

Processing for Relevance

Marjolein Groefsema
University of Hertfordshire

ABSTRACT

One of the basic claims of relevance theory is that because a communicator asks for the attention of the addressee, the addressee is entitled to assume that the communicator is trying to be relevant, and interpret an utterance according to this expectation. This paper addresses the question of what the consequences of this claim are for a model on on-line input processing, and proposes a model in which addressees interpret utterances in accordance with the principle of relevance, by building anticipatory hypotheses about the overall propositional form of the utterance. It is shown how this proposal can account for psycholinguistic findings concerning disambiguation, and for the interpretation of centre embedded sentences.

1. Introduction

In most existing psycholinguistic models of sentence input processing little or no account is taken of the fact that the use of natural language is a process of interaction between communicator and addressee. Many psycholinguists have tried to explain input processing purely in terms of the addressee's activity, without taking into account that the communicator is responsible for the input. However, relevance theory (Sperber & Wilson, 1986, 1995) shows that natural language use is more than the use of a 'context-neutral' language with the context added: the choice of a particular utterance is a consequence of the context in which it is uttered. Sperber & Wilson argue that communicators intend their audience to believe that they are worth paying attention to, and that addressees only pay

attention to information which seems relevant to them, so that a communicator will aim at optimal relevance (i.e. the communicator will aim at causing the addressee to compute an adequate range of contextual effects for a minimal amount of processing effort). It follows from this that a communicator will try to keep processing cost down by accommodating her/his choice of linguistic output to the processing needs of the addressee. This view of communicator/addressee interaction has consequences for a model on input 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.

Based on this view, I will present an account of some of the mental mechanisms and processes that take the addressee from a linguistic input to the full pragmatic interpretation of that input. In section (2) the proposals made by relevance theory with regard to on-line interpretation are discussed, and it is shown how the proposal that during on-line interpretation we make anticipatory hypotheses concerning the propositional form of an utterance can be made more precise. On the basis of this, section (3) introduces an account of on-line processing which takes as its starting point that the interpretation an addressee selects is a consequence of expectations of relevance. A refinement of the notion of 'processing effort' is introduced and it is shown how on the basis of this we can account for experimental findings concerning disambiguation, and for the various degrees of difficulty experienced in processing multiple centre-embedded sentences.

2. Relevance and on-line processing

The psycholinguistic literature on (post-lexical) processing is primarily concerned with how linguistic syntactic representations are constructed on-line. However, linguistic syntactic representations are at most half-way stations in the interpretation process. In order to account for how people recover the full semantic/pragmatic meaning of an utterance, we need, at least, an account of how an addressee computes conceptual representations on the basis of the linguistic input. There is, as yet, little psycholinguistic theorisation of pragmatic aspects of interpretation. This may be partly due to the fact that pragmatics has been mostly concerned with implicatures, rather than with working out the propositional content of an utterance (e.g. Grice, 1975; Levinson, 1983). However, Wilson & Sperber (1981) and Sperber & Wilson (1986, 1995) show that pragmatics does not only play a role in recovering what is implicated, but also in recovering what is explicitly communicated. Processes such as reference assignment, disambiguation, and recovery of elliptical material involve interaction between input and context, guided by pragmatics.

The basic claim of relevance theory is that in processing information, people try to achieve the greatest cognitive effect for the smallest amount of processing effort, i.e. that human cognition is driven by relevance and the maximisation of relevance, where relevance to an individual can be defined as follows:

"Extent condition 1: an assumption is relevant to an individual to the extent that the contextual effects achieved when it is optimally processed are large.

Extent condition 2: an assumption is relevant to an individual to the extent that the effort required to process it optimally is small." (Sperber & Wilson, 1986: 145).

Sperber & Wilson propose that there are three types of contextual effects: strengthenings, where new information causes a person to have more confidence in an assumption already (weakly) held; contradictions, where new information contradicts an existing assumption, causing the existing assumption to be eliminated; and contextual implications, which are implications derived from a new assumption together with assumptions in the context.

To account for how this general notion of relevance helps explain ostensive communication, Sperber & Wilson propose that utterance interpretation is governed by the Principle of Relevance, which states that:

"Every act of ostensive communication communicates the presumption of its own optimal relevance." (158).

An utterance is optimally relevant when it enables the addressee to derive an adequate number of contextual effects for no unjustifiable processing effort.

In order to derive contextual effects, a context has to be found against which the utterance is to be processed. Because relevance theory proposes that an utterance communicates the presumption of its own optimal relevance, the addressee can assume that the relevance of an utterance is given, and therefore need not be assessed. The task of the addressee is rather to select a context which bears out this guarantee of relevance of the utterance. How is this context selected? Sperber & Wilson propose that at the start of processing some new item of information there is an initial context consisting of assumptions left over in the memory of the deductive device from the immediately preceding deductive process. This initial context can be extended in different ways during the interpretation process. One way of extending the context is to add assumptions used or derived in previous deductive processes. A second way is to add assumptions stored under the encyclopaedic entries of concepts already present in the context or in the assumption being processed (see Groefsema (1995a)). A third way of extending the context is to add to it information about the immediately observable environment. However, extending the context involves processing effort, which means that an addressee cannot freely access all kinds of different extensions, because this would diminish the overall relevance of the assumption being processed.

