In: H. Schnelle N. O. Bernsen (eds.): Logic and Linguistics: Research Directions in Cognitive Science – European Perspectives, Vol. 2, Lawrence Erlbaum, p.171-183

Natural Language Systems: Some Research Trends

Wolfgang Wahlster

Computer Science Department,
University of Saarbrücken, 6600 Saarbrücken 11,
Federal Republic of Germany

1. THE IMPORTANCE OF NATURAL LANGUAGE SYSTEMS

In an economy based on the generation and dissemination of information and knowledge, basic research in natural language (NL) understanding and generation can have important positive impacts (cf. Waltz, 1983):

1. It can make computer applications available to segments of the population that are unable or unwilling to learn a formal language.
2. It can increase knowledge productivity in providing automatic means for manipulating knowledge expressed in natural language.

Natural language processing (cf. Allen, 1987) is a prerequisite for advanced knowledge-based systems because the ability to acquire, retrieve, exploit, and present knowledge critically depends on natural language comprehension and production. Natural language concepts guide the interpretation of what we see, hear, read, or experience with other senses.

The knowledge base of a natural language system includes both *linguistic* (e.g. lexicon, grammar, dialogue rules) and *nonlinguistic* (e.g. a description of the objects in the domain of discourse) subparts. Whereas, ideally, the construction of the *nonlinguistic* part of the knowledge base is based on joint research of computer scientists and application specialists, the design and implementation of the linguistic parts rely on cooperation among computer scientists and linguists. For centuries linguists have gathered knowledge about various natural languages. In most cases, unfortunately, this knowledge cannot be used directly in natural language systems because it is represented in computationally untractable formats or because it is not detailed enough for transformation into algorithmic systems. Thus, it is often a collaborative effort among linguists as "experts for language" and computer scientists as "experts for the formal representation of knowledge" to construct linguistic knowledge sources for natural language systems.

A piece of software is called a *natural language system*, if

1. A subset of the input and/or output of the system is coded in a natural language.
2. The processing of the input and/or the generation of the output is based on knowledge about syntactic, semantic, and/or pragmatic aspects of a natural language.

3. KNOWLEDGE SOURCES FOR NATURAL LANGUAGE SYSTEMS

Let us use a simple example to demonstrate the role of natural language in advanced information systems.

Usually, a clerk at an information desk in a train station uses a time-table and a price list to respond to questions of an information-seeking customer. As formatted mass data, these tables and lists are contained in manuals or in a database system. As external data, they are not a part of his internalised knowledge. Because for the clerk, the access to these tables and lists is of critical importance, it is clear that he must be familiar with the organisation of the manuals or the database.

For adequate consultation however, the clerk must, aside from such *access knowledge*, activate other areas of knowledge as well:

1. If the customer says "My son must be on the train to Saarbrücken on Monday. Does it have intercity connections?", he shortens his second sentence by using the pronoun "it" instead of "the train". The clerk's *linguistic knowledge* helps him to select the correct referent for the pronoun. He is able to rule out "My son", "Saarbrücken", and "Monday" as possible antecedents for "it". Furthermore, the clerk can use his linguistic knowledge about speech acts to recognise that the client does not just expect a yes/no answer, but also the departure times of suitable intercity trains to Saarbrücken.

2. If the customer asks "What is the difference between a sleeping-car and a couchette car" the clerk cannot find the answer in his timetable or the price list. But the clerk can use his *conceptual knowledge* to compare both concepts and to identify the distinguishing features of both alternatives for spending the night on a train.

3. If the customer says “I am going to Greece on an excursion together with my professor and our archaeological seminar. Can I use a Eurorail ticket?", the clerk exploits his *inferential knowledge* referring to a rule like "If the client is a student and is under 27 years old, then he can buy a Eurorail ticket". He can apply this rule in a backward chaining mode, which means, that he has to test the if-part of the rule. Using an inference rule like "If someone is attending a university course, then he is a student", the clerk can infer that the customer is a student. Then the clerk can respond with "Yes" if the customer gives an affirmative answer to the clerk's question "Are you under 27?"

4. If the customer says "A return ticket to the Hanover Fair, please" the clerk will most probably offer him a first-class ticket. This response is based on a *user model* which the clerk derives from the assumption that a visitor to a professional fair is on a business trip. With the user class "businessman" the clerk associates certain stereotypical knowledge, e.g. that travel costs are usually reimbursed for business trips. Thus the clerk assumes that a first-class ticket will be preferred by the customer.

