RELEVANCE: PROCESSING IMPLICATIONS

Marjolein GROEFSEMA

Abstract

In Relevance, Sperber and Wilson (1986) argue that an addressee builds anticipatory logical hypotheses during the comprehension process. These hypotheses are needed for disambiguation and reference assignment. If this is the case, then a psycholinguistic model of input processing should account for how these hypotheses are built. According to the 'Relevance driven' processing model proposed in this paper the addressee does not build logical hypotheses on the basis of the natural language syntax, but rather uses information stored in the logical entries of the concepts associated with the lexical items encountered in the utterance, guided by the principle of Relevance. Relevance Theory claims that a communicator aiming at optimal relevance will try to make her/his utterance worth the addressee's while, not only on the message level, but also by keeping processing cost down, by accomodating semantic, syntactic and lexical choices to the processing needs of the addressee. The model proposed here views the interpretation process as a process in which the addressee expects the communicator to be aiming at optimal relevance, and in which s/he will interpret an utterance according to this expectation.

1. Introduction

A theory of natural language interpretation should account for how people compute the full pragmatic meaning of an utterance. An account has to be given of how disambiguation and reference assignment take place, how the addressee restores elliptical material, how implicit information is recovered, and how an addressee works out what the illocutionary force of an utterance is. There is, as yet, little psycholinguistic theorisation of these aspects of interpretation. This may be partly due to the fact that, until quite recently, there was only one serious theory of pragmatics, that developed by Grice (1975). He proposed that conversation is governed by a Co-operative Principle, consisting of a number of conversational maxims. A problem with this theory is that it is not explicit enough to provide a basis for a psycholinguistic account of the role of pragmatics in real-time utterance interpretation. (For critiques, see Sperber and Wilson (1986); Blakemore (1987)). With Relevance Theory, Sperber and Wilson (1986) present a

new approach to pragmatics. Because this theory makes explicit claims about how natural language is interpreted in real time, and gives an explanation of how questions like disambiguation, reference assignment, etc. can be accounted for, it should have consequences for a psycholinguistic model of processing.

The basic claim that Sperber and Wilson make is that in processing information people try to achieve the greatest possible cognitive effect for the smallest possible amount of processing cost. As a consequence of this people will only pay attention to information they think is relevant, or more relevant than any other information they could be attending to at that moment. Sperber and Wilson say that "(...) ordinary utterance comprehension is almost instantaneous, and however much evidence might have been taken into account, however many hypotheses might have been considered, in practice the only evidence and hypotheses considered are those that are immediately acessible." (Sperber and Wilson (1986), pp.66-67). Furthermore, they say that "the fact that verbal comprehension is almost instantaneous, and is achieved with the active help of the speaker, makes the hearers choice of a context from the whole of conceptual memory more amenable to study." (op. cit., p.67). They postulate that human cognition and communication is driven by relevance, and the maximisation of relevance, which accounts for how some background information and context becomes immediately accessible, rather than others; people process utterances in a context that maximises their relevance. They capture these ideas in the principle of Relevance, which says: "Every act of ostensive communication communicates the presumption of its own relevance.", where relevance is defined as:

a) the greater the contextual effects, the greater the relevance,

b) the smaller the processing effort, the greater the relevance;

and where the assumption of optimal relevance is defined as:

"a) The set of assumptions {I} which the communicator intends to make manifest to the addressee is relevant enough to make it worth the addressee's while to process the ostensive stimulus.

b) The ostensive stimulus is the most relevant one the communicator could have used to communicate {I}." (op. cit., p.158).

It follows from this that a communicator aiming at optimal relevance will not only try to make her/his utterance worth the addressee's while on the message

level, but will also try to keep processing cost down, by accomodating semantic, syntactic and lexical choices to the processing needs of the addressee.

In most processing models little or no account is taken of the fact that the use of natural language is a process of interaction between communicator and addressee, in which both parties are involved in establishing succesful communication. Work on 'context-neutral' languages for computers and on Artificial Intelligence seems to have influenced psycholinguists in trying to explain input processing phenomena purely in terms of hearers'/readers' activity, without taking into account that speakers/writers are responsible for the input:

"*It is natural to think of perception and comprehension as including analogues of the parsing operations of formal grammars, and so to view A.I. parsing schemes as potential models of (portions of) some mental processes.*" (Karttunen and Zwicky (1985), p.9).

Relevance theory, however, forces us to review this conception, by showing us that natural language use is more than the use of a 'context-neutral' language with the context added: the choice of a particular utterance by a communicator aiming at optimal relevance is a consequence of the context in which it is uttered, and the addressee processes the utterance based on this.

