Tilburg University

Language understanding by computer

Bunt, H.C.

Publication date:
1990

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):

Bunt, H. C. (1990). *Language understanding by computer: Developments on the theoretical side*. (ITK research report / Institute for Language Technology and Artificial Intelligence; Vol. 14). Unknown Publisher.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E1

CBM
CBM
R
8409
1990 14
14
UNIVERSITY
LEKE
UNIVERSITEIT
BRABANT

ITK

RESEARCH
REPORT

ITK Research Report No. 14

March 12, 1990

Language Understanding
by Computer:
Developments on the
Theoretical Side

Harry C. Bunt

To appear in: Research Directions in Cognitive Science:
European Perspectives. Vol.5: Artificial Intelligence,
N.Berndsen and D.Sleeman (eds.).

ISSN 0924-7807

©1990. Institute for Language Technology and Artificial Intelligence,
Tilburg University, The Netherlands

LANGUAGE UNDERSTANDING BY COMPUTER: DEVELOPMENTS ON THE THEORETICAL SIDE

Harry Bunt

ITK, Institute for Language Technology and AI
Tilburg, The Netherlands

1 Introduction

This paper consists of three parts. In the first part I discuss the notion of *language understanding* and how it relates to *Artificial Intelligence*. In the second part I review some of the more important recent work on the theoretical side in the design of computer systems that are intended to understand natural language. In the third part I present a view on directions in the computational modelling of language understanding that seem most important for the near future.

2 Understanding language

2.1 Human and artificial language understanding

Until two decades ago, the only type of language understander was the human understander; '*language understanding*' was synonymous with human language understanding, and the study of language understanding was part of cognitive psychology and psycholinguistics. In the sixties, Chomsky

pointed out the theoretical importance of the fact that humans are able to understand infinite varieties of natural-language expressions in spite of finite information-processing resources; the implication being that meaning is encoded in natural language in systematic ways, describable by finite sets of grammatical rules and principles in combination with lexical knowledge.

Since computers are able to store and effectively apply lexicons and large sets of rules in complex tasks, the human understander is no longer the only conceivable kind of language understander. When undertaking the design of a language understanding system, we have to face the question what it is exactly that has to happen inside the system in order to speak of ‘understanding’. In other words, what exactly should be the result of an understanding process. This question does not arise so urgently in the case of human language understanding, since it is usually rather obvious whether someone understands something or not. But the computer case is different, as the classical ELIZA program testifies: even when a computer responds in a seemingly intelligent fashion to natural language inputs, it is far from certain that the system actually understands. To examine this, we should consider the system’s *potential* responses to potential inputs, rather than its actual responses to particular inputs. As a system’s potential responses are determined by the internal state that is created by the processing of an input, the internal state is what we should look at. One of the major attractions of the study of computational language understanding is precisely this: we can directly inspect internal states of the system, in contrast with the case of a human language understander.

2.2 Understanding and meaning representation

In practice, language understanding systems are designed so as to produce symbolic structures supposed to represent meaning, so-called ‘meaning representations’. This raises the question what makes these structures representations of *meaning*; what criteria do we have to determine whether the construction of these structures indeed amounts to *understanding*?

For one thing, these structures should themselves have well-defined meanings, otherwise they can hardly explain much about the meanings of natural-language expressions. For this reason we shall in this paper only consider approaches to computational language understanding that use rig-

orously defined meaning representations. Another requirement is that these structures should have the logical properties necessary to explain the semantic relations between natural-language expressions. Early work in language understanding by computer has been unsatisfactory in these respects. For instance, the structures produced by the interpretation module in Winograd's pioneering system SHRDLU (Winograd, 1972) consist of pieces of LISP code which do not permit general reasoning. The system is therefore unable to answer or even interpret a question like '*What is the colour of a red block*' if it cannot find a red block. Schank's 'conceptual dependency' structures have no formal definitions at all, let alone the required logical properties.

The use of explicit meaning representations may be compared to the role of semantic representations in linguistic semantic theories. For a long time, linguistic theories have remained vague about the nature of semantic representations. The development of Montague grammar in the early seventies changed this dramatically, however. Montague's work has popularized the use of expressions in a formal, logical language as semantic representations. In Montague's view, the use of these representations is merely a matter of convenience: in principle, the rules for associating them with natural-language expressions could be combined with the rules for their evaluation, giving rise to more complex rules that assign semantic values directly to natural-language expressions. In this approach, meanings ('semantic values') are abstract, non-symbolic objects, modelled by mathematical constructs like functions from possible worlds and timepoints to sets of individuals. This is the most popular view on meanings and meaning representations in linguistic semantic theories. A notable exception is Kamp's Discourse Representation Theory (Kamp, 1981), which regards both semantic representations and abstract model-theoretic constructs as essential ingredients of a theory of linguistic meaning.

2.3 Meaning representations and meanings

Can a computer understand language without computing meanings, by only computing meaning *representations*? The answer is *no*, as the following example shows.

Suppose a computer is asked the question: *Which flights from Canada*

arrive on Monday? Let's assume that a meaning representation indicates that we are dealing with a question, and that its 'content' is the set of flights from Canada arriving on Monday:¹

- (1) $\langle \text{QUESTION},$
 $\{x \in \text{FLIGHTS} \mid \text{FROM}(x, \text{canada}) \ \& \ \text{ARRIVE}(x, \text{monday})\}$

In order to answer this question, the machine has to identify those objects which are flights, which depart from somewhere in Canada, and which arrive on Monday. Doing so is applying the semantic definition of the representation language and *evaluating* the expression which represents the semantic, model-theoretic content. This is in fact the same as computing the meaning of the question. To answer a question, we might say, the machine has to compute its meaning. This is not as obvious as it may seem, however. What is obvious is that a question has to be *understood* before it can be answered, but understanding could conceivably consist of constructing a meaning representation, rather than computing the meaning. The example makes clear, however, that a computer can only be said to understand a question if, in addition to constructing a meaning representation, there is also the *ability to evaluate* that representation: to compute the meaning.

Incidentally, we see here the one and only appropriate criterion for deciding the adequacy of an alleged semantic representation: when applying the semantic definition of the representation language, it should yield² the semantic object that actually constitutes the meaning of the original natural-language expression. (Such a 'semantic object' may be a real-world object, a number, a truth value, a set of any of these, a relation, an intension, etc.) The only reliable test for whether a computer has understood something is thus to see what happens when it computes the value of the meaning representation that it has constructed.

