL dialects provide two types of terms, namely *concepts* (unary predicates) and *roles* (binary predicates), but they differ with respect to the term-forming operators they contain. Common concept-forming operators are: conjunction ($c_1 \sqcap c_2$), disjunction ($c_1 \sqcup c_2$), and negation ($\neg$ c), as well as quantified restrictions [Quantz 92] such as value restrictions ($\forall$r:c), which stipulate that all fillers for a role r must be of type c, or number restricitions ($\geq n$ r:c or $\leq n$ r:c), stipulating that there are at least or at most $n$ role-fillers of type c for r. Role-forming operators are, besides conjunction, disjunction, and negation, role composition ($r_1.r_2$), transitive closure ($r^+$), inverse roles ($r^-$) and domain or range restrictions ($_c|r$ or $r|_c$). In a description, an object is described as being an instance of a concept (o :: c), or as being related to another object by a role ($o_1$ :: r:$o_2$). Rules have the form $c_1 \Rightarrow c_2$ and stipulate that each instance of the concept $c_1$ is also an instance of the concept $c_2$.

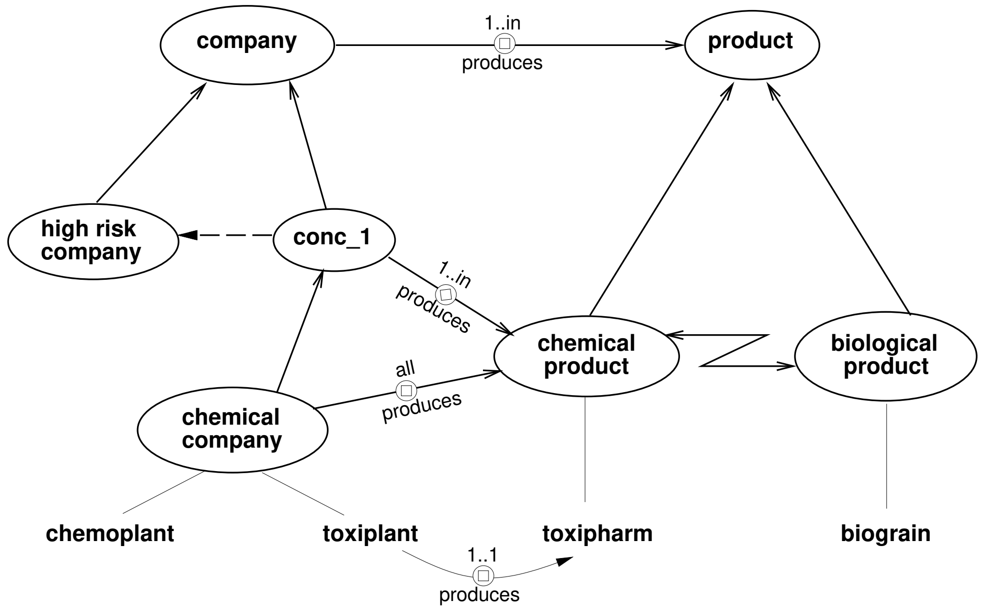


Figure 10: The net representation of the sample modeling. 'conc_1' is the concept some(produces,chemical product).

The example in Figure 9, whose net representation is shown in Figure 10 illustrates the most important aspects of a DL modeling. One role and five concepts are defined, out of which four are primitive (only necessary, but no sufficient conditions are given). Furthermore, the modeling contains one rule and four object descriptions.

In DL, such a modeling is regarded as a set of formulae $\Gamma$. Given the formal semantics of a DL, such a set of formulae will entail other formulae, i.e., there is an entailment relation $\Gamma \models \gamma$. Now the service provided by DL systems is basically to answer queries whether some formula $\gamma$ is entailed by a modeling $\Gamma$. The following list contains examples for the types of queries that can be answered by a DL system:

- $\Gamma \models \mathrm{t}_1 \sqsubseteq \mathrm{t}_2$
  Is a term $\mathrm{t}_1$ more specific than a term $\mathrm{t}_2$, i.e., is $\mathrm{t}_1$ *subsumed* by $\mathrm{t}_2$? In the sample modeling, the concept 'chemical company' is subsumed by 'high risk company', i.e., every chemical company is a high risk company.

- $\Gamma \models \mathrm{t}_1 \sqcap \mathrm{t}_2 \sqsubseteq \bot$

Are two terms $t_1$ and $t_2$ incompatible or disjoint? In the sample modeling, the concepts 'chemical product' and 'biological product' are disjoint, i.e., no object can be both a chemical and a biological product.

- $\Gamma \models o :: c$
  Is an object o an instance of concept c (object classification)? In the sample modeling, 'toxiplant' is recognized as a 'chemical company'.

- $\Gamma \models o_1 :: r{:}o_2$
  Are two objects $o_1$,$o_2$ related by a role r, i.e., is $o_2$ a role-filler for r at $o_1$? In the sample modeling, 'toxipharm' is a role-filler for the role 'produces' at 'toxiplant'.

- $\Gamma \models X :: c$
  Which objects are instances of a concept c (retrieval)? In the sample modeling, 'chemoplant' and 'toxiplant' are retrieved as instances of the concept 'high risk plant'.

- $\Gamma \cup \{\alpha\} \models \perp$
  Is a description $\alpha$ inconsistent with the modeling (consistency check)? The description chemoplant :: produces:biograin is inconsistent, wrt the sample modeling, i.e., 'biograin' cannot be produced by 'chemoplant'.

In Section 7 we briefly show how such queries can be used to support disambiguation.

## 6.2 The Conceptual Hierarchy

In this section we briefly describe the main parts of the conceptual hierarchy, namely the *sort hierarchy*, the *entitity hierarchy*, and the *situation hierarchy*. The contents of the conceptual hierarchy are largely determined by the lexemes contained in the demonstrator word list.

### The Sort Hierarchy

The sort hierarchy shown in Figure 11 comprises the basic categories of the domain model and thus reflects the basic conceptual distinctions. Sortal information of this general level will be used already in Syntax and Semantic Construction,

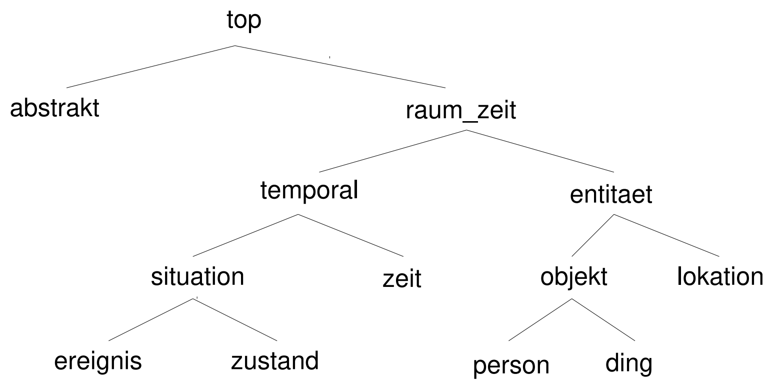


Figure 11: The sort hierarchy.

whereas the more specific conceptual distinctions will be only used by Semantic Evaluation.[27]

It should be noted that on this general level the subsorts of a sort are mutually disjoint. One motivation for this particular sort hierarchy was therefore to introduce sorts comprising those subsuorts between which types shifts occur frequently.