2. The 'relevance' approach to hypothesis formation

Sperber and Wilson (1986) suggest that, during real time utterance interpretation, the addressee builds anticipatory logical hypotheses about the overall structure of the proposition expressed by the utterance. These hypotheses are needed by the addressee to enable her/him to resolve potential ambiguities and ambivalences. Sperber and Wilson assume that these logical hypotheses are built on the basis of anticipatory syntactic hypotheses, unless the addressee already has an incomplete logical form available, which can function as an anticipatory logical hypothesis, as is the case with, for example, answers to questions. They assume that logical forms are "trees of labelled nodes (...) which should be a set of logical categories, perhaps from a fixed range which is part of human mental equipment, which might be regarded as variables over conceptual representations of different types." (op. cit., p.205). They use the proforms of English to represent these variables, e.g., SOMEONE as a variable over conceptual representations of people. Sperber and Wilson assume that after recognizing the first word of an utterance the addressee assigns it to a syntactic category, (e.g. NP), and then makes the anticipatory

syntactic hypothesis that it will be followed by another syntactic category (e.g., NP -> VP). By variable substitution this then yields an anticipatory logical hypothesis, e.g., (NP(John) VP) -> 'John DO/BE/HAVE SOMETHING'. Sperber and Wilson say that "on this approach, there is a clear sense in which the logical category labels correspond to, and are indeed semantic representations of, syntactic category labels of natural language (though there need not be a one to one correspondence)." (op. cit., p.206). However, it does not become clear from their account on what basis and how the syntactic hypotheses are formed, whether these are built for each word, for example, or only for major phrases. Sperber and Wilson say in respect to this that the hearer "might not only identify each word and tentatively assign it to a syntactic category, but use his knowledge of its lexical properties and syntactic co-occurence restrictions to predict the syntactic categories of following words or phrases." (op. cit., p.205). Sperber and Wilson support their claim that syntactic hypotheses are built by referring to Johnson-Laird's chapter 'Parsing and Performance' in Mental Models (1983). What Johnson-Laird actually says about syntactic hypothesis formation is the following:

"*There are in principle several types of prediction that a parser might make. It could predict that the next constituent will be of a particular category. (...) It could predict that a particular constituent must definitely occur at some later point in the sentence, though not necessarily as the next constituent. (...) Finally, the parser could make either sort of prediction with respect to optional constituents. (...) The point about making predictions is to increase efficiency. If too many predictions are made, however, the system will collapse under its own processing load. It would thus be folly to design a system that made top down predictions about, say, occurences of conjunctions, or adverbs like 'only' or any other constituent that is ubiquitous. Such predictions would have to be made after almost every word in a sentence and mostly fail to be fulfilled.*" (Johnson-Laird (1983), pp.320-321).

Unfortunately, Johnson-Laird does not make it clear what principles or strategies the parser uses to decide when to make hypotheses and when to refrain from making them. For example, it does not become clear what the parser will do when it encounters a constituent which gives rise to different expansions, as is the case with 'admitted' in 'Jennifer admitted stealing' (example from Sperber and Wilson (1986), p.206). When the hearer recognizes 'admitted', this gives rise to a range of possibilities of expansion, i.e., S, S', PP (S') and NP (PP). Johnson-Laird's account does not make it clear whether hypotheses would be built for all these possibilities. If so, this would give rise to an extended range of anticipatory logical

hypotheses since there is not necessarily a one-to-one correspondence between the logical category labels and the syntactic category labels; in turn this would increase processing load. If, at the other hand, hypothesis building were postponed until the next word was recognized, to keep processing cost down, this would mean that no anticipatory logical hypotheses could be built, which would create problems for Sperber and Wilson's approach to utterance interpretation, and especially disambiguation.

Another problem with this view on hypothesis building is that it is not clear which logical variables will be substituted for which syntactic categories. Sperber and Wilson give as anticipatory hypotheses built on the syntactic hypothesis (admitted -> NP):

(62) (a) Jennifer let someone in.
(b) Jennifer confessed to something.
(op. cit., p.207).

Here (62a) has SOMEONE, a variable over people, whereas in (62b) SOMETHING is used, a variable over things. If we presume that any NP gives rise to two logical hypotheses, one with SOMEONE and one with SOMETHING, this would have as an undesirable consequence that (admitted -> NP) would yield the logical hypothesis 'Jennifer admitted (confessed to) SOMEONE.' (in the sense of: (?) the thing/proposition that Jennifer admitted to was someone), which would be ruled out on grounds of semantic incompatability. Furthermore, NP is not the only category that could be substituted by more than one logical variable. Jackendoff (1983) identifies and justifies several basic 'conceptual' categories; at least four of these, PLACE, MANNER, TIME and PROPERTY can be realized by a PP. These would give us four different anticipatory logical hypotheses, when a PP is encountered in a syntactic hypothesis. For example, the hypothesis (N -> PP) would give us, by variable-substitution:

"Jennifer SOME PROPERTY..."
"Jennifer SOMEWHERE..."
"Jennifer SOME MANNER..."
"Jennifer SOME TIME..."

Although the first and, arguably, the second of these can be realized as natural language NPs, TIME and MANNER are problematic, because they typically modify events, so that 'Jennifer SOME TIME' and 'Jennifer SOME MANNER' cannot be realised by natural language NPs. This means that, unless one wants to say that there are constraints on which logical variables can be substituted for which syntactic category in which position in the syntactic tree, one ends up with anticipatory logical hypotheses which are unrealizable, and therefore superfluous.

Basing logical hypotheses on syntactic hypotheses may actually lead to false predictions in some cases. Consider the following example:

A: I went to the pictures last night.
B: Me too.