In the 1970s, there were many computer science projects which tried to replace the clerk by an information system. In such a scenario, the customer formulated his request in a query language and a DBMS evaluated the query and retrieved the relevant data from a database in which the timetable and the price list were stored – the result of which being unacceptable service quality for most customers. For someone who uses a train twice a year, it is unreasonable to have to learn a formal query language (even if the query language is very simple like pushing a combination of four buttons out of a menu of 50) for getting information on railway connections. Even if the customer would spend the time to learn the query language, the lack of expressive power of current database query languages compared with natural language, e.g. the absence of *indirect speech acts* and *anaphoric devices* as exemplified in 1, makes such a dialogue with an information system rather frustrating.

In our example scenario, a knowledge-based consultation system could provide an increased consultative capacity by combining a database system with a knowledge base, containing a formal representation of the linguistic, conceptual, and inferential knowledge of the clerk, as well as stereotypes for user modelling.

Although presently for each problem mentioned in 1–4 at least one experimental knowledge-based access system which can adequately handle this type of question exists, it is, of course, a long way from implementing the broad-based universal communication capabilities of a clerk into an integrated and robust real-time system.

A particular problem, far from being solved in any system, is the permanent *knowledge acquisition* of humans. In our example, one usually assumes that a clerk reads newspapers or watches TV news programmes. Therefore, he may know that this year the Davis Cup final takes place in Munich. Thus, if a customer asks for a return ticket to the Davis Cup final, the clerk will probably be able to offer him a return ticket to Munich. For a knowledge-based system today, it would be unrealistic to assume that the system updates its knowledge base daily with information potentially relevant for consultation purposes.

4. RESEARCH OPPORTUNITIES

Certainly, natural language (NL) processing has a broad range of possible application areas. But today, many existing NL systems break down when we begin to scale up to larger discourse domains where inherent limitations of the domain are no longer adequate to resolve ambiguities or to control inference. This means that more *breadth* for a NL system cannot be gained without more *depth* of representation and processing. For many possible applications the systems must have a very large vocabulary, deal with a wide range of linguistic constructions, and have a great variety of meaning representation constructions.

Further improvements of the semantic coverage, the richness of the discourse domains and the broadening of the conversational context in the next generation of NL systems will lead to an explosion in the number of semantic interpretations that the systems will have to process. Many ambiguities which are ignored in current systems will be dealt with. To cut down on the processing of spurious readings, future systems have to check a great variety of constraints imposed by user models and discourse models at an early stage of processing. This means that the improved functionality of future NL systems will lead to a considerable increase in the amount of processing resources required.

One possible solution for avoiding a decrease in response time, is to run the NL systems on suitable parallel machines. Today, machines like BBN's Butterfly or Thinking Machines' Connection Machine, both running parallel Common Lisp, seem to be most appropriate. A prerequisite for the success of such an approach is that *parallelisability* becomes a major design criterion.

There will have to be major breakthroughs in cognitive science in order to obtain general and principled solutions for the following research topics.

3.1 Integrated NL Interfaces

Today, a decision maker often has to access many different software systems in order to solve his problem. The next generation of NL systems should *interface intelligently to multiple underlying systems*. Such an integrated interface would allow the user to concentrate on decision making instead of spending his time on the details of which software system offers the information needed, how a problem should be devided to make use of the various systems, how to translate his query into the input language of the selected system, or how to combine the results from several systems into the desired answer.

For such interfaces meta-knowledge and a reasoning component are needed for determining which underlying system or which combination of systems can best fulfil a user's request. In order to choose the correct subset of systems, the interface has to exploit a model of the capabilities of the underlying systems together with a model of the user goals and intentions in the current conversational context.

An example of current research on this topic is JANUS, being developed by BBN in the Strategic Computing Program. JANUS will be an integrated NL interface to three systems:

1. A database which contains information about ships (IDB).
2. A graphics system which can display ships on maps (OSGP).
3. An expert system for force requirements (FRESH).