The basic distinction is the one between RAUM_ZEIT and ABSTRAKT. The former one covers objects occurring within time and/or space whereas the latter one covers abstract objects as numbers or sciences.

The sort RAUM_ZEIT splits into TEMPORAL and ENTITAET, i.e. into more temporal and more spatial objects. The sort TEMPORAL can be used to account for the systematic type shift from SITUATION to ZEIT; the sort ENTITAET for the shift from OBJEKT to LOKATION.

The distinction between PERSON and DING is used to distinguish between objects occurring as agents, i.e. humans, institutions, and animals and non-agents. The distinction between EREIGNIS and ZUSTAND captures the distinction between dynamic and static situations.

### The ENTITAET Hierarchy

The sort ENTITAET denotes all entities which are material, i.e. spatial objects and substances. The subsorts of ENTITAET are LOKATION, SUBSTANZ and OBJEKT.

[27]The design of the sort hierarchy has been based on discussions with Walter Kasper, Scott McGlashan, and Sebastian Millies.

The LOKATION describes the spatial information of objects which is relevant for the disambiguation of prepositions, e.g. the number of dimensions, information about fuzzy or strict boundaries.

The distinction between LOKATION and OBJEKT is motivated by the systematic type shifts occurring between these entities. To handle these type shifts all "concrete" subsorts like GARTEN (garden) , MENSCH (human being) etc. are underspecified wrt to this distinction. Nevertheless there are fully specified concepts for every subsort i.e. for every underspecified <NAME> there is a sort <NAME>_OBJ, which is subsumed by OBJEKT and a sort <NAME>_LOK, subsumed by LOKATION. Technically this is achieved by using a set of macros, expanded by the preprocessor *m4*. The use of an external macroprocessor is necessary, since BACK and FLEX only support macro definitions in a rather restricted way.

During disambiguation an object of the sort <NAME> can be specialized to <SORT>_LOK in case of the location reading, or to <NAME>_OBJ. This specialisation can be triggered on demand, e.g. if it is necessary for the disambiguation of prepositions, thereby allowing a *variable depth of analysis*. Furthermore, if the object is specialized to _OBJ, the information about its spatial attributes is not lost, because the different specializations (i.e. the _LOK, _OBJ, _SUBS concepts) are linked by specific roles.

### The SITUATION Hierarchy

Situations are entities, which are temporarily and spatially determined. Thus all roles, which are applicable on the concepts RAUM_ZEIT and TEMPORAL are also applicable on situation concepts. Lexemes expressing these concepts are either verbs, (most) adjectives, nominalizations of verbs and nouns denoting events. The "situation"-model provides the domain knowledge for these lexemes,[28] while other aspects of meaning, such as aspect/aktionsart or speech acts, are modeled at other places. Since all these aspects contribute to the overall meaning, the different aspects should not be inconsistent with each other.

At present the representation of SITUATION distinguishes four different points of view onto this concept, namely the "site of the situation", the kind of process, the distinction between event and state,[29] and the perspective from which the situation is seen.

[28]We would like to thank Rita Nübel for the requirements on verb concepts from the point of view of transfer and some very valuable hints.

[29]This distinction may in a later stage of the project rely on the calculation of aspekt/aktionsart.

The concept for the "site of the situation" (PSYCH_SIT) is subcategorized into mental and nonmental situations, where mental situations provide a further distinction into cognitive, perceptive, and communicative situations. Note also that cognitive situations include—atleast at present—emotional situations.[30] The concept of mental situations is partly motivated by the fact that not only communicative processes but also cognitive and perceptive processes are related, atleast in the VERBMOBIL setting, to speech acts.

Viewed from the process perspective, a situation either describes that something is in motion or that some entity is related to some other entity. Both process concepts are furtheron subdivided into

**motion in general:**

- activity (HANDLUNG), where the actor is not changing its place, but may make objects move around,
- movement (FORTBEWEGUNG), where the actor is involved in a change of place,
- interaction (INTERAKTION), where atleast two actors are involved,[31]

**relational:**

- relating an entity to a spatio-temporal location (REL_LOK),
- relating an entity to another entity to define or identify it (REL_CLASS),
- expressing possession respective non-possession in a general sense (REL_POSS).

Finally, the perspective from which the situation is seen describes a situation as being agent centered, affected-object centered, or process centered. This distinction may be helpful for analysing causation phenomena. Figure 12 shows the upper structure of situations and indicates whether a role is applicable ('+') or not ('–').

The concepts described so far constitute an upper structure from which via cross-classification more specific concepts are defined, which may then be related

[30] The role EMOTIONAL with role fillers *yes* or *no* provides this distinction.

[31] It should be noted, that the represention of interactions as a direct subconcept of motion in general is a provisional one, since it can also be seen as an activity with two actor (e.g. to make an agreement) or as a movement (e.g. when the two actors meet). Thus it may be reasonable to subdivide HANDLUNG and FORTBEWEGUNG into one-actor and multi-actor processes.

| | situa-tion-type | pro-cess-type | psych | ment-al-sit-type | emoti-onal | non-sta-tive | change of place | ac-tors | rel-type |
|---|---|---|---|---|---|---|---|---|---|
| psych_sit | + | + | + | - | - | - | - | + | - |
| process | + | + | + | - | - | + | - | + | - |
| ereignis | eventP | + | + | - | - | - | - | + | - |
| zustand | stateP | + | + | - | - | - | - | + | - |
| agtc-process | + | ag_cent | + | - | - | - | - | + | - |
| affc-process | + | aff_cent | + | - | - | - | - | + | - |
| proc-process | + | proc_cent | + | - | - | - | - | + | - |
| mental-sit | + | + | ment. | + | - | - | - | + | - |
| nonmen-tal-sit | + | + | non-ment. | - | - | - | - | + | - |
| perzept. | + | + | ment. | perz | - | - | - | + | - |
| kognit. | + | + | ment. | intern. | no | - | - | + | - |
| kommu-nikat. | + | + | ment. | kom | - | - | - | + | - |
| emotion | + | + | ment. | intern. | yes | - | - | + | - |
| beweg. | + | + | + | - | - | yes | + | + | - |
| relation. | + | + | + | - | - | no | - | + | + |
| interakt. | + | + | + | - | - | yes | no | $>1$ | - |
| handlung | + | + | + | - | - | yes | no | 1 | - |
| fortbeweg. | + | + | + | - | - | yes | yes | 1 | - |
| rel_lok | + | + | + | - | - | no | - | + | lok |
| rel_class | + | + | + | - | - | no | - | + | class |
| rel_poss | + | + | + | - | - | no | - | + | poss |