How can the value of a meaning representation like (1) be computed, when that value is built up from abstract entities (after all, a flight is an abstract concept)? One way is by consulting a data base which contains descriptions of these objects. A data base is commonly viewed as a

¹The argument that follows is independent of the particular, somewhat naive meaning representation used here.

²The computational perspective in fact adds a second criterion of adequacy: tractability (see section 3.1).

representation of the abstract or concrete concepts that make up a certain universe of discourse. Instead of dealing with abstract or real-world objects the machine deals with descriptions of them, and so computer understanding of language after all does not leave the level of symbolic representation and computation.

Another possibility is to view the computing system, including its data base, as an information processing system which is in a certain state, depending on the contents of all its memory registers. Understanding a natural-language expression would change the state of the system; for instance, understanding the above question should bring the system into a state where it knows that the speaker wants to know the value of the corresponding semantic representation. On this view, the meaning representation (1) is a description of the fact that the machine's state changes in that particular way.³

2.4 Sentence understanding and knowledge structures

According to standard linguistic semantic theory, the meaning of a sentence is determined by the meanings of its words plus its syntactic structure. The meaning of a word, moreover, is given by a relation to a real-world or abstract entity. (This applies to *content* words, at least.) The entities involved in the meanings of content words make up the objects in terms of which our knowledge of the world is expressed. The understanding of a word thus amounts to relating it to entities in the interpreter's total knowledge of the world, and understanding a sentence can be characterized as assigning it a semantic structure in terms of the world knowledge accessed by word meanings. This means that the study of language understanding by computers involves both linguistic semantics, as the study of the linguistic encoding of semantic structures, and Artificial Intelligence, as the study of the representation and application of knowledge.

³Below, in section 4, we shall return to this view in the more specific form where meaning is equated with state-changing potential.

2.5 Language understanding and understanding language use

When *studying* language, one may be dealing with *sentences*: word sequences with certain grammatical properties; when *using* language, by contrast, one deals with *utterances*: sentences produced by someone and addressed to someone in a certain context, with a certain intonation or punctuation, and meant to serve a certain purpose. This means that understanding language 'in action' is not 'just' a matter of sentence semantics, but also involves understanding the purpose of its use in a given context. The full meaning of an utterance has, besides a semantic part, a pragmatic part which describes its communicative function (cf. (1) above).

'Context' should be taken in a double sense here: linguistic and nonlinguistic. The linguistic context is the discourse of which the utterance forms part. 'Nonlinguistic context' refers to a variety of factors, including those concerning the *setting* in which language is used and those concerning the *states* of the agents involved, which include their information, plans, goals, hopes, fears, etc. Both linguistic and nonlinguistic context play an essential part in establishing the meaning of an utterance.

Let us look at an example of actual language use to see this. The text (2) is a transcription of a telephone conversation with the information service at Schiphol, Amsterdam Airport (*C* = client, *I* = information service).

- (2)
- 1 *I*: Schiphol information
 - 2 *C*: Good afternoon. This is Van I. in Eindhoven.
I would like to have some information about flights
to Munich. When can I fly there between
now and ... next Sunday
 - 3 *I*: Let me have a look. Just a moment
 - 4 *C*: Yes
 - 5 *I*: O.K., there are ... three flights every day, one at
nine fifty
 - 6 *C*: Yes
 - 7 *I*: one at one forty ... and one at six twenty-five
 - 8 *C*: Six twenty-five ... These all go to Munich

- 9 I: These all go to Munich
 10 C: And that's on Saturday too
 11 I: That's on Saturday too, yes
 12 C: Right ... Do you also have information about the
 connections to Schiphol by train?
 13 I: Yes, I do
 14 C: Do you know how long the train ride takes to Schiphol?
 15 I: You are travelling from Eindhoven?
 16 C: That's right
 17 I: It's nearly two hours to Amsterdam ...
 You change there and then it's another fifteen minutes,
 so you should count on some two and a half hours
 18 C: O.K. ... Thank you
 19 I: You're welcome
 20 C: Good afternoon
 21 I: Good afternoon

The first three sentences: 'Schiphol information', 'Good afternoon', 'This is Van I. in Eindhoven' can clearly not be understood without considering their role and place in the dialogue (and its setting). In fact, their role and place *are* their meaning in this case: their primary function is to assess the contact established, checking the most elementary conditions for starting a dialogue. (Compare the 'Good afternoon' in turn 2 with those in turns 20 and 21!). Likewise, the 'Yes's in 4, 6, 11, and 13 can of course be interpreted only in relation to the previous discourse.

Other cases where contextual information is needed to establish the communicative function are the sentences 'These all go to Munich' and 'And that's on Saturday too', in turns 8 and 10. The transcriptions of these sentences do not contain question marks, since the utterances did not have an obvious 'interrogative' intonation. A context-independent analysis of these sentences cannot possibly reveal that they function as questions. Note, in particular, that the same sentence occurs in turn 8 as a verification and in turn 9 as a confirmation. The question as to what is the function of an utterance in the dialogue arises in fact for *every* utterance; therefore, the semantic analysis of the sentences should all the time be supplemented with a pragmatic one which relates the sentence to the current context.

The example also illustrates the necessity to extend interpretation over sentence boundaries in order to establish their *semantic* content. (3) lists those sentences in the dialogue that deserve a serious semantic analysis.

- (3) 1 I would like to have some information about flights to Munich.
2 When can I fly there between now and ... next Sunday
3 There are ... three flights every day, one
at nine fifty one at one forty ... and one
at six twenty-five
4 These all go to Munich
5 That's on Saturday too
6 Do you also have information about the connections
to Schiphol by train?
7 Do you know how long the train ride takes
to Schiphol?
8 You are travelling from Eindhoven?
9 It's nearly two hours to Amsterdam
10 You change there and then
it's another fifteen minutes, so you should
count on some two and a half hours

Sentence 1 can be interpreted in isolation, but cannot get the *intended* interpretation, where the quantification domain is restricted to flights from Amsterdam. Sentence 2 suffers from the same problem. Sentence 3 even more so; interpreted in isolation, it represents an obviously *false* assertion, where it is true in its intended interpretation! Similar cases of incompleteness, or *ellipsis*, occur in 7, 9, and 10.

Linguistic context is, by definition, required for the interpretation of anaphoric expressions, expressions that refer to something mentioned elsewhere in the discourse. Clear cases in this text are 'that' in 4 and 5, and 'there' in 10. Less clear cases are 'too' in 5 and 'also' in 6. Other context-dependent expressions are, for instance, 'I' in (1), 'now' in (2), and 'you' in (6), (7) and (10).

