anales de psicología
2000, vol . 16, nº 2, 227-249

Murcia (España). ISSN: 0212-9728

Elaborative Inferences

Mauricio Iza

Departamento de Psicología Básica & ILCLI
Universidad de Málaga. Spain

Jesús Ezquerro

Institute for Logic, Cognition, Language, and Information (ILCLI)
Universidad del País Vasco/EHU. Spain

Abstract: According to psychological experimentation, there is evidence of automatic elaborative inferential activity, as well as of incompleteness in backward inference making (anaphora, causal structuring). McKoon & Ratcliff (1992) have argued for a "minimalist" theory of inference in reading, contrasting it with "constructionist" theories, including theories based on mental or situation model. But, minimalism has mischaracterized text understanding, failing to notice that the extent to which an inferential activity is relevant (in focus) affects the speed and easiness of performing inferences, as it happens with anaphora resolution. Although a theory of discourse processing should be both constructionist and minimalist, we discuss the relationship between a computational theory of inference making and a description of the mechanisms underlying our inferential abilities.
Key words: Inferential activity, reading, minimalism, text understanding, discourse processing, computational theory.

Título: Inferencias elaborativas.
Resumen: De acuerdo con los resultados experimentales, existe evidencia tanto de la actividad inferencial elaborativa automática, como de la incompletud en la realización de inferencias puente o hacia atrás (anáfora, estructuración causal). McKoon y Ratcliff (1992) han argumentado a favor de una teoría "minimalista" para las inferencias durante la lectura, contrastándola con las teorías "constructivistas", incluyendo las teorías basadas en modelos mentales o situacionales. Sin embargo, el minimalismo no ha caracterizado apropiadamente la comprensión de textos, al no atender al hecho de que una actividad inferencial relevante (enfocada) afecta a la velocidad y facilidad de la realización de inferencias, tal como sucede en la resolución anafórica. Aunque una teoría de procesamiento del discurso debería ser tanto constructivista como minimalista, se discute la relación entre una teoría computacional de la realización de inferencias y una descripción de los mecanismos subyacentes a nuestras habilidades inferenciales.
Palabras clave: Actividad inferencial, lectura, minimalismo, comprensión de textos, procesamiento del discurso, teoría computacional.

It is plain that, along the course of reading a text, world knowledge is often required in order to establish coherent links between sentences. Therefore, the content grasped from a text turns out to be strongly dependent upon the reader's additional knowledge that allows a coherent interpretation of the text as a whole.

Researchers in Artificial Intelligence have realised long time ago the importance of knowledge about the world for designing models of natural language processing. Similarly, in the field of Psycholinguistics there are several approaches that try to account for the different ways in which world knowledge can be used in tasks of language understanding. Two of them are of special importance. First, knowledge of the context as a whole works as a constraint for interpreting the sentences in a text. Second, knowledge about specific facts, and knowledge about how things usually turn out to be, is normally used to fill in many details that do not figure explicitly in the text.

One way of characterising this additional component of text meaning is doing it in terms of those inferences that a competent reader must undertake to combine the meanings of the different sentences in a suitable way. Since readers have no direct access to speaker's intended meaning, usually have to rely on a process of inference to arrive to an interpretation of the sentences, or well to establish the connections between them. These inferences appear to be of different kinds. It may be the case that agents are able of deriving some specific conclusion from some premises through

*** Dirección para correspondencia:** Mauricio Iza. Departamento de Psicología Básica. Universidad de Málaga. Campus de Teatinos. 29071 Málaga (España).
E-mail: (1) M. Iza: iza@uma.es;
(2) J. Ezquerro: ylpezmaj@st.ehu.es

deductive inference, but they are rarely asked to do it in everyday discourse.

In this concern, it has been proposed that such a range of textual inferences fall into four categories. The first one is *lexical inference*. This sort of inference is required to solve problems of lexical ambiguity or nominal reference (anaphoric reference). The second one includes *inferences of space and time*. For understanding a narrative text, readers have to be able of anchoring the events and episodes described in some spatio-temporal framework. The third kind corresponds to *extrapolative inferences*. Readers have to extrapolate often beyond two events to find some link between them. Readers must infer the sequence of intervening events linking two sentences (events). In doing so, they extrapolate beyond what is literally given in the text. Finally, there are *evaluative inferences*. This type of inferences arises because the meaning of an event depends on the context in which is presented, and readers have to determine this frequently to understand what the text is about (1).

As we can see, inferences can be classified in a number of ways. Another interesting classification is the one between *bridging* and *elaborative* inferences. Bridging inferences *must* be made if a text is to be coherently interpreted; but there are other inferences that, even though their conclusions are normally true or highly likely, are *merely elaborative*, being thus considered unnecessary to link the sentences in a text (2). Both kinds of inferences go forward in opposite direction: whereas the former goes backward, the latter go forward. In addition, it is worth mentioning that there are inferences about the situation described in a text and inferences about the topic of a text or the motives of its writer (3).

Summarising, inferences flesh out a text with additional information retrieved from long-term memory, in such a way that what is finally stored is the information in the input *plus* information that was not explicitly stated but inferred. However, there are two main problems that text comprehension researchers ask about inference making. The first is to determine what (and when) inferences are made as a text is read, and whether they are encoded into a representation of its content. The second concerns the organisation of knowledge in long-term memory, and how cues in a text access that knowledge. In this paper we will try to examine the above mentioned forward inferences within this framework of problems.

Determining the inferences to be made

There are lots of experimental data confirming that readers must perform inferences for interpreting a text. However, it turns out that the notion of inference usually appealed to is rather general, so that it helps us very little to advance in the solution of our first problem, namely, to specify the type of inferences that readers perform.

A tentative answer to this problem is to assume that, in normal cases, bridging inferences are made on-line. In order to collect data on-line, experimental studies have used self-paced reading technique. For example, Haviland & Clark (1974) measure the time spent by subjects when reading a given sentence in two different contexts.

Haviland & Clark (1974) show that identifying referents for definite noun phrases is a highly inferential activity. These authors found that to determine the referent for *the beer* in (2b) took readers significantly longer than in (1b).

(1) a. Mary got some beer out of the car.
b. The beer was warm.

(2) a. Mary got some picnic supplies out of the car.
b. The beer was warm.

This result is explained in terms of a particular aspect of the inferential process described as forming a bridging assumption. The bridging assumption required between (2a) and (2b) is that shown in (2c).

(2c) The picnic supplies mentioned include some beer.

It takes time to make up this type of bridging assumption. In this way, the difference in comprehension times noted between (1b) and (2b) is accounted for. Thus, the consequence to draw from these research findings is that inference takes time.

In a nutshell, information in (2c) can be seen as the missing link required to establish an explicit connection between (2a) and (2b). Nevertheless, some inferences can be also performed without taking additional time. For instance, if the missing link is represented in some stereotypic knowledge format (i.e., frames, schemas, or something like that), it can be automatically activated (or being easily accessible) when the knowledge chunk is activated.

