Incrementality and Locality of Language Comprehension: The Pivotal Role of Semantic Interpretation Schemata

Udo Hahn **Martin Romacker**
Text Understanding Lab
Computational Linguistics Division, Freiburg University
D-79085 Freiburg, Germany
`http://www.coling.uni-freiburg.de`

Abstract

We introduce a computational model of language comprehension that combines locality of syntactic and semantic analysis with incrementality of processing. As the model incorporates inheritance-based abstraction mechanisms we are able to specify a parsimonious inventory of abstract, simple and domain-independent semantic interpretation schemata.

Introduction

Despite a large body of experimental evidence for the incrementality of human language comprehension (e.g., Tyler & Marslen-Wilson (1977), Thibadeau et al. (1982), Garrod & Sanford (1985)), this issue has not been given comparable attention in computational approaches to natural language processing. While in the past only few studies were concerned with the incremental aspects of semantic interpretation (for an overview, cf. Haddock (1989)), interest in this area has recently increased, especially in the cognitive science community (Paredes-Frigolett & Strube, 1996; Hahn & Strube, 1996; Lombardo et al., 1998).

Our contribution to this discussion lies in a model of incremental semantic interpretation whose specifications are very compact. This allows us to treat a variety of linguistic phenomena by only few and general interpretation schemata. In essence, these schemata address structural *configurations* within dependency graphs rather than hook on particular language phenomena or single rules. The grammar model, as well as the domain model make extensive use of inheritance-based abstraction mechanisms. By interfacing the description of semantic interpretation schemata to these inheritance hierarchies, we are able to specify a parsimonious and domain-independent semantic interpretation system.

We start from a lexicalized grammar model based on dependency relations. *Locality* in this framework has a dual reading. On the one hand, it refers to the reachability of syntactically related content words within "minimal" dependency graphs. On the other hand, these minimal dependency graphs have a direct conceptual interpretation. *Incrementality* comes in as local readings are combined on the fly to form larger units of analysis as comprehension unfolds.

Framework for Incremental Interpretation

Knowledge Sources. We supply grammar and domain knowledge, as well as interpretation schemata mediating between these two knowledge sources. *Grammatical knowledge* for syntactic analysis is based on a fully lexicalized dependency grammar (Hahn et al., 1994). Lexeme specifications form the leaf nodes of a lexicon tree, which are further abstracted in terms of word class specifications at different levels of generality. This leads to a word class hierarchy, which consists of word class names $\mathcal{W} := \{\textsc{VerbTrans}, \textsc{Verb}, \textsc{Article}, \textsc{Det}, \ldots\}$ and a subsumption relation $isa_{\mathcal{W}} = \{(\textsc{VerbTrans}, \textsc{Verb}), (\textsc{Article}, \textsc{Det}), \ldots\} \subset \mathcal{W} \times \mathcal{W}$. Inheritance of grammatical knowledge is based on the idea that constraints are attached to the most general word class to which they apply.

A dependency grammar captures binary valency constraints between a syntactic head (e.g., a noun) and one of its possible modifiers (e.g., a determiner or an adjective). In order to establish a dependency relation $\delta \in \mathcal{D} := \{specifier, subject, direct\text{-}object, \ldots\}$ between a head and a modifier, the corresponding constraints on word order, compatibility of morphosyntactic features, as well as semantic criteria have to be fulfilled. Fig. 1 depicts a sample dependency graph in which word nodes are given in bold face and dependency relations between them are indicated by labeled edges.

Conceptual knowledge about the domain is expressed in a KL-ONE-like representation language (Woods & Schmolze, 1992). It consists of concept names $\mathcal{F} := \{\textsc{Sell}, \textsc{Hard-Disk}, \ldots\}$ and a subsumption relation on concepts $isa_{\mathcal{F}} = \{(\textsc{Sell}, \textsc{Action}), (\textsc{Hard-Disk}, \textsc{Product}), \ldots\} \subset \mathcal{F} \times \mathcal{F}$. The relation names $\mathcal{R} := \{\textsc{sell-agent}, \textsc{has-hard-disk}, \ldots\}$ denote conceptual relations also organized in a subsumption hierarchy $isa_{\mathcal{R}} = \{(\textsc{has-hard-disk}, \textsc{has-physical-part}), (\textsc{has-physical-part}, \textsc{has-part}), \ldots\}$.

Linkages between concepts via conceptual relations are determined by dependency relations linking lexical items in the dependency graph directly or in a mediated way (via a series of dependency relations). *Semantic interpretation rules* mediate between both description levels in a way as abstract and general as possible.