User-friendly natural language access systems must be able to deal *with extragrammatical language* and *metalanguage*. Users invariably commit errors of orthography, switch word order, violate agreement, omit function words, insert spurious words, or use incorrect punctuation (cf. Carbonell & Hayes, 1983). They often do not notice their errors, so that the system has to recover from sloppy user input, e.g. by exploiting its task and domain knowledge. Metalanguage, i.e. utterances about other utterances, also occur with some regularity in NL interactions, e.g. "when I say 'copy' I mean 'output on the laser printer'."

Thus, basic research on a theory of understanding metalanguage and extragrammatical language should be encouraged.

3.2 Multimodal Communication

In face-to-face conversation humans frequently use *deictic gestures* (e.g. the index finger points at something) in parallel to verbal descriptions for referent identification. Such a *multimodal* mode of communication can improve human interaction with machines, as it simplifies and speeds up reference to objects in a visual world.

The basic technical prerequisites for the *integration of pointing and natural language* are fulfilled (high-resolution, bit-mapped displays and window systems for the presentation of visual information, various pointing devices such as a light-pen, mouse, or touch-sensitive screen for deictic input). But the remaining problem for cognitive science is that explicit meanings must be given to natural pointing behaviour in terms of a formal semantics of the visual world.

Unlike the usual semantics of mouse clicks in direct manipulation environments, in human conversation the region at which the user points (the *demonstratum*) is not necessarily identical with the region which he intends to refer to (the *referent*). In conventional systems there exists a simple one-to-one mapping of a demonstratum onto a referent, and the reference resolution process does not depend on the situational context. Moreover, the user is not able to control the granularity of a pointing gesture, as the size of the predefined mouse-sensitive region specifies the granularity.

Compared to that, natural pointing behaviour is much more flexible, but also possibly ambiguous or vague. Without a careful analysis of the *discourse context* of a gesture there would be a high risk of reference failure, as a deictic operation does not cause visual feedback from the referent (e.g. inverse video or blinking as in direct manipulation systems).

Although the "common visual world" of the user and the system could be any graphics or image, current projects combining pointing and natural language focus on forms or geographic maps.

For example, the TACTILUS subcomponent of our XTRA (cf. Kobsa et al., 1986) system handles a variety of *tactile gestures*, including different granularities, inexact pointing gestures, and *pars-pro-toto deixis*. In the latter case, the user points at an embedded region when actually intending to refer to a superordinated region. XTRA provides NL access to an expert system, which assists the user in filling out a tax form. During the dialogue, the relevant page of the tax form is displayed on one window of the screen, so that the user can refer to regions of the form by tactile gestures. The syntax and semantics of the tax form is represented as a directed acyclic graph (including relations such as "geometrically embedded" or "conceptual part of"), which contains links to concepts in a KL-ONE knowledge base.

The deixis analyser of XTRA is realised as a *constraint propagation* process over these networks. In addition, TACTILUS uses various other knowledge sources of XTRA (e.g. the semantics of the accompanying verbal description, case frame information, the dialogue memory) for the interpretation of the pointing gesture.

While the simultaneous exploitation of both verbal and non-verbal channels provides maximal efficiency, most of the current prototypes do not use truly parallel input techniques, as they combine *typed* NL and pointing. In these systems the user's hands move frequently back-and-forth from the keyboard to the pointing device. Note, however, that multi-modal input makes even NL interfaces without speech input more acceptable (less keystrokes) and that the research on typed NL forms the basis for the ultimate speech understanding system.

Another restriction of current prototypes is that the presented visual material is fixed and finite, so that the system builder can encode its semantics into the knowledge base. While some of the recent NL interfaces respond to queries by generating graphics, they are not able to analyse and answer follow-up questions about the form and content of this graphics, as they do not have an appropriate representation of its syntax and semantics. Here one of the challenging problems is the *automatic formalisation of synthetic visual information* as a basis for the interpretation of gestural input.

Some of the open questions which have to be solved by future research in cognitive science are:

1. How can *non-verbally communicated information* be included in a formal semantic representation of discourse?
2. What is an *adequate architecture* of parsers and generators for multimodal communication?
3. How could a generator decide whether to use a pointing gesture, a verbal description or a combination of both for referent identification (*knowledge-based media choice*)?
4. What are the temporal interdependences of verbal and non-verbal output in deictic expressions (synchronisation of speech and gesture)?
5. How can we cope with complex pointing actions, e.g. a continuous movement of the index finger (drawing a circle around a group of objects, underlining something, specifying a direction or a path) or a quick repetition of discrete pointing acts (emphatic pointing, multiple reference)?

