Approaches to natural language discourse processing*

Paul Mc Kevitt[†]
Derek Partridge[‡]
Yorick Wilks[§]

Keywords: natural-language processing, natural-language discourse, semantics, intention, coherence, context, OSCON.

Abstract

One of the most difficult problems within the field of Artificial Intelligence (AI) is that of processing language by computer, or *natural-language processing*. A major problem in natural-language processing is to build theories and models of how individual utterances cling together into a coherent discourse. The problem is important because, to properly understand natural language, a computer should have some sense of what it means for a discourse to be coherent and rational. Theories, models and implementations of natural-language processing argue for a measure of coherence based on three themes: meaning, structure, and intention. Most approaches stress one theme over all the others. Their future lies in the integration of components of all approaches. A theory of intention analysis solves, in part, the problem of natural-language dialogue processing. A central principle of the theory is that coherence of natural-language dialogue can be modelled by analysing sequences of intention. The theory of intention analysis has been incorporated within a computational model, called Operating System CONsultant (OSCON), implemented in Quintus Prolog, which understands, and answers in English, English questions about computer operating systems. Theories and implementations of discourse processing will not only enable people to communicate better with computers, but also enable computers to better communicate with people.

[0]Department of Computer Science, University of Sheffield, PO Box 600, Mappin Street, GB- S1 4DU, Sheffield, England, EC, E-mail: p.mckevitt@dcs.shef.ac.uk.
[0]Department of Computer Science, University of Exeter, GB- EX4 4PT, Exeter, England, EC, E-mail: derek@dcs.exeter.ac.uk.
[0]Computing Research Laboratory, Dept. 3CRL, Box 30001, New Mexico State University, Las Cruces, NM 88003-0001, USA, E-mail: yorick@nmsu.edu.
*This work has been funded, in part, by a Research Fellowship from the Computing Research Laboratory (CRL), New Mexico State University, Las Cruces, New Mexico, USA, a grant from U S WEST Advanced Technologies, Denver, Colorado, USA, and a Visiting Fellowship from The School of Electronic Engineering, Dublin City University, Dublin, Ireland, EC.

1 Introduction

One of the most difficult problems within the field of Artificial Intelligence (AI) is that of processing language by computer, or *natural-language processing* (see Partridge 1991). Much of the work in natural-language processing has concentrated on modelling the structure, meaning, and usage of individual utterances[1]. However, it is not often that natural-language utterances occur on their own, independent of some context or other. Hence, one of the problems in natural-language processing is to build theories and models of how individual utterances cling together into a coherent discourse. It can be argued that to understand a discourse properly, a computer should have some sense of what it means for a discourse to be coherent. Current theories and models of natural-language processing argue that the coherence of a discourse can be measured in terms of three main themes: meaning, structure, and intention. Most of the approaches stress one theme over all the others. We shall compare the different approaches and briefly describe how a theory of intention analysis for natural-language discourse processing has been incorporated within a computational model called the Operating System CONsultant (OSCON). OSCON is implemented in Quintus Prolog and understands, and answers in English, English questions about computer operating systems.

We shall only discuss theories and computational models within the symbolic processing framework of artificial intelligence. Those interested in connectionist approaches can look at Barnden and Pollack (1991) and Reilly and Sharkey (1991). First, we shall introduce some of the problems which occur in processing natural-language discourse.

2 Some problems in processing discourse

The set of problems which occur in the processing of natural-language discourse is enormous. Before we go on to look at some of the problems in more detail, we shall discuss the phenomenon of *indexicality*[2] which prevails during any discussion of such problems. Bar-Hillel (1954) introduces indexicality in language where he speculates that 90% of declarative sentences include implicit references of the speaker, addressee, time and/or place of utterance in expressions such as first and second pronouns (*I*, *you*), demonstratives (e.g. *this*), tenses, and adverbs like *here, now, yesterday.* He argued that indexicality is an inherent and unavoidable property of natural language. The argument is basically that utterances like those in (1) below cannot be understood in terms of truth values unless we take their context into account. Thus, with respect to the Gulf crisis in 1990 (1a) would seem true if it was uttered by say, the leader of Japan, but would seem false if uttered by say, the president of the USA. Hence, the speaker of an utterance, or *who* says the utterance, is of importance. It is possible, in general, to map utterances like (1a) into some representation,

[1] The term *utterance* will be used to refer to any unit of natural language, whether it be a word, phrase, sentence, or exclamation. An utterance may be well-formed or ill-formed. Reference will usually be made to written, rather than spoken utterances, unless indicated otherwise.

[2] Indexicals are utterances whose meanings are functions of their context.

but it is not possible to say whether that representation is true, or false. A full account of the context of the utterance would provide such information. (1b) is interesting by the very fact that *when* it occurred is important. If it was uttered before late 1990 then it would seem true, yet soon after that period it would seem false. Finally, the last utterance is of interest because of where it is uttered. For example, if uttered aloud in front of Mrs. Thatcher, while she is leader of the Conservative party, it may sound quite strange, unless it was meant in humour.

(1) a We will not send soldiers into the Gulf.
b The USA sold arms to Iraq.
c I am secretly interested in contesting the Leadership
of the Conservative Party.

Hence, what we see here, is that the truth of utterances depends on who, when, and where these utterances are being said.

There are a number of problems which occur in natural-language dialogue due to the ability of people to use language in context. To give the reader a feel for such problems, we shall discuss some of them.

2.1 Tense

Tense enables sentence tokens to be used where complete meaning of the sentence cannot be derived without taking context into account. Time is usually important in these cases. Take for example the utterances in (2) below.

(2) a The European Community (EC) is going to start a war in the Gulf.
b The European Community (EC) is starting a war in the Gulf.

The complete meaning of (2a) is not determinable, due to an underspecification of time, whereas in (2b) the time is more precisely determinable. The reason for this is that in (2a) it is not clear if the war will be started in days, weeks, or months, whereas in (2b) the time of the start of the war is more precise. A study on the temporal relationships of utterances through verbs, from a logical point of view, is given in Galton (1984). Also, the beliefs[3] and intentions[4] of the speaker come into play, because if utterance (2a) was made by the president of the European Community,

[3]The term *belief* is used to refer to the points of view, or attitudes, an agent has towards some objects, or other agents. There has been much work on the processing of beliefs in natural-language processing. See, for example, the work of Ballim and Wilks (1990, 1991), Barnden et al. (1991), Wilks and Bien (1983), and Wilks and Ballim (1987).

[4]It is difficult to provide a definition of intention, because it means many different things to people from different research backgrounds. In this work, by intention, we mean the purpose of the speaker. This sense of intention will involve the goals, plans, and beliefs of an agent, who has an intention, and how such intentions are indicated to a computer by means of natural language. Such goals could be "to find out where Sainsbury's is", "to eradicate the Poll Tax" or "to print a file on the printer", with the respective plans, [use town-map], [Non-payment/obstruction], [use UNIX-command, 'lpr'].

and it was just at a point where it looked likely that the EC would attack Iraq, then the time reference would be on the order of days rather than say, years.

2.2 Anaphoric reference

In natural language discourse it is possible to use words such as *me*, *here*, and *then* to refer to entities mentioned in other parts of the discourse. This phenomenon is called *anaphoric reference.* When pronouns like *he*, *her*, and *their* are used to refer to other entities such pronouns are called *anaphoric references.*

Some researchers believe that reference is indexical in that it is determined by the *linguistic* context of an utterance, but we believe that reference is indexical also in that it is determined by speaker *intention.* This is shown in (3) below. We argue that the only way the referent of "he" can be determined is by monitoring the intentions of the participants in the dialogue, as "he" could refer to either Mr. Howe, or "Heseltine."

(3) Mr. Howe told Heseltine he should resign because of her attitude to Europe.

The use of *pronouns*, where words like *he, she, and it*, are used to refer to people or things, is a common type of reference. The determination of the the person or thing being referred to,is important. Such entities are called referents. For example, in (4) the utterance depends on who "I" is. If the person saying this utterance is the president of the USA then that is fine, otherwise it might be a case for committing someone to an institution.

(4) I am president of the USA.

The referent of the pronoun "I" is usually considered to be the utterer of an utterance. The use of "he" is more difficult to resolve as the pronoun could be referring to anyone in the discourse under discussion. The speaker's intention can be useful in determining the correct referent. However, note that pronouns do not always have referents. For example, the referent of "his" is unclear in "No man is older than his father." A referent is partially, but not uniquely, determined by the temporal ordering of utterances. Much pronoun reference is resolved by determining the beliefs and intentions of the speaker and hearer that enable the hearer to determine the correct referent.

2.3 Natural-language generation

The context of an utterance will affect how the utterance should be generated[5] in a dialogue. For example, looking at the sample dialogue in (5)[6], between a computer user and a system, it is noted that the system changes its behaviour if it remembers it has answered a related question.

[5] By *generation* we refer to *natural-language generation*, or the generation of natural language from some internal representation. Such representations can take many forms.

[6] In this, and any succeeding dialogues, U refers to a computer user and S to a computer system.

(5) a U: How can I delete a file?
S: 'rm <filename>' will remove filename from your current directory.
For example, 'rm paul.1' will remove paul from your current directory.
b U: How can I delete a directory?
S: Use 'rmdir <directoryname>'.

For query (5a) the system gives an example of how to use the command in question. However, no example is given for query (5b). The user should infer from the context how to use the command. In fact, it may be quite boring for the user to have examples of the use of commands displayed all the time. The point here is that the system should be aware of what it has said to the user to reduce the redundancy of information presented. Certainly, if a user asks the same query one hundred times the system should notice this, and should not just blindly keep answering the same query.

2.4 User modelling

Any useful natural-language interface must build a model of the user for effective cooperative interaction. This is particularly important for the generation of useful natural-language responses. There has been much work in Artificial Intelligence (see Ballim and Wilks 1990, 1991, Carberry 1989, Chin 1988, Kobsa and Wahlster 1988, Mc Kevitt 1991b, and Wilensky et al. 1988), and human-computer interaction (see Benyon 1984, 1987, and Benyon and Murray 1988) on the modelling of users' levels of expertise. Below, there are two example dialogues exemplifying different degrees of user expertise.