Local Computations: Basic Parsing Protocol. The lexicalized grammar and the associated parser we use are embedded in an object-oriented computation model. Dependency relations are computed by lexical objects, so-called *word actors*, through strictly local message passing, only involving the lexical items they represent. The basic protocol for incremental parsing can be sketched as follows (Hahn et al., 1994):

- After a word has been read from textual input, its associated lexeme (specified in the lexicon tree) is identified and a corresponding word actor gets initialized. As most lexemes (verbs, nouns and adjectives) are directly linked to

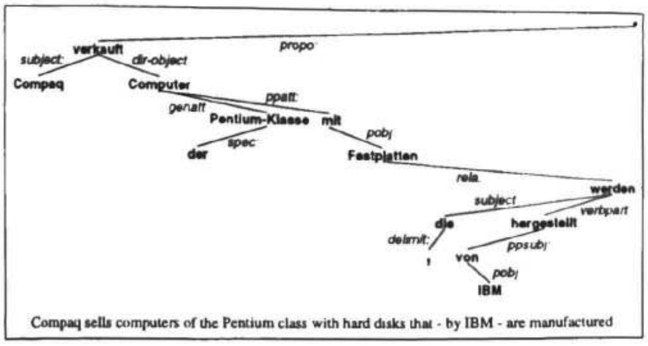

Figure 1: A Sample Dependency Graph

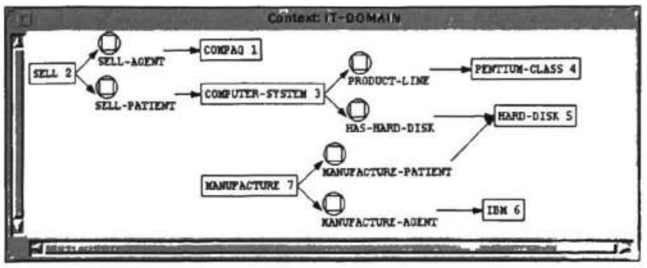

Figure 2: Corresponding Concept Graph

the conceptual system, each lexical item w that has a conceptual correlate C in the domain knowledge base, $w.C \in \mathcal{F}$, gets instantiated in the knowledge base, such that for any instance I_w, initially, $type(\mathrm{I}_w) =: w.C$ (e.g., w = *Festplatte*, $w.C$ = HARD-DISK, and I_w = HARD-DISK.5).

- For integration in the parse tree, the newly created word actor searches its head (alternatively, its modifier) by sending a request for dependential government to its left context. The search space is restricted, since this request is only propagated upwards along the "right shoulder" of the dependency graph constructed so far. All word actors which are addressed this way check, in parallel, whether their valency constraints are met by the requesting word actor.
- If all grammatical constraints are fulfilled by one of the targeted word actors, a tentative semantic interpretation is performed incorporating constraints from the currently checked dependency relation. If a valid result is computed, i.e., only if both grammatical and semantic-conceptual integrity are guaranteed, the acknowledged syntactic head h sends an acceptance message to the dependent modifier m and the screened dependency relation is finally established.

Incremental Semantic Interpretation

In the dependency parse tree from Fig. 1, we distinguish lexical nodes that have a conceptual correlate (e.g., *"Compaq"*, *"verkauft" (sells)*) from others that do not have such a correlate (e.g., *"mit" (with)*, *"von" (by)*). This is reflected in the two basic configurational settings for semantic interpretation:

- **Direct Linkage:** If two lexical nodes with conceptual correlates are linked by a *single* dependency relation, a *direct* linkage is given, which can immediately be interpreted in terms of a corresponding conceptual relation. This is illustrated in Fig. 1 by the direct linkage between *"Compaq"* and *"verkauft" (sells)* via the *subject* relation, which gets mapped to the SELL-AGENT conceptual role linking instances of corresponding conceptual correlates, *viz.* COMPAQ.1 and SELL.2, respectively (cf. Fig. 2). This interpretation uses knowledge about the conceptual correlates and the linking dependency relation, only.
- **Mediated Linkage:** If two lexical nodes with conceptual correlates are linked by a *series* of dependency relations and none of the intervening nodes have a conceptual correlate, a *mediated* linkage is given. Such a "minimal" subgraph can only be interpreted indirectly in terms of a conceptual relation. We include lexical information from intervening nodes in addition to the knowledge about the conceptual correlates and dependency relations. In Fig. 1 this is illustrated by the syntactic linkage between *"Computer"* and *"Festplatten" (hard disks)* via the intervening node *"mit" (with)* and the *ppatt/pobj* relations. This leads to a conceptual linkage between COMPUTER-SYSTEM.3 and HARD-DISK.5 via the relation HAS-HARD-DISK in Fig. 2.