Ellipsis and anaphora occur all over the place in dialogues of the kind

that one might like to have with a computer, and are only two of the wide variety of phenomena that require information from both the linguistic and nonlinguistic context.

3 Language understanding and AI systems

From the above analysis we can conclude that language understanding in machines is a process where formal representations of meaning are constructed and evaluated against a body of linguistic and nonlinguistic knowledge. The main theoretical issues in designing language understanding systems therefore relate to:

1. the definition of adequate and computationally tractable meaning representations, taking both semantic and pragmatic meaning aspects into account;
2. the use of pragmatic and discourse-contextual knowledge in constructing meaning representations; and
3. the use of nonlinguistic knowledge in constructing meaning representations, and its representation for their evaluation.

Focusing on these three issues, I consider in this section some contributions to the study of language understanding from recent work in Artificial Intelligence, in particular in Europe.

3.1 Semantic representation

The design of semantic representation formalisms is difficult because, on the one hand, natural language allows the expression of a very rich variety of semantic structures, whose representation calls for a highly expressive representation language. On the other hand, the more expressive the representation language, the greater the danger that its logical properties become

so complex that the evaluation of its expressions becomes computationally intractable.⁴

For instance, it is well known that many semantic phenomena in natural language can only be handled in an intensional framework. Such frameworks have become well-established in linguistic semantics through the work of Montague, who designed an intensional semantic representation language with a formal semantics, called *IL* (Intensional Logic), which permitted a successful attack on a variety of semantic problems. The computation of the values of *IL* expressions, however, presupposes the explicit availability of specifications of all possible combinations of facts in the domain of discourse; though mathematically elegant, for a semantic domain of realistic size this is computationally intractable.

For similar reasons, most of the AI-work on language understanding uses semantic representations in only limited extensions of *first-order* logic, thus prohibiting things like predicates of sets of individuals, functions from predicates to predicates, etc. This is equally unsatisfactory as the lack of intensions, for instance for treating adverbs. In most semantic representation systems the computation of the value of the representation is a combination of consulting a world model and applying postulates that express dependencies between predicates. The application of these postulates requires an implemented deduction system, and computationally tractable deduction systems are not available for unrestricted higher-order logics.

One of the extensions to first-order logic that is gaining wide acceptance consists in making the logic *many-sorted*. This means that the individual objects inhabiting the semantic domain are subcategorized into a variety of 'sorts', and the expressions referring to individuals are *typed* in order to indicate what sort of individual is denoted. Argument positions of predicates can then be labelled with sorts, to be able to check that predicates are applied to appropriate types of arguments. This has been implemented in several of the more ambitious projects in Europe that aim at the construction of language understanding systems.

One case in point is the LILOG (Linguistics and Logic) project in Germany, carried out by IBM Heidelberg and a group of universities (see Herzog

⁴An interesting attempt to get around this difficulty can be found in Moore(1980), where possible-worlds variables are quantified over in an extensional first-order object language.

et al., 1986). In this project, a semantic representation language has been defined called *LLILOG*, which is basically that of many-sorted first-order predicate logic, with a few notable extensions (see Beierle et al., 1988).

In *LLILOG*, a more complex notion of sorts than the standard one is defined, in that (1) the collection of sorts is partially ordered, and (2) sorts can be structured objects, rather than atomic (just as syntactic categories are nowadays usually viewed as feature bundles, rather than just names; this has in fact inspired the *LLILOG* design). Two ways of defining structured sorts are illustrated by the following examples:

- (4) *red-building* == *building*(*colour* : {*red*})
 person == (*father* : *person*(*name* : *C* : *string*),
 name : *C* : *string*)

The first defines the sort *red – building* as the sort *building* restricted by the feature *colour* : {*red*}. The denotation of an expression of the sort *red – building* would thus be a building with a red colour. The second defines the sort *person* recursively as the property of having a father and a name, where the father is again of the sort *person* and the name of the father is identical to that of the son.

In *LLILOG* sorts are treated as part of the language; thus, a sort declaration like:

- (5) *John* ∈ *student*

where *student* is a sort name, is an *LLILOG* expression saying that John is a student.

Sorts also play a role in the TENDUM project in the Netherlands (see Bunt et al., 1985), where a family of semantic representation languages has been designed called the *EL* family, for *Ensemble-theoretical Language* family. The semantics of these languages is based on ensemble theory, an extension of Zermelo-Fraenkel set theory with objects that have a part-whole structure much like sets, but that have no elements. This extension was designed in order to deal with the semantics of mass terms (Bunt, 1985). For

the rest, the EL languages are fully typed, many-sorted and higher-order; their design is based on the semantic representation languages developed earlier in the PHLIQA question-answering project at Philips Research (see Bronnenberg et al., 1980; Scha, 1983). A sophisticated system of complex types is used in these languages, which are expressed in separate members of the EL family: the EL type languages.

The EL languages are also used in the SPICOS project, a joint project of Siemens and Philips aiming at the development of a documentary data base consultation system with natural language input and output in spoken form (see Thurmair, 1987).

In NATTIE, a natural language project in Britain carried out by SRI Cambridge and Cambridge University, a system is built for translating English sentences into formal representations of their literal meanings; this system is called the Core Language Engine (CLE). In the CLE representation language, many-sorted first-order logic is again the point of departure; lambda abstraction and several other features have been added to this.

The CLE system uses two intermediate levels of linguistic analysis between the natural language sentence and its meaning representation. These are a syntactic analysis and a 'logical form', which represents the literal sentence meaning except that the scopes of quantifiers and the references of anaphoric expressions are unspecified. For instance, the sentence *A bishop visited every college* is represented at logical form level by the expression:

(6) *quant(exists, A, [event, A],*
[visit1, A, qterm(a1, B, [bishop1, B]),
qterm(every1, C, [college1, C])])

The *qterm* expressions here represent unscoped quantifiers. A quantifier scoping algorithm, applied in the stage of constructing the full-fledged meaning representations, finds two possible scope assignments in this example:

(7) *quant(exists, B, [bishop1, B],
 quant(forall, C, [college1, C],
 quant(exists, A, [event, A], [visit1, A, B, C])))*

*quant(forall, C, [college1, C],
 quant(exists, B, [bishop1, B],
 quant(exists, A, [event, A], [visit1, A, B, C])))*

It is tempting to consider the ‘logical forms’, like (6), as a kind of meaning representations which are underspecified in certain respects, namely quantifier scopes. However, the language in which logical forms are represented in CLE only has a syntactic definition and its semantics is by no means obvious; therefore, these do as yet not qualify as proper meaning representations.