Inferences as missing-links

The types of *generally true* missing links are normally presented in terms of a connection between the verb of one sentence or clause, and the definite noun phrase of another, as it can be seen in the following examples:

(1) a. It was dark and stormy the night the millionaire was murdered.
b. The killer left no clues for the police to trace.
(Carpenter & Just, 1977)

(2) a. Mary dressed the baby.
b. The clothes were made of pink wool.
(Sanford & Garrod, 1981)

In (1), we would have a missing link such as "Murdering involves a killer"; and, in (2), a link such as "Dressing involves clothes" would be needed.

This last example is used in a controlled experiment by Sanford & Garrod to test whether or not the type of missing link involved requires the additional processing time noted by Haviland & Clark (1974) concerning the picnic supplies-beer example; nevertheless, there is no time difference.

Sanford & Garrod have suggested that, when the missing link is already part of the text, no additional processing is required to understand the subsequent reference to another element in that knowledge representation. Their claim is that because dressing activates *clothes* in our representation, any subsequent mention of *the clothes* is understood as quickly as it would be if *the clothes* had already been explicitly mentioned. However, since *picnic supplies* did not automatically activate beer in the knowledge representations of Haviland & Clark's subjects, they had to make a bridging assumption that took additional processing.

Thus, as it seems, we have (at least) two categories of missing links: one kind is automatically activated without requiring additional processing time; the other, instead, is not automatic, but the result of a bridging inference, and so requires additional processing time.

In this work, we will focus on the type of representation required by a connectionist system to perform such automatic inferences (4). In order to achieve this, both of the following problems should be solved. First, we have to determine exactly which elements will be automatically activated via the reader's preexisting knowledge representation (5). Secondly, there is a problem with the automatic connection via background knowledge. It is the assumption that the connection can be described in terms of a decomposition of lexical meaning. Chafe (1972) suggests that this may be a reasonable approach, and Sanford & Garrod make the point in processing terms: "when a verb like dress is encountered, this will evoke from memory a representation which contains slots for a variety of entities implied in the meaning of the verb, such as *clothing*" (Sanford & Garrod, 1981, p. 108) (6). But, if this really is the case, then there would be an extremely large, and massively redundant, representation which would be unlikely to lead to the automatic connection type of processing, indicated in their experimental findings. Finally, we will discuss the possibility of using an attentional focus to solve both these problems in order to perform forward inferences in a hybrid system.

Elaborative inferencing

The types of inferences have been classified in two general categories: necessary and elaborative. In the former case, readers perform necessary inferences, such as those needed to maintain referential coherence (e.g., Haviland & Clark, 1974; McKoon & Ratcliff, 1986; O'Brien, Duffy & Myers, 1986), or to establish causal relations (e.g., Myers, Shinjo & Duffy, 1987; van den Broek, 1990, 1995, 1996). In the latter case, readers draw elaborative inferences to predict upcoming consequences or information (e.g., McKoon & Ratcliff, 1986; O'Brien, Shank, Myers & Rayner, 1988; Garrod, O'Brien, Morris & Rayner, 1990). For instance, perhaps as soon as comprehenders read that "The actress fell from the 14th floor", they predict that the actress died. If so, they would be generating what is called a predictive or forward inference. However, experimental data show that this inference is only partially encoded into the mental representation of a text (McKoon & Ratcliff, 1986).

Similarly, during anaphora processing, its interpretation can depend upon some arbitrary part of general knowledge not easily accessible. So, anaphora needs inference processes, based on world knowledge or, on context, that must be performed very quickly. As we shall see later, some of these inference processes can be elaborative inferencing processes that help in the interpretation of a subsequent anaphora.

When elaborative inferences are made

Haviland & Clark's (1974) experiment showed that inferences necessary to establish the coherence of a text are made during the course of reading. However, there is currently some controversy in psychology about elaborative inferences. A text can support indefinitely many inferences of this kind, and it is obvious that not all of them can be drawn during reading. In the early 1970s it was considered both in psychology and in AI that are performed just the most probable ones, although evidence suggesting this conclusion came from memory experiments, being thus indirect (Garnham, 1985).

The theory that assumes that inferences *are* encoded into memory representations as texts are read has been called by Garnham (1982) the "immediate inference" theory. An alternative theory, that appears more suitable in some respects, is the *deferred inference* theory. Its plausibility comes from the fact that indefinitely many inferences can be made from a sentence, or set of sentences. If all of them were encoded, the mental representation of even a single sentence would turn out infinitely large. In any event, the major part of the inferred facts would not be useful, being thus inefficient to encode them. The deferred inference theory claims that only necessary inferences are drawn at the time sentences are encoded, given that such inferences are necessary for a text to be coherently interpreted. Therefore, elaborative inferences are made just in case they are needed, for instance to answer questions or solving problems of interpretation. If they can be drawn from the text, then they can also be derived from an adequate memory representation of it.

Experiments on elaborative inference making carried out by Corbett & Dosher (1978) and Singer (1979, 1980, 1981), where both probable and improbable explicit instruments of verbs were presented, agree with Thorndyke (1976) on this point, namely, that inferences are made just when they are necessary for comprehension or question answering (the deferred inference theory) (7).

Garnham (1982) has shown a way to reconcile data supporting immediate and deferred inference theories. The immediate inference theory claims that, for example, implicit and highly probable explicit case fillers are treated in the same way: both are encoded. The *omission theory* (Garnham, 1982) suggests that neither of two are encoded into a representation of content, though explicit case fillers are encoded into a relatively short-lived representation of surface form. Highly probable fillers are reconstructed, for answering questions, by

means of a process similar to the one required in the immediate inference theory for inferring implicit case fillers at the time of encoding.

The omission theory explains why implicit and highly probable explicit case fillers are confused in long-term memory: neither are encoded. It also explains why, at short delays, explicit case fillers are easy to verify (unlike implicit case fillers), they are present in a representation of surface form (8). However, more recent experimental evidence suggests that some implicit case filler can be immediately encoded in different degrees.

Psychological evidence for elaborative inferencing

As O'Brien *et al.* (1988) argued there are two major sources of difficulty in the study of elaborative inferences. First, in so far as they are not necessary for comprehension, it is difficult to predict exactly whether or not a reader will perform some inference, and when will do it, as well as to know exactly what that inference would consist in. Second, readers normally generate elaborative inferences only in limited situations, preferring to delay any inferential process until it becomes necessary.

A type of elaborative inference that depends on knowledge about empirical regularities is *instantiation*. Instantiation occurs when a general term, such as *fish*, takes a more concrete interpretation from context.

For instance, Garnham (1981b) offered evidence for on-line instantiation. He found that:

The shark swam rapidly through the water.

was read slowly after:

The fish avoided the swimmer.

since there is nothing in that sentence to indicate that the fish is a shark. However, with a context that cued an instantiation:

The fish attacked the swimmer.

the use of *shark* in the following sentence did not produce any difficulty. This suggests that, when people read that a fish attacked a swimmer, they represent the fish as a shark or something like that.