In order to increase the generality and to preserve the simplicity of semantic interpretation we introduce a generalization of the notion of dependency relation such that it incorporates direct as well as mediated linkage: Two content words (nouns, adjectives, adverbs or full verbs) stand in a *mediated syntactic relation*, if one can pass from one word to the other along the connecting edges of the dependency graph without traversing word nodes other than prepositions, modal or auxiliary verbs (i.e., elements of closed word classes). In Fig. 1, e.g., the tuples *("Compaq", "verkauft")*, *("verkauft", "Computer")*, *("Computer", "Festplatten")*, *("Festplatten", "hergestellt")* stand in a mediated syntactic relation, whereas, e.g., the tuple *("verkauft", "Festplatten")* does not, since *"Computer"* is an intervening content word node.

We then call a series of contiguous words in a sentence S that stand in a mediated syntactic relation a *semantically interpretable subgraph* of the dependency graph of S. Semantic interpretation will be started whenever two word nodes with associated conceptual correlates are dependentially connected so that they form a semantically interpretable subgraph. As we will see, in some cases the dependency structures we encounter will have no constraining effect on the kind of conceptual relations we check (e.g., genitives). There are other cases (e.g., prepositional phrases), however, where constraints on possible interpretations can be derived from dependency structures and the lexical material they embody.

Basic semantic interpretation schema. Semantic interpretation is executed via a search in the domain knowledge base by combining two sorts of knowledge — first, concept pairs for which connecting relation paths have to be determined; second, constraints on the kinds of permitted or excluded conceptual relations when connecting relations are being computed. Concept pairs represent the content words linked at the dependency level within the semantically interpretable subgraph, while constraints on relations account for the dependency relation(s) between them. Schema (1) describes a mapping from the conceptual correlates, $h.C_{from}$ and $m.C_{to}$, in $\mathcal{F}$ of two syntactically linked lexical items, h and m, respectively, to connected relation paths R_{con}. A relation path $rel_{con} \in R_{con}$ composed of n relations, $(r_1, ..., r_n)$, is called *connected*, if for all its n constituent relations the concept type of the domain of relation r_{i+1} subsumes the concept type of the range of relation r_i.

$$si : \begin{cases} \mathcal{F} \;\times\; 2^{\mathcal{R}} \times 2^{\mathcal{R}} \;\times\; \mathcal{F} & \rightarrow \;\; 2^{R_{con}} \\ (C_{from}\,,\; R_+\,,\, R_-\;,\; C_{to}) & \mapsto \;\; \widetilde{R_{con}} \end{cases} \quad (1)$$

As an additional filter, si is constrained by all conceptual relations $R_+ \subset \mathcal{R}$ a priori permitted for semantic interpretation, as well as all relations $R_- \subset \mathcal{R}$ a priori excluded from semantic interpretation (concrete examples will be discussed below). Thus, $rel \in \widetilde{R_{con}}$ holds, if rel is a connected relation path from C_{from} to C_{to}, obeying the restrictions imposed by R_+ and R_-. For ease of specification, R_+ and R_- consist of the most general conceptual relations possible. Prior to semantic processing we expand them into their transitive closures, incorporating all their subrelations in the relation hierarchy. Hence, $R_+^* := \{\, r^* \in \mathcal{R} \mid \exists\, r \in R_+ : r^* \; isa_{\mathcal{R}} \; r \,\}$ (correspondingly, R_-^* is dealt with).

We also define the function $get\text{-}roles(C) =: CR$, which extracts the set of all conceptual relations CR associated with a concept C. Applying $get\text{-}roles$ to C_{from} extracts the roles that are used as starting points for the path search according to the defined restrictions. R_+ restricts the search to relations contained in $CR \cap R_+^*$, *iff* R_+ is not empty (otherwise, CR is taken as it is), R_- allows only for relations in $CR \setminus R_-^*$.

If the function si returns the empty set (i.e., no valid interpretation can be computed), no dependency relation will be established. Otherwise, for all resulting relations $\text{REL}_i \in \widetilde{R_{con}}$ an assertional axiom is added to the knowledge base by asserting the proposition $(h.C_{from}\ \text{REL}_i\ m.C_{to})$, where REL_i denotes the i^{th} reading.