There are good reasons for being interested in formal meaning representations which are underspecified in certain respects, since the use of formal representation languages sometimes has the undesirable effect that one is forced to be more articulate than natural language warrants. Consider, for instance, the sentence *The boys carried the boxes upstairs*. The sentence leaves open whether individual boys carried individual boxes, or piles of boxes, or groups of boys carried individual boxes, etc. Since the sentence asserts something about *the boys* and *the boxes*, rather than *some (of the) boys/boxes*, a semantic representation should somehow quantify universally over the sets of boys and boxes under consideration. But should these quantifiers relate to individual boys and boxes, or to sets of them? A predicate-logic representation language forces us to make specific choices in these matters, and thus to treat the sentence as multiply ambiguous, although intuitively it isn’t ambiguous at all; rather, it is *vague* in some respects. Only if this vagueness is taken over in the formal representation, one may truly claim to represent what the sentence expresses. (See further Bunt, 1984.)

An interesting aspect of the CLE representation language, illustrated by (5), is that state/event variables are used as arguments of predicates. A verb is not treated as expressing a relation between its subject and its complements, but, following Davidson (1967), as describing an event (or state) where that relation holds. This opens the possibility of treating

optional verb phrase modifiers as predications of events, which in turn permits a uniform treatment of prepositional phrases, independent of whether they modify nominal or verb phrases. This approach is also followed in the representation language CML (Conceptual Modelling Language) of the LOQUI system, a natural language interface to databases developed as part of the ESPRIT project LOKI (see Imlah, 1987), and in the representation language used in the ACORD project, a mainly Scottish-German ESPRIT project (see Calder, Klein, Moens and Zeevat, 1987).

However, some verb phrase modifiers cannot be handled adequately in this way. For instance, fast swimming goes much slower than fast running; therefore, one would like to interpret the adverb *fast* relative to the action involved. For this purpose, Pulman (1987) has proposed the addition of higher-order predicates applicable to state/event variables. For example, the sentence *Mary is swimming fast* would be represented as:

$$(8) \exists e : SWIM(mary, e) \ \& \ FAST(e, \lambda e' : \exists x : SWIM(x, e'))$$

which can be read as: 'There is a swimming by Mary and, by the standards of those events where someone swims, that event was fast'.

This extension has not been implemented so far.

Pulman suggests that "in an actual implementation, we can recast the effects of such a function (like *fast*, HB) in terms of meaning postulates, or rules of inference, at least in the cases where the property in question is capable of being dealt with in a quantitative way. So we might encode something like 'tall' as:

$$(9) GREATER - THAN(height - of(x), 100ft) \rightarrow TALL(x, tree)"$$

This could indeed be done for adjectives like *tall*, but it is unclear how this could work for adverbs like *fast* where it seems wrong to transpose the speed of the event to the speed of its subject. (For example, when *Mary moves her hand fast* it is the object of the verb rather than the subject to which a qualification of speed applies.)

Most of the representation languages mentioned so far take first-order logic as point of departure. This is by far the most common approach in A.I.

circles, but things are changing. The representation languages in the TENDUM project are based on higher-order type logic, and make use in their interpretation of ideas akin to situation semantics (see Barwise & Perry, 1983; Fenstad et al., 1987). In the ACORD project the representation language is based on discourse representation theory; situation semantics is the basis of work carried out at the university of Oslo (see Vestre, 1987; Colban & Fenstad, 1987). These approaches offer better perspectives for integrating sentence semantics with discourse semantics (see Guenther, 1988; Fenstad, 1988).

3.2 Pragmatics and discourse

The importance of both linguistic and nonlinguistic context for understanding an utterance in natural language was stressed in 2.5. In fact, contexts and utterances are related in a similar way as hens and eggs, since contexts give rise to utterances and utterances create contexts. Utterances create *linguistic* contexts by definition; less obvious is the fact that utterances also create *nonlinguistic* contexts.

Linguistic communication occurs because some agent wants to achieve certain goals. These goals may vary from specific, well-defined ones such as wanting to know what time it is, to very general ones like creating a pleasant atmosphere. An agents' goals and beliefs form the primary driving forces behind his utterances, and they form part of the nonlinguistic context. It is also this part of nonlinguistic context which is largely *created by* utterances, since the utterances change the agents' beliefs, goals and other considerations; in other words, the utterances *change the context*. The new context which an utterance creates will determine the continuation. Utterances thus produce new contexts, and contexts produce new utterances.

The study of language in relation to context, in particular to nonlinguistic context, takes place in the branch of linguistics called *pragmatics*; the study of language in relation to linguistic context, in particular as far as interpretation is concerned, is called *discourse semantics*. Discourse semantics has in recent years been an area of highly active research by linguists, philosophers, cognitive scientists as well as A.I. researchers; for an

up-to-date review see the chapter by Guenther in volume 2 of the present series.

The A.I. work in pragmatics which is most relevant to consider here is that where functional aspects of language use are studied in relation to formal and computational representation of mental states and dialogue organization.

When studying the functions of natural-language utterances, it is inviting to adopt the concepts of speech act theory. Speech act theory views the use of language as the performance of acts of communication. Central in this approach is the notion of *illocutionary acts*, being the actions performed in using language. For instance, when we say that a declarative sentence is used to make a request, the request is the illocutionary act performed. Other examples of illocutionary acts are promises, threats, verifications, lies and warnings.

Speech act theory also distinguishes a *perlocutionary* dimension of language use, which concerns the effects that a linguistic act altogether achieves. For instance, convincing someone is not an illocutionary act, but may be achieved indirectly as the effect of one.

Speech act theory has traditionally focused on the development of taxonomies of illocutionary acts and the conditions for their correct performance. This does not immediately bear fruit for the design of language understanding systems. More fruitful is the approach taken by Allwood (1976), where utterances are viewed as actions that signal certain aspects of the speaker's mental state. Such an approach also underlies versions of speech act theory developed in AI by Perrault, Cohen, Allen, Levesque, Bunt, Ostler, and others. In some of these versions, illocutionary acts play an explicit role, in that utterances are assigned illocutionary act labels as part of the representation of their meaning (as in (1) above). In others their role is only instrumental in designing a process that explicates which aspects of a speaker's mental state are signalled by which kinds of utterances (see e.g. Cohen & Levesque, 1985). The most sophisticated versions which can be found in the literature have so far not been implemented; less sophisticated versions are used in the LOQUI system (Wachtel, 1987) and in the TENDUM system (Bunt, 1986).