Therefore, context can be used to elaborate a representation of the object. Such elaboration is based on knowledge about the world, not on lexical semantics. This is consistent with the claim that the representation of the content of a sentence -its mental model- is not linguistic in nature. Its components would represent objects, not word meanings (9).

Experimental research on inference processes that take place during reading is usually intended to confirm whether or not some particular type of information is inferred. Recently, McKoon & Ratcliff (1986) have proposed an alternative framework for research in this area. Their main claim is that inferences are not necessarily encoded into the memory representation of a text in an all-or-none fashion, but instead they can be encoded in different degrees. The degree of encoding ranges over not encoded at all, to partially encoded, to exemplars encoded. For McKoon & Ratcliff, a partially encoded inference is a set of meaning features that does not completely instantiate the inference (10).

Minimal inference processing framework. Inferences that occur during reading can differ regarding the degree in which they are encoded. Within this framework, an inference can minimally represent some set of features or propositions that do not instantiate the whole inference. Different types of inferences can be encoded with different amounts or strengths of information. If the strength of encoding is relatively high, then the effects of inference will appear under a variety of retrieval conditions. Nevertheless, if the strength is low, then the effects of inference could only appear under optimum retrieval conditions.

Minimal-inference approach suggests comparisons between different types of inference under different types of retrieval conditions (as a way of mapping information involved in the

mental representation of a text). Examples of inferences with different degrees of encoding have been provided in several experiments designed to examine the encoding degree of inferences about predictable events and contextually defined categories (11).

McKoon & Ratcliff (1989a) showed that inference performance does not depend on textual information, but on semantic-associative information. For these authors, there are at least two ways in which associative semantic information can interact with new information during reading:

1) Semantic association may provide relations between two pieces of explicitly stated information, and these relations may contribute to the construction of inferred connections.

In previous research for explicitly stated information, the usual finding is that the higher the semantic association between words in a context, the easier is the processing. Corbett (1984) has presented data that suggest that the difficulty of interpretation of a category name used as an anaphora depends on the typicality of the alternative referents, and Roth & Shoben (1983) have shown the dependence of relative typicality on context.

2) Semantic information can support inferences during reading, thus contributing to elaborative inferences. Elaborative inferences do not connect information explicitly stated in the text; rather they are inferences that add new, no yet stated, information.

There are few data about the effects of semantic association in the case of elaborative inferences. Some previous work suggests that inferences about the instruments of verbs depend on the degree of association between the instrument and the verb (McKoon & Ratcliff, 1981). However, the instrument of the verb was explicitly presented, in such a way that all inference processes had to do, was to connect an instrument stated at one point in the text to the appropriate verb stated at a later point in the text.

McKoon & Ratcliff's (1989a) Experiment 1 provides evidence that semantic association affects inference processes even when information to-be-inferred has not been stated in the text (12). Inferences about the most typical examples of such categories are encoded into a high level during reading, and the content of the inference is made up of information relating properties of the most typical exemplar to textual information.

This emphasis on semantic association determines the degree of specificity of inferences in text processing. According to the degree of specificity inferences establish mainly local coherence with minimal encoding of other kinds of inferences (McKoon & Ratcliff, 1988b). McKoon & Ratcliff (1986) distinguish inferences that are necessary to connect propositions by argument repetition, from inferences not necessary to achieve coherence. In the first category are included inferences that establish connections between two instances of the same concept (McKoon & Ratcliff, 1980; Ratcliff & McKoon, 1978), and inferences that fix the referent of an anaphor (Haviland & Clark, 1974; Corbett, 1984; Corbett & Chang, 1983; Dell, McKoon & Ratcliff, 1983; McKoon & Ratcliff, 1980). There is a great amount of evidence confirming that these inferences are generated quickly and automatically. McKoon & Ratcliff interpret this evidence as favouring the hypothesis that information in the working memory representation of a text is, like direct semantic-associative information, easily available (Kintsch & Vipond, 1979).

In the second category are included inferences that are not necessary for coherence, such as elaborative inferences studied by McKoon & Ratcliff (1986) and Singer & Ferreira (1983), inferences that fill in schema information (Alba & Hasher, 1983; Seifert, McKoon, Abelson & Ratcliff, 1986), and inferences about the instruments of verbs (Corbett & Dosher, 1978; McKoon & Ratcliff, 1981). Also in this category there can be inferences about the global structure of a text (see, McKoon & Ratcliff, 1992). For all of these elaborative inferences, there is some evidence

that they can be automatically processed during reading. Nevertheless, without the support of directly available information, explicit encoding of these inferences requires an excess of computation.

McKoon & Ratcliff (1989b) modify in the following way their general framework of minimal inference processing: if a specific inference is allowed by easily available general knowledge from long-term memory, then it will be constructed even if it is not required for coherence. However, this modification has only application for the most easily available information. Other inferences such as those about predictable events, about default values in schema representations, and about the instruments of verbs, are usually encoded partially or not at all.

For instance, in Sanford & Garrod's (1981) model, it is misleading to say that given "Jane unlocked the door" the inference "She used a key" will be drawn. Instead, for Sanford (1990), such *knowledge* becomes *accessible*; nevertheless, in the case of instantiation, context can be restrictive enough to define the nature of the role.

O'Brien *et al.* (1988) reported that readers generated elaborative inferences only when a text contained characteristics (a strong biasing context or a demand sentence) that made it easy to predict the specific inference that a reader would draw, and virtually eliminated the possibility of the inference being disconfirmed. Garrod *et al.* (1990), however, argued that the two text characteristics manipulated might have produced different types of elaborative inferencing: biasing contexts result in a passive form of elaborative inferencing.

Passive elaborative inferencing. As has been just said, O'Brien *et al.* (1988) reported experiments about elaborative inferences using texts containing characteristics that made it easy to predict the specific inference that a reader would draw. This occurred only in presence of a strong biasing context before the point at which the inference was expected (see text 1 in Table 1), or when the text invited an inference by means of a demand sentence (see text 2 in Table 1). As soon as these characteristics were relaxed, evidence for the establishment of the inference disappeared.

Table 1: Example passages used in O'Brien, Shank, Myers, & Rayner (1988)

Text 1
All the mugger wanted was to steal the woman's money. But when she screamed, he <stabbed> <assaulted> her with his (knife/weapon) in an attempt to quiet her. He looked to see if anyone had seen him. He threw the *knife* into the bushes, took her money, and ran away.
Text 2
Joan was delighted when Jim gave her a ring with a (diamond/large stone) in it. <He had asked her to marry him, and now they were officially engaged.> <He often bought her expensive and unusual gifts, and this was no exception.> She went to show her father. *He asked what kind of gem it was*. She excitedly told him that it was a *diamond* from her boyfriend.

Note: The high-context version was created by including the text in the first set of brackets; the low-context version was created by using the text in the second set of brackets. The antecedent is in parentheses, with the explicit and implicit versions to the left and right of the slash, respectively. The target anaphora is in Italics, and the demand sentence in the text 2 is in Italics.