This proposal gives the following picture of the comprehension process. The linguistic input maps onto an (incomplete) conceptual representation (a non-linguistic logical form). At points where the logical form is incomplete (for example, where reference has to be assigned), this logical form has to be enriched into a fully propositional form. This logical form does not have to be recovered completely before any enrichments can take place.


Processes such as reference assignment, disambiguation, and recovery of elliptical material occur on-line, in accordance with the principle of relevance. Given that each cycle of context extension involves additional processing effort, the addressee will select the most accessible value, consistent with the principle of relevance, in each case. It follows from this that the first interpretation found consistent with the principle of relevance is the only interpretation consistent with the principle of relevance: any further interpretation which would yield contextual effects would automatically falsify the second part of the definition of relevance, because it would involve more processing effort.

Sperber & Wilson propose that, in order to save a hearer from going through a lot of fruitless processing, a speaker aiming at optimal relevance phrases her/his utterance in such a way as to facilitate early, and correct disambiguation, reference assignment, and enrichment. This raises the question of how a speaker can anticipate the way in which a hearer is going to interpret an utterance. Sperber & Wilson propose that the hearer makes anticipatory hypotheses about the overall logical form of the utterance, and that s/he resolves potential ambiguities and indeterminacies on the basis of these hypotheses. By virtue of being a hearer as well as a speaker, the speaker can anticipate what hypotheses the hearer is going to make, and structure his/her utterance accordingly.¹

Sperber & Wilson 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.*” (Sperber & Wilson, 1986: 205).

As a simplified illustration of the logical form of a sentence such as (1a) they give (1b):

(1)(a) John invited Lucy



Sperber & Wilson argue that, viewed in this way, the logical form of (1a) carries the information that John did something, that John invited someone, etc.; in other words, a number of analytic implications of (1a) can be read off the logical form representation directly.

The question then arises how hearers construct these anticipatory logical hypotheses. One possibility (put forward by Sperber & Wilson) is to postulate that the hearer builds these anticipatory logical hypotheses on the basis of anticipatory syntactic hypotheses, by variable substitution. However, this proposal faces a number of problems. In the first place, it is not clear on what basis and how the syntactic hypotheses are formed, whether these are built for

each word, for example, or only for major phrases (see Groefsema (1992) for an extended evaluation of this problem). In the second place, it is not clear which variable will be substituted for which syntactic categories. For example, on encountering (2) we can build the syntactic hypothesis that *put* can be followed by an NP and a PP:

(2) Mary put the book on the table.

However, a PP can express concepts of different sorts of conceptual categories (cf. Jackendoff, 1983, 1990), such as a property 'with big ears,' a time 'at five o'clock,' a manner 'by car,' and a place 'on the table.' In principle, these different categories could give rise to different logical hypotheses. The problem with this is that often these hypotheses would be superfluous, in that they would never actually be realised, and moreover, they miss the point that by virtue of its meaning *put* involves a place where something is put.

Moreover, the linguistic form of an utterance often underdetermines the conceptual representation the addressee constructs on the basis of the utterance (e.g. Blakemore, 1987; Carston, 1988; Groefsema, 1992, 1995a; Sperber & Wilson, 1986, 1995), as in the following examples (Carston, 1988: 165, 167), where we understand the linguistic input in (3a) and (4a) along the lines of (3b) and (4b) respectively:

(3)(a) He ran to the edge of the cliff and jumped.

(b) He ran to the edge of the cliff and jumped over the cliff.

(4)(a) Susan's performance isn't good enough.

(b) The way in which Susan performs something is not good enough for something.

Nor is the structure of the natural language input and the structure of the conceptual representation always isomorphic, as in examples (5) and (6) (Pustejovski, 1991: 424, 426):

(5) John began a novel.

(= John began to read a novel,
or: John began to write a novel)

(6) John knows the plane's arrival time.

(= John knows what time the plane will arrive)


The proposal that anticipatory logical hypotheses are based on anticipatory syntactic hypotheses cannot account for how we formulate the right logical hypotheses for any of these examples.