3.3 Spatial Descriptions

When seeing a series of TV pictures showing a part of a highway where several hundred vehicles are lined up one behind the other, each one moving forward only at a snail's pace, we can sum up the scene with the expression "traffic-jam". This is a typical example of a large class of situations that can be described with the statement "one word says more than a thousand pictures" – a reversal of the classical saying.

One of the major goals of cognitive science for the years to come is to gain a better understanding of how perception interacts with language production. A connection to the real world via perception is an optimal starting point for the investigation of *referential semantics* in natural language systems. Traditionally, referential semantics is the study of how phrases in a sentence connect to objects and events in the real world. One of the goals of language-oriented AI research is to attain a completely operational, extreme form of referential semantics that reaches down to the sensoric level.

However, because so far most NL systems have no access to sensory data, in NL research referential semantics often considers only the relationship of phrases to terms in the knowledge representation language. This means that the result of the referential analysis of a nounphrase like "the red car" is simply an identifier like "CAR123", which may be interpreted as an individual constant in a logic-based representation, a node in a semantic network, or a frame instance, depending on the particular knowledge representation language underlying the system. Moreover, a sentence like "The red car stopped now" is mapped onto some simple event representation, e.g. named "EVENT07" which links the referents for the nounphrase and the temporal adverb together. In the model-theoretic sense, for these systems, the knowledge base itself plays the role of the model, so that each syntactic constituent corresponds to an object in the knowledge base.

Such an approach obviously is inadequate for explaining the detailed semantics of spatial prepositions, locomotion verbs or temporal adverbs. Taking referential semantics seriously means that *tactile, acoustic* and *visual perception processes* must be coupled with language analysis and generation processes, so that their mutual dependencies can be studied. If an Al system is ever to use verbs like "hit", "push" and "touch" correctly, taking the subtle semantic differences between them into account, it has to rely on lexical entries for these verbs which establish explicit links to sensory information.

Some authors claim that for symbols to have meaning to an AI system, there must be *sensory symbo*ls to which the *nonsensory symbols* are related by some computational formalism (e.g. Woods, 1983). Even if one does not adopt this extreme position, it is clear that the semantic objects onto which spatial descriptions are mapped must be elements of a domain with a rich geometric structure. Unlike most other semantic theories, the model theory of *situation semantics* (cf. Barwise & Perry, 1983) explicitly refers to a sort of locations L, which consists of connected regions of space-time and various structural relations between them, so that it has much promise for application in natural language scene description (cf. Fenstad, 1988).

Whereas for many years there was little interaction between computer vision and language researchers, in the last few years this situation is changing.

A great practical advantage of natural language scene description is the possibility of the application-specific selection of *varying degrees of condensation* of visual information. The vast amount of visual data accumulated in medical technology, remote sensing and traffic control, for example, can only be handled by a machine. As opposed to a representation of the results of processing digitised image sequences in the form of graphical output, a natural language description of an image sequence can provide the user with more information in less time. If a system is capable of describing the results of interpreting an image sequence in a medical context as "stricture of the left kidney artery", the doctor can first classify this description according to the diagnostic context and later go back to specific segments of single relevant images.

Numerous open questions concerning the formal reconstruction of the interplay between "seeing" and "speaking" must be further explored and resolved.

A problem which is generally left unsolved is one where not only the course of a trajectory in time and space is of decisive importance for the selection of an adequate description. For example, even if all temporal and spatial requirements for the description of the observed trajectory of a moving vehicle are met, a description such as (1) might still be felt to be inadequate:

1. The car is parking in front of the traffic lights.

Only by taking the *intention behind an action* into consideration (cf. Retz-Schmidt, 1986) can an adequate description of the same trajectory be given as in (2):

2. The car is waiting in front of the traffic lights.

One criterion for the choice of soccer as a domain of discourse in the VITRA project (cf. Wahlster, 1988b) was the fact that the influence of the agents' assumed intentions on the description is particularly obvious here. Thus (3) and (4) describe the same process in time-space but imply different team membership for player Meyer:

3. Meyer kicked the ball out of play next to the goal.
4. Meyer barely missed the goal.

In (3), the player has no intention of getting the ball into the goal, but deliberately kicks it out of play. In (4), by comparison, the player's kick was clearly aimed at the goal as expressed in the verb "miss".