Figure 12: The upper structure of the SITUATION hierarchy.

in some way or other to lexemes. To model such specific concepts some additional roles are of course necessary. A first set of these roles are the thematic relations and their more transfer-specific subrelations, which are discussed in section 6.3. Thus besides their role in the syntactic/semantic interface they get an additional function in the conceptual modeling. Reusing thematic relations at this level may be too straightforward, since they entail some strong commitments. But these commitments don't seem to be too severe. Furtheron, it is a matter of convenience not to have an additional mapping process.

In the following list additional roles are described, which are necessary to model situative concepts:

- The role HANDLUNG_DIR characterizes a process as being either directed towards the actor or away from it.
- The role INTENDIERT states whether a situation is accidental or not, while CAUSATIV and CAUSATIV_FUER denotes the causer resp. the causee of a situation.
- The role PUNKTUELL distinguishes between punctual and extended situations.
- Finally the roles INCHOATIV and KONKLUSIV characterize the starting resp. the ending phase of a situation.

Additionally, it seems worth to consider, whether situative concepts should be characterized by their typical duration, because such an information may support some decisions in the area of aspect and aktionsart. But since just typical amounts of time can be expressed, the modeling of this feature should be done with defaults.

To put some flesh to the modeling of situations, some examples are explained in the following. At first some conceptualisations for cognitive situations are shown:

wissen := kognition and zustand and process_typ:ag_cent
and emotional:no and the(agent,person) and the(theme,anything)
and sicher_p:sicher

denken := kognition and ereignis and process_typ:ag_cent
and emotional:no and the(agent,person) and the(theme,anything)
and intendiert:yes

einfallen := kognition and ereignis and process_typ:aff_cent

and emotional:no and the(experiencer,person) and the(c_o_s,anything)
and c_o_s_type:kreativ and punktuell:yes

WISSEN is a agent-centered state with SICHER_P indicating that the THEME is for the actor sure knowledge as opposed to e.g. the concept GLAUBEN. The cognitive event conceptualisations DENKEN und EINFALLEN differ wrt. the actor which is in one case intentive and an agent, while in the other case it is an experiencer and the event is punctual. They also differ wrt. to THEME being an INCREMENTAL_THEME in EINFALLEN.

As an additional example WARTEN and TREFFEN are both non-mental situations and activities (HANDLUNG) as well as agent-centered, but they differ wrt. to the situation-type and TREFFEN not necessarily having a theme:

warten :< handlung and nonmental_sit and process_typ:ag_cent
and the(agent,person) and the(c_o_s,raum_zeit)
and c_o_s_type:def_n_subinterval

treffen :< handlung and ereignis and nonmental_sit
and process_typ:ag_cent and intendiert:yes
and the(local,lokation) and all(agent,person)
and atleast(2,agent) and handlung_dir:zu_ag
and the(teil_von,interaktion).

It should be kept in mind that the main purpose of the domain model is to support the discrimination of word sense for transfer and semantic evalution and not to give a general definition of what a lexeme may mean. A consequence of this approach to modelling world knowledge is that it can not be entailed from the non-existence of a thematic relation in the conceptualisation of a situation that a phrase with the lexeme related to this conceptualisation should not have this thematic relation as an adjunct. An examination whether such an entailment is sensible is among others subject of further research.

## 6.3 Thematic Relations

Though thematic relations have been used in numerous approaches and an abundant number of articles has been written on this subject[32] there is still no consensus on their nature and their definition in sight. But their repeated use shows the

[32]To mention just a few of these publications, either for their influence or for giving an overview, the reader is referred to [Gruber 65], [Fillmore 68], [Jackendoff 83, Jackendoff 90], [Rauh 88], [Dowty 91].

very necessity of a notion of such a kind. A closer look at these numerous approaches, allows one to mainly distinguish two different research directions. On the one hand thematic relations are seen together with some primitive operators such as GO, CAUSE, etc. as building stone of conceptual structure - a position, most prominently represented by R. Jackendoff. Despite a lot of attractive features Jackendoff's approach includes some very strong commitments, most notably the direct mapping between the linguistic surface and the conceptual structure.

On the other hand, they are seen as part of the syntactic/semantic interface in order to state semantic generalizations over grammatical functions. Thus their task is either to group together different grammatical function, as is the case with the subject of an active sentence and the *by*-phrase of a passive one, or to make a more fine-grained distinction of one grammatical function, where e.g. the thematic relation of the subject is an AGENT, EXPERIENCER, or THEME depending on the type of the verb.

In this approach thematic relations are ascribed to arguments and to adjuncts, while sentential modifiers are excluded. Nevertheless the border between adjuncts and sentential modifiers is not clear-cut, since some adjuncts may also serve as sentential modifiers. But this would add another complication to the already exisiting ones.

One problem that attracts much interest, possibly due to its frequency, is the clear distinction between AGENT, PATIENT, and THEME. Usually, the AGENT performs the action, the PATIENT is affected by it, and the THEME is the entity in motion or undergoing a change of state. The distinction between the first two relations seems to be obvious, but both may also be classified as THEME, which may lead to unexpected results depending on the ordering of the classification statements.

A solution to this probleme, which at least partially gives up the "one actant, one relation"-principle, is the introduction of an additional *causal* dimension as proposed in [Jackendoff 90] or [Grimshaw 92], which consists of AGENT and PATIENT, while THEME remains in the original one. The above mentioned principle is thus weakened to a "one actant, one relation from a dimension"-principle, thus allowing THEME to be coindexed either with AGENT or under certain conditions (see below) as PATIENT. This coindexing, which is explictly triggered by some features, will prevent the classification of the AGENT as a PATIENT.