(6) a U: How can I delete a file?
S: Use 'rm'.
For example, to remove the file named foo, type 'rm foo'.
b U: What does ruptime do?
S: ruptime is used to list the uptime of all machines
on the network, list the number of all users on them,
and list their load and average.

(7) a U: How can I find out the inode of a file?
S: Use ls -i.

b U: What does uptime do?
S: uptime is like ruptime, except ruptime is for all
machines on the network.

These example dialogues taken from (Chin 1988) show how the Unix Consultant (UC) system behaves depending on the level of expertise of the user. When the user asks (6a) it would seem that he/she doesn't know very much about operating systems. The user is asking a query about a simple topic: how to delete a file. Therefore, the system answers with the right command and an example of its use. When the user asks about ruptime, the system gives a detailed answer as the user is less knowledgeable.

In the second dialogue, (7), the user is asking about finding the inode of a file, which is a more complex query. Chin (1988) argues that such a query indicates that the user might have some expert knowledge of operating systems. Therefore, the answer is given in a short format, which presumes that the user knows how to specify the command syntax. Also, the system assumes that the user is an expert. Then, in (7b) when the user asks about "uptime" the system assumes that "ruptime" is already known to the user, as he/she is considered somewhat of an expert. Hence, it is important to incorporate a user-modeling component into a system for dialogue processing so that the system has a good feel for the type of user being dealt with. Also, the type of user will affect the way the system generates answers to queries. The relation between a dialogue model and a user-model has been the subject of much debate (see Schuster et al. 1988, and Kobsa and Wahlster 1988). A problem with Chin's (1988) approach is that it concentrates too much on the domain specific knowledge of the user whereas we should try to develop more universal aspects of user-models.

In summary, four pertinent subproblems in relation to the problem of natural-language discourse processing have been discussed: (1) the processing of verbs indicating temporality, (2) anaphor resolution, (3) response generation, and (4) user-modelling. Many of the theories of natural-language discourse processing claim to solve problems such as these. The common theme of all such theories is that they incorporate a model of the coherence of discourse. Now we shall move on to discuss the phenomenon of coherence in more detail.

3 Coherence of natural-language discourse

One of the features of natural-language discourse, when used by people, is that it is usually coherent to some extent. If you speak to your car mechanic about what is wrong with your car you will usually try to make your dialogue coherent, so that the mechanic will understand what you are trying to convey. Of course, if the mechanic does not understand your description then the continuing coherence of the discourse may not be obvious to him. Hence, it seems that understanding is inextricably linked with coherence, and the less coherent a conversation, the harder it is to understand, and the less one understands a conversation the less it will appear to be coherent.

The idea of coherence in discourse is tied up with rationality. If we speak to someone we expect him/her to behave rationally, and we recognise that rationality by the coherence of what he/she says. So, if a policeman asks you, "Where were you on the 17th March 1991?," and you reply to the policeman, "Where were you on 17th March 1991?," it might seem rather odd to him. However, the policeman will, most likely, not assume that you are behaving irrationally. He will first try to work out why you gave the answer that you did, and may decide that you did so because you wanted to avoid answering the question. When James Joyce uses the sentence,

> The great fall of the offwall entailed at such short notice the pftjschute of Finnegan, erse solid man, that the humptyhillhead of himself prumptly sends an unquiring one well to the west in quest of his tumptytumtoes: and their upturnpikepointandplace is at the knock out in the park where oranges have been laid to rust upon the green since devlinsfirst loved livvy.
> — Joyce (1939, p. 4)

in his book, Finnegans Wake, it seems ridiculous. The reader of Finnegans Wake does not assume that Joyce is irrational; Joyce probably intended to make a point while using such a strange sentence. Likewise, if someone says, "Mr. Brooke is going to solve the Northern Ireland problem," and you reply, "And pigs will fly," then he does not assume that you are irrational; he works out that you probably said something silly because you wanted to imply the silliness of his initial statement.

However, there seem to be certain rules about how far one can go, so as not to appear irrational. For example, if I uttered, "same please," and I was first in queue at the restaurant, it would not make sense. Also, consider the situation where two colleagues step up to a bar counter, and one orders a rum and coke from one bartender, while the other says, "same please," to another bartender, where the second bartender has not overheard the first request. The second bartender will be understandably confused. In the latter two examples there seems no obvious explanation of why speakers uttered what they did, whereas there is, at least, some explanation in the former examples.

The point is that, under *normal* circumstances, people try as much as possible to assume a speaker is rational, until they cannot do so any more. We assume people are rational, and that they cooperate when they use language to communicate with us. This issue has received much attention in the philosophy of language, linguistics, and artificial intelligence. A philosopher of language, Grice (see Grice 1957, 1969, 1971, 1975) has termed this phenomenon the Cooperative Principle.

As coherence of natural language is an indication of rationality, this coherence is very important. It is important because how we interact with people is dependent on their rationality, and if someone seems irrational then we may decide to terminate communication with them. Likewise, we may decide to terminate natural-language dialogue communication with a computer if the dialogue appears incoherent, and hence irrational. And the computer may do the same if we appear irrational. Hence, if a computer is to conduct natural-language dialogue communication then it is important that it has some sense of what it means for a dialogue to be coherent.

Halliday and Hasan (1976) have argued that coherence in discourse is a linguistic property of sentences. They point out that it is (1) anaphoric expressions, and (2) lexical items appearing across sentences, which establish coherence. In the former case, the links between references, and their previous occurrence, is considered, and in the latter case words are considered. Halliday and Hasan used the word *cohesion* to describe the linguistic properties of text that contribute to coherence, and argued that contributions to coherence are not a function of the reference of the words in a discourse, but of the forms themselves, and their histories of occurrence.

Another set of theories (Rumelhart 1975, and van Dijk 1972, 1977) treat coherence as reflecting a defining property for discourse, as comparable to grammars for sentences. They argue that coherence of discourse is due to grammars of coherence, and they have specified these grammars in detail. A good survey of such grammars occurs in Andersen and Slator (1990).

A more useful possibility of defining the coherence of discourse might be one which considers Grice's Cooperative Principle. Each utterance is considered to say something necessary, true, and relevant to accomplishing some objective, in which it is believed that the speaker and listeners are mutually interested. A coherent discourse would be one where an interpreter can reconstruct the speaker's plan from his/her utterances, and make inferences about them to understand the dialogue at hand. Hence, coherence here is defined not by the discourse components themselves, but by the effort needed to construct a plan to attribute to the speaker involved in producing the text. This will depend on how hard, or easy, it is to take each sentence as representing a true, necessary, and relevant contribution to the plan. It seems to be the case that the coherence of text is based, not only on properties of the text itself, but also on the inferences made by discourse hearers. Now, that we have discussed the phenomenon of coherence in discourse we will move on to discuss the various theories of natural-language discourse processing.

4 Theories of discourse processing

There are a number of theories of natural-language discourse processing, and computational models of these theories exist. Theories concentrate on the themes of semantics, structure, and intention. A common principle of all approaches is that they provide a model of the coherence of discourse. Semantic theories argue that the coherence of a discourse is a feature of its meaning and that if you model the meaning, the coherence falls out of that. For example, there are theories regarding the coherence of semantics in discourse such as those proposed by Fass (1988), Schank (1972, 1973, 1975), Schank and Abelson (1977), and Wilks (1973, 1975a, 1975b, 1975c).

Structure-based theories argue that a discourse can be modelled in terms of structural units which can be recognised and marked out in the discourse. These theories have given names like *topic* and *focus* to such units. Examples of such theories are proposed in Alshawi (1987), Dale (1988, 1989), Grosz (1983), Grosz et al. (1983), Grosz and Sidner (1986), Sidner (1983, 1985), and Webber (1978).

Finally, other theories model the coherence of discourse from the point of view of intention, or the goals, plans and beliefs of the participants in the discourse. These approaches argue that people's intentions underlie their use of language, and that by modelling these intentions one can model language. The approaches do not argue that intentions in people's brains can be seen, but that people's intentions can be recognised, and inferred from the utterances they use. Examples of such approaches are given in Allen (1983), Appelt (1981, 1985), Carberry (1989), Cohen et al. (1982), Hinkelman and Allen (1989), Hobbs (1979), Litman and Allen (1984), Schank and Abelson (1977), and Wilensky (1983). Ideas here emerge from claims made by philosophers of language such as Austin (1962) and Searle (1969). Such philosophers argue that the motivation for people to use language is to achieve their intentions.

Various approaches to modelling the coherence of dialogue shall now be investigated, and particularly the latter approach on the analysis of intentions. There has been much research conducted on developing theories, and designing computational models for understanding natural language discourse. Much of the work has been on developing models for text and speech processing. Some work has concentrated more specifically on processing natural language dialogue. The next three sections serve as an introduction to the area of discourse modelling and the various points of view on how it should be accomplished. Many of the approaches have common themes, while there are differing points of view between them. The one common theme of all the models is that of coherence of discourse. First, we will concentrate on semantic-based approaches to modelling discourse.

5 Discourse modelling through semantics

Semantic-based approaches argue that the coherence of a discourse can be analysed from the point of view of its meaning or semantics. Approaches centre on techniques for recognising and representing the semantics of the discourse. There are many methods for conducting semantic processing, but most techniques represent semantics using some form of meaning representation. Formal languages for the representation of meaning of natural language utterances include propositional logic, predicate calculus, semantic networks (see Quillian 1969), Conceptual Dependency (see Schank 1972, 1973, 1975, Schank and Abelson 1977), and Preference Semantics (see Fass 1988 and Wilks 1973, 1975a, 1975b, 1975c).