One advantage of this domain of discourse is that, given the position of players, their team membership and the distribution of roles in standard situations (e.g. penalties and corners), stereotypical intentions can be assumed for each situation. Given the current state of *plan recognition* research, then, the chances of successfully reaching the described research goal are better than in other less schematised situations.

A second problem arises from the fact that a system, in order to generate communicatively adequate descriptions, must construct a model of the visual conceptualisations which the system's utterance elicits in the hearer's mind. Such a user model (cf. Wahlster & Kobsa, 1986) can become relevant for the decision during sentence generation as to whether, instead of a definite description, a pronoun might also be understandable for the hearer.

Let us suppose that the system has just generated the following text as a description of an observed situation:

5. In the left half, Jones is running toward the goal with the ball. Meyer is chasing him and trying to attack him. But Meyer is too slow.

Because it is not possible for the hearer to visually follow the action on the field, s/he can only form a rough idea of the spatial setting. It is imperative that the system be able to put itself into the hearer's place and take the hearer's possibly imagined conceptualisation into account before continuing to generate sentences.

If the system is planning to generate sentence (6), it must decide, in order to conform with the conversational maxim of cooperativity, whether the referent of the pronoun "him" can be unambiguously determined by the hearer:

6. Now only the goal keeper is between him and the goal.

Only "Jones" and "Meyer" in the preceding text are possible referents for the pronoun. Because Meyer was mentioned last, this referent is the first to suggest itself to the reader in purely textual terms. In the sense of an *anticipation feedback loop* (cf. Jameson & Wahlster, 1982), however, the system could recognize that this resolution of the anaphora is inconsistent with the assumed spatial conceptualisation in the hearer's mind by accessing the *imagination component* of the user model. Therefore, "Jones" is the only unambiguous referent for the pronoun that is compatible with the user model. Only after such a successful understanding process has been anticipated should the planned sentence be generated. Otherwise, the system would not be able to employ pronominalisation to shorten its sentences but would have to resort to the use of proper nouns, for example.

3.4 Non-monotonicity in Understanding

Due to the serial nature of utterances and dialogues, non-monotonicity is pervasive in natural language processing. When reading text left-to-right, *default assumptions* made early in the sentence must often be withdrawn as reading proceeds (cf. Zernik & Brown, 1988). A similar problem occurs when following a dialogue, as assumptions about the dialogue partner, derived from his dialogue behaviour at an early stage of the conversation, often have to be retracted in a later phase of the dialogue (cf. Wahlster, 1988a).

During the last years, non-monotonic inference was identified as a central part of *common sense reasoning*. A number of formalisms for non-monotonic logic have been proposed, none of them completely successful (cf. Ginsberg, 1987). On the practical side there has been better success at integrating so-called *reason-maintenance systems* into AI systems for belief revision.

However, the role of default reasoning in natural language processing is an important research question, which has barely been addressed.

Zernik and Brown (1988) show that even during the various processing stages for an apparently simple sentence [see (7a) and (7b) below] interpretations are asserted and retracted dynamically. Sentence (7a) yields the initial hypothesis that "John got the battery".

7a. John needed a new battery. He took it

This interpretation is based on two assumptions: Unless otherwise observed, a generic word like "take" indexes the generic meaning of a physical transfer and the pronoun "it" refers to the last physical object mentioned in the discourse. However, as reading proceeds,

7b. John needed a battery. He took it up with dad

the initial interpretation must be revised, because a more specific lexical entry is found. The idiomatic sense of "He took it up with his dad" leads to an interpretation in which "it" refers to John's goal of getting a battery and "take up" is understood as "raising an issue". However, the initial reading must be recovered when reading the end of the sentence.

7c. John needed a battery. He took it up with dad from the basement.

The *Non-Monotonic Grammar* (NMG) described by Zernik and Brown is a first attempt to use a reason maintenance package for enhancing a parser's capabilities. NMG uses dependency directed backtracking, so that unlike other current parsers it does not have to recompute the initial interpretation after retracting the idiomatic reading. Processing garden path sentences and parsing in the presence of lexical gaps are other tasks highlighting the role of default reasoning in text comprehension.