But nevertheless there are still border cases, which make it hard to decide whether an argument is an AGENT or PATIENT. A recent approach by D. Dowty in [Dowty 91] breaks down agent-hood as well as patient-hood into a number of features or proto-relations, where their existence or non-existence contributes to

| *Relation* | *Super Role* | *Domain* | *Range* |
|---|---|---|---|
| agent | them_rel | situation | raum_zeit |
| experiencer | agent | situation | person |
| patient | them_rel | situation | anything |
| theme | them_rel | situation | anything |
| way_rel | them_rel | event | anything |
| source | way_rel | event | anything |
| path | way_rel | event | anything |
| goal | way_rel | event | anything |
| local | lokalisiert them_rel | situation | raum_zeit |
| temporal_rel | them_rel | situation | temporal |
| instrumental | them_rel | event | objekt |
| benefactive | them_rel | event | person |
| concomitant | them_rel | event | ding |
| final | them_rel | situation | event |
| causal | them_rel | situation | event |
| modal | them_rel | situation | anything |
| comitative | modal | situation | person |
| incremental_theme | theme patient | event | anything |

Figure 13: Thematic relations in Version 1.0.

the formation of an AGENT- or PATIENT-relation. This approach seems to be quite promising, but this still has to be worked out in detail.

Another observation in [Dowty 91], namely the identification of INCREMENTAL_THEME as a subtype of THEME which contributes to patient-hood, is also worth being introduced as a thematic relation. An INCREMENTAL_THEME denotes objects which are the traditional "effected" or "destroyed" objects or objects undergoing a *definite* change of state. Examples of an INCREMENTAL_THEME are "to write an article" or "to determine the next meeting".

Note that the distiction between AGENT and EXPERIENCER is explicitly modeled by a rule stating that the experiencer has to be a non-volitional agent.

Based on the up-to-now announced requirements to describe verbs and prepositions for the purposes of semantic evaluation and transfer[33] a list of thematic relations has been compiled, which may still be incomplete. Figure 13 contains the thematic relations contained in the current domain model. For each relation we also list

- its domain, i.e. the concept on which the relation is applicable,
- its range, i.e. the concept to which fillers of the relation have to belong,
- its super roles, i.e. the relations which are specialized by it.

The following list gives very provisional descriptions of the thematic roles in order to get a common understanding and from which project-internal definition may evolve. But possibly in the short run the accompanying examples, in which the relevant phrase is emphasized, will do a better job.

**agent** is the active force in a situation.

(32) *Er* schreibt den Termin ins Notizbuch.

**experiencer** is a person being an active, but non-volitional force in a situation.

(33) Das passt gut *bei mir*.
Wann wär's *ihnen* denn recht.

**theme** is the entity which is moving or undergoes a change of state.

(34) Er schreibt *den Termin* ins Notizbuch.
Er trifft *sie*.

**patient** is the entity which undergoes a change of state induced by the action of the agent.

(35) Er lädt *ihn* ein/aus
Er trifft *sie*.

**way_rel** subsumes the following three relations ascribed to motions. Thus this relation may be used as an underspecification in cases where just one of them is necessary. Furtheron these relations are applicable not only to motion in space.

[33] We would like to thank Bianka Buschbeck, Rita Nuebel, Markus Egg for their contributions. However any errors or probably still existing misconceptions fall into our responsibility.

**source** is the location from which a motion event starts

(36) Er kommt *aus Hamburg*.
Bei mir geht's *von neun Uhr an*.

**path** is an intermediary location during a motion event.

(37) ... und kommen dann *über den Fasanengarten* rüber zum Gebäude.

**goal** is the location where a motion event ends.

(38) die ganze Sache *auf Montag* verlegen
*in die Schweiz* fahren

**local** is the location at which a situation takes place.

(39) Ich sitze gerade *in meinem Büro* und sehe *in meinem Kalender*, daß...

**temporal_rel** places either a punctual or an extended situation on the time axis.

(40) *In dieser Woche* kann ich nicht.
*Morgens* bin ich im Büro.

**instrumental** is the object with which an event is primarily performed.

(41) Er schreibt den Termin *mit dem Bleistift*.
Er fährt *mit dem Zug*.

**benefactive** is a person benefitting from the event.

(42) Notieren Sie *mir/für mich* den Termin auf.

**concomitant** is an object which accompanies the performance of an event.

(43) wir treffen uns dann in der Eingangshalle *mit den Unterlagen*.

**final** is the aim of a situation.

(44) den Raum *für die Tagung besorgen*

**causal** is the reason why a state exists or an event occurs.

| *Feature* | *Range* | *Example* |
|---|---|---|
| in_jahr | Number | in_jahr:1994 |
| in_monat | januar..dezember | in_monat:september |
| in_woche | 1..52 | in_woche:14 |
| am_wochentag | montag..sonntag | am_wochentag:dienstag |
| am_datum | 1..31 | am_datum:6 |
| am_tageszeit | morgen..abend | am_tageszeit:nachmittag |
| um_uhrzeit | 0..2359 | um_uhrzeit:1655 |

Figure 14: Major features used to model calendar information.

(45) Ich rufe mal an ... *wegen der zweitägigen Arbeitssitzung*.

**modal** is the manner in which a situation exists .

(46) Er kam *in Eile* ins Büro.
Ich kann *notfalls* einen Raum besorgen.

**comitative** is the accompanying person in a situation.

(47) Ich möchte einen Termin *mit ihnen* ausmachen.

For a transfer, which uses these relations, these relations are not specific enough to get the appropriate wording in the target language. In the case of english as target language one has e.g. for INSTRUMENTAL to distinguish between INSTRUMENTAL_PROPER (with) and INSTRUMENTAL_TRANSPORT (by). These more fine-grained thematic relations, called *relational concepts* are refinements of thematic relations just in the same way as GOAL refines the WAY_REL-relation.

## 6.4 Calendar Information

The domain model contains concepts and roles which can be used to describe the temporal entities referred to in the dialogues. Since this involves mainly dates and names of days or months, we call this information *calendar information*. Calendar information can be used

1. as a basis for Transfer/Generation (according to the hypothesis that in many cases the speech-event type and the date described is the equivalent of translation;