There has been much debate with respect to how syntactic and semantic processing within natural language systems should be integrated. Some argue for systems where syntactic analysis is conducted first, and is then followed by semantic analysis. Some argue that syntax analysis should be conducted hand-in-hand with semantic analysis. Others argue that semantic analysis should come first, with syntactic analysis subservient. The latter approach has been taken most notably by Schank and Riesbeck (see Schank 1973, and Schank and Riesbeck 1981). The semantics-first camp argue that people still understand and can build representations of ill-formed syntactic sentences like, "Paul Japan go to" with ease, even if such utterances are strange. Hence, they argue, syntactic analysis is not central to natural-language processing. We will discuss one example of semantic-based approaches to natural language discourse processing, that of Wilks. Detailed discussions of other examples of the semantics approach are found in Mc Kevitt et al. (1992a).

5.1 Wilks

Wilks (see Wilks 1973, Wilks 1975a, 1975b, 1975c) built a computer program for text processing which handled paragraph-length input. It used a system of preferential choice, or *Preference Semantics*, between deep semantic patterns based on *semantic density*. The program was applied to the problem of machine translation.

Wilks' technique uses a fragmentation routine to fragment text and then represent it in an interlingual structure called a *semantic block*. This block consists of *templates* bound together by *paraplates* and *common sense inferences*. Each of these three items consists of formulas, with predicates and functions ranging over the formulas and subformulas. The formulas, in turn, consist of *elements*. Semantic items represent text items as shown in Table 1.

Items in semantic representation	*Corresponding text items*
formula	English word sense
template	English clause or simple surface item
semantic block	English paragraph or text

Table 1: Wilks' Semantic representation of text items

The paraplates and common-sense inferences combine to cause bound templates in the semantic block. Semantic elements exist as primitives out of which all the above complex items are made up.

Elements are primitive semantic units used to express the semantic entities, states, qualities, and actions about which humans speak and write. Examples are (a) entities: MAN(human being), (b) actions: FORCE(compels), and (c) cases: TO(direction). In addition to these primitive elements, there are *class* elements whose names begin with an asterisk. Examples are *ANI, for the class of animate elements, MAN, BEAST, and FOLK. *HUM is the class of human elements MAN and FOLK.

Semantic formulas are constructed from elements. They express the sense of English words. One formula exists for each sense. The most important element is always the rightmost, which is called the *head*, and it expresses the most general category under which the word sense in question falls. The formula for the action sense of *drink* is shown below, and the *cases* and *values* of the subformula shown in Table 2.

"drink" (action) –>

((*ANI SUBJ) (((FLOW STUFF) OBJE)
liquidity
((SELF IN) (((*ANI THRU PART)) TO)
aperture
(BE CAUSE)))))

Subformula	*Case(Act)*	*Value*
(*ANI SUBJ)	SUBJ	*ANI
((FLOW STUFF) OBJE)	OBJE	(FLOW stuff)
(SELF IN)	IN	SELF
(((*ANI (THRU PART)) TO)	TO	(*ANI (THRU PART))
(BE CAUSE)	CAUSE	BE

Table 2: *Cases* and *Values* for Wilks' subformula representation

The explanation of the subformula is that:

- the preferred agent is animate
- the preferred object is liquid
- the container is the self, the subject
- the direction of the action is a human aperture (i.e. the mouth)
- the action is of causing to be (somewhere else)

Drink is an action preferably done by animate things to liquids, or flowing substances, causing the liquid to be in the animate thing, and via a particular aperture of the animate thing, the mouth.

Templates are networks of formulas giving connectivity between agent, action, and object formulas. Each clause, phrase or primitive sentence of text (called fragments) is seen by a program as *senses* or strings of formulas drawn from a dictionary. There is one formula for each text word.

Bare templates exist as sequences of agent, action, and object heads into which semantic formulas can slot. If there exists a series of formulas whose heads are identical to a bare template of elements then the sequence of formulas is a template for that fragment together with any other formulas that depend on the main three. The program tries to locate one, or more, templates in each string of formulas by looking at head elements and seeking proper sequences of heads. It is an hypothesis of the work that an inventory of bare templates can be built up for the analysis of ordinary language.

Wilks later introduced what he called *Pseudo-Texts* into his system (see Wilks 1975c). He points out that sentences which break preferences are the norm in language use, and not the exception. He notes that in the sentence, "My car drinks gasoline," none of the senses for car are animate, and so the system simply accepts what it is given. However, this is not enough as far as

[(MAN)	inject	liquid	]
[(MAN)	(USE)	#tube	] insertion of fuel
[ICengine	(USE)	#liquid	] petrol moving engine
[*	(MOVE)	[]	] moves the car

Figure 1: A Pseudo-Text for car

understanding the utterance goes. He discovered the need for representations of objects such as "car". The program should have access to a sufficiently rich knowledge structure for "car" and be able to notice that cars consume particular types of fluid, that is of the same semantic structure as the relation in which a drinker normally stands to a liquid to be drunk.

Pseudo-Texts are structures of factual and functional information about a concept, or item, and are similar to frames in the sense of Minsky (see Minsky 1975), Charniak (1977), and Schank (see Schank and Abelson 1977). An example of a Pseudo-Text for car is given in Figure 1. A Pseudo-Text is a semantic block consisting of a linear sequence of template forms tied by case ties. The linear order of the templates is taken to indicate normal time sequence. The program accesses the Pseudo-Text for car and seeks the template in it with the closest match to the source form template. Wilks uses a notion of *projection* to describe the process of placing *use* in place of *drink* in the formula for drink.

In summary, Wilks' theory of Preference Semantics is one which assumes that the coherence of a discourse can be computed bottom up from the semantics of individual sentences. The coherence is computed in terms of the preferences of templates for various formulas. There is no notion of determining structural properties of a discourse, or of processing intention within Wilks' theory.

6 Modelling discourse from structure

Structure-based approaches to discourse modelling tackle discourse coherence from the point of view of discourse structure. These approaches argue that there are structures in discourse, whether explicit, or implicit, and those structures can be recognised and represented. Much of the work centres on strategies for recognising and representing the *topic* and *focus* of the discourse. Examples of structure-based approaches are described in Alshawi (1987), Dale (1988, 1989), Grosz (1983), Grosz et al. (1983), Grosz and Sidner (1986), Sidner (1983, 1985), and Webber (1978). There are also formal structure-based approaches to discourse understanding discussed in Barwise and Perry (1983), Heim (1982), and Kamp (1981). We will discuss the approach of Grosz and Sidner (1986) here. A detailed description of other approaches is given in Mc Kevitt et al. (1992a).

6.1 Grosz and Sidner

Grosz and Sidner (1986) propose a new theory of discourse structure. Structure consists of three subcomponents: (1) structure of sequences of utterances or *linguistic structure*, (2) structure of purposes or *intentional structure*, and (3) state of focus of attention or *attentional state*. It could be argued that Grosz and Sidner's approach should be taken as being general, but we do not do that because, although they stress the importance of intention just as much as structure, they represent intentional entities inside structural concepts. Linguistic structure is the natural aggregation of segments of discourse. Utterances are the actual saying or writing of particular sequences of phrases and clauses and are the linguistic structure's basic elements. Intentional structure captures discourse-relevant purposes expressed in the segments as well as relationships among them. Attentional state is focus of attention of participants as discourse unfolds. This records objects, propositional relations and intentions that are important at each point in the discourse.

They point out that the subcomponent distinction is important for explaining cue phrases, referring expressions and interruptions. Their basic elements of a computational theory of discourse structure simplify and expand upon previous work. The work involves the integration of two lines of research: (1) focusing in discourse (Grosz 1978a, 1978b, 1981), and (2) intention recognition in discourse (Allen 1983, Litman 1985, Pollack 1986, and Sidner 1983, 1985). However, the former structure-based research is predominant in their theory. Grosz and Sidner concentrate on generalising the task-oriented elements that constitute discourses. Discourses are task-oriented but the tasks are varied. Tasks are physical, mental and linguistic. By distinguishing subcomponents they simplify computational mechanisms of related work. They present an abstract model of discourse structure and state that they do not recognise the need for the specifics of a computer program, or for a psychological model.

Structure, intention and attention allow the Conversational Participant (CP) to determine how an individual utterance fits the rest of the discourse, i.e. why it was said, and what it means. Expectations of what is to come are also integrated in generation and interpretation. Abstraction of the focus of attention of participants summarises information from previous utterances and hence there is no need for a complete history of the discourse.

With regard to *linguistic structure* utterances are aggregated into discourse segments. The utterances in a segment are relevant to that segment, and the segment is relevant to the discourse. Two consecutive utterances may be in different, or the same segments. Nonconsecutive utterances may be in the same segment. Segmentation of discourses has been illustrated by Grosz (1978b) for task-oriented dialogues, Linde (1979) for arguments and Reichman-Adar (1984) in informal debates, explanation, and therapeutic discourse. There has been a lot of debate on segmentation boundaries. There are few psychological studies, although Mann (1975) has shown that when people segment dialogues they do so in similar ways.

Linguistic expressions are the primary indicators of discourse segment boundaries. Cues like "in the first place" are dividers that function to indicate these boundaries. Reichman (1985) calls these, *clue words* and Grosz and Sidner, *cue phrases*. Boundaries can indicate changes in (1) attention, or (2) intention. Grosz and Sidner point out that discourse segmentation is useful in defining referring expression boundaries. The constraints for intra-segment reference are different from inter-segment reference.

With regard to *intentional structure* participants have purposes in dialogue. Some of the purposes will distinguish coherent, from incoherent, discourse. They define foundational purpose as discourse purpose (DP) which is the intention that underlies the discourse engagement. Discourse Segment Purpose (DSP) is the intention of the segment. Typically, the Initiating Conversational Participant (ICP) will have a number of different kinds of intentions associated with the use of discourse. The recognition of DP or DSP is essential to achieving purpose. ICP's use discourse for impressing or teaching. Private intentions are involved here too. Here is an example of a type of intention that can serve as DP/DSP which is intended to have some agent perform some physical task:

> Intend that Ruth intend to fix the flat tire.