The view that logical hypotheses are conceptual structures rather than linguistic structures allows for an alternative proposal. Postulating that the mind computes structured conceptual representations entails that information about the syntax of these representations

has to be stated somewhere. Sperber & Wilson propose that concepts consist of conceptual addresses which give access to different types of information, lexical, logical and encyclopaedic. This view of concepts allows for the proposal that one type of information that is stored under a conceptual address is information concerning what type of concepts the concept has to or can combine with to yield a well-formed conceptual structure, as a logical selection frame. Postulating this has as a consequence that we can account for how anticipatory logical hypotheses are built without having to appeal to anticipatory syntactic hypotheses: the phonological form of a lexical item maps onto a conceptual address (or different conceptual addresses in the case of ambiguity), which gives access to the logical entry of that concept. From this the addressee can recover the selection frame of the concept, and use this selection frame as an anticipatory hypothesis about the logical form of the utterance.

Although Sperber & Wilson postulate that there may be a set of conceptual categories which is part of basic human equipment, they do not specify what these conceptual categories are. Jackendoff (1983, 1990) proposes and defends a set of basic conceptual categories: THING, EVENT, STATE, ACTION, PLACE, PATH, PROPERTY, AMOUNT, and MANNER. He argues that we need to distinguish these categories as entities at the conceptual level because we can quantify over instances of them, point at instances of them, and question instances of them. Jackendoff leaves it open whether a category PREDICATE should be distinguished. However, Kempson (1988) shows that we need to postulate a logical variable corresponding to a VP to account for VP anaphora (see also Groefsema (1995b) for arguments that we need a category PREDICATE). Adopting these categories gives us a basis for making Sperber & Wilson's proposals more precise.²

Phonological forms of lexical items then map onto conceptual addresses which give access to different types of information, including selection information. For example, the conceptual address for the verb *open* gives access to the selection frame which specifies what the concept can combined with to form a complex concept:^{3, 4}



One can open a thing in a certain manner, for example, you can open a door quickly, which means that the PREDICATE can contain a MANNER to yield a complex PREDICATE, although it does not have to. Moreover, the event expressed by *open* can have a PROPERTY, such as that the opening of the door was done with a key, which means that the EVENT can contain a PROPERTY although it does not have to.⁵

The conceptual address for *open* also gives access to a selection frame which specifies what the concept has to combine with to form a well-formed conceptual structure (i.e. a propositional form):

$$\left[\begin{array}{l} \text{[PLACE][TIME][EVENT_]} \\ \text{PROPOSITION} \end{array} \right]$$

Similarly, the conceptual address for *girl* will give access to the following selection frame:

$$\left[\begin{array}{l} \text{GIRL} \\ \{\{\text{PROPERTY}\}\} \\ \text{THING} \end{array} \right] \text{a}$$

This frame tells us that the concept GIRL belongs to the category THING, which has to be individuated in some way to occur as a constituent in conceptual structure (indicated by the index a). Moreover, it tells us that GIRL can occur with a PROPERTY, although it does not have to. PROPERTIES modifying THINGS play a different role from, for example, PROPERTIES modifying EVENTS. Whereas PROPERTIES modifying EVENTS change the nature of the EVENT involved, PROPERTIES modifying THINGS help individuate the instance of the THING at stake. PROPERTIES modifying THINGS can be either concepts of the conceptual category PROPERTY, or they can be propositions, expressing an assumption about the THING at stake.

The conceptual address for *girl* will also give access to a second selection frame which specifies what the concept has to combine with to form a well-formed conceptual structure (i.e. a propositional form):

$$\left[\begin{array}{l} \text{[PLACE][TIME][EVENT/STATE][PREDICATE][THING_]} \\ \text{PROPOSITION} \end{array} \right]$$

We can make a distinction between concepts that occur in conceptual structure and concepts that do not. This second type of concept has the function of constraining the interpretation of the utterance in some way. Within this second group of concepts we can distinguish different sorts of concepts. What they have in common is that none of them gives access to encyclopaedic information: their 'meaning' is exhausted by how they constrain the interpretation process. This means that they do not have a further role to play in general cognitive processes, and consequently they do not occur in logical forms. For example, complementizer *that* signals that a proposition is following, by which it constrains the possible logical hypothesis that can be built. In other words, *that* maps onto a conceptual address which gives access to the selection frames:

$$\left[\text{PROPOSITION} \right]$$

$$\left[\begin{array}{l} \text{[PLACE][TIME][EVENT/STATE][PREDICATE][PROPOSITION_]} \\ \text{PROPOSITION} \end{array} \right]$$

The first frame simply says that a PROPOSITION is following, and the second frame specifies that it has to combine with certain categories to yield a well-formed conceptual structure.