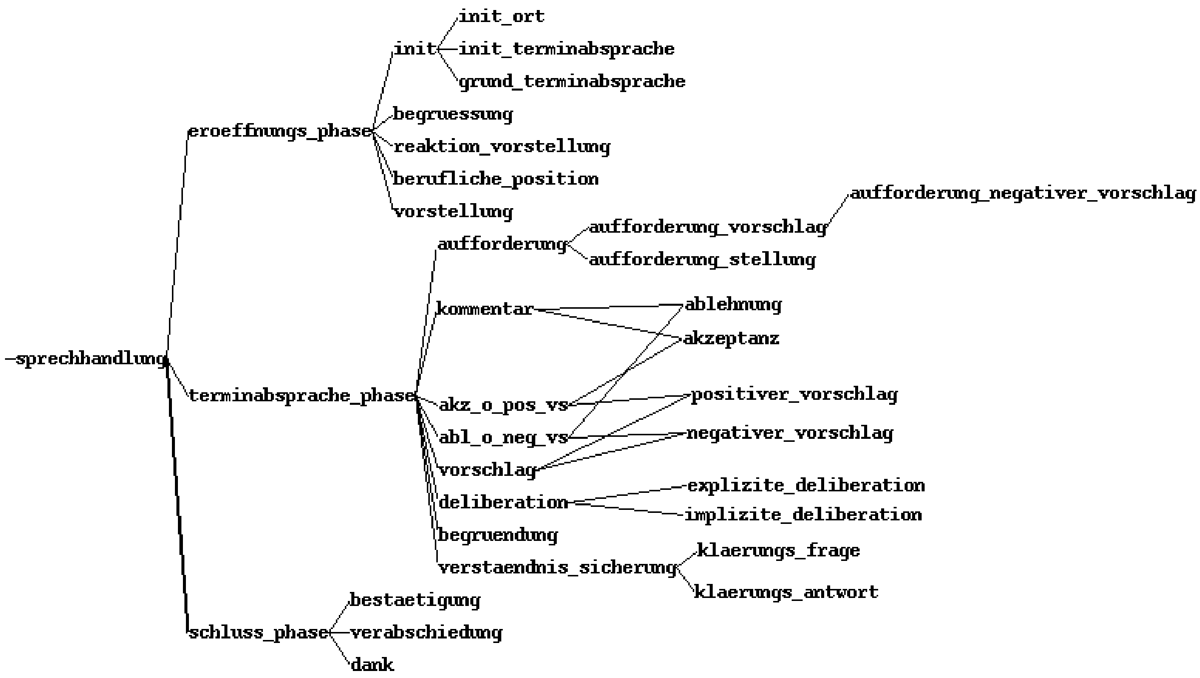


Figure 15: The speech-event hierarchy.

2. for macro-structural preference rules in the determination of a speech-event type (e.g. by checking whether a temporal description can be a refinement of a previously proposed date, or whether it has to be a description of a new alternative);

3. to detect errors resulting from speech recognition (e.g. ‘31. Juni’ is not a valid date).

The main features used for describing calendar entities together with range restrictions and exemplary values are shown in Figure 14.

## 6.5 The Speech-Event Hierarchy

Figure 15 shows the hierarchy of speech-event types as modeled in the BACK system. Each speech-event type corresponds to a BACK concept. In addition to that we introduced some “artificial” concepts that cannot be regarded as speech-event types—these concepts play a crucial role in the process of automatically assigning speech-event types to utterances. They are used to store already computed information or to structure the concepts in order to allow for the formulation of preference rules operating on a higher level. There is, for example, a speech-event type 'akz_o_pos_vs' representing that the type is either an acceptance or a positive proposal. The speech-event types are described in detail [Schmitz, Jekat-Rommel 94].

# 7 How To Use the Domain Model

There are several ways to use the domain model, which we will briefly illustrate in this section. In principle one can distinguish between

1. subsumption/disjointness checking
2. retrieval of sub and super concepts
3. retrieval of role information
4. default-based reasoning

Note that this section is mainly meant to illustrate possible uses of the domain model and does not describe how the domain model is actually used in the demonstrator. This depends on the design decisions taken in the other modules.

### Subsumption and Disjointness Checking

The most obvious use of a conceptual hierarchy is to check satisfaction of selectional restriction by testing subsumption. If we have, for example, a selectional restriction requiring the argument of a preposition to be a noun phrase referring to a location, we can check whether the type of the referent of the actual noun phrase is subsumed by LOKATION.

Subsumption checks can be implemented by using the BACK predicate

```
backask(subsumes(C1,C2))
```

which succeeds iff the concept 'C1' subsumes the concept 'C2'. Note that this use of the domain model implicitly assumes a closed world. In principle it would be more appropriate to check whether the referent could be an instance of LOKATION instead of asking whether it is already known to be an instance of LOKATION. Thus we would have to call[34]

[34] '\+' is the Quintus-Prolog predicate for negation as failure.

\+(backask(disjoint(C1,C2)))

which succeeds if 'C1' and 'C2' are not disjoint, i.e. can be unified. However, we can argue that there are ontologically essential categories for which subsumption checking is adequate, since it yields equivalent results as disjointness checking. Note that this equivalence holds if

subsumes(c1,c2) fails iff disjoint(c1,c2) succeeds

which in turn holds if the sub concepts of each concept are mutually disjoint.

Whereas this correspondence between subsumption and disjointness holds for the sort hierarchy, it is not generally true for the whole domain model. Given the open-world assumption underlying Description Logics, selectional restrictions should be in general checked by testing disjointness. The restriction is violated only if the concept required by the restriction and the concept actually occurring as an argument are disjoint.

### Retrieval of Sub and Super Concepts

Retrieval of sub and super concepts is another functionality provided by the domain model. This may be used, for example, to find more specific or more general concepts in case a concept is not lexicalized in English. The FLEX system provides a *filter mechanism* which allows limited second-order predications, i.e. the assignment of properties to concepts and roles. Thus we can mark all concepts in the hierarchy which are lexicalized in English by adding a filter, e.g. 'eng_lex'. The query

direct_supers(Conc,[eng_lex],Supers)

then retrieves the most specific concepts subsuming 'Conc' and satisfying the filter 'eng_lex'.

In the BACK system (direct) super and sub concepts are backtracked by

tboxget(direct_super_concept(Conc,Super))
tboxget(super_concept(Conc,Super))
tboxget(direct_sub_concept(Conc,Sub))
tboxget(sub_concept(Conc,Sub))

In addition to these predicates which provide information about concepts, the predicate 'msc' provides information about objects. Given an object, this predicate returns the *most specific concepts* it instantiates.

### Retrieval of Role Information

The programming interface of BACK allows the retrieval of information about a particular role at an object or a concept. The predicates

tboxget(number_restriction(Conc,Role,Min,Max))
tboxget(value_restriction(Conc,Role,VR))
tboxget(fillers(Conc,Role,Fillers))

can be used to retrieve the the *cardinality*, the *value restriction*, and the *fillers* of a role.