After recognizing 'me' the hearer assigns it to the syntactic category N or NP and can then, because of its case, build the syntactic hypothesis (NP -> conj. NP), e.g., 'Me and John...'. By variable substitution this will yield the anticipatory logical hypothesis 'speaker and SOMEONE...'. Alternatively, 'me' can be recognized as a displaced object, so that the syntactic hypothesis can be built that it will be followed by a sentence with an empty 'object-slot'. The hearer has no place in either hypothesis to fit 'too' into, which will render 'me too' uninterpretable. This would imply that a speaker aiming at optimal relevance would not utter 'me too'; but surely one would want to say that B's response is more relevant than the alternative utterance "I went to the pictures last night too.", since, although both utterances have the same contextual effects, the effort required to process 'me too' is (intuitively) small, whereas processing "I went to the pictures last night too." takes more effort, which is not set off by any extra contextual effects. In general, it seems that elliptical sentences present a problem as far as syntactic hypothesis building is concerned, because they induce the wrong hypothesis to be built, e.g.,

"Hope to see you soon."
"Police rescued from gang."

This can also happen with sentences with dislocated constituents, e.g.,

"The question of financing the project we discussed last week."
"That film we saw on the ferry coming to Britain."

There is nothing about these dislocated NPs that distinguishes them from ordinary subject-NPs, so that the wrong hypothesis (NP -> VP) can be built on recognising them. Alternatively, we would be forced to adopt the view that for every first encountered NP in an utterance the hypothesis that it is dislocated would have to be built, as well as (NP -> VP), something which is clearly undesirable.

3. An alternative approach to logical hypothesis building.

Sperber and Wilson see anticipatory logical hypotheses "as playing a crucial role in disambiguation and reference assignment." (Op. cit., p.206). They note that "the experimental literature on disambiguation suggests that disambiguation and reference assignment are (...) to some extent 'top-down' processes: that the hearer makes anticipatory hypotheses about the overall logical structure of the utterance and resolves potential ambiguities on the basis of these." (op. cit., p.205). However, as I argued in the previous section, their account of how anticipatory logical hypotheses are constructed comes up against a range of problems, although these problems do not concern the feasibility of having anticipatory logical hypotheses as such. This then raises the question of whether it is necessary to postulate that anticipatory logical hypotheses are based on anticipatory syntactic hypotheses, and indeed, whether postulating this is a prerequisite for the way in which Sperber and Wilson view the interpretation process, and a consequence of Relevance Theory.

Within the framework of relevance theory, anticipatory logical hypotheses are hypotheses about the structure of the logical form of the utterance. For Sperber and Wilson "a logical form is a well-formed formula, a structured set of constituents, which undergoes formal logical operations determined by its structure. (...) when a natural-language sentence is uttered, the linguistic input system automatically decodes it into its logical form (or in the case of an ambiguous sentence into a set of logical forms), which the hearer is normally expected to complete into the fully propositional form that the speaker intended to convey." (op. cit., pp.72-73). This notion of logical form is different from the level of LF within the GB-framework, where LF is defined as the level of grammar at which quantifier scope and other properties are directly represented (cf. Chomsky, 1981, 1986). In GB this representation is completely syntactic, whereas in Relevance Theory logical form is a representation of an inner 'Language of Thought' (Cf. Fodor, 1981). Kempson (1986) says about this difference:

"*Relevance theory assumptions [about logical form, mg.] differ from GB assumptions in particular with respect to the vocabulary in which LF is constructed (...), there is a mapping onto incomplete expressions of the language of inference from s-structure as part of the grammar, a mapping which defines the level of logical form. Such logical forms are mixed representations. They involve expressions of the language of inference which may be incomplete. Points at which they are not complete are specified as metavariables, themselves place-holders for the*

value to be assigned, together with associated restrictions on the value that metavariable may take. These metavariables are then assigned same language-of-inference representations as value by pragmatic principles and not by any algorithmic device associated with the grammar of the input language." (Kempson (1986), p.8).

As both Sperber & Wilson and Kempson point out, the logical forms that the addressee builds on the basis of linguistic information have to be completed into the fully propositional forms, which were intended to be conveyed. In order to do this, the addressee will have to build further anticipatory logical hypotheses (as Kempson points out: "points at which they are not complete are specified as metavariables ...."). Since one of the central claims of relevance theory is that the linguistic content of an utterance underdetermines its propositional content, it follows from Relevance Theory itself that logical hypotheses cannot only be based on anticipatory syntactic hypotheses by variable-substitution. This leaves us with two possibilities. One possibility is that the addressee builds anticipatory logical hypotheses partly based on syntactic hypotheses and partly in a different way. If this is the case then an account will have to be given of how the processor deals with the problems that we have encountered in the previous section, and also what it is that complements the syntactic processor in building hypotheses, and how this is done. The second possibility is that anticipatory logical hypotheses are built in a different way altogether. I will argue that adopting this possibility will open the way to a better account of how people process language in real time.

Within a GB-framework, LFs are composed of lexical items and natural language syntax. Chomsky (1986) says that what the lexicon contains is:

"*(...) for each lexical item, its (abstract) phonological form and whatever semantic properties are associated with it. Among these will be the 'selectional properties' of heads of constructions (...) Let us call these properties 'semantic selection' (s-selection) (...)*" (Chomsky (1986), p.86).