Integration of Knowledge Levels

In the course of the interpretation process, we apply a number of specializations of the general semantic interpretation schema (1). One major source from which specializations arise are positive lists, D_+^{lexval}, and negative lists, D_-^{lexval}, of syntactic dependency relations, from which corresponding conceptual constraints, R_+ and R_-, can be directly derived. Knowledge about D_+^{lexval} and D_-^{lexval} is encoded at the level of *word classes* $\mathcal{W}$, such that $lexval \in \mathcal{W} \times \mathcal{D}$, and, thereby, inherited by all their subsumed lexemes. For instance, the word class of transitive verbs, VERBTRANS $\in \mathcal{W}$, defines for a subject valency $D_+^{\langle verbtrans,\ subject\rangle} := \{subject\}$ and $D_-^{\langle verbtrans,\ subject\rangle} := \emptyset$. We distinguish three basic cases:

1. Knowledge available from syntax *determines* the semantic interpretation, if $D_+^{lexval} \neq \emptyset$ and $D_-^{lexval} = \emptyset$ (e.g., the subject of a verb).
2. Knowledge available from syntax *restricts* the semantic interpretation, if $D_+^{lexval} = \emptyset$ and $D_-^{lexval} \neq \emptyset$ (e.g., most prepositions).
3. If $D_+^{lexval} = \emptyset$ and $D_-^{lexval} = \emptyset$, no syntactic constraints apply and semantic interpretation proceeds *entirely concept-driven*, i.e., it relies on the domain knowledge only (e.g., for genitive attributes).[1]

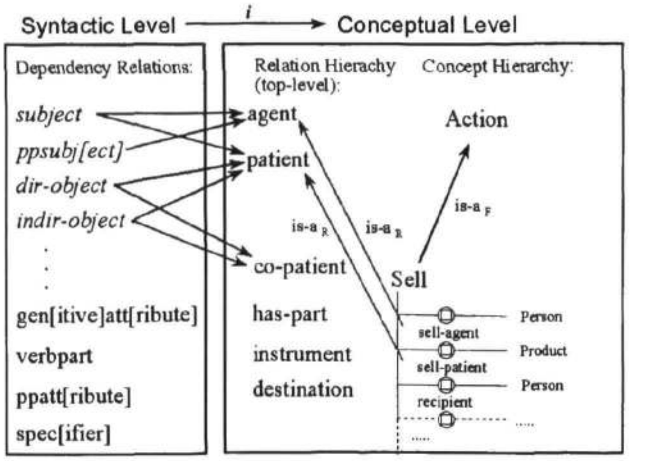

Figure 3: Relations between Knowledge Levels

For syntactic constraints to be propagated to the conceptual level, we define a mapping, i: $\mathcal{D} \rightarrow 2^{\mathcal{R}}$, from dependency relations to sets of conceptual relations. R_+ and R_- must be computed from D_+^{lexval} and D_-^{lexval}, respectively, by applying the interpretation function i to each element of D_+^{lexval} and D_-^{lexval}. So, $R_+ := \{y \mid x \in D_+^{lexval} \wedge y \in i(x)\}$ and $R_- := \{y \mid x \in D_-^{lexval} \wedge y \in i(x)\}$. An illustration is given in Fig. 3. On the left side, at the syntactic level proper, a subset of the dependency relations contained in $\mathcal{D}$ are depicted. Those that have associated conceptual relations are shown in italics. *dir[ect]-object*, e.g., must conceptually be interpreted in terms of PATIENT or CO-PATIENT; *gen[itive]att[ribute]*, however, has no direct conceptual correlation as this dependency relation does not restrict conceptual interpretation at all.

At the conceptual level, two orthogonal taxonomic hierarchies exist, one for relations, the other for concepts (cf. Fig. 3, right side). Both are organized in terms of subsumption hierarchies ($isa_{\mathcal{F}}$ and $isa_{\mathcal{R}}$). Also, both hierarchies interact, since relations are used to define concepts. The concept SELL is a subconcept of ACTION. It has a role SELL-PATIENT whose filler's type must be a PRODUCT. SELL-PATIENT itself is subsumed by the more general relation PATIENT.

Sample Analyses

In the examples we discuss now, we start from the interpretation of direct linkage, and then turn to the interpretation of mediated linkages considering increasingly complex configurations in the dependency graph as given by prepositional phrases and passives in relative clauses.

Interpreting direct linkage. When the first word in our sample sentence, *"Compaq"*, is read, its conceptual correlate COMPAQ.1 is instantiated immediately. The next word, *"verkauft" (sells)*, also leads to the creation of an associated instance (SELL.2). The word actor for *"verkauft"* then attempts to bind *"Compaq"* as its syntactic subject. Since we encounter a direct linkage, the semantic interpretation

---

[1] We have currently no empirical evidence for the fourth possible case, where $D_+^{lexval} \neq \emptyset$ and $D_-^{lexval} \neq \emptyset$.