In O'Brien *et al.*'s (1988) experiments, reader's eye movements were monitored as they were reading the passages of text. In these studies the context was manipulated to modify the degree to which it restricted the antecedent for a subsequent reference. For example, con-

sider text 1 in Table 1. What was subsequently to be identified by the reader as *the knife* in the last sentence was introduced in a restricting context sentence such as "He stabbed her with his weapon", or by a less restricting context such as "He assaulted her with his weapon". In a half of the conditions, the antecedent entity was lexically specified as the target in the context sentence (i.e., *his weapon* was related with *his knife*). O'Brien *et al.* were thus able to compare gaze duration (the amount of time the eye remained on a word before moving on to another word) on the target anaphor (e.g., *knife*) to determine whether the reader used the context to elaborate on the original referent. The hypothesis was that if there was no difference in gaze duration between conditions with a lexical specified antecedent *versus* those ones without, then it could be assumed that the reader must have inferred the specification from the surrounding context.

In two of the studies (Experiments 1 and 2) just such an effect was observed. There was a reliable interaction between context (restricting *versus* non restricting) and lexical specification (explicit *versus* implicit introduction of the antecedent). Readers meeting with strong restricting contexts seemed to infer a more specific referent for the antecedent than would be licensed by its initial description. In Experiment 1, the antecedent was restricted with a very strong immediate context of introduction (as in the *knife* example above); in Experiment 2, the introducing context was much weaker, but a demand sentence that focused attention on the exact nature of the antecedent occurred before the critical target reference was found. The text 2 in Table 1 shows a typical passage from Experiment 2. The sentence in italics in the example is the demand sentence that focuses the reader's attention on the target concept and seems to force the elaborative inference that the large stone is in fact a diamond.

However, Garrod *et al.* (1990) reconsidered the above data and offered two refinements to O'Brien *et al.*'s conclusions. First, the two text features manipulated by O'Brien *et al.* (a strong biasing context or a demand sentence) might have produced different types of elaborative inferencing. Garrod *et al.* argued that biasing contexts result in a passive form of elaborative inferencing, involving setting up a context of interpretation, whereas the presence of a demand sentence invites the reader to actively predict a subsequent expression. Secondly, clear evidence for either type of inference will be apparent only with truly anaphoric materials.

In this way, Garrod *et al.* (1990) argue that, in the previous case of *the actress*, with a sufficient constraining context, there would have been encoded much more semantic features defining *dead*, leading to the activation of a predicted inference. Therefore, they conclude that it is plain that conditions under which evidence for elaborative inferences will be evident are limited. One such condition is a highly restricting context (13).

Summarising remarks. As we have seen, comprehenders may draw inferences to predict upcoming consequences. For instance, perhaps as soon as comprehenders read the sentence "The actress fell from the 14th floor", they predict that the actress died (14). In doing it, they would be generating what is called a predictive inference. Inferences can also be drawn to improve a story. If so, comprehenders would be generating what is called an elaborative inference.

In contrast to predictive or elaborative inferences, some inferences are drawn to resolve inconsistency. Coherence inferences (bridging inferences) fill in missing information to resolve a contradiction between previous and current passages. In this way, coherence inferences bring together events of a story and thereby improve the story's cohesion.

Coherence inferences are, in general, more likely to be drawn than predictive or elaborative inferences (McKoon & Ratcliff, 1986). In fact, comprehenders are just as fast at verifying some types of information they assumed through coherence inferences, as they are at verifying information that was explicitly stated; in contrast, comprehenders go much more

slower at verifying information assumed through predictive or elaborative inferences (Singer, 1980, 1996; Singer & Ferreira, 1983; Singer, Graesser & Trabasso, 1994; Trabasso, 1991).

However, a distressing question for theories of inference is how to characterise the inferences people perform, and the circumstances under which they do it. According to minimalism theory (15), inferences that are merely elaborative are not prompted from the beginning, but only if they are subsequently required. Answering either a question posed by an experimenter, or one arising in the mind of the reader, would be examples of this point. Likewise, only inferences that establish local cohesive links are made automatically.

Hence, the most important determinants of inference making are, on the one hand, the establishment of local coherence; on the other, ready availability of knowledge, which can underlie elaboration. However, a complete representation of a text involves elaborative (minimalist or non-minimalist) inferences. The reason is that texts do not describe situations completely (16).

Here several examples: there can be performed some global inferences, given they are necessary for a correct interpretation of a text. Gernsbacher & Robertson (1992) showed that fictional characters' emotional states are readily inferred. Inferences about emotional states could be based on readily available knowledge, but the mundane nature of the story does not grant that information will be readily available. This hedge generalizes the widely accepted idea that elaborative inferences can be made subsequently to the initial reading of a text.

There is a wide agreement on the existence of several general kinds of inferences. Some are required to maintain textual coherence, e.g., referential coherence and causal coherence. The claim that such necessary inferences are routinely made during encoding of a text, is widely shared in research (cf. Garrod et al., 1990; McKoon & Ratcliff, 1986; Potts, Keenan & Golding, 1988). There are other inferences, however, that are not required to establish text coherence, inferences that can be considered as elaborative. Elaborative inferences include instrumental inferences, inferences that would supply a typical instrument for a verb (e.g., inferring *hammer* for *pounding a nail*); semantic inferences, that add contextually appropriate shades of meaning to a concept (e.g., emphasising *round* for tomato in *the girl rolled a tomato*); and predictive inferences, which would produce information about the likely outcome of a described event. The extent to which readers draw elaborative inferences, and the conditions that prompt them, are still much at issue (17).

There are very different views about how frequently forward inferencing goes on, and many mental models' theorists seem to believe that there is a lot of forward inferencing (18). However the views on this issue tend toward minimalism, in the sense that they don't assume that many forward inferences are made as a matter of fact (Garnham & Oakhill, 1992; Stevenson, 1993). Likewise, the role of implicit knowledge doesn't need to be necessarily treated as inference (*accessibility*).

Elaborative inferences can be made during comprehension, but only very slowly, and without being fully represented (if at all) in a mental model. Likewise, elaborative inferences use operational processes on premises based on evidence in the input, and on inferences retrieved from long-term memory (Stevenson, 1993). This way, researchers are considering parallel architectures as computational models of this processing. One of the reasons is that these architectures execute just the relevant inferences, because a connectionist system constrains inferences through its interconnected network of excitatory and inhibitory links (Kintsch, 1988). However, as Sanford (1990) pointed out, relevance itself seems to be very difficult to define, and thus to establish a procedure to treat it. In the case of forward inferencing, a step could be made in strongly biasing contexts (Garrod *et al.*, 1990), where inferences are constrained to those that are relevant to the given context.

Finally, inferences of all types (including elaborative activity) will be more frequent with

main characters than with secondary ones (see Garrod & Sanford, 1988). This amounts to mapping onto background knowledge that serves the intended perspective of the main character. A knowledge based account of understanding with this kind of constraint, or preference, could be to recruit to single-perspective knowledge, which we would normally acquire as a result of our experience. Also, this raises the interesting question of when and how a new entity is considered more topical than the previous topic entity.