We also find concepts which do not themselves occur in conceptual structures, but whose 'meaning' is partly procedural, i.e. they give a procedure for the interpretation of some conceptual category (see e.g. Wilson & Sperber, 1993). An example of this is the indefinite article. What the indefinite article does is signal that an instance of a THING is at stake, i.e. it has associated with it a procedure for setting up a concept of an instance of some THING. In line with the proposals made here, we can say that the conceptual address of the indefinite article will give access to the following selection frames:

[_{THING}]
 (Procedure: create instance)

[[_{PLACE}] [_{TIME}] [_{EVENT/STATE}] [_{PREDICATE}] [_(THING _)]]
 PROPOSITION

Similarly, Blakemore (1987) shows that discourse connectives should be analyzed as encoding procedures for the pragmatic computations a proposition may enter into. For example, she argues that the meaning of *so* is an instruction to interpret the proposition it introduces as a logical consequence of some other proposition. We can view discourse connectives as mapping onto concepts which give access to selection frames, which in turn give access to the procedure associated with the concept.

3. A relevance-driven account of input processing

I propose an account of on-line comprehension, in which the interpretation process is driven by the principle of relevance, and which takes into account that the input for comprehension is never 'neutral', but is produced by a communicator aiming at optimal relevance.

In this model a linguistic processor is responsible for recognizing the phonological (or orthographic) form of a lexical item. The phonological (or orthographic) forms give access to the concepts associated with them, which in turn give access to the selection frames associated with the concepts. Based on these selection frames, anticipatory hypotheses are made regarding the logical form of the utterance. These 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.

3.1. Processing effort

As we have seen above relevance theory predicts that an addressee selects the interpretation which yields adequate contextual effects for as little processing effort as possible. The

different sorts of contextual effect are clearly defined by Sperber & Wilson. The question then is what constitutes processing effort. A partial answer to this question is provided by the relevance theory notion of accessible contexts. Relevance theory proposes that the interpretation of a linguistic input depends on its interaction with assumptions (the context) which are either already held in memory, or can be constructed from assumption schemas, to yield contextual effects. These assumptions and assumption schemas have to be retrieved from memory. The more accessible an assumption is, the easier it is to recall from memory. Relevance theory proposes that contexts are ordered according to accessibility, and that extending a context is a cyclic process, i.e. only some extended contexts are accessible from the immediate context, but these extended contexts make further extensions accessible, which in turn make further extensions accessible, etc. Actually accessing a context involves processing effort, and each step of context extension involves more processing effort, so that from this we can conclude that the fewer steps of context extension needed to yield an interpretation with adequate contextual effects, the smaller the processing effort.

However, this is not all that needs to be said about processing effort. In many psycholinguistic experiments, ambiguous sentences such as in (7), and 'garden-path sentences' such as in (8), are presented to subjects in isolation, i.e. the 'null context':

- (7)(a) Joyce said Tom left yesterday.
 (b) John hit the girl with a book.

- (8)(a) Since Jay always jogs a mile seems like a short distance to him.
 (b) The horse raced past the barn fell.

When these sentences are presented to subjects in isolation, they are not genuine acts of ostensive communication, i.e. they are not used to communicate propositions to the subjects. This means that selecting one rather than another interpretation is not going to give the subject more contextual effects, for example, knowing that Joyce said something yesterday does not give the subject more contextual effects than knowing that Tom left yesterday, if the subject does not know who Joyce and Tom are. Moreover, by formulating these sentences the researchers are not aiming at optimal relevance, i.e. they are not accommodating the choice of sentence to the processing needs of the subjects. A communicator actually uttering one of the sentences in (7) and (8) would only do so if the utterance is the most relevant s/he could use in order to communicate the proposition s/he wants to convey, i.e. when it accommodates the processing needs of the addressee best.. This means that experiments in which sentences are processed in isolation do not necessarily reflect normal utterance interpretation: the initial context in which the sentence is processed does not contain any assumptions that have a bearing on the interpretation, and the interpretation of the sentence will yield no or very little contextual effects. Even so, it has been found that ambiguous sentences, as in (7), have a favoured interpretation when processed in isolation, i.e. those in (9) (e.g. Ferreira & Clifton, 1986; Frazier, 1987):

- (9)(a) Joyce said (Tom left yesterday).
 (b) John (hit (the girl)(with a book)).

And the 'garden-path' sentences, as in (8), are called so because they cause subjects to select the wrong analysis. If we cannot account for this in terms of contextual effects, we will have to account for it in terms of processing effort. However, the difference in processing effort between the different analyses cannot be explained in terms of numbers of context extensions involved in the different interpretations, because extending the context does not make one interpretation more relevant than the other.

We can explain this difference in processing effort by postulating that 'accessibility' applies to conceptual slots in logical hypotheses as well as to assumptions in the context. A conceptual slot is the most accessible (i.e. immediately accessible) when it is being inserted in a logical hypothesis, when it is being filled, and when a constraint is put on its interpretation. When new conceptual slots are inserted into the logical hypothesis they in turn become more accessible. We can think of accessibility in terms of activation: when a concept or conceptual slot is accessed it receives activation, thereby becoming more accessible, but when new concepts or conceptual slots are accessed activation on the earlier ones fades away. This proposal does not mean that we can equate accessibility with 'having been postulated most recently' *per se*, because two conceptual slots of the same type may have the same level of activation if they are further up-stream in a logical hypothesis, even though one will have been postulated more recently than the other.