The FLEX system supports *term-valued* features, i.e. features taking concepts or roles instead of objects as fillers. This allows the modeling of rules for *conceptual disambiguation*, as

```
pred:anbieten1 and the(ref comp theme_rel,termin_c)
               => conc:termin_vorschlagen_c
pred:vergessen and the(ref comp theme_rel,zeit_c)
               => conc:nicht_passen_c
pred:gut and the(ref,passen_c)
               => conc:gut_passen_c
pred:unguenstig and the(ref,zeit_c)
               => conc:schlechter_termin_c
```

Given the above predicate we can then retrieve the filler for the feature 'conc' at a particular object.

### Default-Based Reasoning

Finally, the FLEX system contains a nonmonotonic extension based on *weigthed defaults* [Quantz, Suska 94]. In the demonstrator weighted defaults will be used to support the automatic assignment of speech-event types to utterances.

The implementation of weighted defaults in FLEX is based on *situated descriptions*. Thus instead of saying that an object o is an instance of a concept c, such a description can be restricted to a particular situation. The situated description 'o :: c in s' thus expresses that o is an instance of c in situation s. Situations can also be *extended*: if situation $s_2$ extends situation $s_1$ all descriptions valid in $s_1$ also hold in $s_1$.

Given the information modeled with weighted defaults, we can assign default extensions to each situation. The weights of the defaults are then used to score these default situations and the situation with the minimal negative score is the preferred default situation. Evaluating defaults thus yields a table which assigns default situations to strict situations. On the basis of this table default information can be retrieved by using the standard retrieval predicates.

# 8 Future Work

The experiences we made during our work on this first version of the domain model will have considerable impact on our future work. In this final section we briefly sketch some of the issues we want to address in future versions of the domain model.

As has been pointed out, the domain model has to take into account requirements arising from rather different tasks (e.g. ordering of syntactic parses, conceptual disambiguation, determination of speech-event types). One of the major problems thus was the integration of these heterogeneous requirements into a homogeneous domain model. We plan to use two different mechanisms to facilitate this integration:

1. the specification of requirements should include examples showing how the information is actually intended to be used;
2. the domain model should be more transparent wrt the use of the information it contains.

One way of realizing the second item is to explictly represent the users of each concept and role in the domain model, e.g. by means of the filter mechanism in FLEX.

In general, we plan to consider the use of additional tools to support the "distributive" development of the domain model. This involves, for example,

- the compilation of (parts of) the domain model from FLEX to STUF;
- the computation of differences between different versions of the domain model.

In addition to these organizational issues, we also envisage to reconsider certain parts of the domain model. For one thing, we will probably distinguish more carefully between thematic and conceptual relations. Thematic relations will then be used on the linguistic level, i.e. the role 'agent' will be a role between two linguistic signs, whereas conceptual relations will be used on the referential/conceptual level, e.g. the role ‘proposer’ will be a role between two discourse referents.

Finally, we will carefully evaluate, to which degree type shifting is relevant in the VERBMOBIL demonstrator. Though the current domain model provides the basis for both a variable depth of analysis and a straightforward treatment of systematic type shifts, we still have to see, how this is actually used in the demonstrator.

## Acknowledgments

Our work on domain modeling has been influenced by discussions with partners at various workshops and project meetings. In the above presentation we have tried to acknowledge all substantial input we have received in the course of the domain modeling.

We would like to thank all colleagues who have helped to produce this version of the domain model. In particular we would like to thank Bianka Buschbeck, Markus Egg, and Rita Nübel for providing specific requirements which served as a basis for large parts of the conceptual hierarchy.

# References

[Allen 87] J. Allen, *Natural Language Understanding*, Menlo Park: Benjamin/Cummings, 1987

[Almuallim et al. 94] H. Allmuallim, Y. Akiba, T. Yamazaki, A. Yokoo, S. Kaneda, “Two Methods for Learning ALT-J/E Translation Rules from Examples and a Semantic Hierarchy”, *COLING-94*, 57–63

[Austin 62] J.L. Austin, *How To Do Things with Words*, Oxford, 1962

[Bach, Harms 68] E. Bach, R.T. Harms (eds), *Universals in Linguistic Theory*, London: Holt, Rinehart and Winston, 1968

[Bade et al. 94] U. Bade, S. Heizmann, S. Jekat-Rommel, S. Kameyama, D. Krause, I. Maleck, B. Prahl, W. Preuß, *Wizard-of-Oz-Experimente mit dem VERBMOBIL-Simulator*, VM-Memo 24, Juni 1994

[Barwise 89] J. Barwise, *The Situation in Logic*, Stanford: CSLI Lecture Notes 17, 1989

[Barwise, Perry 83] J. Barwise, J. Perry, *Situations and Attitudes*, Cambridge (Mass.): MIT Press, 1983

[Bateman 92] J. Bateman, "The Theoretical Status of Ontologies in Natural Language Processing", in [Preuß, Schmitz 92], 50–99

[Benthem 81] J.v. Benthem, "Why Is Semantics What?", in [Groenendijk et al. 81], 29–49

[Bierwisch 83] M. Bierwisch, "Semantische und konzeptuelle Repräsentation lexikalischer Einheiten", *Studia grammatica* **XII**, 61–99

[Bosch 85] P. Bosch, "Lexical Meaning Contextualized", in [Hoppenbrouwers et al. 85], 251–258

[Bosch 93] P. Bosch, "Lexical Meaning and Conceptual Representation", Draft, March 1993

[Brachman, Schmolze 85] R.J. Brachman, J.G. Schmolze, "An Overview of the KL-ONE Knowledge Representation System", *Cognitive Science* **9**, 171–216, 1985

[Brachman et al. 91] R. Brachman, D.L. McGuiness, P.F. Patel-Schneider, L. Alperin Resnick, A. Borgida, "Living with CLASSIC: When and How to Use a KL-ONE-like Language", in [Sowa 91], 401–456

[Brown, Yule 83] G. Brown, G. Yule, *Discourse Analysis*, Cambridge: Cambridge University Press, 1983

[Coulmas 81] F. Coulmas, *Routine im Gespräch*, Wiesbaden: Athenaion, 1981

[Cresswell 78] M. J. Cresswell, "Semantic Competence", in [Cresswell 88], 12–33

[Cresswell 88] M. J. Cresswell, *Semantical Essays*, Dordrecht: Kluwer, 1988

[Cruse 86] D.A. Cruse, *Lexical Semantics*, Cambridge: Cambridge University Press, 1986

[Davidson 67] D. Davidson, "Truth and Meaning", in [Davidson 84], 17–36

[Davidson 84] D. Davidson, *Inquiries into Truth and Interpretation*, Oxford: Clarendon Press, 1984

[Davidson, Harman 72] D. Davidson, G. Harman (eds), *Semantics of Natural Language*, Dordrecht: Reidel, 1972