Grosz and Sidner define structural relations that play an important role in discourse structure. A dominance relation is defined as an action that satisfies one intention, e.g. DSP1 may be intended to provide part of the satisfaction of another, say e.g. DSP2.

DSP1 contributes to DSP2
DSP2 dominates DSP1

A dominance relation places a partial-ordering on DSPs called a *dominance hierarchy*. A second structural relation is *satisfaction-precedes* which defines the case where one DSP must be satisfied before another. For some discourses, including task-oriented ones, the order in which the DSPs are satisfied may be significant.

Attentional state is the abstraction of focus of attention and saliency at different points in the discourse. It is a property of the discourse and not the participants. It is inherently dynamic, recording objects, properties and relations that are salient at each point. The following structures are defined:

(1) Focus space: a set of these models attention.

(2) Transition rules: model changes in attention state. These specify conditions for adding/deleting focus spaces.

(3) Focusing structure: the collection of focus spaces available at any one time that are salient at each point.

(4) Focusing: the process of manipulating focus spaces.

The focusing process associates focus spaces with each discourse segment. The space contains entities that are salient because of (1) saliency in the segment, or (2) producing or comprehending utterances in the segment (see Grosz 1978b). Also the focus space includes the DSP. The fact that the focus space contains the DSP shows us that there is a commitment to structure/topic over intention. Focus spaces are maintained in a data structure called a *stack*.

Unlike Wilks' semantics-based approach described in the previous section, Grosz and Sidner do not concern themselves with meaning, but the role of DP/DSP in structure. We query Grosz and Sidner's use of the stack as a data structure representing the focus of attention of the discourse. A stack is a data structure where information is piled on top of previous information. The structure has the inherent quality that data first-in is last-out and data last-in is first-out. Terms like PUSH have been coined to describe placing data onto the top of the stack and POP for removing the data off the top of the stack. The motivation for the use of a stack as opposed to a more complex data structure such as a network is not elabourated upon in the work of Grosz and Sidner. If the stack is motivated as a mechanically convenient, and simple method of representing a discourse then there are no arguments against it. However, if the stack structure is motivated by more theoretical considerations then we must object. The principle of a stack would mean that they believe dialogue to be inherently structured like a stack. They would be claiming that people in conversation bring up topics, discuss them, and then leave the topics again just like one would PUSH data and POP data from the stack. Although, we have no access to empirical data which would argue against Grosz and Sidner's use of a stack we find it hard to believe that dialogue will always be so structured.

In summary, Grosz and Sidner's approach to modelling the coherence of discourse is motivated by modelling the topic and focus of the discourse. This is integrated with models of the intentions and linguistic structure of the discourse. The main mechanism of their computational model is a stack which stores information about context by PUSHing focus spaces onto the stack. However, this has an inherent principle of locality about it. Also, the determining factor for focus spaces is syntactic, with the use of conversational cues, with no regard for the semantic content of utterances.

7 Modelling discourse using intention

Intention-based approaches to discourse modelling argue that the coherence of a discourse can be modelled from the point of view of the intentions (or goals, plans, or beliefs) of discourse participants. The approaches centre on the recognition and representation of goals, plans or beliefs of participants in a discourse. The models developed here are based on the assumption that it is the intentions of the participants in a discourse that lend coherence to the discourse.

Examples of research incorporating intention analysis include Allen (1983), Appelt (1981, 1985), Carberry (1989), Cohen et al. (1982), Hinkelman and Allen (1989), Hobbs (1979), Litman and Allen (1984), Schank and Abelson (1977), and Wilensky (1983). There are also plan-based approaches to text processing such as those by Alterman (1985) and Lehnert (1978, 1982). While detailed discussions of all of these are given in Mc Kevitt et al. (1992a) we discuss here the approach of Cohen et al. (1982). First, we start with a discussion of some philosophical work on intention.

7.1 Philosophical views of intention

Speech acts, or communicative acts, are a means of describing intention in discourse or dialogue. Austin (1962) studied a class of natural language utterances which he called *performatives* that do not refer to states in the world, but to acts such as *promising*, *threatening* and *naming*. Austin argued that it does not make sense to take the traditional approach of understanding utterances like, "I name this ship the 'Beer Engine"', in terms of truth and falsity, but to consider whether they are *felicitious*, or whether they are appropriate in the context where they are used.

Austin defined an *illocutionary* act as the act of something being done by performing the *locutionary* act of uttering. An illocutionary act may be questioning, commanding or describing. Also, a speaker will, through performing a given illocutionary act, perform other actions. These actions which can be intended, or unintended, are called *perlocutionary* acts. If a person issues a warning, then the purpose will probably be to warn the hearer. Likewise, if an utterance is made, issuing a request for a confirmation, the effect may be that the hearer will give that confirmation.

An utterance may contain linguistic expressions which serve to indicate the illocutionary force of the utterance. In (1a) below, "warn" is a *warning* that something has gone wrong, and in (1b) "order" indicates an *order* for something to be done.

(1) a I warn you, be careful.
b I order you to take charge.

Warning and *ordering* are caused by uttering specific words. Austin termed verbs like these *performatives*. Only certain verbs act as performatives. For example, in (2a) and (2b) nothing is performed by the verbs therein.

(2) a I know what I know.
b I love all things sweet.

The terms *speech act* or *communicative act* have become common in the literature as alternative terms for *illocutionary act*. Traditionally, however, speech act and communicative act cover a wider range of phenomena than illocutionary act. *Explaining, intimating, confirming,* and *predicating* are all examples of speech acts.

Searle (1969) extended Austin's work by formalising the structure of felicity conditions associated with a variety of speech acts. He classified all speech acts as incorporating one of five *illocutionary points*: Assertives, Directives, Commissives, Expressives, and Declarations. The important distinction of Searle's work is that he covered *all* utterances, and not just those having explicit performative verbs such as, *promise*, or *declare*. For example, the statement, "I'll be there," can be a promise event though it's in the form of a statement. There is an enormous amount of research in the area of illocutionary force, and Gazdar (1979) provides a useful bibliography.

In discourse, people tend to follow certain pragmatic rules of behaviour of exchange. The philosopher of language Grice (1975), proposed that conversational exchange and rational behaviour are governed by what he called the *Cooperative Principle* [7]:

> Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged. — Grice (1975, p. 45).

Grice described four categories, or special cases, of this principle, and gave examples of language and non-language applications. Grice called the four categories *maxims*. He pointed out that as long as participants in a mutual conversation assume that each other is adhering to the Cooperative Principle, meanings can follow as inferences from some maxim not being violated. The meanings passed from speaker to speaker do not have to expressed explicitly. The maxims are as follows:

I Quantity: (1) Make your contribution as informative as is required (for the current purposes of the exchange).

(2) Do not make your contribution more informative than is required.

II Quality: Try to make your contribution one that is true.

(1) Do not say what you believe to be false.

(2) Do not say that for which you lack adequate evidence.

III Relation: Be relevant.

IV Manner: Be Perspicuous.

(1) Avoid obscurity of expression.

(2) Avoid ambiguity.

(3) Be brief (avoid unnecessary prolixity).

(4) Be orderly.

Grice pointed out that the maxim of Quality was much more important than other maxims. Violating *Quality* is a moral offence, whereas violating the others is, at most, inconsiderate or rude. It can be possible that situations arise where the speaker cannot conform to all the maxims at once. It may not be possible to say as much as is necessary in some situation, without saying things with no adequate evidence.

[7] It is this principle, which enables us to explain the irrationality of the examples on restaurants, and bars, described earlier.

A speaker may make an utterance in discourse which disregards a maxim (e.g. Relation) even when the speaker wishes to convey a relevant statement. In situations where speakers do not obey the maxims, hearers are still prepared to believe that speakers are abiding by the Cooperative Principle. Otherwise, the hearer must assume that the speaker is behaving irrationally. Many people have misunderstood what Grice meant by his maxims. He did not strive to say that people do, or should, use discourse to conform to the maxims, but that assuming the maxims exist, we can explain utterances that seem illogical or irrational. This assumption demonstrates that people cooperate in dialogue as much as possible.

Some examples of the use of the conversational maxims are in conversations like the following shown in (3). In (3b) the speaker uses rhetoric to show that what speaker1 says is of course true. On the surface (3b) does not seem relevant, and therefore would break the maxim of relation. However, deeper insight through analysis of the fact that (3b) is a question which is obviously true, shows that (3b) is meant to convey some obvious rhetorical fact that renders (3a) obviously true.

(3) a Speaker1: Is Margaret Thatcher in favour of military action?
b Speaker2: Is the Pope catholic?

The use of the Cooperative Principle can lead to irony and humour in discourse. Consider the dialogue in (4)[8]. In (4a) an air-hostess, or steward, Julie is walking through the cabin asking the passengers if they would like some tea. However, in (4b) the senior flight attendant indicates that Julie should/could be more polite about how she is doing the asking. This is indicated by speech intonation and stress. The senior attendant is using the Cooperative Principle to indicate her intention indirectly. In (4c) Julie turns the tables, pretends to take the senior attendant literally, and tells her that, yes, she would like a cup of tea.

(4) [The context here is an airplane where a Flight Attendant (Steward) is asking passengers whether they would like some tea as she walks through the cabin]
a Flight Attendant (Julie): Tea, Tea, Tea....
b Senior Flight Attendant: Julie, "Would you like some tea?"
c Flight Attendant (Julie): Yes, please.

As speakers expect hearers to accept the Cooperative Principle, the speaker can exploit it, and speak in a way that his/her behaviour must be interpreted according to the Principle. If the speaker's utterance seems irrelevant, the hearer will try to construct an inference sequence

[8]This example is taken from a documentary about the life of Flight Attendants (Stewards) shown on British Television (Channel 4) during April 1991.

which will force the utterance to be relevant, or cooperative. The exploitation of the maxims is a technique by which utterances can be used to convey more than is literally indicated. Grice termed this phenomenon, *implicature*.