User modelling is another important research area in natural language processing where non-monotonic reasoning plays a crucial role. Consider the following dialogue (U = User, S = System) with a tutoring system (cf. Wahlster, 1988a):

8. **S**: Tell me about California.
9. **U**: San Francisco is the capital of California.
10. **S**: No, that's wrong.
11. **U**: I see. So, that's not the capital.
12. **U**: Then, what is the capital?
13. **S**: Sacramento.
14. **S**: Now, tell me why you mentioned San Francisco first, when you began to talk about California.

A simple consequence of the user's response (9) is an entry in the system's user model, which represents the fact, that the system believes that the user believes (B1). After (10), and certainly after (11), the model should contain (B 1'):

B1. capital (CaIifornia, San-Francisco).
B1'. not (capital (California, San-Francisco).
B2. capitaI (California, Sacramento).

This means that the user-modelling component has to remove (B1) from the user model [in a reason maintenance system this causes (B1) to be added to the set of beliefs, which are currently "out"]. After (13) the user's belief (B2) should be added to the system's user model. If the *a priori* user model contains “For each state there exists one and only one capital" as a mutual believed fact, then the user-modelling component can also remove (B1’) after adding (B2). The GUMS system is a first attempt to integrate a belief revision component into a dialogue-based user-modelling component (cf. Finin & Drager, 1986).

3.5 Transmutable Systems

A general-purpose natural language dialogue system should be adaptable to applications that differ not only with respect to the domain of discourse, but also to dialogue type, user type, and intended system behaviour. In Wahlster and Kobsa (1986), we call such Systems, which are transportable and adaptable to diverse conversational settings, *transmutable systems*. A first attempt to build a transmutable system was our design of the experimental dialogue system HAM-ANS (see Hoeppner et al., 1983), whose dialogue behaviour can be switched from a "cooperative" mode (e.g. the system answers questions about a traffic scene) to a "interest-based" mode (e.g. the system tries to persuade the user to book a room in a particular hotel).

When people communicate, they do so for a purpose specific to the conversational situation. On the other hand, most of the systems developed so far have no interest beyond providing the information-seeking user with relevant data. In the long run, natural language systems as components of advanced knowledge-based systems must perform a greater variety of illocutionary and perlocutionary acts: they may teach, consult, or persuade the user, inspire him to action or argue with him (see Bates & Bobrow, 1984; Wahlster, 1984; Webber, 1986; Woods, 1984). The major problem builders of transmutable systems are confronted with, is the lack of representational vocabulary for the declarative description of the relationship between the system and the user, the *system’s intended dialogue behaviour* and the associated *conversational tactics*.

3.6 Text Generation

The task of text generation involves translating knowledge represented in a formal language in a computer’s memory into natural language. For example, information encoded in a knowledge representation language like KL-TWO regarding the use or repair of a technical system could be used as the basis for the automatic generation of instruction manuals in a *variety of natural languages*. Moreover, from the same representation, different manuals could be generated for different audiences, such as beginners, expert users, or maintenance personnel. Ultimately, techniques of user modelling (cf. Wahlster & Kobsa, 1986) can tailor the documents to the background of each particular individual, making the text more understandable and generating the correct level of detail. Especially in a tutorial framework, it is important to make the generated text interesting. Heuristics for increasing the tension and fluency of a text must be integrated. In order to speed up the comprehension process the system has to generate *meta-utterances* like "as I have stated before" or "generally speaking" (cf. Zuckermann & Pearl, 1984).

There are two main aspects of generation: (1) deciding what to say, and (2) deciding how to say it. For the first task it is important to treat text generation as a special case of a goal-oriented action, which requires planning and reasoning. Combined with speech act theory the planning approach to text generation promises significant advances, but it presupposes efficient inference systems for reasoning about the beliefs, goals and actions of rational agents.

Another goal which requires considerable basic research is that of matching the NL production capabilities of systems with their comprehension capabilities. This is a prerequisite for building advanced *writer's workbenches*.

Such document preparation aids could detect errors in spelling and grammar, suggest paraphrases of passages of the text to make them more understandable, suggest ways to shorten the text or to restructure the document. This involves integrating work on text processing, document formatting and natural language processing.

3.7 Tools for NL Research

An investment in good tools for NL research and the development of NL systems will pay excellent dividends. Such tools make it possible to test new theories or methods and to build new systems more rapidly, by using *off-the-shelf components* for programs.