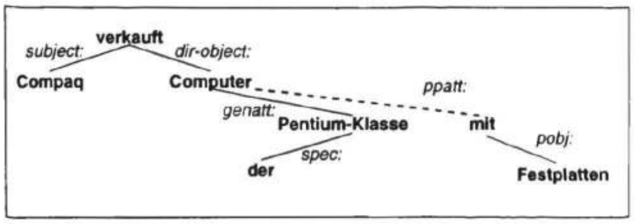

Figure 4: Dependency Graph during PP-Attachment

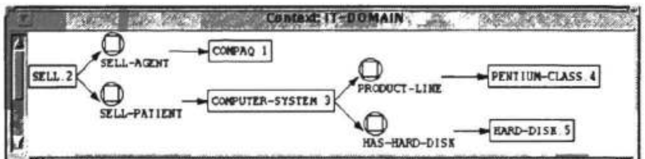

Figure 5: Concept Graph after PP-Attachment

schema si_{dir} is instantiated using the contextual information (the *subject* relation of a transitive verb is to be checked) as actual parameters. We incorporate two types of information — first, grammatical constraints from the word class of the verb *"verkauft"*, *viz.* $D_{+}^{\langle verbtrans,\ subject\rangle} := \{subject\}$ and $D_{-}^{\langle verbtrans,\ subject\rangle} := \emptyset$, which are mapped to {AGENT, PATIENT} by the function i (cf. Fig. 3); second, knowledge about the concept types of COMPAQ.1 and SELL.2 (COMPANY and SELL, respectively). Hence, si_{dir}(SELL, {AGENT, PATIENT}, $\emptyset$, COMPANY). Extracting the role set from SELL (cf. Fig. 3), only SELL-AGENT and SELL-PATIENT are allowed for interpretation as they are included in the transitive closure of AGENT and PATIENT. Checking sortal integrity succeeds only for SELL-AGENT (COMPANY is subsumed by LEGAL-PERSON and by PERSON). In an analogous way, the semantically interpretable subgraph < *"verkauft"* – *dir-object* – *"Computer"*> is dealt with.

When syntactic information does not constrain the semantic interpretation we have to proceed in an entirely concept-driven way. This holds for the third complete subgraph, < *"Computer"* – *genatt* – *"Pentium-Klasse"*>. The actual parameters provided lead to si_{dir}(COMPUTER-SYSTEM, $\emptyset$, $\emptyset$, PENTIUM-CLASS). We extract all roles contained in the concept definition of COMPUTER-SYSTEM and iteratively check for each role whether PENTIUM-CLASS is a legal filler. This is only the case for the relation PRODUCT-LINE. Though various linguistic phenomena (subjects, direct objects and genitives) are covered, a single schema, si_{dir}, is sufficient for semantic interpretation of direct linkage configurations.

Interpreting mediated linkage. For the interpretation of mediated linkage, information supplied by the intervening lexical nodes is incorporated. It is contained in *lexeme*-specific lists, $R^{lex} \subset \mathcal{R}$, since specifications at the word-class level ($D_{+/-}^{lexval}$) turn out to be too general here. This applies to specific lexical exemplars from closed word classes encountered in mediated linkages (e.g., prepositions). So, the number of additional specifications required remain fairly small.

Consider Fig. 4 where a semantically interpretable subgraph consisting of three word nodes (*"Computer"*, *"mit"*, *"Festplatten"*) is currently being processed (indicated by the dashed line). In particular, the word actor for *"mit" (with)* tries to determine its syntactic head.[2] We consider prepositions as relators carrying conceptual constraints for the corresponding instances of their syntactic head and modifier. The "meaning" of a preposition is encoded in a set $R^{Prep} \subset \mathcal{R}$, for each preposition in $Prep$, holding all permitted or excluded relations in terms of high-level conceptual relations. For the preposition *"mit" (with)*, e.g., we have $R_{+}^{mit} :=$ {HAS-PART, INSTRUMENT, HAS-PROPERTY, HAS-DEGREE, ...}. If a syntactic dependency relation between a head and a prepositional modifier has to be checked, the corresponding list R^{Prep} has to be matched against the restrictions imposed by the preposition's syntactic head via $D_{+/-}^{lexval}$.

When *"mit"* attempts to be governed by *"Computer"* the mediated linkage results in the instantiation of the specialized interpretation function si_{prep} which is applied exclusively for all occurrences of PP-attachments. The conceptual entities to be related are denoted by the leftmost and the rightmost node in the actual subgraph (i.e., *"Computer"* and *"Festplatten"*). Since the dependency relation between the head *"Computer"* and its modifier *"mit"*, *ppatt*, does not impose any restrictions at all (cf. Fig. 3), i.e., $D_{+/-}^{lexval} = \emptyset$, semantic interpretation boils down to the evaluation of si_{prep}(COMPUTER-SYSTEM, {HAS-PART, INSTRUMENT, HAS-PROPERTY, HAS-DEGREE, ...}, $\emptyset$, HARD-DISK). By extracting all conceptual roles and checking for sortal consistency, only HAS-HARD-DISK $isa_{\mathcal{R}}$ HAS-PART yields a valid interpretation to directly relate COMPUTER-SYSTEM and HARD-DISK. The state of semantic interpretation after PP-attachment is given in Fig. 5. This also indicates that during each step of the analysis a corresponding conceptual interpretation is already available.