AI work on speech acts has resulted in significantly deeper insights into the nature of speech act concepts. In a critical review of the foundations of

speech act theory, Levinson (1983) suggests that the most promising way of obtaining solid foundations for notions like illocutionary acts consists in viewing these as functions that map contexts into contexts, as suggested earlier e.g. by Gazdar (1981). The crucial question is then, of course, how a notion of context can be defined to make this idea operational. The AI work on speech acts gives an answer to this question: context can be construed as the relevant aspects of the mental states of the participants in the communication. Which aspects of the mental states are 'relevant' is determined by the overall setting in which the communication takes place. For example, in pure information-exchange dialogues the relevant aspects include knowledge and belief with various gradations of certainty, as well as goals and plans, but exclude emotions like hopes and regrets.

3.3 Language interpretation and knowledge processing

We have seen above that language understanding by its very nature is a combination of determining linguistic structure and applying world knowledge, since the understanding of a sentence involves the identification of the knowledge elements that the linguistic expression refers to. One of the characteristics of the study of language understanding by computer is that the connection between linguistic structure and nonlinguistic knowledge is made explicit and brought within the scope of investigation.

Apart from the role that nonlinguistic knowledge plays in providing the roots for the analysis of meaning, nonlinguistic knowledge is also needed to make linguistic interpretation feasible. One of the most striking properties of natural language expressions is their ambiguity, both at word level and at sentence level. By far the majority of words have a variety of possible meanings, and by far the majority of sentences a variety of possible interpretations due not only to lexical ambiguities, but also to different scope assignments to quantifiers, conjunctions, temporal adjuncts, etc.; due to different choices of grammatical function, of attaching prepositional phrases, and so on. For human understanders, most of this ambiguity goes virtually unnoticed because the application of world knowledge at an early stage makes one reading more plausible than the others. The design of artificial understanding systems makes us acutely aware of both the perplexing

degree of ambiguity in natural language and of the absolute necessity of applying world knowledge at an early stage.

From a logical point of view, there are two methods of combining knowledge elements: the *deductive* and the *model-theoretic* method, also known as the *syntactic* and the *semantic* method, respectively. The deductive method considers the facts about the domain of discourse, assumed to be true, as axioms. Rules of inference are applied to these axioms plus the axioms of logic to try to deduce the truth of certain propositions. The model-theoretic approach works by means of recursive evaluation of complex formulae combined with the assignment of semantic values to the constants and variables of the language in which the formulae are expressed.

We have seen above that most language understanding systems in AI, to the extent that they use formal meaning representations, stick to first-order predicate logic or a modest extension of it. In that case, deductive methods can be applied; for more complex representation languages the deductive method runs into serious difficulties because of its computational complexity. As semantic representation languages become more and more powerful, the model-theoretic approach seems to be the only one possible. This approach has other limitations, however (see 4.3).

The interplay of linguistic and nonlinguistic knowledge would obviously be facilitated by expressing both in the same representation formalism. This would open the way to applying deductive or model-theoretic computations to combinations of linguistic and nonlinguistic elements. A step in this direction is taken in the *LLILOG* language, which has inherited for the definition of complex sorts as in (4) constructions from a formalism designed for the representation of syntactic feature structures, the STUF formalism (see Smolka, 1988; Bouma, König & Uszkoreit, 1988). Similarly, in the TENDUM project the design is under way of a language in the *EL* family for describing syntactic structures, feature operations, and syntactic-semantic grammar rules.

4 Future research

In the previous section we have focused on three issues in the computer understanding of natural language:

1. the definition of adequate and computationally tractable meaning representations;
2. the use of pragmatic and discourse-contextual knowledge in constructing meaning representations;
3. the use of nonlinguistic knowledge in constructing and evaluating meaning representations.

The development of artificial language understanding systems requires further research on each of these topics; such research may, on the other hand, be expected to contribute significantly to our understanding of language understanding. The design of language processing systems moreover by necessity involves the *integration* of a wide range of theories concerning the many aspects of language. A decently designed language understanding system should not be based on separate theories of semantics, pragmatics, discourse, knowledge representation and reasoning, but on a combination of such theories, brought together in a unifying framework. Below I briefly consider the future of research on each of the three topics mentioned above as well as the possibility of a unifying framework.

4.1 Semantic representation

Within the wider perspective of language understanding in context, taking discourse context and nonlinguistic knowledge into account, classical semantic sentence analysis may almost seem a futile business. Indeed, it is fair to say that a considerable number of studies of semantic problems at sentence level have been misdirected, since the key to their solution is to look beyond the sentence boundary, at discourse level, and to bring nonlinguistic contextual knowledge to bear. Nonetheless, the adequate representation of the semantic content encoded by syntactic structures remains a task of obvious fundamental importance. And there is still a lot of work to be done in this area, especially in refining representation formalisms in a way that supports the systematic construction of semantic representations for natural-language sentences.

The following example illustrates the kind of problems to be addressed. Suppose we want to construct a semantic representation of the sentence *The people who presented the reports on morphology are with Siemens*. The use of the definite articles indicates that there is a predication ranging over known domains, and that all the people and reports in these domains are involved. All the people involved are said to be with Siemens, so a universal quantification over the domain of the relevant people is in order. However, the quantification over reports on morphology should not be universal, since the relative clause is not meant to select those people who presented *every* report on morphology, but rather those who presented *some* of those reports. The problem is that, like in the notorious ‘donkey sentences’, the same determiner seems to require a universal quantifier in some syntactic positions and an existential one in others. But the problem is more serious: if we use an existential quantification here, and construct a representation like (16), we miss the point that the sentence states something about *all* the reports on morphology.

$$(16) \quad \forall x \in \{y \in PEOPLE \mid \exists z \in MORPHREPORTS : PRESENT(Y, Z)\} \\ WITH(x, Siemens)$$

The representation (16) would express a true proposition if some of the reports were not presented at all. So in order to represent that all the reports in question are involved, we have to do something to the effect of adding the clause (17):

$$(17) \quad \forall x \in MORPHREPORTS : \exists y \in PEOPLE : PRESENT(y, x)$$

The difficulty here is not so much to give a correct semantic representation of the sentence in standard first-order logic, but to arrive at such a representation through a systematic procedure based on the syntactic and lexical information provided by the sentence.