[Dölling 92] J. Dölling, "Flexible Interpretation durch Sortenverschiebung", in [Zimmermann, Strigin 92], 23–62

[Donini et al. 91] F.M. Donini, M. Lenzerini, D. Nardi, W. Nutt, "Tractable Concept Languages" *IJCAI-91*, 458–463

[Dowty 91] D. Dowty, "Thematic Proto-roles and Arguments Selection", *Language* **67**, 547–619, 1991

[Dowty et al. 81] D. R. Dowty, R. E. Wall, S. Peters, *Introduction to Montague Semantics*, Dordrecht: D. Reidel, 1981

[Fillmore 68] C.J. Fillmore, "The Case for Case", in [Bach, Harms 68], 1–88

[Fodor 78] J.A. Fodor, "Tom Swift and His Procedural Grandmother", in [Fodor 81], 204–224

[Fodor 81] J.A. Fodor, *Representations*, Brighton: The Harvester Press, 1981

[Gabbay, Guenthner 83] D. Gabbay, F. Guenthner (eds), *Handbook of Philosophical Logic*, Vol. I: Elements of Classical Logic, Dordrecht: Reidel, 1983

[Galton 88] A. Galton, "Formal Semantics: Is It Relevant to Artificial Intelligence?", *Artificial Intelligence Review* **2**, 151–165, 1988

[Garey, Johnson 79] M.R. Garey, D.S. Johnson, *Computers and Intractability*, New York: Freemann, 1979

[Goodman 68] N. Goodman, *Languages of Art*, Indianapolis: Hackett, 1976[2]

[Goodman 78] N. Goodman, *Ways of Worldmaking*, Indianapolis: Hackett, 1978

[Goodman and Nirenburg 91] Goodman, S. Nirenburg (eds), *The KBMT Project: A case study in Knowledge-Based Machine Translation* San Mateo: Morgan Kaufmann Publishers, 1991

[Grimshaw 92] J. Grimshaw, *Argument Structure*, Cambridge: Cambridge University Press, 1992

[Groenendijk et al. 81] J. Groenendijk, T.M.V. Janssen, M. Stokhof (eds), *Formal Methods in the Study of Language*, Amsterdam: Mathematical Centre Tracts 135, 1981

[Groenendijk et al. 87] J. Groenendijk, D. de Jong, M. Stokhof (eds), *Studies in Discourse Representation Theory and the Theory of Generalized Quantifiers*, Dordrecht: Foris, 1987

[Gruber 65] J.S. Gruber, *Studies in Lexical Relations*, PhD thesis, MIT, Cambridge, 1965

[Halliday 70] M.A.K. Halliday, Language Structure and Language Function. In J. Lyons, editor, *New Horizons in Linguistics*, Middlesex: Penguin, 1970

[Henschel, Bateman 94] R. Henschel, J. Bateman, "The Merged Upper Model: A Linguistic Ontology for German and English", *COLING-94*, 803–809

[Hoppe et al. 93] T. Hoppe, C. Kindermann, J.J. Quantz, A. Schmiedel, M. Fischer, *BACK V5 Tutorial & Manual*, KIT Report 100, Technische Universität Berlin, 1993

[Hoppenbrouwers et al. 85] G. Hoppenbrouwers, A. Weijters, P. Seuren (eds), *Meaning and the Lexicon*, Dordrecht: Foris, 1985

[Hovy, Nirenburg 92] E. Hovy, S. Nirenburg, "Approximating an Interlingua in a Principled Way", in *Proceedings of the DARPA Speech and Natural Language Workshop*, Hawthorne, NY, Feb. 1992

[Jackendoff 83] R. Jackendoff, *Semantics and Cognition*, Cambridge: MIT Press, 1983

[Jackendoff 90] R. Jackendoff, *Semantic Structures*, Cambridge: MIT Press, 1990

[Kamp, Reyle 93] H. Kamp, U. Reyle, *From Discourse to Logic*, Dordrecht: Kluwer, 1993

[Kanger, Öhman 80] S. Kanger, S. Öhman (eds), *Philosophy and Grammar*, Dordrecht: Reidel, 1980

[Kay et al. 91] M. Kay, J.M. Gawron, P. Norvig, *Verbmobil: A Translation System for Face-to-Face Dialog*, August 1991

[Keenan 75] E. Keenan (Ed.), *Formal Semantics of Natural Language*, Cambridge: Cambridge University Press, 1975

[Keil 79] F.C. Keil, *Semantic and Conceptual Development: An Ontological Perspective*, Cambridge (Mass.): Harvard University Press, 1979

[Keil 89] F.C. Keil, *Concepts, Kinds, and Cognitive Development*, Cambridge (Mass): MIT Press, 1989

[Klose 93] G. Klose,*Task-Oriented Modeling for Natural Language Processing Systems*, KIT-Report 109, Technical University Berlin, Computer Science Department, September 1993

[Klose et al 92] G. Klose, E. Lang, Th. Pirlein, *Ontologie und Axiomatik der Wissensbasis von LILOG*, Berlin: Springer, 1992

[Kuhn 62] T. Kuhn, *The Structure of Scientific Revolutions*, Chicago, 1962

[Leblanc 83] H. Leblanc, "Alternatives to Standard First-order Semantics", in [Gabbay, Guenthner 83], 189–274

[Lenat, Guha 90] D. Lenat, R.V. Guha, *Building Large Knowledge-based Systems: Representation and Inference in the Cyc Project*, Addison Wesley, 1990

[Levinson 83] S.C. Levinson, *Pragmatics*, Cambridge: Cambridge University Press, 1983

[Lewis 72] D. Lewis, "General Semantics", in [Davidson, Harman 72], 169–218

[Lyons 77] J. Lyons, *Semantics*, Cambridge: Cambridge University Press, 1977

[Lyons 91] J. Lyons, "Bedeutungstheorien", in [Stechow, Wunderlich 91], 1–24

[Mann 85] W.C. Mann, Janus abstraction structure – draft 1. Draft, USC – Information Science Institute, Marina del Rey, California, 1985

[Mellish 88] C. Mellish, "Implementing Systemic Classification by Unification", *Computational Linguistics* **14**, 40–51, 1988

[Minsky 68] M. Minsky (Ed.), *Semantic Information Processing*, Cambridge (Mass): MIT Press, 1968

[Minsky 75] M. Minsky, "A Framework for Representing Knowledge", in [Winston 75], 211–277

[Montague 74] R. Montague, *Formal Philosophy: Selected Papers of R. Montague*, edited by R. Thomason, New Haven: Yale University Press, 1974