To convey an idea of the generality and flexibility of our approach, we discuss a more complex example. In Fig. 6 the relative clause *"die von IBM hergestellt werden" (that are manufactured by IBM)* has already been analyzed and IBM.6 figures as MANUFACTURE-AGENT of MANUFACTURE.7 (cf. Fig. 2). Since the *subject* valency of the passive auxiliary *"werden"* is occupied by a relative pronoun (*"die"*), the interpretation of that structure must be postponed until the pronoun's referent becomes available. Passive interpretation is performed by another specialization of the general interpretation schema, si_{pass}. As with certain prepositions, constraints come directly from a positive list $R_{+}^{passaux} :=$ {PATIENT, CO-PATIENT}, which resides in the lexeme specification for *"werden"*. The items to be related are contained in the semantically interpretable subgraph spanned by *"die" (that)* and *"hergestellt" (manufactured)*.[3]

The referent of the relative pronoun *"die"* becomes avail-

---

[2] We do not attach prepositions to possible heads immediately. Rather we have a built-in delay mechanism so that the nominal modifier of a prepositional head has to be determined first. With this content item attached (in our example, HARD-DISK.5), compatibility checks with a potential head, as searched by the prepositional modifier, have a reasonable conceptual grounding.

[3] Pronominal reference is resolved when the head of the relative clause (*"werden"*) has determined its own head. As with prepositions, this delay mechanism is justified by the fact that semantic interpretation is only reasonable when a basic conceptual grounding has already been established (i.e., *"werden"* must have been bound to the content word *"hergestellt" (manufactured)*).

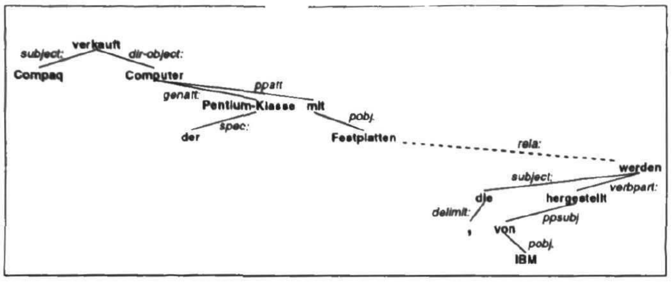

Figure 6: Dependency Graph for Relative Clause

able when the syntactic head of the relative clause (*"werden"*) has determined its head (cf. the dashed line in Fig. 6). Two alternatives arise, namely *"Computer"* and *"Festplatten" (hard disks)*. The choice of the latter leads to those parameters for si_{pass} needed for the *subject* interpretation.[4] si_{pass} inverses the argument structure by taking the leftmost word node (i.e., the pronoun's referent *"Festplatten"*) as C_{to} and the rightmost (*"hergestellt"*) as C_{from}. Hence, si_{pass}(MANUFACTURE, {PATIENT, CO-PATIENT}, $\emptyset$, HARD-DISK). The final interpretation is depicted in Fig. 2 linking MANUFACTURE.7 and HARD-DISK.5 via the role MANUFACTURE-PATIENT. Obviously, integrating intra- and extrasentential anaphorical phenomena (Strube & Hahn, 1995) necessitates a slight extension of our notion of a semantically interpretable subgraph. In case pronouns are involved, such a subgraph is interpretable *iff* all referents are made available.

Evaluation of Semantic Interpretation

In a small-scale evaluation study, we started from a domain ontology that is divided into an upper generic part (composed of 1,100 concepts and relations) and various domain-specific parts. In the study we report on two specialized domains were dealt with — an information technology (IT) model and an ontology covering parts of anatomical medicine (MED) (each domain model, in addition, contributes 1,100 concepts and relations). Corresponding lexeme entries in the lexicon provide linkages to the ontology. We also assume a correct parse to be delivered for the semantic interpretation process.

We then took a random selection of 54 texts (comprising 18,500 words) from our two text corpora. For evaluation purposes, we concentrated on the interpretation of genitives (direct linkage), prepositional phrase attachments and auxiliary as well as modal verbs (both variants of mediated linkage). In the following, we will focus on the discussion of the results from the semantic interpretation of genitives (cf. Table 1).

The chosen texts contained a total of almost 250 genitives, from which about 59%/33% (MED/IT) received an automatic interpretation, with correct ones for 57%/31% (recall). An interpretation was considered *correct* when the conceptual correlates of lexical items in a semantically interpretable subgraph were conceptually related in an adequate way.