Optional slots differ from the other slots in a logical hypothesis in that they do not have to be filled by any conceptual material in order to yield a well-formed representation: they merely show what a concept can combine with to form a complex concept, or, as is the case with PROPERTIES occurring with THINGS, they can help the addressee to pick out the THING at stake. This means that they do not become immediately accessible on being inserted in the hypothesis. Rather, optional slots become accessible either when the addressee recovers a concept of the same type which cannot be inserted into a non-optional slot, or they become accessible because they are activated by something in the context.

Given this account of accessibility of conceptual slots, we can say that filling a conceptual slot involves least processing effort if the conceptual slot is immediately accessible.


3.2. Disambiguation

On the basis of the proposal that processing effort is defined both in terms of the accessibility of assumptions and the accessibility of conceptual slots in logical hypotheses, we can account for why the sentences in (7) and (8) are interpreted the way they are, when processed in isolation. As we saw above, when (7a) is processed in isolation, it is interpreted as (9a) rather than (9c):


- (9)(a) Joyce said (Tom left yesterday).

(c) Joyce said (Tom left) yesterday.


We can account for this preference as follows: when the addressee encounters *Joyce*, this will give access to the selection frames:




Because the addressee has not got a hypothesis yet about the logical form of the utterance, both selection frames need to be accessed. These selection frames yield the hypothesis:



Said gives access to the selection frames:





Since the addressee already has an anticipatory hypothesis available, only the first frame needs to be accessed, and the information can be unified with the existing hypothesis, giving rise to a further hypothesis. This is in accordance with the principle of relevance, since accessing both selection frames would give rise to extra processing effort, which would not be offset with an increase in contextual effects (in fact, the information in the second selection frame cannot be incorporated into the current hypothesis at all). Moreover, the past tense constrains the TIME to a time before time of utterance:




Tom gives access to the selection frames:



Because there is no THING slot available in the hypothesis, the first selection frame can't be incorporated in it, and the second frame needs to be accessed. Because there is a PROPOSITION slot available in the hypothesis, the information in the second frame can be unified with the existing hypothesis, yielding a further hypothesis:⁶



On encountering *left* the PREDICATE slot can be filled by the information in its selection frame, and the past constrains the TIME slot in the subordinate PROPOSITION to a time before the time of utterance. This yields the following hypothesis:






On encountering YESTERDAY, this will be fitted into the most accessible TIME slot. The TIME slot in the subordinate PROPOSITION is the TIME slot postulated most recently, and moreover the tense of LEAVE has put a constraint on its interpretation immediately before YESTERDAY is encountered, so that it is more accessible than the TIME slot in the main clause. Because of this, the processor will not even consider the TIME slot in the main clause. Only if YESTERDAY is incompatible with conceptual material already available, for example in the case of *Joyce said Tom is leaving yesterday*, or if it is incompatible with assumptions in the context, will this assignment be rejected, and the TIME slot in the main PROPOSITION tried. Because rejecting the first interpretation, and selecting the TIME slot in the main PROPOSITION would involved more processing effort, which is not offset by an increase in contextual effects, *Joyce said (Tom left yesterday)* is the most relevant interpretation, and therefore the interpretation that the addressee selects when processing (7a) in isolation.

When (7b), repeated here, is processed in isolation, it is interpreted as (9b) rather than as (9d):


- (9)(b) John (hit (the girl)(with a book)).
- (d) John (hit (the girl with a book)).

Again, we can account for this by looking at how the logical hypothesis for this sentence is built. By the time the addressee encounters *with*, s/he has the following hypothesis available:



When the addressee accesses the selection frame of *with*, s/he will find that it introduces a PROPERTY involving a THING. This then makes both the optional PROPERTY slots (accompanying GIRL and HIT) in the logical hypothesis accessible. Why then would the

PROPOSITION₂



Because the addressee cannot choose between these two hypotheses at this stage, s/he will maintain both. These hypotheses will be processed in parallel, until a choice can be made between them. When the addressee encounters the indefinite article *a*, this will give access to the logical selection frames:

[THING]
 (Procedure: create instance)


[[PLACE] [TIME] [EVENT/STATE] [PREDICATE] ([THING -])]
 PROPOSITION

These selection frames show both that an instance of a THING is following, and what that THING has to combine with to yield a PROPOSITION. Why then does the addressee select the interpretation of the THING as the direct object, rather than as the subject of the second PROPOSITION? As we have seen, the addressee only uses information from both selection frames if s/he does not have a logical hypothesis yet, or if the information in the first selection frame cannot be fitted into the hypothesis. In this case, the first selection frame yields the information that an instance of a THING is following. Although there is no THING slot in hypothesis 1, there is in hypothesis 2 into which the instance of THING can be fitted. Because the indefinite article signals that an overt THING is following, and there is no comma to indicate that the first proposition is complete, there is no reason for the addressee to assume that the THING signalled by the indefinite article should be taken as the beginning of the second proposition rather than as the THING following JOG.