Foregrounded inference contexts

What has been called "implicit inference", seems to consist of a device that builds a *single* mental model on the basis of discourse, its context, and background knowledge (19). Such a knowledge is embodied in the model by default, that is, it is maintained in the model provided there is no subsequent evidence overruling it. No attempt is undertaken to search for an alternative model unless such evidence arises. For this reason the process can be very rapid; it becomes as automatic as any other cognitive skill that calls for no more than a single mental representation at time. Also for this reason, implicit inferences lack the guarantee, the mental imprimatur, associated with explicit deductions.

As we have seen, predictive inferences are a representative case in the general controversy over elaborations, and they have received extensive attention in the literature (e.g., Duffy, 1986; McKoon & Ratcliff, 1986; Singer & Ferreira, 1983; Singer, Graesser & Trabasso, 1994; Trabasso, 1991; van den Broek, 1990, 1995). One reason for the interest in predictive inferences is their significant relation to what is known about the causal structure of narratives. The processing of narratives is facilitated by the logical ordering of events. If readers are going to make an inference to facilitate the assimilation of upcoming text, it seems likely that they would make an inference that is related to the causal coherence of the narrative. So, an appropriate approach should determine the contexts in which readers use elaborative inferences (e.g., McKoon & Ratcliff, 1989b).

Writers use several techniques to emphasize particular concepts. For instance, a term that serves as topic of a sentence, or a term to which a pronominal reference is made, can become the focus of processing (Chafe, 1972; Sanford & Garrod, 1981). Such terms are said to be foregrounded due to their high likelihood of being connected with later elements of text. Typically, terms in the foreground are the ones most likely to be evaluated as the possible referents of a next noun or pronoun.

These terms, therefore, are especially important for coherence. If elaborations are going to be made, they are more likely to be related to foregrounded concepts, because these concepts are likely to be repeated, and become part of the overall representation of the text (Sanford & Garrod, 1981).

In the case of predictive inferences, it seems unlikely that readers would make inferences about the consequences of an action involving a concept that is not in the foreground. However, an action involving a foregrounded concept might lead to a prediction about what is likely to come next in the text. Whitney, Ritchie & Crane's (1992) data revealed just that predictive inferences are generated only about concepts that are foregrounded in the passages.

However, traditional models, such as classical logic, have severe difficulties to account for this type of reasoning. The approximate, evidential and adaptive nature of commonsense reasoning, as well as its spontaneity and speed, demands to look for different formalisms and frameworks. In our opinion, massively parallel connectionist models of approximate rule-based reasoning are more suitable for doing the task.

The temporal synchrony approach proposed by Shastri & Ajjanagadde (1993; Shastri & Grannes, 1995) brings about specific and psychologically meaningful claims on the nature of reflexive reasoning, and it may play a representational role in neural information processes. Concretely, it relates the capacity of

the working memory that underlies reflexive reasoning to biological parameters (20).

The reasoner connectionist network has been extended in order to overcome some of its limitations (21), and to improve its performance. While these extensions have been moved towards a more distributed character of representations, there are also other schemas that employ different techniques for high-level data, or knowledge representation, that are related to rule-based reasoning (22). In particular, the *robust reasoning* model proposed by Sun (1993b; 1994) tries to bring together reasoning and similarity. Next, we will broadly which system could fit better with our purposes, as well as its implications, in order to define the main lines of a cognitive architecture for comprehension (23).

Elaborative implementations

Traditionally, a text has been thought of as a set of connected propositions, with inferences either filling in gaps or elaborating the structure. Either way, it makes sense to say that all inferences will be predicate-argument form. However, one danger with forward inferencing is that there are so many possible inferences that it can go on forever (24).

In general, it is assumed that people probably do not do much forward inferencing. Except as Schank (1986) suggests when they ask themselves questions in order to explain and generalize their experiences (cited by Collins & Michalsky, 1989). In any case, people do some forward inferencing and their guess is that the same pattern occurs. But they do not carry it very far because the certainty of the inference quickly falls below some threshold of plausibility.

It is assumed that inferential activity may depend upon a continuous checking of an indeterminate number of input signals against norms for each signal type (syntactic, semantic, textual, conceptual and probabilistic), where the norms might be determined locally by the text or globally by more general knowledge based expectations (Stevenson, 1993; Suh & Trabasso, 1993). This way of thinking about inference generation presents unsurmountable problems for conventional AI modelling and seems to be naturally explicable within the connectionist frameworks.

Connectionist networks are well suited to everyday common sense reasoning (Sun, 1992). Their ability to simultaneously satisfy multiple soft constraints allows them to select from conflicting information in firing a plausible interpretation of a situation (25).

For instance, the types of elaborations that will be temporarily activated include forward inferences about the likely consequences of an event. This accords well with Kintsch's (1988) construction-integration theory of comprehension. In his framework, the initial stage of comprehension is a construction stage in which a small set of elaborations are activated from general knowledge, elaborations that are most closely associated with the propositions directly corresponding to the linguistic input. These elaborations are constructed without the guidance of the complete context -there is only the hope that some of these elaborations might turn out to be useful. Many of these elaborations do not play a role once the context has been further clarified, and thus, are pruned out of the representation. An inference about highly likely consequences of an event explicated in the text captures the characterization just sketched. It would emanate from general knowledge, would be very closely related to the text input, and might well play an eventual role in comprehension of the text.

Kintsch (1988) proposes a hybrid connectionist model where general knowledge is represented in a connectionist network in which related concepts and propositions have excitatory connections. During the construction stage of the model, elaborative inferences are made. Propositions that are close associates of the concepts in initial propositions are activated. However, this process is context-free and activates a lot of irrelevant material (26).

The role of focus

It is widely argued that only a subset of conceptualisations resulting from discourse is available in attentional *focus*, a kind of working memory containing representations of currently important discourse entities. Focus can be understood in terms of the ease of accessibility of potential antecedents and control of inference patterns (see, for example, Grosz, 1977; Hudson, Tannenhaus & Dell, 1986; Sanford & Garrod, 1981). Thus it will be easier to access representations which are in focus than those which are in long-term memory, and this has implications for ease of reference resolution (27).

So, it seems also necessary to take into account the role played by the different characters in a text. Elaborative activity seems to be more through main characters than through secondary characters. We need both to differentiate the degree of activation of these characters (and its role in controlling inferencing), and, to deal with the question of when and how this difference of activation may be modified. For instance, it is possible that a given elaborative inference could produce a change in the activation of the characters (i.e., a focus shift).

From a psychological point of view, Focus system is considered as a kind of working memory that is subject to severe capacity constraints (28). However, it also plays an important role during inference performance. Murray, Klin & Myers (1993) show that consequences of events are inferred when they are extremely predictable and strongly in focus. They suspect that propositions in both the text and the knowledge base are activated in parallel, resulting in either a backward or a forward inference, depending upon whether the cause of the focal event is to be found in something that has already taken place in the text, or in some future action or event that is highly predictable given the text and the reader's general knowledge.