	MED	IT
# texts	29	25
# words	4,300	14,200
# genitives ...	100	147
... with interpretation	59	49
...... correct	57	46
...... incorrect/none	2	3
... without interpretation	41	98
recall	57%	31%
precision	97%	94%

Table 1: Empirical Results for the Interpretation of Genitives

Slightly more than half of the loss we encountered can be attributed to insufficient coverage in the two *domain models*. The remaining cases can be explained by insufficient coverage of the *generic model* and reference to *other domains*, e.g., fashion or food. Some minor loss is also due to phrases referring to *time* (e.g., "the beginning of the year"), *space* (e.g., "the surface of the storage medium"), and *abstract* notions (e.g., "the acceptance of IT technology"), as well as *evaluative* expressions (e.g., "the advantages of plasma display") and *figurative* language (e.g., "the heart of the notebook").

The concrete values we found, disappointing as they may be for recall (57%/31%), encouraging, however, for precision (97%/94%), can only be interpreted relative to other data still lacking on a broader scale. Judged from the poor figures of our recall data, there is no doubt whatsoever that conceptual coverage of the domain constitutes *the* bottleneck for any knowledge-based approach to NLP. Sublanguage differences are also mirrored in these data, since medical texts adhere more closely to well-established concept taxonomies than magazine articles in the IT domain.

Related Work

Perhaps the most influential paper that treats incremental semantic interpretation from a modern rule-based perspective is the one by Pereira & Pollack (1991). We share their ideas to preserve as much as possible the principle of *compositionality* (evidenced by the smooth integration of interpreted subgraphs into the already constructed partial dependency graph), and to integrate the *discourse context* into the interpretation process as early as possible (evidenced by the resolution of the relative pronoun, though we have not gone into the details of anaphora resolution; cf. Strube & Hahn (1995)). We differ, however, fundamentally with respect to grammar specification. Pereira and Pollack use a PATR-style unification grammar, which comes without any mechanisms for inheritance to support rule specifications (semantic theories in that area do also not go beyond sortal taxonomies for semantic labels, cf. Creary & Pollard (1985)). This approach suffers from "flat" representations that require to enumerate semantic rules for specific grammatical phenomena and so they tend to proliferate.

Milward (1995) already points out a major advantage of lexicalized over rule-based approaches. In processing a sentence incrementally using a lexicalized grammar, we do not have to look at the grammar as a whole (as with rule-based grammars), but only at the grammatical information indexed for each of the words from the input stream, thus allowing for *efficient processing*. When one takes inheritance mechanisms

---

[4] The alternative choice (*"Computer"*) creates a local ambiguity, because COMPUTER-SYSTEM.3 can also be related to MANUFACTURE.7. However, heuristics are applied to select the most plausible reading, which are sensitive to preferential criteria such as the length of role compositions, the types of relations encountered, etc. (Hahn & Markert, 1997).

into account *efficient encoding* becomes another asset to this approach, both for the lexicon and semantic interpretation.

Sondheimer et al. (1984) and Hirst (1988) also propose models of incremental semantic interpretation. Their use of KL-ONE-style representation languages allows them to exploit property inheritance (or typing) built-ins. The main difference to our approach lies in the status of the semantic rules. Sondheimer et al. attach interpretation rules to each *role (filler)* and, hence, have to provide detailed specifications reflecting the idiosyncrasies of each semantically relevant role attachment. Property inheritance comes only into play when the selection of alternative semantic rules is constrained to the one(s) inherited from the most specific case frame. Hirst uses strong typing at the conceptual *object* level only, while we use it simultaneously at the grammar and domain knowledge level for the processing of semantic schemata.

Charniak & Goldman (1988) and Jacobs (1991) specify semantic *rules* in the context of inheritance hierarchies. So, they achieve a similar gain in generality as we do. Unlike our approach, they still provide specific rules for grammatical phenomena (genitives, adjectival noun modifiers, etc.), while we abstract from these phenomena and collapse them in *single* schemata (as with si_{dir}) whenever possible.

The incorporation of inheritance-based abstraction principles within a lexicalized grammar as a basis for a truly cognitive model of on-line comprehension has also been suggested by Jurafsky (1992). While his proposal is focused on architectural issues how to combine linguistic, computational and psychological criteria for efficient comprehension (e.g., by introducing load constraints on STM, lexical salience weights), he is not explicit about the details of grammatical and semantic specifications. The same argument applies to Lombardo et al.'s description of an incremental interpreter based on a lexicalized dependency grammar (Lombardo et al., 1998). Focus is on the procedural aspects of parsing and simultaneous semantic interpretation rather than on a methodology for semantic interpretation. Paredes-Frigolett & Strube (1996) introduce an approach to interfacing an HPSG parser with a powerful logical representation language such that incrementality of parsing and interpretation are preserved. They discuss a sample parse that builds on large amounts of fine-grained knowledge pieces, but they do not provide evidence for the generality and scalability of their approach beyond the example they discuss. Our approach offers descriptional parsimony as required by any reasonable model of cognitively plausible language comprehension. We have also indications that it scales to real-world text understanding applications (Hahn et al., 1999).