Let us now look at why subjects get 'garden-pathed' with sentences such as (8b):

(8)(b) The horse raced past the barn fell.

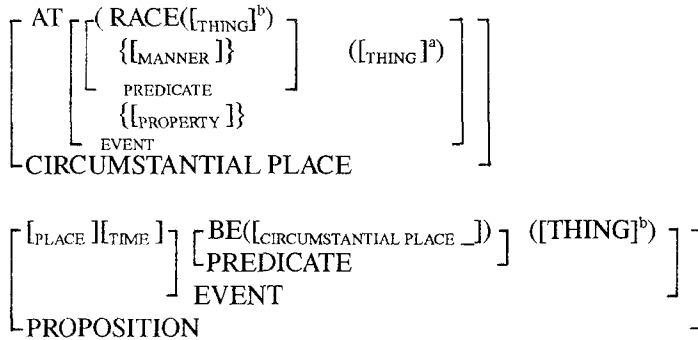
By the time the addressee encounters *raced*, s/he will have the hypothesis:



Raced will map onto a number of different concepts. The past reading of (7b) will give access to the selection frame RACE₁:



The passive reading will give access to the selection frames for RACE₂:



In principle, the addressee has a choice between fitting RACE₁ into the PREDICATE slot in the hypothesis, and fitting RACE₂ into the PROPERTY slot in the hypothesis. However, the PREDICATE slot has to be filled with conceptual material in order to yield a PROPOSITION, while the PROPERTY slot is an optional slot. This means that the PREDICATE slot is accessible, while the PROPERTY slot only becomes accessible if the addressee recovers a concept which cannot be inserted into a slot which has to be filled in order to yield a proposition, or if assumptions in the context make the resulting interpretation relevant. In this case, the first selection frame of RACE₁ can be inserted into a slot which has to be filled in order to yield a proposition, and moreover, there are no assumptions in the context which make inserting RACE₂ in the PROPERTY slot of HORSE

relevant, so that the addressee will choose the assignment of RACE₁ to the PREDICATE slot.

What we see then is that we can define processing effort not only in relation to contexts, but also in relation to logical hypothesis formation and completion. In relation to context extension we saw that the fewer steps of context extension needed to yield an interpretation with adequate contextual effects, the smaller the processing effort; in relation to filling conceptual slots in a logical hypothesis we saw that the more accessible a conceptual slot, the smaller the processing effort involved in recovering it. In actual utterance interpretation, the amount of processing effort involved in computing an interpretation will depend on the interaction of these different processes.

3.3. Multiple centre-embedded sentences

Different explanations have been proposed to account for the difficulty of processing multiple centre-embedded sentences, such as:

- (10) 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 complementizers can make sentences difficult to understand: the complementizer signals that a new phrasal node should be started. However, the difficulty in processing sentences like (10) does not seem to be due to the absence of complementizers. Newmeyer (1983) points out that multiple centre-embedded constructions which complementisers are unacceptable:

- (11) The cheese that the rat that the cat chased ate was rotten.

Newmeyer argues that we can explain the unacceptability of multiple centre-embedded sentences by combining the competence model of generative grammar with a model of immediate memory storage. However, this proposal does not give us an explanation of why the sentence in (10) and (11) are unacceptable, while Smith's (1989) example in (12) is fine, even though it has the same multiple centre-embedded structure:

- (12) The game those boys I met invented resembles chess.

Moreover, it does not give us an explanation of why there is a gradient of unacceptability, so that Smith's (1989) example in (13) is more unacceptable than (10) and (11) even though it involves fewer words:

- (13) Oysters oysters oysters split split split.

We can explain this phenomenon by looking at what is involved in the interpretation of these different sentences, i.e. how the addressee recovers 'who did what to whom'. As we shall see it turns out that the concepts referred to in the different sentences constrain the interpretation to a greater or lesser extent. Let us consider (10). On encountering *The* the addressee recovers the selection frame for THE:

[THING]
 (Procedure: accessible instance)

$$\left[\begin{array}{l} \text{[PLACE] [TIME] [EVENT/STATE] [PREDICATE] ([THING _])} \\ \text{PROPOSITION} \end{array} \right]$$


On the basis of this the addressee can set up a hypothesis, into which RAT can be fitted next:

$$\left[\begin{array}{l} \text{[PLACE] [TIME]} \\ \text{PROPOSITION} \end{array} \left[\begin{array}{l} \text{[PREDICATE]} \\ \text{EVENT/STATE} \end{array} \left(\begin{array}{l} \text{[RAT} \\ \text{[PROPERTY]} \\ \text{THING} \end{array} \right)^a \right] \right]$$