Following on from the work of Grice, it is possible to assume that if people are acting in accordance with the Cooperative Principle then that provides a basis for their dialogue being coherent. Intention-based approaches to discourse processing are, in a large part, rooted in the philosophical works of Austin, Searle and Grice. Let us move on now to discuss an example theory of natural-language dialogue processing which utilises intention.

7.2 Cohen, Perrault and Allen

Cohen et al. (1982) view natural language question-answering as a special case of natural language dialogue, and they develop a theory of how the user's intent can be inferred. They observe that usually people do not say what they mean, and people expect hearers to infer intentions that they are meant to infer. On analysis of dialogues with a real natural-language system called PLANES they noted that (1) users expect the system to infer and respond to their apparent but unstated goals, (2) users' subsequent utterances depend upon the system's responses, and (3) the system should always keep track of what is communicated. They point out that users (1) expect more than just answers to isolated questions, (2) expect to engage in a conversation with a system where coherence is manifested in the interdependence of their unstated plans and goals, (3) expect that the system incorporate its own responses into its analysis of their utterances, (4) maintain these expectations even when it is obvious that the system is not a competent conversationalist, and (5) expect the system to react to unstated goals. Also, users expect the system to use indirect speech acts, and users make complicated requests using several utterances. Users expect the system to be aware of their reference failures, the failures of their presuppositions, and ensure that they are not misled. The system should expect that the users are dependent on it.

They argue that (1) intentions are identified with plans, and (2) that utterances are planned by speakers to achieve effects on hearers. To solve these problems Cohen et al. argue that the system must (1) engage in reasoning about how an utterance is being used (user's intentions), decide (2) what communication actions it should perform, and (3) how they should be performed. They point out that there is no direct mapping between utterance form and the action it is being used to perform. Also, they make the assumption that the speaker is a rational agent. There are two processes a system must embody to provide a basis for solutions to intention-discrimination problems: (1) Plan Construction, and (2) Plan Recognition.

Plan construction is a process where an agent can examine the consequences of sequences of future actions before executing them. Plan recognition is the observation of actions performed by an agent and the predicting of subsequent actions. Inferring a plan which the user may be following is Plan Recognition and involves beliefs about (1) the agents' beliefs, (2) conditions that are true at the end of an action, and (3) likely plans and goals of the agent.

They introduce the notion of Intended Plan Recognition which is where the system cannot simply infer and act upon what the user wants but must act/infer on what the user wants the system to "think" that the system wants. Intended Plan Recognition is conducted by the following rules:

SBUW(Do S A):	System believes that the user wants it to do A.
SBUW(SBUW(Do S A)):	System believes that the user wants it to think he/she wants it to do A.
SBUW(SBUW(A)) –> SBUW(SBUW(B)):	Inferring other goals such as B from A.

A means of controlling Intended Plan Recognition is based on (1) a heuristic that terminates inference chains that lead to mutually exclusive alternatives, and (2) assuming the speaker is a rational agent.

Indirect speech act recognition is the discovery of goals inferred during Intended Plan Recognition. Intended Plan Recognition produces goals the system is supposed to attribute to the user. For imperatives the hearer will believe that the speaker wants the hearer to do some act A. The system believes that the user wants it to do A. Plan recognition is the process of inferring the plan an agent may be following and comes from:

(1) Observing an utterance of a sentence.

(2) Assuming the agent wanted to do it.

(3) Inferring the agent wanted the typical act effect.

(4) Characterising the effects of uttering of sentences to be hearer beliefs about what the speaker wants.

For example, to demonstrate plan recognition we can take the utterance, "Do you know where the Enterprise is?" (1) From the syntax and semantics of the utterance the system recognises that the user intends it to believe that the user wants to know whether the system knows where the Enterprise is. (2) The system can infer that the user wants to know whether the system knows where the Enterprise is, then adopt the user's knowing whether the system knows as a goal. (3) then the system can satisfy the goal by telling the user whether it knows or not.

Therefore, a system should (1) reason about others' actions, (2) plan utterances to achieve specific communicative goals from beliefs and intentions of the user, and (3) recognise the user's utterances as part of larger plans over several utterances. The recognition of user's utterances as intentions within sequences is a major goal of the work described in the next section.

Cohen, Perrault and Allen discuss two implemented prototype systems: (1) a simple question-answering framework, and (2) a decision-support framework. The former system is a simple program for understanding, acting as an information clerk at a train station, and which expects users to want to board, or meet, trains. The latter system engages in dialogues about a display screen. Each system has the ability to (1) distinguish beliefs and wants of the user from its own, and (2) to process indirect speech acts. We shall discuss the question-answering system here.

The question-answering system acts as an information clerk at a train station. The system expects users to want to board or meet trains. In the following dialogue fragment the system answers (1a) literally, but also provides gate information. The system deduces that the user does not know the gate information but needs it in order to achieve his/her goal(s).

(1) a U: When does the Montreal train leave?
b S: 3:15 at gate 7.

The system can also infer user intentions based on sentence information. For example, to give the reply in (2b) below the system uses its expectations to infer that the user has a goal of boarding the 3:15 train to Windsor, and that the user also needs the gate information in order to do so.

(2) a U: The 3:15 train to Windsor?
b S: Gate 10.

In the following dialogue the system must generate a question to disambiguate trains *to* Windsor and trains *from* Windsor.

(3) a U: When is the Windsor train?
b S: Do you want to go to Windsor?
c U: Yes.
d S: 3:15.

The system has the capability of analysing a wide range of requests including conventional forms such as the following:

(4) Do you know when the Windsor train leaves?
I want to know when.....
Tell me when...
Can you tell me when....

The system can aso handle the following nonconventional forms:

(5) John asked me to ask you when the next train to Windsor leaves.
John wants to know when the next Windsor train leaves.

The program which processes these examples consists of the following components: (1) a parser using syntactic and semantic information to produce a literal interpretation of the input, (2) a plan-recognition component having expectations about goals, and when given a set of these, and observed actions from (1), it infers a plan, (3) an obstacle-detection component which analyses a plan for steps that the user cannot perform easily, and (4) a plan construction component, which given a goal, plans a course of action. Plan recognition was effected by search through a space of pairs of plan fragments.

In summary, Cohen, Perrault and Allen's work argues that the coherence of a discourse can be modelled from the point of view of recognition of the intentions of the participants in the discourse. An important feature of their work is the principle that intention sequencing is a fundamental contributor to discourse coherence.

Now that we have discussed all three major approaches to discourse processing we briefly discuss a theory of intention analysis which has been incorporated within a computational model.

8 A theory of intention analysis

A theory of intention analysis (see Mc Kevitt 1991b) has been proposed as a model, in part, of the coherence of natural-language dialogue. A central principle of the theory is that coherence of natural-language dialogue can be modelled by analysing sequences of intention. The theory has been incorporated within a computational model in the form of a computer program called the Operating System CONsultant (OSCON) (see Guthrie et al. 1989, Mc Kevitt 1986, 1991a, 1991b, Mc Kevitt and Wilks 1987, and Mc Kevitt et al. 1992b, 1992c, 1992d). OSCON, which is written in Quintus Prolog, understands, and answers in English, English queries about computer operating systems.

The computational model has the ability to analyse sequences of intention. The *analysis of intention* has at least two properties: (1) that it is possible to recognise intention, and (2) that it is possible to represent intention. The syntax, semantics and pragmatics of natural-language utterances can be used for intention recognition. Intention sequences in natural-language dialogue can be represented by what we call *intention graphs*. Intention graphs represent frequencies of occurrence of intention pairs in a given natural-language dialogue. An ordering of intentions based on *satisfaction* exists, and when used in conjunction with intention sequences, indicates the *local*[9] and *global* degree of expertise of a speaker in a dialogue.

[9]By *local* expertise we wish to stress the fact that sometimes experts can act as novices on areas of a domain which they do not know well.

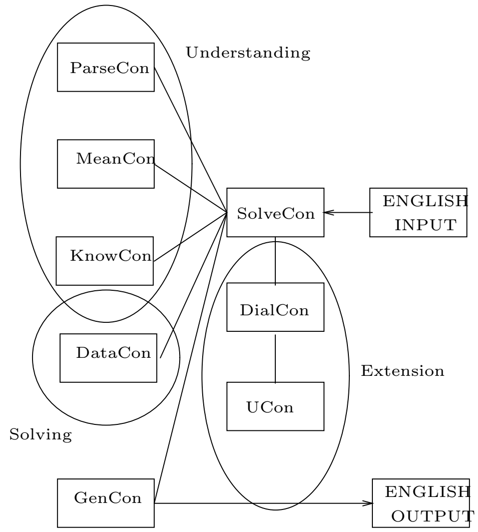

Figure 2: Architecture of the Operating System CONsultant (OSCON) system

The architecture of the OSCON system consists of six basic modules and two extension modules. There are at least two arguments for modularising any system: (1) it is much easier to update the system at any point, and (2) it is easier to map the system over to another domain. The six basic modules in OSCON are as follows: (1) ParseCon: natural-language syntactic grammar parser which detects query-type, (2) MeanCon: a natural-language semantic grammar (see Brown et al. 1975, and Burton 1976) which determines query meaning, (3) KnowCon: a knowledge representation, containing information on natural-language verbs, for understanding, (4) DataCon: a knowledge representation for containing information about operating system commands, (5) SolveCon: a solver for resolving query representations against knowledge base representations, and (6) GenCon: a natural-language generator for generating answers in English. These six modules are satisfactory if user queries are treated independently, or in a context-free manner. However, the following two extension modules are necessary for dialogue-modelling and user-modelling: (1) DialCon: a dialogue modelling component which uses an intention matrix to track intention sequences in a dialogue, and (2) UCon: a user-modeller which computes levels of user-satisfaction from the intention matrix and provides information for both context-sensitive and user-sensitive natural-language generation. A diagram of OSCON's architecture is shown in Figure 2.