Sperber and Wilson, however, follow Fodor (1983) in assuming that the language faculty is a modular input system that translates 'lower level' sensory representations into 'higher level' conceptual representations, which are of the same format as 'higher level' representations from other input mechanisms. Logical forms (in various states of completion) are well formed formulae containing concepts and logical variables, rather than lexical items. Sperber and Wilson see concepts as psychological objects, which contain three distinct types of information: logical, encyclopaedic and lexical. They say that:

"*on this approach a conceptual address is thus a point of access to the logical, encyclopaedic and linguistic information which may be needed in the processing of logical forms containing that address.*" (Sperber and Wilson (1986), p.86).

This division into different types of information raises the question of where s-selection properties (of whatever format) are stored. Whether they are contained within the lexical entry applying to natural language syntax, which is the only possibility within the Chomskyan view, or whether they apply to logical form rather than natural language and therefore are part of the logical entry of a concept. Let us consider the verb 'eat'. 'Eat' can be used transitively or intransitively, which gives us either two s-selection frames, i.e.,

(Agent(_)) and (Agent(_ theme))

or one s-selection frame with an optional theme:

(Agent(_ (theme))).

If s-selection frames are stored in the mental lexicon, then the parser will have to make a choice, e.g., (Agent(_)) will be chosen in the case of 'John was eating'. However, 'John was eating' implies that John was eating SOMETHING (even though the speaker emphasizes the act of eating as opposed to what was eaten, by choosing this particular utterance), as is illustrated by the following conversation:

A: John was eating when I phoned him.
B: Oh, what was he eating ?
A: I don't know. It doesn't matter, does it ?
B: He always cooks such wonderful meals !

What A wants B to recover when she utters "John was eating..." is the logical form 'John was eating SOMETHING...', intending the information conveyed in this logical form to be relevant for B in its own right, i.e., although A implies that John was eating something, she intends the focus of her utterance to be on the act of eating. However, B processes this utterance in a context (which B explicates in her second utterance), which makes it relevant for her to know what concept should be inserted in place of the logical variable SOMETHING. If we take s-selection frames to be specifications about what the logical form of an utterance is going to be, rather than its lexical form, then we can account for the difference in emphasis between the transitive and intransitive use of 'eat' without having to postulate that a choice has to be made between different s-selection frames, or different options within one s-selection frame; we can instead say that 'eat' just has one 'logical variable selection

frame':

(SOMEONE(_ SOMETHING))

which is stored in the logical entry of the concept. Depending on what proposition the communicator wishes to convey, s/he can choose whether to lexicalize the logical variable SOMETHING, or not. This has as a consequence that it becomes superfluous to postulate that s-selection frames are stored in the mental lexicon as well: all the addressee has to do is recognize the phonological form of a word and then access the logical entry of the concept that that word refers to, which will give her/him the (logical) s-selection frame associated with that concept, as illustrated in figure 1):

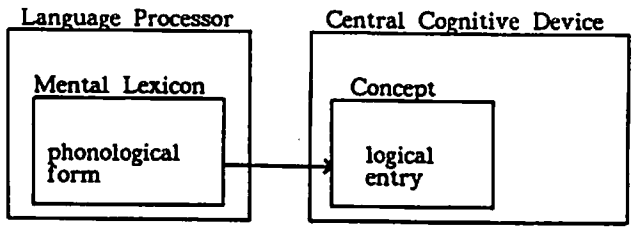

Figure 1): Recovering the s-selection frame.

Viewed in this way, s-selection frames show what their associated concepts can combine with in order to yield a well-formed logical form. This means that one does not have to say that the logical structure of an utterance is specified by the argument structure associated with the lexical items in a given s-structure. Rather, viewing s-selection frames in this way enables us to say that the thematic roles and relations that we can observe in natural language are a 'translation' of the logical form specification of concepts.

Postulating that s-selection frames (as I will continue to call them) are contained in the logical entry of a concept has as a consequence that we can show how anticipatory logical hypotheses are built without having to appeal to anticipatory syntactic hypotheses: when the addressee accesses the lexical entry of a concept, this will give her/him the s-selection frame, which tells her/him what logical variable(s) the concept needs to combine with in order to yield a well formed logical formula. The addressee can then use this s-selection frame as an anticipatory logical hypothesis. E.g., on recognizing 'a' the addressee can build the anticipatory logical hypothesis:

(((_)(SOME N'))(SOME PREDICATE))
SOME PROPOSITION

In this hypothesis 'a' works as a 'semantic pointer' telling the addressee that the identity of whatever SOME N' refers to has not yet been established. A question to ask here is whether these s-selection frames consist of the same information as Chomsky envisages for s-selection frames as stored in the mental lexicon. I want to suggest that these logical s-selection frames are in fact richer. In the first place, the s-selection frame will show what the concept has to combine with in order to yield a well formed formula. For example, GIRL will have the specification:

(((SOME)(_))(SOME PREDICATE))
SOME PROPOSITION

where SOME is a logical variable which has to be filled by a specifier. All the variables in this specification have to be filled by concepts, either derived from the utterance that the addressee is interpreting, or from the context in which it is interpreted. If no specific concept is available then the addressee may insert the basic category concept in place of the logical variable that is derived from it. In the second place, the s-selection frame will show what the concept can combine with, although it does not have to. E.g., GIRL will have as a further specification:

((_)(SOME PROPERTY))
SOME N'

Since one of the relevance theory assumptions is that the interpretation process is an on-line process, this means that as soon as a first anticipatory logical hypothesis is formed, it can undergo pragmatic processes. This means that logical variables like SOME PROPERTY above can, in a sense, be activated during the interpretation process. If, for example, it becomes important to an addressee to know which girl is meant, while processing "The girl...", this will 'activate' the logical variable SOME PROPERTY. If the communicator does not forward the information in her/his utterance, this will cause the addressee to ask "Which girl ?", if this is relevant enough for her/him. The s-selection frames will thus have different tiers, which contain these different kinds of information. It is not always necessary to access the whole s-selection frame of a concept. For example, when an addressee, processing "The girl...", accesses the s-selection frame of THE, this will yield the anticipatory logical hypothesis:

((THE (SOME N'))(SOME PREDICATE))
SOME PROPOSITION

so that, when 'girl' is recognized, only the tier of the s-selection frame of GIRL which gives the further specification needs to be accessed: the addressee already has an anticipatory hypothesis available. In the third place, an s-selection frame may contain restrictions on what the value of a concept that can be assigned to a variable may be. For example, BUY will have the (partial) s-selection frame:

(SOMEONE (__ SOMETHING FOR SOMEONE))
SOME PROPOSITION

with an associated restriction on FOR SOMEONE that, unless the communicator states for whom, or this is immediately recoverable from the context, this variable refers to the person who does the buying. Furthermore, any utterance or series of connected utterances refers to a certain time at which the proposition expressed took, takes or will take place. This time can be denoted by the tense of the verb used, it can be made more precise by the communicator by using a time-adverb, prepositional phrase, or subordinate clause, or it can be made more precise by the addressee from accessing the context. This suggests that any logical hypothesis has a logical variable SOME TIME associated with it.

Similarly, any utterance or series of connected utterances that refers to an EVENT (in the sense of Jackendoff 1983) implies that the event took place somewhere, even though the focus of the utterance may be on what took place, rather than where it happened. This seems to point at the presence of a logical variable SOMEWHERE in the s-selection frames of concepts connected to verbs that refer to an event.

Let us consider an example from the previous section:

A: I went to the pictures last night.
B: Me too.

When the addressee recognizes 'me', s/he can access the logical entry of the associated concept and build the logical hypothesis:

((X)(SOME PROPERTY)) speaker(X)
SOME N'
(($_{N'}$ _)(SOME PREDICATE)(SOME TIME))

where (speaker(X)) is a restriction on the value of X. This yields:

(((B)(SOME PROPERTY))(SOME PREDICATE)(SOME TIME)).

When the addressee recognizes 'too', s/he can take this as a 'semantic pointer' or instruction, to fill in the variable slots with the preceding values that are most accessible, i.e., the values from the immediately preceding utterance. This yields the logical form:

((B (SOME PROPERTY))(went to the pictures)(last night)).

An objection can be made here that 'me' does not have nominative case, so that it cannot function as the subject-argument of a proposition. I think that this objection would be valid if the mapping from natural language onto logical form was a process of translating, of variable-substitution. Since I am claiming that the structure of the logical form is not (directly) dependent on the syntactic form of the utterance, but that it is a discrete system, it follows that it is not necessary to postulate that (explicit) case plays a role in logical form. In fact I want to suggest that case is a property of natural language and not of the language of thought, and that it only comes into play in language production. Consider the following example:

Child: Him is nice !
Father: You should say "he is nice", not "him is nice".

If logical form formation was a process of variable-substitution, then it would be difficult for the father in the above exchange to interpret his child's utterance at all, but in actual fact it is not. This follows straightforwardly if we take it that the father accesses the logical entry for HIM and finds:

((Y)(SOME PROPERTY)) male(Y)
SOME N' +pronominal
(($_{N'}$)(SOME PREDICATE)(SOME TIME))

where (+pronominal) can be taken as a restriction on how to find the referent of the expression used. Another apparent problem may be that if we specify the s-selection frame of HIM as above, we run into trouble when 'him' is encountered in direct or indirect object position, e.g. 'I saw him yesterday'. This, however, is not a real problem, since by the time the addressee encounters the object, s/he has already formed specific logical hypotheses, so that there is no need to access the whole logical entry of the concept.

4. Processing implications.

I want to propose a processing model in which the interpretation process is driven by the principle of relevance, a model that takes into account that the input for comprehension is never 'neutral', but is produced by a communicator aiming at optimal relevance. This means that the communicator will not only try to make her/his utterance worth the addressee's while on the message level, but will also try to keep processing cost down, by accomodating semantic, syntactic and lexical choices to the processing needs of the addressee.