It is all too well known that classical first-order logic is not powerful enough for the semantic representation of natural language; the representation of

adverbials considered above illustrates this (see also Allen, 1988). The introduction of higher-order predicates or intensions leads to representation languages which are computationally intractable, however. Since the correct representation of semantic structures is impossible in very simple languages, there is really no alternative but to introduce more powerful languages and be very selective in their use in inferential processing. That is, one should not want to use the representations in these languages in general deductive processes of the kind of mechanical theorem provers. The use of appropriate semantic representation formalisms is thus intimately connected with issues of knowledge representation and application, taken up below.

From a linguistic semantic point of view, the hunt should go on for formal representation systems which are powerful enough to represent the semantic distinctions that can be expressed in natural language, without thrusting distinctions upon us which may be relevant from a logical point of view but are not justified on linguistic grounds. Moreover, a useful representation system should allow systematic derivation of representations from linguistic structures.

4.2 Pragmatics and discourse semantics

Compared to semantic interpretation at sentence level, pragmatic interpretation at utterance level and semantic interpretation at discourse level are still poorly understood.

Semantic interpretation at discourse level concerns the way in which linguistic context contributes to the determination of sentence semantics. In (3) above we have seen the crucial and pervasive nature of this contribution illustrated in a sample of natural discourse. Anaphora, ellipsis, and the contextual determination of the domain of reference of noun phrases are some of the notorious manifestations of this phenomenon. Ambiguity is another one. For instance, the classical ambiguity concerning the attachment of the preposition phrase in the sentence *He saw a girl with the telescope* vanishes if the sentence is preceded by: *John looked down the dusky street, wondering whether there was somebody at the far end. It was impossible to see. He took the telescope he'd received for his birthday, placed it on the window-sill and looked again.*

Pragmatic interpretation at utterance level concerns the determination of the communicative functions (or 'illocutionary forces') of utterances. The AI work on speech act theory, mentioned in 3.2, has so far paid little attention to the question *how* the communicative functions of utterances can actually be determined. It seems rather obvious that, just like the semantic content, the communicative function of an utterance is partly determined by properties of the utterance itself, partly by the linguistic context, and partly by the nonlinguistic context. Exactly which properties of utterances are involved (syntactic structure, lexical choices, special particles, prosodic features, punctuation,..); how these interact; what contextual features are involved, and how all these factors interact, is still only poorly understood. Some interesting studies in this area, which contribute to the empirical foundations of the kind of speech act theory emerging from AI, are Beun (1988; 1989) and Saebo (1988).

Another aspect of the pragmatic part of language understanding which is still underdeveloped, but beginning to receive attention in AI, concerns the perlocutionary dimension of speech acts. This dimension presents very complex problems, since the perlocutionary effects which an utterance may bring about are hard to predict. They are not 'controlled' by the utterance, but depend on properties of the mental states that they interact with. But this does not mean that nothing systematic can be said about perlocutionary effects. For instance, in the setting of an information-exchange dialogue like (2), where one of the participants is considered to be the expert on the topic of conversation, the information on the topic supplied by the expert will in general be taken over by the other participant *unless* it would conflict with some other information available to that participant. This means that the adoption of beliefs should take place via *defeasible* rules. Recently, the first steps in developing such rules within a framework of default reasoning have been made by Perrault (1988) and Appelt & Konolige (1988). The formulation of perlocutionary rules and the design of appropriate reasoning mechanisms for predicting perlocutionary effects constitute a task of fundamental importance for the development of models of language understanding and language understanding systems, since speech acts are typically performed with the intention to produce a particular *perlocutionary* effect.

4.3 Linguistic interpretation and knowledge processing

We have seen above that the explicit formal representation of linguistic semantic structures requires a powerful representation language, which poses a problem for effective application of world knowledge - at least when the deductive method is followed. The obvious consequence would seem to be that deductive methods should be avoided; the model-theoretic approach is, by its top-down recursive nature, more directed than the deductive method. However, the model-theoretic method cannot achieve everything through recursive evaluation.

The reason for this is that the model-theoretic approach assumes the value of a complex expression to depend only on the values of its atomic constituents, given the syntactic structure of the expression. But this is not always the case. Suppose we want to express the knowledge that every airplane has an engine, even though we are unable to specify the engine(s) of any particular plane. In this case we know that the formula (18) is true, even though recursive evaluation would fail:

$$(18) \quad \forall x : PLANE(x) \rightarrow \exists y : ENGINE(y) \& HAS(x, y)$$

Recursive evaluation would say that this formula is true in case the expression to the right of the arrow would evaluate to *true* for every value of x which belongs to the set of planes. That expression would in turn evaluate to *true* if, for any value a of x , there is a value b of y such that b is an engine which a has. However, we assumed that we didn't know any engine b of any air plane a , so this expression does not come out true. The point is that our knowledge of the truth of the proposition (18) is *independent* of our knowledge of its constituents.

In model-theoretic semantics, this is solved by adding postulates to the knowledge base, which are considered as constraints that the semantic values of the constants and variables of the language should satisfy. But bringing the knowledge expressed by these postulates to bear amounts to performing deductions, and brings us back to the deductive method. All in all, hybrid systems which combine syntactic and semantic methods of inferencing will probably have to be developed.

It was noticed above that the interplay of linguistic and nonlinguistic knowledge would be facilitated by expressing both in the same representation formalism. But it seems unlikely that this is always possible. 'Visual' knowledge, such as what a rose looks like, or other sensory knowledge such as how a banana tastes or how honeysuckle smells, does not seem to be representable in symbolic form. Is this kind of knowledge relevant to language understanding?

Until very recently, there was little reason to worry about such matters when building a language understanding system. However, a recent trend discussed in Wahlster's chapter in vol. 2 of the present series, is to make such systems *multimodal*. One may, for instance, ask questions about information acquired through a visual channel, which means that semantic representations will contain symbols that refer to picture elements. Research on such systems may be expected to bring some insights into this largely unexplored domain.

4.4 A unifying framework

One of the most important reasons why the design of language understanding systems has a special contribution to make to our understanding of language understanding, is that in such a system many pieces must fit together: syntax, morphology, semantics, pragmatics, discourse semantics, and nonlinguistic knowledge. Moreover, they have to fit together in a computationally feasible way. Within linguistic theory, frameworks have to some extent been constructed where syntax, morphology and semantics fit together; the integration of pragmatics, discourse semantics and nonlinguistic knowledge into these frameworks has not yet been achieved in satisfactory ways. An important task on the theoretical side of designing language understanding systems is the development of a unifying framework in which these pieces fit together at a theoretical level.