ParseCon consists of a set of Prolog predicates which read natural-language input and determine the type of query being asked, or intention type presented, by the user. For each type of query there are tests for characteristic ways that people might utter that query. ParseCon uses a semantic grammar, in the Definite Clause Grammar (DCG) [10] formalism of Prolog.

MeanCon consists of predicates which check queries for important information. There are predicates which check for mentioned (1) command names (e.g. "ls", "more"), (2) command-

[10] Definite Clause Grammars (DCG's) were first developed by Pereira and Warren (1980) as a tool to be used in Prolog for natural-language processing.

effect specifications (e.g "see a file"), and (3) concepts, or objects (e.g. "file", "directory"). In case (2) there are specific types of information searched for: (1) **verb** specifying action (e.g. "see", "remove"), (2) **object** of action (e.g. "file"), (3) **modifier** of object (e.g. "contents"), and (4) **location** of object (e.g. "screen"). MeanCon also checks for option verbs (e.g "number") and option verb objects (e.g. "lines"). MeanCon contains a dictionary of English words that define categories such as "person", "modifier", "article", "quantifier" and "prepositions".

KnowCon consists of a set of data files to represent knowledge about the domain language used for understanding English queries. Data files here contain information about English verbs which denote types of command or action. Examples of categories of action are: (1) creating, (2) screenlisting, (3) printerlisting, (4) sending, (5) transferring, and (6) removing. KnowCon also contains grammar rules for operating system objects like "date", "file" and "directory". The grammar rules encode characteristic ways in which people talk about the objects in English.

DataCon consists of a set of data files defining detailed information about operating system commands. This information is stored for the UNIX and MS-DOS Operating Systems. The data for UNIX is split among seven files about commands: (1) preconditions, (2) effects, (3) syntax, (4) names, (5) precondition options, (6) effect options, and (7) name options. The first four files contain basic data about commands while the last three contain data for options. For MS-DOS, data is contained in just four files which are similar, in spirit, to the first four here.

SolveCon is a solver which constructs and matches representations of user queries (called *Formal Queries*) against the knowledge base, DataCon, and produces an instantiated Formal Query which serves as an answer for the query. SolveCon is the heart, or driver, of the OSCON program because it contains the information for mapping English sentences into instantiated formal queries. It contains a set of complex rules which call other OSCON modules to determine (1) query type, (2) intention type, and (3) the instantiated Formal Query for that query. The determination of intention type is a two stage process where natural-language queries are first mapped into query types, and then into intention types. SolveCon also checks for repetitions by comparing the propositional content, or topic, of the current intention against that of the previous.

GenCon is the natural-language generator for OSCON and maps instantiated information from SolveCon into English answers. Here, there are algorithms for printing out (1) preconditions, (2) effects (or postconditions), and (3) syntax of commands. Also, there are predicates for printing out examples of the use of commands and command compositions. The type of query asked by the user determines the information presented in English to the user.

DialCon is the dialogue modeller for OSCON which updates the *intention matrix* representing intention pair frequencies in the dialogue. Matrix update is conducted by locating the relevant cell in the matrix which needs to be updated, and increases its count by 1. DialCon indexes the cell in the matrix by pairing the current intention type with the previous.

UCon is the user-modelling component of OSCON. UCon derives a binary measure of user expertise, *expert* and *novice*. UCon applies a user-modelling function to the intention matrix to determine levels of user *satisfaction* and *dissatisfaction*. Initially, the user is assumed to be an expert. Subsequent changes in the levels of satisfaction and dissatisfaction will result in changes in the level of user expertise. Such information is used by GenCon to generate context-sensitive and user-sensitive natural-language responses. In Appendix A we see how the system modifies its responses based on the level of user satisfaction. A detailed analysis of how the system can modify its natural language responses is given elsewhere (see Mc Kevitt 1991b). We will not discuss details of processing within components of the OSCON system. These can be found in Mc Kevitt (1991b).

The OSCON system is used to test the what we call the *Intention-Computer* hypothesis: that the analysis of intention facilitates effective natural-language dialogue between different types of people and a computer. OSCON provides positive evidence for this hypothesis (see Mc Kevitt 1991b, Mc Kevitt et al. 1992b, 1992c, 1992d).

9 Conclusion

That completes our discussion of the various approaches to processing natural language discourse. There are three main approaches to processing discourse: modelling semantics, structure or intention. The different approaches, emphasise some of these phenomena more than others. However, a central theme of much, if not all, of the work is that natural-language discourse processing can be conducted by modelling the coherence of a discourse. The motivation for a concentration on discourse coherence is the fact that rational speakers usually use coherent discourse and that essential information will otherwise be lost.

Semantic-based approaches concentrate on recognising and representing the meaning of a discourse. This is completed by representing the semantics of individual utterances in the discourse and linking these representations together. The semantics-based approach attempts to infer relationships between utterances, a process called inferencing, so that implicit links between utterances can be discovered. Such inferencing attempts to maximise the coherence of the discourse. Semantics-based approaches argue that by modelling the semantics of a discourse other information will fall out of that.

Structure-based approaches model the discourse from the point of recognising and representing explicit and implicit structures in the discourse. A number of different names are given to entities which represent discourse structure. Many of the theories define explicit spaces which represent implicit spaces in the discourse. Much of the work involves recognising and representing these spaces and the relationships between them. Various names have been given to such spaces. Grosz and Sidner (1986) use the terms *topic* and *focus*, while Reichman (1985) uses the term *context factor*. Alshawi (1987) uses *focus spaces* for representing information in discourse. All of the approaches argue that discourse structure can be recognised, in part, by syntactic markers called *conversational cues* or *clue-words*. Although many of the structure-based approaches mention other elements of the discourse such as semantics and intention they consider them secondary to structure. For example, Grosz and Sidner and Reichman consider intention as being subordinate to topic.

Intention-based approaches model the discourse from the point of view of the intentions (or goals, plans or beliefs) of the participants in the discourse. Intention-based approaches are rooted in the philosophical works of Austin (1962), Searle (1969), and Grice (1957, 1969, 1971, 1975). The coherence of a discourse is determined by the coherence of the intentions of the participants. Much of the work here involves recognising and representing the goals, plans or beliefs of the speaker and the relationships between them. Again, there are a number of names for structures representing intention. Pustejovsky claims to provide an integrated theory and computational model of discourse processing, which models semantics, structure and intention. This general approach would seem to be the best solution for modelling discourse. Future work will involve incorporating aspects of all theories into computational models and then testing such models.

Finally, we have described briefly how a theory of intention analysis which models, in part, the coherence of natural-language discourse, is incorporated within a computational model called OSCON. OSCON has the capability of recognising and representing intentions in natural language dialogue. A model of the user can be constructed from intention sequences in a dialogue. In turn, the user-model is used to dynamically modify OSCON's English answers in dialogue. OSCON provides evidence for the Intention-Computer hypothesis that the analysis of intention facilitates effective natural-language dialogue between different types of people and a computer.

There are a large number of open questions in discourse processing and we see the integration and testing of different approaches a fruitful path to follow. Work on natural-language discourse processing will not only enable computers to better understand people, but also people to better understand computers.

10 Acknowledgements

We would like to thank David Benyon and Antony Galton and an anonymous referee for comments on this research.

11 References

Allen, James F. (1983) Recognising intentions from natural language utterances. In *Computational Models of Discourse*, M. Brady and R.C. Berwick (Eds.), 107-166. Cambridge, MA: MIT Press.

Alshawi, Hiyan (1987) *Memory and context for language interpretation.* Cambridge, United Kingdom: Cambridge University Press.

Alterman, R. (1985) A dictionary based on concept coherence. In *Artificial Intelligence*, 25, 153-186.

Andersen, S. and B.M. Slator (1990) Requiem for a theory: the 'story grammar' story. In *Journal of Experimental and Theoretical Artificial Intelligence (JETAI)*, 2,3, 253-275, July-September.

Appelt, Douglas E. (1981) *Planning natural-language utterances to satisfy multiple goals.* Technical Note 259, SRI International, 333 Ravenswood Avenue, Menlo Park, California 94025, USA.

Appelt, Douglas E. (1985) *Planning English sentences*. Cambridge, United Kingdom: Cambridge University Press.

Austin, J.L. (1962) *How to do things with words (Second Edition)*. J.O. Urmson and Marina Sbisà (Eds.). Cambridge, MA: Harvard University Press.

Ballim, Afzal and Yorick Wilks (1990) Stereotypical belief and dynamic agent modeling. In *Proceedings of the Second International Conference on User Modelling*, University of Hawaii at Manoa, Honolulu, Hawaii, USA.

Ballim, Afzal and Yorick Wilks (1991) *Artificial Believers*. Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Barnden, John A. and Jordan B. Pollack (1991) (Eds.) *Advances in connectionist and neural computation theory, Vol. I: high level connectionist models*. Norwood, N.J.: Ablex Publishing Corporation.

Barnden, John A., Sylvia Candelaria de Ram, David Farwell, Louise Guthrie, Stephen Helmreich, Paul Mc Kevitt, and Yorick Wilks (1991) The relevance of beliefs to natural language communication. In *Intercultural Communication Studies*, 1(1), 237-272. Also as, "The need for belief modelling in natural language processing," In *Proceedings of the International Conference on Cross-Cultural Communication (ICC-CC-89)*, Trinity University, San Antonio, Texas, USA, March, 1989.

Barwise, J. and J. Perry (1983) *Situations and attitudes*. Cambridge, M.A.: Bradford Books, MIT Press.

Bar-Hillel, Y. (1954) Indexical expressions. In *Mind*, 63, 359-379.

Benyon, David (1984) MONITOR – a self-adaptive user interface. In *Proceedings of Interact '84, First IFIP Conference on Human-Computer Interaction, B. Shackel (Ed.)*, Amsterdam: Elsevier/North-Holland.