[Morris 38] C. Morris, "Foundations of the Theory of Signs", in [Morris 71], 13–71

[Morris 71] C. Morris, *Writings on the General Theory of Signs*, The Hague: Mouton, 1971

[Nunberg 78] G. Nunberg, *The Pragmatics of Reference*, Bloomington: Indiana University Linguistics Club, 1978

[Partee 80] B.H. Partee, "Montague Grammar, Mental Representation, and Reality", in [Kanger, Öhman 80], 59–78

[Partee 87] B.H. Partee, "Noun Phrase Interpretation and Type-Shifting Principles", in [Groenendijk et al. 87], 115–143

[Partee 92] B.H. Partee, "Semantic Categories and Semantic Type", in [Rosner, Johnson 92], 97–126

[Penman 89] The Penman Natural Language Group. The Penman User Guide. Information Sciences Institute, Marina del Rey, California, December 1988

[Pinkal 85] M. Pinkal, *Logik und Lexikon*, Berlin: de Gruyter, 1985

[Pinkal 88] M. Pinkal, *Neuere Semantikmodelle für die Verarbeitung natürlicher Sprache*, Universität Hamburg: FBI-HH-M-164/88, Dezember 1988

[Pollard, Sag 87] C. Pollard, I.A. Sag, *An Information Based Syntax and Semantics*, Vol. I Fundamentals, Stanford: CSLI Lecture Notes 13, 1987

[Potts 75] T.C. Potts, "Model Theory and Linguistics", in [Keenan 75], 241–250

[Preuß, Schmitz 92] S. Preuß, B. Schmitz (eds) , *Workshop on Text Representation and Domain Modelling*, KIT-Report 97, Technische Universität Berlin, 1992

[Pustejovsky 89] J. Pustejovsky, "Current Issues in Computational Lexical Semantics", *EACL-89*, xvii–xxv

[Pustejovsky 93a] J. Pustejovsky, "Type Coercion and Lexical Selection", in [Pustejovsky 93b], 73–94

[Pustejovsky 93b] J. Pustejovsky (Ed.), *Semantics and the Lexicon*, Dordrecht: Kluwer, 1993

[Putnam 80] H. Putnam, "Models and Reality", in [Putnam 83], 1–25

[Putnam 83] H. Putnam, *Realism and Reason*, Philosophical Papers, Vol. 3, Cambridge: Cambridge University Press, 1983

[Putnam 87] H. Putnam, *The Many Faces of Realism*, La Salle: Open Court, 1987

[Quantz 92] J.J. Quantz, "How to Fit Generalized Quantifiers into Terminological

Logics", *ECAI-92*, 543–547

[Quantz 95] J.J. Quantz, *Preferential Disambiguation in Natural Language Processing*, KIT Report in preparation

[Quantz, Schmitz 94] J.J. Quantz, B. Schmitz, "Knowledge-Based Disambiguation for Machine Translation", *Minds and Machines* **4**, 39–57, 1994

[Quantz, Suska 94] J.J. Quantz, S. Suska, "Weighted Defaults in Description Logics—Formal Properties and Proof Theory", in B. Nebel, L. Dreschler-Fischer (eds), *KI-94: Advances in Artificial Intelligence*, Berlin: Springer, 1994, 178–189

[Quantz et al. 94] J.J. Quantz, G. Dunker, V. Royer, "Flexible Inference Strategies for DL Systems", in F. Baader, M. Lenzerini, W. Nutt, P. F. Patel-Schneider (eds), *International Workshop on Description Logics*, DFKI Report D-94-10, DFKI Saarbrücken, 27–30, 1994

[Quillian 68] M.R, Quillian, "Semantic Memory", in [Minsky 68], 216–270

[Quine 51] W. v.O. Quine, "Two Dogmas of Empiricism", in [Quine 53]

[Quine 53] W. v.O. Quine, *From a Logical Point of View*, Cambridge (Mass.): Harvard University Press, 1980[2]

[Rauh 88] G. Rauh, *Tiefenkasus, thematische Relationen und Theta-Rollen*, Tübingen: Gunter Narr Verlag, 1988

[Rosner, Johnson 92] M. Rosner, R. Johnson (eds), *Computational Linguistics and Formal Semantics*, Cambridge: Cambridge University Press, 1992

[Schmitz 94] B. Schmitz, "Die Verbindung von mikro- und makrostruktureller Information für die Übersetzung von Verben im Verbmobilszenario", VM-Memo 30, Technische Universität Berlin, Juni 1994

[Schmitz, Jekat-Rommel 94] B. Schmitz, S. Jekat-Rommel, "Eine zyklische Approximation an Sprechhandlungstypen– zur Annotierung von Äußerungen in Dialogen", Verbmobil Report 28, September 1994

[Searle 80] J.R. Searle, "Minds, Brains, and Programs", *The Behavioral and Brain Sciences* **3**, 417–457 1980

[Searle 83] J.R. Searle, *Intentionality*, Cambridge: Cambridge University Press, 1983

[Skirbekk 77] G. Skirbekk (Ed.), *Wahrheitstheorien*, Frankfurt: Suhrkamp, 1977

[Sommers 65] F. Sommers, "Predicability", in M. Black (Ed), TitPhilosophy in America, London: George Allen & Unwin Ltd, 1965, 262–281

[Sommers 71] F. Sommers, "Structural Ontology", *Philosophia*, 21–42, 1971

[Sowa 91] J. Sowa (Ed.), *Principles of Semantic Networks: Explorations in the Representation of Knowledge*, San Mateo: Morgan Kaufmann, 1991

[Stechow, Wunderlich 91] A. von Stechow, D. Wunderlich (eds), *Semantik—Semantics*, Berlin: de Gruyter, 1991

[Sundholm 83] G. Sundholm, "Systems of Deduction", in [Gabbay, Guenthner 83], 133–188

[Tarski 44] A. Tarski, "Die semantische Konzeption der Wahrheit und die Grundlagen der Semantik", in [Skirbekk 77], 140–188

[Thomason 74] R. Thomason, "Introduction", in [Montague 74]

[Wahlster 94] W. Wahlster, "Nach der ADL-Evaluation: Bemerkungen zur Situation der KI in Deutschland", *KI*, 54–56, Juli 1994

[Winston 75] P.H. Winston (Ed.), *The Psychology of Computer Vision*, New York: McGraw-Hill, 1975

[Wittgenstein 56] L. Wittgenstein, *Philosophische Untersuchungen*, Frankfurt: Suhrkamp,

[Zimmermann, Strigin 92] I. Zimmermann, A. Strigin (eds), *Fügungspotenzen*, Berlin: Akademie Verlag, 1992