Thus, Duffy (1986) had passages in which, for example, a protagonist has just been served soup when the train "screeched to a stop". In this example, subjects were tested with the probe word "spill" and no evidence was found that the inference (that the soup spilled) was active. However, in Duffy's material, the sentence focusing on the soup and the sentence stating that the train suddenly stopped were separated by a sentence in which the protagonist reached for the salt. So, the soup was no longer in focus when the train halted and the failure to activate the targetted inference is not surprising (Murray *et al*., 1993, p. 471).

For instance, Grosz's (1986) attentional state contains information about the objects, properties, relations,and discourse intentions that are most salient at any given point. It essentially summarizes information from previous utterances crucial for processing subsequent ones thus obviating the need for keeping a complete history of the discourse. It is inherently dynamic, recording the objects, properties, and relations that are salient at each point in the discourse (29).

A different representation of each argument and predicate together with an explicit encoding of inferential chain, between argument predicates, seems essential if a system is wanted to be applied to a large number of dynamic binding that results from these rule applications.

Similarly, Noordman & Vonk (1992) have shown that the reader's knowledge of the world is an important factor in controlling inferences. In their work, the role of the reader's knowledge with respect to the information to be inferred is investigated by varying the materials in terms of their familiarity to the reader (novice vs. expert).

Also, with the aid of some additional control mechanism, such as focus, it may be possible to design a system that combines forward as well as backward reasoning and admits evidential rules. Such a system will be capable of (i) representing incoming information and making predictions based on this information by using its long-term knowledge, and (ii) generating explanations for, and testing

the consistency of, incoming information by referring to its long term knowledge.

However, according to Sanford (1990), relevance cannot itself be used as a procedural criterion for a system to use in controlling inferential activity. Basically, inferential activity should be a function of the structure of the text, of the choice of words, and of the topicalization devices used by the writer (30). It should not surprise anyone that elaborative inferences can be made; the questions are how the state of an inference relates to focus and how focus relates to language input.

Discussion

Comprehension of sentences is located somewhere along a continuum between perception and problem solving processes. It is rarely equated with perception, because perceptual researchers are mainly concerned with more elementary processes, being content leaving complexities of comprehension to others (e.g., to the study of reading). On the other hand, comprehension processes are frequently subsumed under general problem solving. However, as Kintsch (1992, 1994) argues, comprehension is a domain *sui generis* in which it is indeed useful to consider much of what has been regarded as problem solving from the perspective of comprehension (see, e.g., Mannes & Kintsch, 1991; Mannes, 1994).

The meaning of a text results from the sum of its propositional content -information given in sentences determines what can be considered as the propositional content expressed by them-, together with what is inferred from the text in the course of reading. Therefore, comprehension results in a representation of the linguistic input (a propositional representation) that is then converted into a mental model of the text through the use of inferences based on non-linguistic knowledge (Stevenson, 1993, p. 244).

In this context, according to Garrod, Freudenthal & Boyle (1994), the question is how and when text inferences come during reading, as well as the extent to which to establish the full meaning of a sentence is independent and secondary in relation to the task of deriving its propositional content. Some authors (Kintsch, 1974; Kintsch & van Dijk, 1978) have suggested that establishing the propositional content of the sentence is the initial goal of interpretation, and takes place before the full interpretation.

However, in determining the propositional content expressed by a sentence, readers have to take into account information advanced previously in the discourse. We have shown this point in the case of sentences containing indexical expressions such as anaphoric noun-phrases, or tensed verbs, whose semantic interpretation depends upon foregoing temporal information in the text. Another issue is that many non-indexical expressions take very different interpretations in different contexts. Therefore, it seems inefficient in terms of the system's computational costs do not execute contextually appropriate decisions just at time, say, when the relevant expressions in the sentence are encountered. Otherwise, the system should have to reinterpret the whole sentence afterwards an initial propositional representation has been established.

We have examined whether such a contextually significant representation is built up immediately and continuously during comprehension, or it results instead from subsequent processes of integration and inference. In fact, we have found some evidence favouring the former view. For instance, in the research developed by Sanford and Garrod, readers not only fix discourse reference on-line, but they seem to have also immediate access to much more extensive information about the likely state of the referent given the preceding context. This means, in terms of inference, that readers infer much more about the referent of the noun phrase than its identity (see, e.g., Anderson, Garrod & Sanford, 1983).

However, as we have seen, such an account is by no means new. Theorists such as Just & Carpenter (1980, 1992) have argued strongly for what they call the *immediacy assumption* dur-

ing written language processing. These authors claimed that a reader tries to interpret every content word as soon as it is met, and at all levels of processing. The problem is to account for such apparent power in the processor. After all, as Kintsch, among others, has stressed, we know from other investigations of human intellectual function that we are processors of very limited immediate capacity, and it would seem unlikely that all inferences could come about in the limited time spent by a fluent reader to fix any word or phrase of text. In particular, we will suggest that this issue depends on the sort of representational architecture assumed in a theory of comprehension.

Limitations of Construction-Integration theory

A common objection raised against any account of language processing that assumes continuous use of knowledge based inference, is that drawing inferences is computationally costly and hence must be minimised, given we want to grant that the processor will operate as fast as it seems to do. If all possible inferences were drawn all the time, many of these would turn out inappropriate by the time the end of the sentence is reached, hence the argument for carrying out less costly computation first, and minimising the use of inference.

We have defended in other places that focus mechanism can perform this double function, mainly in relation with inference control and anaphora resolution. At the same time, this mechanism would give more consistency to the model. Nevertheless, the working memory proposed by Kintsch is very different from the attentional mechanism we have talked about (31). The use of context (32) to disambiguate between alternative results of memory retrieval, and to select between the results of the construction processes, is closely related to the notion of a current "focus". In the framework for language interpretation that we have been developing, while context represents information about the salience of knowledge, focus specifies the entity, or group of entities, that are currently most salient. In other words, we view focus as a derived notion determined by the available contextual information. Just as it happens with context, keeping track of focus has efficiency as consequence. This is because memory retrieval searches can begin with the items currently in focus (Grosz, 1977). Moreover, as we will see below, it can be useful to differentiate between explicit and implicit focus during text comprehension.

In this sense, a connectionist approach seems to be more close to the *real* nature of the construction process. Connectionist models proposed for reflexive and commonsense reasoning try to perform this process automatically and within some biologically motivated values (capacity bounds). For instance, Rohwer (1992, 1993) presents an apparently novel hybrid representation of phasic (Shastri & Ajjanagadde, 1993; Shastri & Grannes, 1995) and tensor-product representations which retains the desirable properties of each.

Likewise, in the case of learning, connectionist models for script-based text processing deal naturally with the interaction between episodic and long-term memory. After all, readers make predictions about what will happen next. And they do this both on the basis of particular episodes they have found in memory, and on generalizations drawn from similar experiences (Schank, 1982). In this sense, we will examine whether distributed representations would bring about some insights to define a more *complete* connectionist architecture for comprehension. Nevertheless, we will concentrate, for the moment, on the nature of focus.