Benyon, David (1987) User models –what's the purpose?. In *Proceedings of the Second Intelligent Interface Meeting*, M. Cooper and D. Dodson (Eds.). Alvey Knowledge-Based Systems Club, Intelligent Interfaces Special Interest Group, 28-29 May, London. London: The Alvey Directorate.

Benyon, David and Dianne Murray (1988) Experience with adaptive interfaces. In *The Computer Journal*, Vol. 31, No. 5, 465-474.

Brown, John Seely, Richard R. Burton and Alan G. Bell (1975) SOPHIE: a step towards creating a reactive learning environment. In *International Journal of Man-Machine Studies*, 7, 675-696.

Burton, R. (1976) *Semantic grammar: An engineering technique for constructing natural-language understanding systems*. BBN Report No. 3453, Bolt, Beranek and Newman, Cambridge, MA, USA.

Carberry, Sandra (1989) A pragmatics-based approach to ellipsis resolution. In *Computational Linguistics*, Vol. 15(2), June, 75-96.

Charniak, Eugene (1977) A framed painting: the representation of a commonsense knowledge fragment. In *Cognitive Science*, 1(4).

Chin, David (1988) Exploiting user expertise in answer expression. In *Proceedings of the Seventh National American Conference on Artificial Intelligence (AAAI-88)*, Vol. 2, 756-760, St. Paul, Minnesota, USA, August.

Cohen, Phil R., C. Raymond Perrault and James F. Allen (1982) Beyond question answering. In *Strategies for natural language processing*, Wendy G. Lehnert and Martin H. Ringle (Eds.), 245-274. Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Dale, Robert (1988) The generation of subsequent referring expressions in structured discourses. In *Advances in Natural Language Generation, An interdisciplinary perspective*, Communication in Artificial Intelligence Series, Michale Zock and Sabah, Gérard (Eds.), 58-75. London: Pinter Publishers.

Dale, Robert (1989) Cooking up reference expressions. In *Proceedings of the 27th Annual Meeting of the Association for Computational Linguistics, 68-75*, University of British Columbia, Vancouver, British Columbia, Canada, June.

Fass, Dan (1988) *Collative semantics: a semantics for natural-language processing*. Ph.D. Dissertation, Department of Computer Science, University of Essex. Also, as Memorandum in Computer and Cognitive Science, MCCS-88-118, Computing Research Laboratory, Dept. 3CRL, Box 30001, New Mexico State University, Las Cruces, NM 88003-0001, USA.

Galton, Antony (1984) *The logic of aspect: an axiomatic approach*. Oxford: Claredon Press.

Gazdar, Gerald (1979) *Pragmatics, implicature, presupposition, and logical form*. New York: Academic Press.

Grice, H.P. (1957) Meaning. In *Philosophical Review*, LXVI, No. 3, 377-388, October. Also in, *Semantics*, D. Steinberg and L. Jakobvits (Eds.), New York: Cambridge University Press, 1971.

Grice, H.P. (1969) Utterer's meaning and intentions. In *Philosophical Review*, LXXVIII, No. 2, 147-177, April.

Grice, H.P. (1971) Utterer's meaning, sentence-meaning, and word-meaning. In *Oxford readings in philosophy*, J.R. Searle (Ed.), 54-70. Oxford, United Kingdom: Oxford University Press.

Grice, H.P. (1975) Logic and conversation. In *Syntax and Semantics, Vol. 3, Speech Acts*, P. Cole and J.L. Morgan (Eds.), 41-58. New York: Academic Press.

Grosz, Barbara J. (1978a) Focusing in dialog. In *Proceedings of the Second Workshop on Theoretical Issues in Natural Language Processing (TINLAP-2)*, 96-103, University of Illinois, Urbana-Champaign, July. Also in, *Proceedings of the 1978 Meeting of the Association for Computational Linguistics, American Journal of Computational Linguistics*, 3, 1978.

Grosz, Barbara J. (1978b) Understanding spoken language. In *Discourse Analysis*, Walker, D. (Ed.), Chapter IX, 235-268. New York: Elsevier/North-Holland.

Grosz, Barbara J. (1981) Focusing and description in natural language dialogues. In *Elements of discourse understanding*, A. Joshi, B. Webber and I. Sag (Eds.), 84-105. Cambridge: Cambridge University Press. Also as, Technical Note 185, Artificial Intelligence Center, SRI International, Menlo Park, California, USA, July 1978.

Grosz, Barbara J. (1983) TEAM, a transportable natural language interface system. In *Proceedings of the 1983 Conference on Applied Natural Language Processing*, 39-45, Santa Monica, California, USA, February.

Grosz, Barbara, Aravind Joshi and S. Weinstein (1983) Providing a unified account of definite noun phrases in discourse. In *Proceedings of the 21st Annual Meeting of the Association for Computational Linguistics*, 44-50, Cambridge, Massachusetts, USA. Also in, *Proceedings of the International Joint Conference on Artificial Intelligence*, Cambridge, MA.

Grosz, B.J. and C.L. Sidner (1986) Attention, intentions and the structure of discourse. In *Computational Linguistics*, Vol. 12, 3, 175-204.

Guthrie, Louise, Paul Mc Kevitt and Yorick Wilks (1989) OSCON: An operating system consultant. In *Proceedings of the Fourth Annual Rocky Mountain Conference on Artificial Intelligence (RMCAI-89), Subtitled, 'Augmenting Human Intellect By Computer'*, 103-113, Registry Hotel, Denver, Colorado, USA, June.

Halliday, M.A.K. and R. Hasan (1976) *Cohesion in English*. London: Longman.

Hinkelman, Elizabeth and James Allen (1989) Two constraints on speech act ambiguity. In *Proceedings of the 27th Annual Meeting of the Association for Computational Linguistics*, 212-219, University of British Columbia, Vancouver, British Columbia, Canada, June.

Hobbs, Jerry (1979) Coherence and coreference. In *Cognitive Science*, 3, 1: 67-90.

Joyce, James (1939) *Finnegans wake*. London: Faber and Faber.

Kamp, H. (1981) A theory of truth and semantic representation. In *Formal methods in the study of language, Part I*, J. Groenendijik, Th. Janssen, and M. Stokoff (Eds.), 277-322. Amsterdam, The Netherlands: MC TRACT 135.

Heim, I. (1982) *The semantics of definite and indefinite noun phrases*. Unpublished Ph.D. dissertation, University of Massachusetts, Amherst, Amherst, Massachusetts, USA.

Kobsa, Alfred and Wolfgang Wahlster (1988) *User models in dialog systems*. Berlin: Springer-Verlag.

Linde, C. (1979) Focus of attention and the choice of pronouns in discourse. In *Syntax and Semantics, Vol. 12, Discourse and Syntax*, Givon, T., (Ed.). New York, New York: Academic Press.

Lehnert, Wendy (1978) *The process of question answering*. Hillsdale, N.J.: Lawrence Erlbaum Associates.

Lehnert, Wendy (1982) Plot Units: a narrative summarization strategy. In *Strategies for natural language processing*, W.G. Lehnert and M.H. Ringle (Eds.). Hillsdale, N.J.: Lawrence Erlbaum Associates.

Litman, Diane (1985) *Plan recognition and discourse analysis: an integrated approach for understanding dialogues*. Ph.D. Dissertation and Technical Report TR-170, University of Rochester, Rochester, New York, USA.

Litman, Diane and James Allen (1984) A plan recognition model for clarification subdialogues. In *Proceedings of the COLING-84*, 302-311.

Mann, William C. (1975) *Why things are so bad for the computer naive user*. Technical Report RR-75-32, Information Sciences Institute, University of Southern California, California, USA, March.

Mc Kevitt, Paul (1986) *Building embedded representations of queries about UNIX*. Memorandum in Computer and Cognitive Science, MCCS-86-72, Computing Research Laboratory, Dept. 3CRL, Box 30001, New Mexico State University, Las Cruces, NM 88003-0001, USA.

Mc Kevitt, Paul (1990) Acquiring user models for natural language dialogue systems through Wizard-of-Oz techniques. In *Proceedings of the Second International Workshop on User Modeling*, University of Hawaii at Manoa, Honolulu, Hawaii, USA, March.

Mc Kevitt, Paul (1991a) Principles and practice in an operating system consultant. In *Artificial Intelligence and Software Engineering, Vol. 1*, Chapter on 'AI Mechanisms and techniques in practical software', Derek Partridge (Ed.). New York: Ablex Publishing Corporation.

Mc Kevitt, Paul (1991b) *Analysing coherence of intention in natural-language dialogue*. Ph.D. Thesis, Department of Computer Science, University of Exeter, Exeter, England, EC, September.

Mc Kevitt, Paul (1992) The OSCON operating system consultant. In *Intelligent Help Systems for UNIX – Case Studies in Artificial Intelligence*, Springer-Verlag Symbolic Computation Series, Peter Norvig, Wolfgang Wahlster and Robert Wilensky (Eds.). Berlin: Springer-Verlag.

Mc Kevitt, Paul and Yorick Wilks (1987) Transfer Semantics in an Operating System Consultant: the formalization of actions involving object transfer. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence (IJCAI-87)*, Vol. 1, 569-575, Milan, Italy, EC, August.

Mc Kevitt, Paul, Derek Partridge and Yorick Wilks (1992a) *A survey of approaches to natural language discourse processing.* Technical Report 235, Department of Computer Science, University of Exeter, GB-EX4 4PT, Exeter, England, EC.

Mc Kevitt, Paul, Derek Partridge and Yorick Wilks (1992b) *Analysing coherence of intention in natural language dialogue.* Technical Report 227, Department of Computer Science, University of Exeter, GB- EX4 4PT, Exeter, England, EC.