In order to speed up the development of large lexicons and grammars, and to ensure their well-formedness and consistency, specialised software tools must be developed, much like the structured editors and programming environments that improve programmer productivity. Utilities to trace the application of the lexicon and the grammar to a set of examples and to display the processing graphically can improve the debugging and the quality assurance processes.

We do need to have these tools well-documented, portable, reliable and widely distributed as public domain software. The sharing of tools should be encouraged by funding the development and maintenance of research tools, for example morphological analysers, parser generators, knowledge representation systems, planning and inference components, and language generators.

REFERENCES

Allen, J. (1987). *Natural language understanding*. Menlo Park: Benjamin/Cummings.

Barwise, J. & Perry, J. (1983). *Situations and attitudes*. Cambridge, Mass.: MIT Press.

Bates, M. & Bobrow, R.J. (1984). Natural language interfaces: What's here, what's coming, and who needs it. In W. Reitman, (Ed.); *Artificial intelligence applications for business*, pp.179-194. Norwood: Ablex.

Carbonell, J.G. & Hayes, P.J. (1983). Recovery strategies for parsing extragrammatical language. *Journal of Computational Linguistics*, 9, (3-4), 123-146.

Fenstad, J.E. (1988). Natural language systems. In R. Nossum, (Ed.), *Advanced course on AI.* Heidelberg: Springer.

Finin, T. W. & Drager, D. (1986). GUMS: A general user modeling system. In *Proc. of the 6th Canadian Conf. on AI*, pp.24-29. Montreal.

Ginsberg, M.L. (Ed.) (1987). *Readings in nonmonotonic reasoning*. Los Altos, Calif.: Kaufmann.

Hoeppner, W., Christaller, Th., Marburger, H., Morik, K., Nebel, B., O'Leary, M. & Wahlster, W. (1983). Beyond domain-independence: Experience with the development of a German language access system to highly diverse background systems. In *Proc. of IJCAI-83*, pp.588-594. Karlsruhe.

Jameson, A. & Wahlster, W. (1982). User modelling in anaphora generation: Ellipsis and definite descriptions. In *Proc. of 1st ECAI*, pp.222-227. Orsay.

Kobsa, A., Allgayer, J., Reddig, C., Reithinger, N., Schmauks, D., Harbusch, K., & Wahlster, W. (1986). Combining deictic gestures and natural language for referent identification. In *Proc. of COLING-86*, pp.356-361.

Retz-Schmidt, G. (1986). Script-based generation and evaluation of expectations in traffic scenes. In H. Stoyan, (Ed.), *GWAI-85. 9th German Workshop on Artificial Intelligence*. Heidelberg: Springer.

Wahlster, W. (1984). Cooperative access systems. *Future Generation Computer Systems*, *1,* (2), 103-111.

Wahlster, W. (1988a). Distinguishing user models from discourse models. In A. Kobsa & W. Wahlster, (Eds.), *Special issue on user modeling*. Journal of Computational Linguistics.

Wahlster, W. (1988b). One word says more than a thousand pictures. *Computers and Artificial Intelligence,7*.

Wahlster, W. & Kobsa, A. (1986). Dialog-based user models. In G. Ferrari (Ed.), *Special issue on natural language processing. IEEE Proceedings, 74* (7), 948-960.

Waltz, D. (1983). Artificial intelligence: An assessment of the state-of-the-art and recommendation of future directions. *AI Magazine*, Fall, 55-67.

Webber, B.L. (1986). Questions, answers and responses: Interacting with knowledge base systems. In: M. Brodie & J. Mylopoulos, (Eds.), *On knowledge base management systems*. New York: Springer.

Woods, W. A. (1983). Under what conditions can a machine use symbols with meaning? In *Proc. of IJCAI-83*, pp. 47-48. Karlsruhe.

Woods, W. A. (1984). Natural language communication with machines: An ongoing goal. In W. Reitman, (Ed.), *Artificial intelligence applications for business*, pp. 195-209. Norwood: Ablex.

Zernik, U. & Brown, A. (1988). Default reasoning in natural language processing: A preliminary report. To appear in *Proc. of COLING'88*, Budapest.

Zuckermann, I. & Pearl, J. (1984). *Listener model for the generation of meta-technical utterances in math tutoring*. Technical Report, UCLA, Cognitive Science Lab.