On encountering *the* again, the addressee recovers that an accessible instance of a THING is following. Because there is no THING slot in the hypothesis, the addressee has to access the second selection frame of THE, which shows a PROPOSITION can be formed. Because RAT can take a PROPOSITION as a PROPERTY, this information can be fitted into the hypothesis. Because this PROPOSITION contains a THING slot, CAT can be fitted into it. This process is repeated on encountering *the dog*, while BITE will be fitted into the most recently postulated PREDICATE slot. The THING slot following BITE in the hypothesis can be filled by CAT, as the CAT is the most accessible THING, by virtue of the PROPOSITION being a PROPERTY of the CAT, yielding the following hypothesis:

$$\left[\begin{array}{l} \text{[PLACE] [TIME]} \\ \text{PROPOSITION}_1 \end{array} \left[\begin{array}{l} \text{[PREDICATE]} \\ \text{EVENT/STATE} \end{array} \left(\begin{array}{l} \text{[RAT} \\ \text{[PROPOSITION}_2 \text{]} \\ \text{THING} \end{array} \right)^a \right] \right]$$

$$\left[\begin{array}{l} \text{[PLACE] [TIME]} \\ \text{PROPOSITION}_2 \end{array} \left[\begin{array}{l} \text{[PREDICATE]} \\ \text{EVENT/STATE} \end{array} \left(\begin{array}{l} \text{[CAT} \\ \text{[PROPOSITION}_3 \text{]} \\ \text{THING} \end{array} \right)^b \right] \right]$$



This logical hypothesis presents the addressee with a number of problems. The addressee is faced with finding or setting up referents for the 'accessible instances' of RAT, CAT and DOG. The PROPOSITION modifying RAT should help the addressee pick out the intended referent, but since that PROPOSITION itself contains an accessible instance of CAT, finding a referent for RAT becomes dependent on finding a referent for CAT. This is repeated in that the PROPOSITION modifying CAT should help the addressee pick out the intended referent, but since that PROPOSITION itself contains an accessible instance of DOG, finding a referent for CAT is made dependent on finding a referent for DOG. There is no dog in the context, nor does accessing the encyclopaedic entries of RAT, CAT and DOG yield any accessible dog, so that the addressee cannot find any referent.

On encountering *chased* the addressee has two PREDICATE slots available into which the concept could be fitted, both of which are equally accessible. Encyclopaedic entries of the concepts encountered do not constrain the interpretation: it could equally well be the cat as the rat who did the chasing, and it could even be the dog, performing a number of different actions. Moreover, the cat, the rat and the dog could all be chased. When the addressee encounters *ran away*, this could again apply equally well to the cat as the rat, and again encyclopaedic entries of the concepts involved do not constrain the interpretation. What we see then is that it is not due to the repetition of structurally undifferentiated phrases that the human processor is too confused to cope, as Smith (1989) proposes, but rather that this is due to there not being any basis on which to decide what reference should be assigned to the different animals, and to there not being any basis on which to decide what concepts should go where in the logical hypothesis.

When we compare this with the interpretation of (12), a different picture emerges:

(12) The game those boys I met invented resembles chess.

Building a logical hypothesis for (12) occurs along the same lines as for (10). The addressee is faced with similar problems as in processing (10), in that s/he has to find or set up a referent for GAME. The PROPOSITION modifying GAME should help the addressee pick out the intended referent, but since that PROPOSITION itself contains BOYS, finding a referent for GAME becomes dependent on finding referents for BOYS. However, from this stage the interpretation of (12) differs from that of (10). The PROPOSITION modifying BOYS should help the addressee pick out the intended referent, and in this case the PROPOSITION does constrain the interpretation of BOYS. In the first place, *I* maps onto

a THING which is constrained by it being the communicator. This means that even when (12) is processed in isolation, the addressee can set up a (partial) conceptual representation for the THING. Moreover, *meet* can be easily disambiguated to MEET THING. Although again the THING could be either GAME or BOYS, the addressee can recover that games are not usually met by anyone, but that people are, so that the indeterminacy can be resolved in favour of BOYS. This means that the addressee now has a constraint on the interpretation of BOYS, namely that the BOYS at stake are boys that the communicator met. When INVENT THING is encountered, there are two PREDICATE slots available into which it could be fitted, and a filler for THING has to be found. However, the addressee can easily recover that people are not usually invented, but that they do invent things or ideas, so that the indeterminacy can be resolved in favour of BOYS INVENT GAME. This in turn gives the addressee a constraint on the interpretation of GAME, namely that the GAME at stake is a game invented by the boys that the communicator met. When RESEMBLE THING is recovered, there is only one PREDICATE slot left into which it can be fitted, which yields GAME_i RESEMBLES THING, and this hypothesis is confirmed by CHESS, since chess is a game as well.