This conception together with the notion of logical hypothesis building, as I have developed it in the last section, has consequences for a model of processing, because it forces us to look at the interpretation process as a process in which the addressee expects the communicator to be aiming at optimal relevance, and in which s/he will interpret an utterance according to this expectation. In this paper, I will not have anything to say about the role of the prosodic processor, although I assume that prosodic properties have an effect on the choice of context in which an utterance will be processed, and, at times, on the building of logical forms. Likewise, I will not have anything to say about the way in which the phonological form (or the orthographic form) of a lexical item is recognized during the interpretation process. My general picture of the language processing system is given in figure 2 on the next page.

In this model, the 'linguistic processor' is responsible for recognizing the phonological (or orthographic) form of a lexical item. It does not, however, assign it to a syntactic category, nor does it assign a syntactic structure to the string of lexical items that make up the utterance. Postulating that all that an addressee needs in the interpretation process is to recover the phonological (or orthographic) form of a word from the mental lexicon and then access the logical entry of the concept associated with that form, in order to recover the s-selection frame, has as a consequence that there is no role left for a parser. By accessing the s-selection frame in the logical entry of the concepts, the logical processor can determine the logical category of a concept. The concepts are assigned a logical structure by the logical processor, in the form of logical hypotheses, as set out in the previous section. These logical hypotheses can be enriched by accessing the context and the encyclopaedic entries of concepts, in accordance with the principle of relevance, until a fully propositional form has been built.

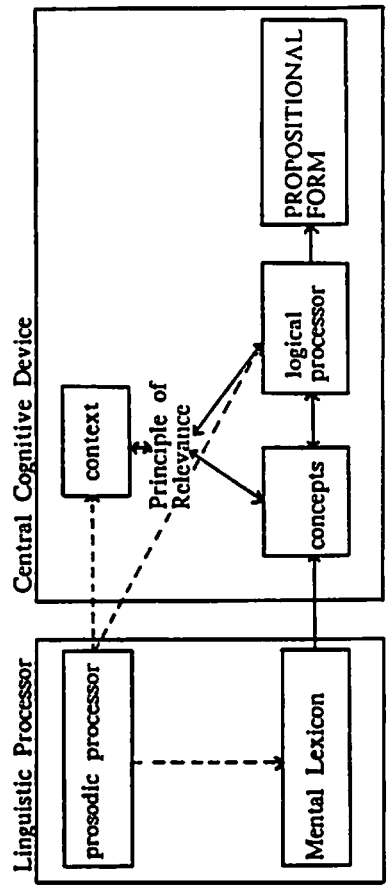

Figure 2: A 'relevance driven' input processing model.

Different parsing strategies have been proposed to account for the difficulty of processing sentences like:

1) The rat the cat the dog bit chased ran away.

For example, Kimball (1973) proposed seven parsing principles, one of which, the principle of New Nodes, explained why deletion of the complementizers can make sentences difficult to understand: the complementizer signals that a new phrasal node should be started. Although complementizers have this function in general, it does not seem to me that inserting them in sentence (1) makes it easier to understand (or only marginally so):

1') The rat that the cat that the dog bit chased ran away.

This shows up even more clearly in the following sentences:

2) Oysters oysters oysters split split split.

2') Oysters that oysters that oysters split split split.

Another of Kimball's principles, the principle of Two Sentences, restricts the parser to analysing daughters of only two S-nodes at one time. This principle explicitly states that the parser cannot handle multiple centre-embedded sentences; but this would predict that a sentence like (3) is difficult to understand, on a par with sentences (1) and (2):

3) The game those boys I met invented resembles chess.

Intuitively though, one would want to say that this sentence can be understood reasonably easily. Frazier and Fodor (1978) proposed that the parser works in two stages in order to account for problems like the above. The first, the PPP or sausage machine, could only handle six words at a time, the second, the SSS (Sentence Structure Supervisor), assembled the overall phrase-marker out of the 'packets' handed to it by the PPP. Both stages were governed by the principle of Minimal Attachment, which says that one does not postulate any (potentially) unnecessary nodes. This principle would lead the PPP to interpret 'The rat the cat the dog' as a conjunction, so that the parser would run into problems when 'bit chased ran away' was encountered. Although this explains why sentence (1) is difficult to process, it does not work for sentences like (2), because there are only six words in the whole sentence and therefore the PPP should be able to handle it. Frazier (1987) argues for the "Garden Path" model, which says that the parser just operates on two principles, the principle of Minimal Attachment (see

above), and the principle of Late Closure, which says that "If grammatically permissable, attach new nodes into the clause or phrase currently being processed (i.e. the phrase or clause postulated most recently.)". (Frazier (1987), p.4).