Such a framework can be constructed by combining the ideas underlying the AI-approach to speech act theory discussed in 3.2 (including the empirical and perlocutionary extensions mentioned in 4.2) with those underlying current developments in discourse semantics and 'dynamic' sentence semantics.

The latter approach, put forward recently by Groenendijk and Stokhof

(1987; 1988), applies the interpretation method of Dynamic Logic, developed in the theory of programming languages, to the semantics of classical logical languages like those of predicate logic and intensional logic. When these languages are used for semantic representation, the method is indirectly applied to natural-language semantics.

The idea of Dynamic Logic (see Harel, 1984) is that the meaning of a statement in a programming language is the way the state of a machine changes when the program is run. A computer program is treated as a formal-language expression which contains variables, whose values characterize the state of the machine. Running the program changes the state of the machine as the values of variables change. Thus, if g is the function describing the values of the variables before the program P is run, and h the function describing this afterwards, the pair $\langle g, h \rangle$ represents a state transition that forms part of the meaning of P . The full meaning of P is the set of all state transitions which, depending on the initial state, may be brought about by P .

This method has been applied to the semantics of first-order logic (Groenendijk & Stokhof, 1987) and Montague's intensional logic (Groenendijk & Stokhof, 1988). Instead of considering the meaning of a formula ϕ to be its truth value, relative to a model M and an assignment of values to variables g , the meaning of ϕ is taken to be the set of all assignment-changes $\langle g, h \rangle$ such that ϕ is true relative to M and h . The following example illustrates this:

$$(19) \quad \|\exists x : \psi\| = \{ \langle g, h \rangle \mid h[x]g \text{ and } h(x) \in \|\psi\| \}$$

where the notation ' $h[x]g$ ' means that h is equal to g except, possibly, for the values assigned to x .

Groenendijk & Stokhof give a very elegant application of this approach to the treatment of anaphora. When interpreting the discourse *A man is walking in the park. He whistles*, it is possible to first construct a representation for the first sentence, and subsequently add a representation for the second sentence and conjoin this to the representation of the first:

$$(20) \quad (\exists x : [MAN(x) \& WITP(x)]) ; WHISTLE(x)$$

The ‘;’ in this formula represents a noncommutative *and*, whose semantics is defined as follows:

$$(21) \quad \|\phi; \psi\| = \{ \langle g, h \rangle \mid \text{there is an assignment } k \text{ such that} \\ \langle g, k \rangle \in \|\phi\| \text{ and } \langle k, h \rangle \in \|\psi\| \}$$

The combination of (19) and (21) has the amazing effect that the free variable x in *WHISTLE*(x) is bound by the existential quantifier, even though it is outside its scope. Intuitively, this is precisely how anaphora seem to work in natural language.

Whether this dynamic form of model-theoretic semantics can be applied successfully to other discourse phenomena than anaphora still remains to be seen. The underlying idea that “*the meaning of a sentence does not lie in its truth conditions, but rather in the way it changes the (representation of the) information of the interpreter*” is quite similar to the common starting point of discourse-semantic theories, which view the interpretation of multi-sentence discourse as a process that grows and updates representation structures (see Guenther, 1988). It is also closely akin to the idea that the pragmatic meaning of an utterance consists in the way the semantic content changes the mental states of the participants involved.

The combination of these ideas, put forward first in Bunt (1988a; 1988b; 1990), leads to a dynamic semantic/pragmatic interpretation theory which is not confined to the sentence level, as follows.

First, the representation structures grown incrementally according to discourse theories, are construed as *representations of the mental states* of the agents involved. These representations are highly complex structures, having components that consist of information relative to a certain propositional attitude and a certain agent (including multiple nestings and mutual beliefs). Second, the pragmatic part of the meaning of an utterance (its communicative function) is construed as identifying the components of the current representation structures that the utterance addresses. Third, and finally, the semantic part of the meaning of an utterance is construed as indicating how the information within the components, identified by the pragmatic meaning, is to change.

A perlocutionary dimension can be added to this framework as the way in which further changes occur in the contents of the components involved

of the dynamic representations.

Such a framework comprises the common ideas of incremental, dynamic interpretation in currently developing theories of sentence and discourse semantics, and in A.I. in work on pragmatics and default reasoning. The further development of a unifying framework along these lines holds the perspective of a theory of natural language understanding and communication in which all the pieces fit together that are now often scattered over different partial theories. Bringing them together in one coherent theory may be expected to contribute to our understanding of language understanding as well as to the establishment of the foundations of language understanding systems.

References

- Allen, J. (1988) *Natural language understanding*. Benjamin/Cummings, Menlo Park, CA.
- Allwood, J. (1976) *Linguistic communication as action and cooperation*. Gothenburg Monographs in Linguistics 1, Gothenburg University.
- Alshawi, H., Moore, R.C., Moran, D.B. & Pulman, S.G. (1987) Research Program in Natural-language Processing, annual report. SRI International, Cambridge Computer Science Research Centre, July 1987.
- Appelt, D. & Konolige, K. (1988) A practical nonmonotonic theory for reasoning about speech acts. SRI Technical Note 432, April 1988.
- Austin, J. (1962) *How to do things with words*. Clarendon Press, Oxford.
- Barwise, J. & Perry, J. (1983) *Situations and Attitudes*. Bradford Books/MIT Press, London & Cambridge, Mass.
- Beierle, C., Dörre, J., Pletat, U., Rollinger, C., Schmitt, P. and Studer, R. (1988) The knowledge representation language *LLILOG*. LILOG Report 41, IBM Deutschland, Stuttgart.