What we see then is that we can account for the differences in processing difficulty of multiple centre-embedded sentences by looking at how the kinds of concepts involved together with encyclopaedic information concerning those concepts help or hinder completion of the logical hypotheses into a propositional form.

Conclusion

Sperber & Wilson (1995: 206) say that they see anticipatory logical hypotheses "*as playing a crucial role in disambiguation and reference assignment.*" In this paper it has been shown that proposals made by relevance theory can indeed be translated into an account of on-line sentence processing. It has been shown how logical hypotheses can be built on-line on the basis of information concerning the ways in which concepts can combine to form well-formed propositional forms. As we have seen, this account makes the right predictions with regard to experimental findings concerning ambiguous and 'garden-path' sentences, and with regard to multiple centre-embedded sentences.

Notes

1. This does not mean that a speaker always will be optimally relevant, for example a speaker may make the wrong assessment of what assumptions are immediately accessible to the addressee.

2. Although Jackendoff's Conceptual Semantics has been criticised for not being able to deal with such issues as quantification and scope phenomena, Zwarts & Verkuyl (1994) show that Jackendoff's system of representation can be given a model-theoretic treatment. They say about this that: "... it turned out that CS can be interpreted in a way that gives it the same expressive

power as a first-order logic with other sorts of entities besides individuals." (26).

In this paper I will make use of informal representations.

3. Jackendoff gives verb meanings in terms of lexical decompositions, while relevance theory proposes that concepts are simple and unanalyzable. One can adopt Jackendoff's conceptual categories without thereby committing oneself to a decompositional view of concepts, as will be done here. For a discussion of concepts in relevance theory, see Groefsema (1997).

4. Square brackets indicate conceptual constituents, while the label indicates the type of conceptual category. Curly brackets are used to indicate optional conceptual constituents.

5. The kind of property that can occur with a particular concept can be specified using 'assumption schemas' (Sperber & Wilson, 1986).

6. Because of space constraints, I represent the subordinate PROPOSITION separately.

Works Cited

- Altmann, G. & Steedman, M. "Interaction with context during human sentence processing." *Cognition* 30 (1988): 191-238.
- Blakemore, D. *Semantic Constraints on Relevance*. Oxford: Blackwell, 1987.
- Carston, R. "Explicature, implicature and truth theoretic semantics." *Mental Representations: The Interface between Language and Reality*. Ed. R. Kempson. Cambridge: CUP, 1988.
- Ferreira, F. and C. Clifton. "The independence of syntactic processing." *Journal of Memory and Language*, 25 (1986): 348-368.
- Frazier, L. "Theories of sentence processing." *Modularity in Knowledge Representation and Natural-Language Understanding*. Ed. J.L. Garfield. Cambridge, Mass.: MIT Press, 1987.
- Grice, H.P. "Logic and conversation." *Syntax and Semantics*, 3. Eds. P. Cole and J. Morgan. New York: Academic Press.
- Groefsema, M. *Processing for Relevance: A Pragmatically based Account of How we Process Natural Language*. Ph.D. Thesis, University College London, 1992.
- _____. "Understood arguments: a semantic/pragmatic approach." To appear in *Lingua* 96 (1995a).
- _____. "On 'benter' and 'succeive': a new perspective on argument structure." Paper presented at the 7th European summer school in logic language and information, Barcelona, 1995b.
- _____. "Concepts and word meaning." *Proceedings of the University of Hertfordshire Relevance Theory Workshop*. Ed. M. Groefsema. Chelmsford: Peter Thomas and Associates.
- Jackendoff, R. *Semantics and Cognition*. Cambridge, Mass: MIT Press, 1983.
- _____. *Semantic Structures*. Cambridge, Mass.: MIT Press, 1990.
- Kempson, R. "Logical form: the language and cognition interface." *Journal of Linguistics* 24 (1988): 398-432.
- Levinson, S. *Pragmatics*. Cambridge: CUP, 1983.
- Pustejovski, J. "The generative lexicon." *Computational linguistics* 17 (1991): 409-441.
- Sperber, D. & D. Wilson. *Relevance: Communication and Cognition*. Oxford: Blackwell, 1986.
- _____. *Relevance: Communication and Cognition*. 2nd edition. Oxford: Blackwell, 1995.
- Wilson, D. and D. Sperber. "On Grice's theory of conversation." *Conversation and Discourse*. Ed. P. Werth. London: Croom Helm, 1981.

- _____. Linguistic form and relevance. *Lingua* 90 (1993): 1-25.
- Zwarts, J. and H. Verkuyl. "An algebra of conceptual structure: an investigation into Jackendoff's conceptual semantics." *Linguistics and philosophy* 17 (1994): 1-28.