Unfortunately, this still does not account for the difference in difficulty of understanding sentences (1) and (2), as opposed to sentence (3). I want to suggest that no speaker aiming at optimal relevance would ever produce an utterance like (1) or (2), because the processing cost can be reduced by using a different utterance expressing the same proposition, e.g., 'The rat that was chased by the cat that was bitten by the dog ran away.'. This difference in processing cost, however, is not due to the complexity of the syntactic parsing, but rather to the amount of processing that the addressee is required to do in order to recover directly the logical form of the proposition expressed, i.e. 'who did what to whom'. The difference between sentences (1)/(2) and sentence (3) is that in the first ones there are no restrictions to differentiate between what can and what cannot be taken as the logical subjects and objects of the different predicates, e.g., the cat, the rat and the dog can all be involved in biting, chasing and running away, and also in being bitten and being chased. In sentence (3), however, there are restrictions: the addressee can work out that games do not, generally speaking, meet somebody, or are met; that they do not invent things, but can be invented; that people on the other hand, can meet people and can invent things, but are not usually invented. The addressee can then build logical hypotheses accordingly. When the addressee, while processing an utterance of sentence (1), accesses the logical entry for BITE, this will yield a logical hypothesis, part of which is:

((DOGi)(BIT (SOME X))),

where SOME X can equally well refer to THE RAT as to THE CAT. In contrast, when an addressee, while processing an utterance of sentence (3), accesses the logical entry for MEET, this will yield the partial logical hypothesis:

((X)(MET (SOMEONE))),
speaker(X)

so that the addressee can work out that SOMEONE cannot refer to THE GAME, which makes THOSE BOYS the easiest accessible candidate. Similarly, there are no restrictions on which concepts can be the logical subject and object of CHASE, whereas INVENT will yield the logical hypothesis:

(BOYSj (INVENTED (SOMETHING))),

so that THE GAME is the easiest accessible object. By the time the addressee accesses the logical entry for RESEMBLE, s/he already has recovered the logical forms (SPEAKERi MET BOYSj) and (BOYSj INVENTED GAMEk), so that RESEMBLED can be inserted in a straightforward way into the variable slot SOME PREDICATE connected to SOME GAME.

Minimal Attachment has also been used as an explanation for the 'garden path' effect in sentences like:

4) The horse raced past the barn fell.

5) The boat floated down the river sank.

However, Crain and Steedman (1985) reported that this effect disappears when sentences like (4) and (5) were processed in a suitable context: "The experiments (...) suggest that there may be no intrinsically garden-pathing structures whatever, but rather that, for any given sentence, there are certain contexts (including the null context) which induce garden paths, and certain others that do not." (Crain and Steedman (1985), p.345). In order to account for this, they proposed three principles; the principle of A Priori Plausibility:

"*If a reading is more plausible in terms either of general knowledge about the world, or of specific knowledge about the universe of discourse, then, other things being equal, it will be favored over one that is not.*" (op.cit., p.330),

The principle of Referential Success:

"*If there is a reading that succeeds in referring to an entity already established in the hearer's mental model of the domain of discourse, then it is favored over one that does not.*" (op. cit., p.331),

and the principle of Parsimony:

"*If there is a reading that carries fewer unsatisfied but consistent presuppositions or entailments than any other, then, other criteria being equal, that will be adopted as the most plausible by the hearer, and the presuppositions in question will be incorporated in his or her model.*" (op. cit., p. 333).

A weakness of Crain and Steedman's proposal is that the principles do not follow from any more general theory about the interpretation process; as Frazier (1987) argues:

"*Without knowing what constitutes a decisive bias during the left-to-right ongoing analysis of a sentence, we do not know how the above principles apply (...) Having received and analysed an entire sentence, it is certainly possible to*

identify the locally possible analyses of the sentence and determine the presuppositions carried by each, the relative a priori plausibility of each, and the referential consistency of each with the preceding discourse context. But under what circumstances can this be determined on a word to word basis ?" (Frazier (1987), p. 13).

Relevance theory gives us an insight into how this can be done: in the communication process hearer AND speaker aim at optimal relevance. Crain and Steedman's principles do not have to be proposed as seperate principles, because (pragmatic) plausibility, referential success and parsimony simply follow from the principle of relevance. Let us consider the example Frazier uses to argue against a principle of plausibility as guiding the online analysis in the interpretation process:

"*With respect to plausibility differences, it is often unclear what would count as a sufficient difference in the likelihood of different readings to permit early resolution of ambiguity. For example, in (15), is the difference in the likelihood of "answer" vs. "duck" as the direct object of "knows" sufficiently great that a simple object analysis is selected in (a), but a sentential complement reading in (b) ? And it is unclear how the parser could (even in principle) evaluate the relative plausibility of the direct object and sentential complement reading without yet knowing the identity of the embedded verb ? Clearly it cannot simply assume, say, that the higher a phrase is on some animacy hierarchy, the more plausible it is as a subject; this might work in (15), but not in general (e.g., consider (16)).*

(15) a. John knew the answer (...)
b. John knew the duck (...)

(16) a. John heard the answer (...)
b. John heard the duck (...)".
(op.cit., p.13).

Frazier is perfectly right when she says that we cannot use a principle of a priori plausibility to decide which analysis should be made of the different sentences in (15) and (16). The point is, however, that Frazier is using 'sentences' here, isolated objects, and not 'utterances'. Frazier does not use (15) and (16) to communicate propositions to her addressees. In other words, Frazier is not aiming at optimal relevance as far as the content of these sentences is concerned, she does not accomodate semantic, syntactic and lexical choices to the processing needs of her addressees; in fact, she does exactly the opposite. A communicator actually uttering one of the sentences in (15) and (16) will only do so when s/he thinks the utterance is the most relevant one s/he can use in order to communicate the proposition s/he wants to convey, i.e.,

when it fulfills the processing needs of the addressee best.