The role of attention as controller

The main function of an attentional system is supporting and controlling text inferences. Baddeley (1986), for instance, assumes that the central executive of his model of working memory has attentional capacities, being able to select and drive control processes (p. 71).

The role of attention during text comprehension has been broadly explored within the Focus Framework (Sanford & Garrod, 1981).

Within this framework, there have been proposed two dynamic attentional components, explicit and implicit focus, and they are assumed to play different roles in text inference.

Explicit 'versus' implicit focus. The first way to differentiate the two focus partitions is to consider the type of information involved at any moment. Explicit focus deals with the currently relevant entities, while implicit focus represents the currently relevant scenarios. This allows the whole structure to correspond to a *model* of the current state of the discourse world, distinguishing both the relevant entities in that world (episodic information), and the situations (semantic-pragmatic information) in which they play a role.

A second distinction has to do with their psychological properties. Explicit focus has a limited capacity, and holds only a small number of characters "in foreground" at a given time. On the contrary, the only limitation on *scenarios* available by means of implicit focus comes from constraints on their logical compatibility, and on the mapping possibilities they afford for the entities represented in explicit focus.

The third distinction is for the function played in referential processing. Explicit focus works in the interpretation of pronouns, while definite descriptions are interpreted taking account the distinctions represented in implicit focus (i.e., discourse roles). Thus, any dynamic text representation system should capture these differences in interpretation, as well as its consequences for the control of textual inferences.

Inference Control. This framework assumes that text inferences involve two components: a *topic* of inference (typically, the entity that the inference relates to and is about), and a *content* (the information that is inferred about that *topic* that goes beyond what is explicitly expressed in the text itself). Explicit focus plays a role in constraining the *topic* of the inference, while implicit focus plays a role in constraining its *content*. That is to say, the relevant background information is the source of situational inferences. Once the reader has identified the type of situation portrayed by the text, this will constraint the content of many inferences about the various individuals in the story, and the events in which they take part.

Therefore, while constraining the topic of inference is a very active process, that reflects what the reader is attending to at a given moment, role constraints are more deeply embedded as part of the background knowledge, a knowledge that the process can ask for when it is needed to solve problems of text integration. This way, implicit focus can be seen rather as playing a passive role in interpretation if compared with explicit focus, which in some sense drives the reader's expectations.

The different consequences of assuming a twofold focus system can be seen in the previously seen studies made by Garrod *et al.*'s (1990) and in the O'Brien *et al.*'s (1988), devoted to explore the effects of role restriction constraints on the time taken to interpret subsequent references. A result of these studies is that introducing a demand sentence in the text, that is to say, a sentence explicitly focusing on the nature of the antecedent, promotes it as topic of the inference, but its content comes from the implicit focus constraint.

Extending the reasoner with an attentional module. Along this work we have been looking for a system that could take into account these inferential processes during the construction of a mental model of the text. For instance, connectionist script-based models (also models based on sequential or causal inferences) use case-role representations to solve some role-binding cases in highly restricted contexts (implicit definite descriptions; see, Garrod & Sanford, 1990), by means of stored knowledge.

In this sense, an obvious assumption is to consider that background knowledge is represented within the reasoner module (Sun, 1994). During anaphora resolution, the interaction with the focus shifting mechanism could serve to differentiate between explicitly and implicitly focused entities. For instance, explicit focus should be those entities explicitly mentioned by

the text while, all the rest that has been inferred should constitute the implicit focus.

This way, it would be possible to simulate how both modules could learn to control inferences during text processing through interconnections among the reasoner and the attentional module. By this way, some cases of elaborative inferences, such as conceptual, predictive and instrumental inferences, could be integrated within this compound architecture. Furthermore, incomplete processes in forward inferencing would depend on its distributed representation. According to Sanford (1990), a possible solution for such processes could to let the relation constrained, but uninstantiated (e.g., that only some microfeatures were minimally activated).

However, such an account could seem a simplified solution for the real nature of inference control during text processing, if applied to the sort of inferences and anaphora resolution in highly constrained contexts. In the control of inferential processes, it is reasonable to suppose that a constrained number of inferences are made at any time, and that these inferences are central or relevant to the discourse. Given that, it is necessary to take into account when forward inferencing should be activated, and what kind of inferences should be suitable at a given moment. These two aspects are intimately related with the context at hand, rather than with focus. In fact, the focusing of entities (their degree of activation) depends on the role they play in the situation found in the text (33).

However, the interaction between this contextual module and the reasoner should be different regarding the focusing case. The best thing to do concerning this issue, is to encode possible contextual links among propositions and their degree of cohesion (coherence) within the discourse. As we have seen in the CI model, such a module is also required to integrate contextually the conclusions drawn out from inference performance (e.g., irrelevant knowledge that has been activated).

Sanford (1990) believes that a plausible solution might be to furnish with some specific value, pragmatically guided, to the relationship, so that it could be given an answer to the question about what the relation is. However, taking account that the two control aspects mentioned above do depend on the inputs and outputs that the neural net has to deal with, it seems necessary to develop more psychological studies for forward inferencing. For instance, Stevenson, Crawley & Kleinman (1994) have showed that, in sentence's continuation tasks with two antecedents (each occupying a different thematic role), there is a preference for referring to a particular thematic role (e.g., agent *versus.* patient). These authors have interpreted such preferences as due to focusing on the consequences of the represented event in a mental model of the sentence, thus suggesting that thematic roles may provide a bridge at the level of syntactic relationships of structures.

Final remarks

Global assumptions of discourse comprehension involve coherence, i.e. roles and fillers, and incrementality, namely, that a contextually meaningful representation is built up immediately and continuously during comprehension. In so far as a great number of aspects of the situation are represented during discourse comprehension, there exists the possibility for inference to connect what is currently under interpretation to the mental model of the discourse.

Nevertheless, in spontaneous, or natural, discourse, it is extremely difficult to provide the single set of inferences that an individual reader would perform to reach an interpretation. We could think, as Clark (1977) does, that there is a set of *necessary* inferences, which every reader must perform. As it seems, these necessary inferences are exactly of the type that, on experimental evidence at hand, do not require additional processing time. However, it happens that texts readers normally come upon show a minimal amount of formal cohesion. Furthermore, such texts assume massive amounts of existing background knowledge, and normally require to make whatever infer-

ences they feel germane to reach an understanding of what is being conveyed.

In this concern, Garrod, Freudenthal & Boyle (1994) distinguishes two types of inference processes: (i) *True inference* process in which an inference scheme is applied to a set of discrete (e.g., propositional) premises; (ii) *Pseudo-inference* processes that arise from interpreting expressions against a mental model of the discourse domain. Thus, many text inferences based on knowledge are really pseudo-inferences that become immediately available during the initial processing of the sentence, while true-inferences are barely computed, and do not enter directly into the initial interpretation. The primary processing would represent the first step, and it is characterised by automatic and rapid decision making. The secondary processing, on the other hand, represents processes triggered by failures at the primary level, and it is possibly subject to the reader's control.