- Beun, R.J. (1988) Declarative or interrogative, that is the question. Forthcoming in *Journal of Pragmatics*.
- Beun, R.J. (1989) Declarative question acts. In : M.M. Taylor, F. Néel & D. Bouwhuis (eds.) *The structure of multimodal dialogue*. North-Holland, Amsterdam.
- Bouma, G., König & Uszkoreit, H. (1988) A flexible graph-unification formalism and its application to natural-language processing. *IBM Journal of Research and Development* 32 (2), 170 - 184.
- Bronnenberg, W.J., Bunt, H.C., Landsbergen, S.P.J., Scha, R.J.H., Schoenmakers, W.J., & Utteren, E.P.C. van (1980) The question answering system PHLIQA1. In: L. Bolc (ed.) *Natural language communication with computers*. Hauser, Munich.
- Bunt, H.C. (1981) *The formal semantics of mass terms*. Ph.D. thesis, University of Amsterdam.
- Bunt, H.C. (1984) The treatment of quantificational ambiguity in the TENDUM system. In *Proceedings COLING'84*, Stanford, CA.
- Bunt, H.C. (1985) *Mass terms and model-theoretic semantics*. Cambridge University Press, Cambridge.
- Bunt, H.C. (1986) Information dialogues as communicative action in relation to partner modelling and information processing. In : M.M. Taylor, F. Néel & D. Bouwhuis (eds.) *The structure of multimodal dialogue*. North-Holland, Amsterdam 1989.
- Bunt, H.C. (1988a) On-line interpretation in speech understanding and dialogues systems. In: H. Niemann, M. Lang, & G. Sagerer (eds.) *Recent advances in speech understanding and dialog systems*. Springer, Berlin.
- Bunt, H.C. (1988b) Towards a dynamic interpretation theory of utterances in dialogue. In: H. Bouma & B. Elsendoorn (eds.) *Working models of human perception*. Academic Press, New York.

- Bunt, H.C. (1990) DIT - Dynamic interpretation in text and dialogue. In: L. Kálmán, ed. *Proc. of Second Symposium on Logic and Language*, Hajduszóboszló, September 1989.
- Bunt, H.C., Beun, R.J., Dols, F.J.H., Linden, J.A. van der & Schwartzberg, G.O. thoe (1985) The TENDUM dialogue system and its theoretical basis. IPO Annual Progress Report 19, 105-113.
- Calder, J., Klein, E., Moens, M. & Zeevat, H. (1987) Problems of dialogue parsing. ACORD Research Paper EUCCS/RP-1, Centre for Cognitive Science, University of Edinburgh.
- Cohen, P.R. & Levesque, H.J. (1985) Speech Acts and Rationality. In: *Proc. 23rd Annual Meeting of the ACL*, pp. 49-59.
- Colban, E. & Fenstad, J.E. (1987) Situations and prepositional phrases. In: *Proceedings of ACL Europe Conference*, Copenhagen.
- Davidson, D. (1967) The logical form of action sentences. Reprinted in his *Essays on Actions and Events*. Clarendon Press, Oxford 1980.
- Deemter, C.J. van, Brockhoff, G., Bunt, H.C., Meya, M., & de Vet, J.M. (1985) From TENDUM to SPICOS, or: how flexible is the TENDUM approach to question answering? IPO Annual Progress Report 20, 83 - 90.
- Fenstad, J.E. (1988) Natural Language Systems. In: R.T. Nossurm (ed.) *Advanced Topics in Artificial Intelligence*, Lecture Notes in Computer Science 345, Springer, Berlin.
- Fenstad, J.E., Halvorsen, P.-K., Langholm, T. & van Benthem, J. (1987) *Situations, Language and Logic*, Reidel, Dordrecht.
- Gazdar, G. (1981) Speech act assignment. In: A.K. Joshi, I. Sag & B.L. Webber (eds.) *Elements of discourse understanding*. Cambridge University Press, Cambridge, U.K.
- Groenendijk, J.A.G. & Stokhof, M.B.J. (1987) Dynamic predicate logic. Unpublished paper, ITLI, University of Amsterdam.

- Groenendijk, J.A.G. & Stokhof, M.B.J. (1988) Context and information in dynamic semantics. In: H. Bouma & B. Elsendoorn (eds.) *Working models of human perception*. Academic Press, New York.
- Guenthner, F. (1988) Discourse: understanding in context. In *Cognitive Science in Europe vol. 2* (forthcoming).
- Harel, D. (1984) Dynamic Logic. In: D. Gabbay & F. Guenthner (eds.) *Handbook of Philosophical Logic, vol II*. Reidel, Dordrecht.
- Herzog, O. et al. (1986) LILOG - Linguistic and logic methods for the computational understanding of german. LILOG Report 1b, IBM Germany, Stuttgart.
- Imlah, W.G. (1987) CML in LOQUI. Unpublished LOKI working paper BI-30, Hamburg University.
- Kamp, H. (1981) A theory of truth and semantic interpretation. In: J. Groenendijk & M. Stokhof (eds.) *Formal methods in the study of language*. Mathematisch Centrum, Amsterdam.
- Montague, R. (1974) *Formal Philosophy*. Yale University Press, New Haven.
- Moore, R.C. (1980) Reasoning about knowledge and action. SRI Technical Note 151, SRI International, Menlo Park, Cal.
- Niedermair, G. (1986) Divided and valency-oriented parsing in speech understanding. *Proceedings of COLING'86*, Bonn, 593 - 595.
- Perrault, C.R. (1988) An application of default logic to speech act theory. In : M.M. Taylor, F. Néel & D. Bouwhuis (eds.) *The structure of multimodal dialogue*. North-Holland, Amsterdam.
- Pulman, S. (1987) Events and VP modifiers. In: *Recent developments and applications of natural language understanding*. Unicom, Uxbridge.
- Scha, R.J.H. (1983) *Logical foundations for question answering*. Ph.D. thesis, University of Groningen.

- Saebø, K.J. (1988) A cooperative yes-no query system featuring discourse particles. *Proceedings of COLING'88*, Budapest, 549 - 554.
- Searle, J.R. (1969) *Speech acts*. Cambridge University Press, Cambridge, U.K.
- Smolka, G. (1988) A Feature Logic with Subsorts. LILOG Report 33, IBM Germany, Stuttgart.
- Strawson, P.F. (1964) Intention and convention in speech acts. *Philosophical Review* 73 (4), 439 - 460.
- Thurmair, G. (1987) Semantic processing in speech understanding. In: H. Niemann, M. Lang, & G. Sagerer (eds.) *Recent advances in speech understanding and dialog systems*. Springer, Berlin.
- Vestre, E. (1987) Representasjon av Direkte Spoersmaal. Cand. Scient. Thesis, University of Oslo.
- Wachtel, T. (1987) Discourse structure in LOQUI. In: *Recent developments and applications of natural language understanding*. Unicom, Uxbridge.
- Wahlster, W. (1988) Research trends in natural language systems. In *Cognitive Science in Europe vol. 2*. (forthcoming).
- Winograd, T. (1972) *Understanding Natural Language*. Edinburgh University Press.

Bibliotheek K. U. Brabant

17 000 011 13211 6