Conclusions

We proposed a principled approach to the design of compact, yet highly expressive semantic interpretation schemata. They derive their power from two sources. First, the organization of grammar and domain knowledge, as well as semantic interpretation mechanisms, are based on inheritance principles. Second, grammar and domain knowledge interact closely via a lean interface — the hierarchy of interpretation schemata. So, the incrementality of semantic interpretation can be described in a general, parsimonious and, hopefully, plausible way. Also, semantic interpretation which has recently become a somewhat marginalized topic of NLP research is given a focused and self-contained theoretical foundation.

Acknowledgements. M. Romacker is supported by a grant from DFG (Ha 2097/5-1). We owe special thanks to Katja Markert for fruitful discussions.

References

Charniak, E. & R. Goldman (1988). A logic for semantic interpretation. *Proc. of the 26th Annual Meeting of the ACL*, pp. 87–94.

Creary, L. G. & C. J. Pollard (1985). A computational semantics for natural language. In *Proc. of the 23rd Annual Meeting of the ACL*, pp. 172–179.

Garrod, S. & A. J. Sanford (1985). On the real-time character of interpretation during reading. *Language and Cognitive Processes*, 1(1):43–59.

Haddock, N. J. (1989). Computational models of incremental semantic interpretation. *Language and Cognitive Processes*, 4(3/4):337–368.

Hahn, U. & K. Markert (1997). In support of the equal rights movement for literal and figurative language: a parallel search and preferential choice model. In *Proc. of the 19th Annual Conference of the Cognitive Science Soc.*, pp. 289–294.

Hahn, U., M. Romacker & S. Schulz (1999). How knowledge drives understanding — matching medical ontologies with the needs of medical language processing. *Artificial Intelligence in Medicine*, 15(1):25–51.

Hahn, U., S. Schacht & N. Bröker (1994). Concurrent, object-oriented natural language parsing. *International Journal of Human-Computer Studies*, 41(1/2):179–222.

Hahn, U. & M. Strube (1996). Incremental centering and center ambiguity. In *Proc. of the 18th Annual Conference of the Cognitive Science Society*, pp. 568–573.

Hirst, G. (1988). Semantic interpretation and ambiguity. *Artificial Intelligence*, 34(2):131–177.

Jacobs, P. S. (1991). Integrating language and meaning in structured inheritance networks. In J. Sowa (Ed.), *Principles of Semantic Networks*, pp. 527–542. Morgan Kaufmann.

Jurafsky, D. (1992). An on-line computational model of human sentence interpretation. In *Proc. of the 10th National Conference on Artificial Intelligence*, pp. 302–308.

Lombardo, V., L. Lesmo, L. Ferraris & C. Seidenari (1998). Incremental interpretation and lexicalized grammar. In *Proc. of the 20th Annual Conference of the Cognitive Science Society*, pp. 621–626.

Milward, D. (1995). Incremental interpretation of categorial grammar. In *Proc. of the 7th Conference of the European Chapter of the ACL*, pp. 119–126.

Paredes-Frigolett, H. & G. Strube (1996). Integrating world knowledge with cognitive parsing. In *Proc. of the 18th Annual Conference of the Cognitive Science Society*, pp. 92–97.

Pereira, F. & M. Pollack (1991). Incremental interpretation. *Artificial Intelligence*, 50(1):37–82.

Sondheimer, N., R. Weischedel & R. Bobrow (1984). Semantic interpretation using KL-ONE. In *Proc. of the 10th International Conference on Computational Linguistics & 22nd Annual Meeting of the ACL*, pp. 101–107.

Strube, M. & U. Hahn (1995). PARSETALK about sentence- and text-level anaphora. In *Proc. of the 7th Conference of the European Chapter of the ACL*, pp. 237–244.

Thibadeau, R., M. Just & P. Carpenter (1982). A model of the time course and content of reading. *Cognitive Science*, 6:157–203.

Tyler, L. & W. Marslen-Wilson (1977). The on-line effects of semantic context on syntactic processing. *Journal of Verbal Learning and Verbal Behavior*, 16:683–692.

Woods, W. & J. Schmolze (1992). The KL-ONE family. *Computers & Mathematics with Applications*, 23(2/5):133–177.