Mc Kevitt, Paul, Derek Partridge and Yorick Wilks (1992c) *Why machines should analyse intention in natural language dialogue.* Technical Report 233, Department of Computer Science, University of Exeter, GB-EX4 4PT, Exeter, England, EC.

Mc Kevitt, Paul, Derek Partridge and Yorick Wilks (1992d) *Experimenting with intention in natural language dialogue.* Technical Report 234, Department of Computer Science, University of Exeter, GB- EX4 4PT, Exeter, England, EC.

Minsky, Marvin (1975) A framework for representing knowledge. In *The Psychology of Computer Vision*, P.H. Winston (Ed.), 211-217. New York: McGraw-Hill.

Moore, R.C. (1980) *Reasoning about knowledge and action.* Technical Report No. 191, Artificial Intelligence Center, SRI International, Menlo Park, CA, USA.

Partridge, Derek (1991) *A new guide to Artificial Intelligence.* Norwood, New Jersey: Ablex Publishing Corporation.

Pereira, Fernando and David Warren (1980) Definite clause grammars for language analysis – a survey of the formalism and a comparison with augmented transition networks. In *Artificial Intelligence*, 13, 231-278.

Pollack, Martha E. (1986) *Inferring domain plans in question-answering.* Unpublished Doctoral Dissertation, University of Pennsylvania, Philadelphia, Pennsylvania, USA.

Pustejovsky, James (1987) An integrated theory of discourse analysis. In *Machine Translation: Theoretical and methodological issues*, Studies in Natural Language Processing, Sergei Nirenberg (Ed.), 168-191. Cambridge, United Kingdom: Cambridge University Press.

Quillian, M. R. (1969) The teachable language comprehender: a simulation program and theory of language. In *Communications of the Association for Computing Machinery*, 459-76, 12, 8, August.

Reichman-Adar, R. (1984) Extended person-machine interface. In *Artificial Intelligence*, 22(2), 157-218.

Reichman, R. (1985) *Getting computers to talk like you and me.* Cambridge, MA: MIT Press.

Reilly, Ronan and N.E. Sharkey (Eds.) (1991) *Connectionist approaches to natural language processing, Vol. 1.* Hillsdale, NJ: Lawrence-Erlbaum.

Rumelhart, David (1975) Notes on a schema for stories. In *Representation and Understanding*, D. Bobrow and A. Collins (Eds.), 211-236. New York: Academic Press.

Schuster, Ethel; David Chin, Robin Cohen, Alfred Kobsa, Kathrina Morik, Karen Sparck Jones, Wolfgang Wahlster (1988) Discussion section on the relationship between user models and discourse models. In *Computational Linguistics, Special Issue on User Modeling, Alfred Kobsa and Wolfgang Wahlster (Guest Eds.)*, Vol. 14, No. 3, 79-103, September.

Schank, Roger C. (1972) Conceptual dependency: a theory of natural language understanding. In *Cognitive Psychology*, 3(4): 552-631.

Schank, Roger C. (1973) Identification and conceptualizations underlying natural language. In *Computer Models of Thought and Language*, R. Schank and K. Kolby (Eds.). San Francisco, CA: Wh Freeman and Co..

Schank, Roger C. (1975) *Conceptual information processing.* Fundamental Studies in Computer Science, 3. Amsterdam: North-Holland.

Schank, Roger C. and Robert P. Abelson (1977) *Scripts, plans, goals and understanding: an inquiry into human knowledge structures.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Schank, Roger and Chris Riesbeck (Eds.) (1981) *Inside Computer Understanding: Five Programs Plus Miniatures.* Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Searle, John R. (1969) *Speech acts: An essay in the philosophy of language.* London: Cambridge University Press.

Sidner, Candace L. (1983) What the speaker means: The recognition of speakers' plans in discourse. In *International Journal of Computers and Mathematics, Special Issue in Computational Linguistics*, 9(1):

71-82.
Sidner, Candice L. (1985) Plan parsing for intended response recognition in discourse. In *Computational Intelligence*, 1(1): 1-10, February.
van Dijk, T.A. (1972) *Some aspects of text grammars.* The Hague: Mouton.
van Dijk, T.A. (1977) *Text and context: Explorations in the semantics and pragmatics of discourse.* London: Longmans.
Webber, Bonnie (1978) *A formal approach to discourse anaphora.* Doctoral Dissertation. Also in, BBN Report No. 3761, Bolt, Beranek and Newman, Cambridge, MA, USA.
Wilensky, Robert (1983) *Planning and understanding.* Reading, MA: Addison-Wesley.
Wilensky, Robert, David N. Chin, Marc Luria, James Martin, James Mayfield and Dekai Wu (1988) The Berkeley UNIX Consultant project. In *Computational Linguistics*, 35-84, Vol. 14, No. 4, December.
Wilks, Yorick (1973) An artificial intelligence approach to machine translation. In *Computer Models of Thought and Language*, R. Schank and K. Kolby (Eds.). San Francisco, CA: Wh Freeman and Co..
Wilks (1975a) Preference semantics. In *Formal semantics of natural language*, Keenan, Edward (Ed.). Cambridge, United Kingdom: Cambridge University Press. Also as, Memo AIM-206, Artificial I ntelligence Laboratory, Stanford University, Stanford, California, USA, July 1973.
Wilks (1975b) An intelligent analyzer and understander of English. In *Communications of the ACM*, 18(5), 264-274, May. Also in, *Readings in Natural Language Processing*, Barbara Grosz, Karen Sparck Jones and Bonnie Webber (Eds.), 193-204, Los Altos, California: Morgan Kaufmann, 1986.
Wilks, Yorick (1975c) A preferential, pattern-seeking semantics for natural language inference. In *Artificial Intelligence*, 6, 53-74.
Wilks, Yorick and J.S. Bien (1983) Beliefs, points of view and multiple environments. In *Cognitive Science*, 795-119. Also in, *Sixth International Joint Conference on Artificial Intelligence (IJCAI-)*, 431-453, USA.
Wilks, Yorick and Afzal Ballim (1987) Multiple agents and the heuristic ascription of belief. In *Proceedings of the 10th International Joint Conference on Artificial Intelligence (IJCAI-87)*, 119-124, Milan, Italy, August.

Appendix A

OSCON sample run

This Appendix demonstrates that the OSCON system dynamically modifys its natural language responses dynamically, based on the types of utterances, and hence intention, that appear in the input.

```
| ?- oscon.
OSCON Program 1.0 (Sun-4, SunOS Release 4.1)
Copyright (C) 1988, Computing Research Laboratory. All rights reserved.
Dept. 3CRL, Box 30001, NMSU, Las Cruces, NM (505) 646-5466, USA.

U S WEST Advances Technology

Please input your question ending with with a '?'
Write 'quit.' when you are finished
OSCON can answer queries for UNIX or MSDOS.
Type one of these to the prompt.
Do you want answers for UNIX or MSDOS?
The default is UNIX [unix]: unix

How can I help you?

--> what does ls do?

'ls <directoryname>' will display directory contents on the screen.
```

```
information  : 0  0  0  0  0   0  =  0
description  : 0  0  0  0  0   0  =  0
elaboration  : 0  0  0  0  0   0  =  0         Satisfaction    = 0
explanation  : 0  0  0  0  0   0  =  0         Dissatisfaction = 0
repetition   : 0  0  0  0  0   0  =  0

nointention  : 0  0  0  0  0   0  =  0

                                  == 0
```

OSCON displays the 'intention matrix' and levels of 'satisfaction' and 'dissatisfaction'. Note that OSCON caters for 5 'real' types of intention, and nointentions. The 'intention matrix' is a 2-dimensional, 6 X 6 matrix. The 'intention matrix' shows all intention-sequence counts to be 0 as only one intention has been entered into the program. Totals for each intention type are also shown. 'Satisfaction' and 'dissatisfaction' levels, computed from the matrix, using the user modelling function, are also 0.

```
--> how do i see my file?

'more <filename>' will display file contents on the screen.

information  : 0  0  0  1  0   0  =  1
description  : 0  0  0  0  0   0  =  0
elaboration  : 0  0  0  0  0   0  =  0         Satisfaction    = 1
explanation  : 0  0  0  0  0   0  =  0         Dissatisfaction = 0
repetition   : 0  0  0  0  0   0  =  0

nointention  : 0  0  0  0  0   0  =  0

                                  == 1
```

This 'intention matrix' shows a total intention sequence count of 1, and the intention count for the sequence 'explanation → information' is set to 1. 'Satisfaction' also totals to 1.

```
--> how does he see my file?

'more <filename>' will display file contents on the screen.

information  : 0  0  0  1  0   0  =  1
description  : 0  0  0  0  0   0  =  0
elaboration  : 0  0  0  0  0   0  =  0         Satisfaction    = 2
explanation  : 0  0  0  0  0   0  =  0         Dissatisfaction = 0
repetition   : 1  0  0  0  0   0  =  1

nointention  : 0  0  0  0  0   0  =  0

                                  == 2
```

This 'intention matrix' shows a total sequence count of 2. Now, a 'information → repetition' sequence has occurred. Note that the syntax of this query does not have to be exactly the same as the previous query to denote a repetition. 'Satisfaction' increases to 2.

```
--> how do they see my file?

'more' is a command.
```

```
It has the syntax ‘more <filename>’.
‘more <filename>’ will display file contents on the screen.
‘more’ only takes the argument ‘file’.

information  : 0  0  0  1  0   0  =  1
description  : 0  0  0  0  0   0  =  0
elaboration  : 0  0  0  0  0   0  =  0        Satisfaction    = 2
explanation  : 0  0  0  0  0   0  =  0        Dissatisfaction = 3
repetition   : 1  0  0  0  1   0  =  2

nointention  : 0  0  0  0  0   0  =  0

                                    == 3
```

A ‘repetition → repetition’ intention loop has occurred tilting the level of ‘dissatisfaction’ higher than ‘satisfaction’. As a reaction, more information is returned to the user. Remember that the user modelling function gives intention repetitions which are along diagonals a weight of 3.