The findings of Crain and Steedman concerning 'garden path' sentences follow straightforwardly from the model proposed here: a communicator aiming at optimal relevance will never knowingly lead her/his addressee 'up the garden path', and in order to prevent this, s/he must supply a context in which it is relevant for the addressee to go for the intended interpretation. For example, when processing an utterance of sentence (6), the logical processor will, when accessing the logical entry for HORSE, build the logical hypothesis:

((THE (HORSE (SOME PROPERTY)))(SOME PREDICATE)(SOME TIME)).

If there is nothing in the context that makes it relevant for the addressee to know which horse is referred to (as it would be the case when this sentence is processed in isolation), the optional variable SOME PROPERTY will not be 'activated', and the logical processor will insert RACED in the SOME PREDICATE-slot; on the other hand, if there is something in the context that makes it relevant for the addressee to know which horse is referred to (e.g., because different horses are raced in different places at the same time), this will activate the SOME PROPERTY variable, causing the logical processor to insert RACED in the SOME PROPERTY-slot. This approach does not imply that misinterpretations never happen. As Sperber and Wilson (1986) point out, communicators aiming at optimal relevance may not always succeed. It helps to explain, however, how human beings can achieve successful communication in the face of 'garden path' effects and ambiguity.

5. Conclusions.

In this paper, I have examined the consequences that Relevance theory has for a model of input processing. Relevance theory claims that human communication is aimed at achieving optimal relevance. It follows from this that a communicator will try to make her/his utterance worth the addressee's while, by accomodating semantic, syntactic and lexical choices to the processing needs of the addressee. This has as a consequence that we cannot view the process of interpreting natural language as a 'context-neutral' process, with the context added: the choice of a particular utterance is a consequence of the context in which it is uttered. To account for this I propose a 'relevance driven' model, which explains input processing in terms of the interaction of input and

context, guided by the principle of relevance, in that the addressee builds anticipatory logical hypotheses on the basis of this principle. Frazier (1987) says that : "(...) today there is, in [her, mg.] assessment, overwhelming evidence that perceivers do grammatically structure a linguistic input during comprehension." (Frazier (1987), p.2). In my view, the assumption that language processing can be studied without taking the role of the communicator into account, seems to have had consequences for the way in which much psycholinguistic research has been conducted, for the selection of experimental stimuli, and for the interpretation of research findings. I would argue that the evidence for a syntactic parser is in fact dependent upon this assumption. The 'relevance driven' model shows that one need not postulate an autonomous parser. All an addressee needs to recover is the s-selection frame which is stored in the logical entry of a concept. This s-selection frame, or part of it, can then be used to build a logical hypothesis. The model I propose can account for experimental findings, which have presented problems for processing models with a syntactic parser, e.g., multiple centre embedded sentences. Moreover, pragmatics plays an integral role in this processing model: it is no longer grafted on to the model as a psycholinguistic afterthought.

References

Blakemore, D. (1987) Semantic Constraints on Relevance. Oxford: Blackwell.

Carston, R. (1988) Language and cognition. In F.J. Newmeyer (ed.) Linguistics: The Cambridge Survey, Volume 3. Cambridge: University Press.

Chomsky, N. (1981) Lectures on Government and Binding. Dordrecht: Foris.

Chomsky, N. (1986) Knowledge of Language: Its Nature, Origin and Use. New York: Praeger.

Crain, S. and Steedman, M. (1985) On not being led up the garden-path: The use of context by the psychological parser. In D Dowty, L. Karttunen and A Zwicky (eds.) Natural Language Parsing. Cambridge: University Press.

Fodor, J.A. (1981) Representations. Hassocks: Harvester Press.

Fodor, J.A. (1983) The Modularity of Mind. Cambridge, Mass.: MIT Press.

Frazier, L. (1987) Sentence processing. In M. Coltheart (ed.) Attention and Performance XII. London: Lawrence Erlbaum Associates.

Frazier, L. and Fodor, J.D. (1978) The sausage machine: A new two-stage parsing model. Cognition, 6, pp. 291-325.

Grice, P. (1975) Logic and conversation. In P. Cole and J. Morgan (eds.) Syntax and Semantics 3: Speech Acts. New York: Academic Press.

Groefsema, M. (1988) On the interaction of pragmatic principles and parsing principles: A 'Relevance driven' model. MA thesis, UCL.

Jackendoff, R. (1983) Semantics and Cognition. Cambridge, Mass.: MIT Press.

Johnson-Laird, P. (1983) Mental Models. Cambridge: University Press.

Karttunen, L. and Zwicky, A. (1985) Introduction. In D. Dowty, L. Karttunen and A. Zwicky (eds.) Natural Language Parsing. Cambridge: University Press.

Kempson, R. (1986) Reconstruction and Logical Form. (Ms.).

Kimball, J. (1973) Seven principles of surface structure parsing in natural language. Cognition, 2, pp. 15-47.

Sperber, D. and Wilson, D. (1986) Relevance: Communication and Cognition. Oxford: Blackwell.