According to Stevenson (1993), a representational system, able to support language and thought, requires, firstly, parallel recognition, retrieval processes, and serial operational processes; secondly, domain-specific linguistic knowledge, and knowledge from other domains. Concretely, Stevenson proposes a hybrid system. A system where symbolic units are stored in a connectionist network. This parallel network can be used both for parallel processes of retrieval, and for serial operational processes that manipulate what has been retrieved.

Connectionism devotes more attention to the evolutionary basis of thought than classical AI does. For instance, a recurrent backpropagation net can keep traces of the history, as well as respond to current context. Likewise, it can infer context based on various cues. In addition, PDP-systems can do "naturally" some things that are very difficult to accomplish for traditional AI: tasks such as recognising family-resemblance, or constructing conceptual prototypes, among others. These capabilities incline us to think that a satisfactory simulation of human thought would include mechanisms of connectionist sort. Nevertheless, some human cognitive activities seem to require processes closer to traditional AI (Boden, 1993)

Finding a methodology to enable a connectionist system to perform the kind of high level inference that symbolic AI approaches deploy, has been relayed mainly on localist representations. For instance, almost all connectionist reasoning systems impose the restriction that just one rule can fire at time (Touretzky & Hinton, 1988) (34).

A connectionist architecture for comprehension should take into account, on the one hand, evidence of incomplete processes during the performance of a range of cognitive tasks such as comprehension, learning and reasoning. In addition, it must involve a limited capacity for the temporary holding and manipulation of (pragmatic) information (Baddeley, 1986; Sanford, 1990). On the other hand, a possible way to control the explosion of inferences within this working memory, at least for some tasks, can be to allow attentional focusing constraint inference (e.g., feature enhancement and suppression), yielding a system with knowledge in a more structured form.

A no yet solved problem, however, is how to decide between either to perform a backward, or well a forward inference during reading. A specific case of this problem is, for instance, to allow the *context* module to select the type of reasoning to be performed in a given situation. To the date, this issue has not been seriously addressed in the literature, and it is clear that we will need to combine both forward and backward reasoning.

On the other hand, some inferences cannot be automatically constructed. Surely, what will be required to account for these inferences, are models of strategic, goal-based generation processes. Such processes have been considered as secondary, and used to revise primary inferences (i.e., to perform adjustments in the focusing process: inference due to mismatching). Nevertheless, it is obvious that the use of explicit rules, and the corresponding inferences, turns out unavoidable in text comprehension tasks, in so far as they are straightfor-

wardly displayed. This is a further reason that favours a hybrid approach to comprehension processes (35).

The main difference between the present and other architectures already proposed is that construction is guided by attentional states (focus and topic), in order to avoid the knowledge explosion that the employment of full-blooded symbolic systems would lead to. For instance, Aretoulaki and Tsujii (1994) proposed a hybrid symbolic-connectionist architecture where standard symbolic parsers interact with a back-propagation feedforward net to generate texts' abstracts. As it happens in CI architecture, symbolic parsers compute a set of linguistic and extra-linguistic features. The net then integrates them in order to establish the relative importance of each sentence for abstract organisation. Other additional symbolic modules operate on the list of important sentences and generate its abstract as well. In this way, the advantages of both classical and connectionist approaches are retained. Nevertheless, as a consequence of contrasting this system, it is necessary to incorporate a larger amount of information in the input units of this net. Furthermore, this type of architectures violates, to some extent, the basic principles of connectionism, because either they include some non-connectionist modules, or they send complex symbolic messages trough the links.

To end up, a hybrid model, depending on its requirements, could incorporate new computational mechanisms and assess them, according to their performance and generality, in order to generate texts. Generation is the task of deciding which, among the facts we want to communicate, should be literally expressed, and which ones should be omitted, leaving their grasping to the reader's inferencing capabilities. Recent psychological evidence emphasizes the importance of elaborative inferences for the construction of fluent texts, in opposition to bridge inferences that only become activated when gaps in coherence get up. These inferences create expectations to avoid such breakdowns in the coherence chain, and increase the reading facility by helping to avoid redundancies (Mehl, 1994; Bakunas, 1996). In the process of discourse planning, speakers have to test whether a proposition to be uttered can have been inferred from preceding propositions. However, in many contexts, the performance of such inferences is not plausible. To some extent, this problem is akin to one of the versions of the famous frame problem, because it involves decisions about what will not be affected by a given action.

Bridging in itself covers very different types of inference. Consider those inferences assigning antecedents to anaphoric pronouns: to justify them it must be shown how this process can preserve truth (e.g., Discourse Representation Theory proposed by Hans Kamp). But inferences based on our knowledge of causes and effects must be justified in a different way: one which connects the concept of causation with the one of probability. The important point is that our inferences must yield true –or at least probably true- conclusions, given our premises. If this is not so, we make a lot of mistakes, and understanding and communication cannot be built on mistakes.

Notes

1. The first two classes of inference mentioned are relatively straightforward taken account that it is possible, at least in principle, to describe the conditions under which a reader will have to perform them. In the last two classes, the significance of an event depends on the knowledge of what can happen in a certain context (see Trabasso & Suh, 1993).
2. Here, we will broadly discuss this type of forward inferences as opposed to backward (bridging) inferences.
3. This last type of inference, of which many conversational implicatures are examples, contributes to the pragmatic interpretation of the text.
4. The traditional approach seems inappropriate to deal automatically with this type of inference. We can't manage with these inferences as rule-based operations on symbolic structures.
5. Knowledge-activation is clearly context-dependent for naturally occurring texts. This problem is very similar o those noted with representing background knowledge, that is, how we set the boundaries on these representations.
6. Sanford & Garrod's proposal that automatic connections are made between elements in a text via preexisting knowledge representation could be used as a basis or deciding which missing links are, and which are not, likely to be inferences.
7. In Thorndyke's experiment, subjects falsely acknowledged an inference that was independently judged to be likewise plausible. However, Thorndyke's have rather low probabilities. They are not safe in the sense of what inferences to highly plausible implicit instruments are -explicit information to the contrary would not be expected if they were incorrect. His results may not, therefore, be relevant to whether highly probable inferences are made when they are unnecessary for comprehension.
8. Kintsch (1974) provides further evidence for this interpretation of Singer's results. After 20 minutes, when surface form is no longer available, explicit information does take to be verified just as long as inferable information.
9. Garnham (1979) showed that verbs as well as nouns can be instantiated. *Fried* is a better recall cue than *cooked* for: "The housewife cooked the chips".
10. Notwithstanding, attributes of meaning might likewise be well encoded as propositions.
11. In the experiments, test words expressing possible inferences from texts are presented for recognition within immediate testing, only 250 ms. after the text.
12. Strong semantic associates could support the inference in several ways:
    a) the amount of information encoded for the inference could be greater or more specific, or
    b) the strong associates could make the inference more probable.
13. For Suh & Trabasso (1993), negative findings on elaborative inferences reported by McKoon & Ratcliff (1986) may also have resulted from mismatched probes. For example, the sentence from McKoon & Ratcliff's (1986) study: "The actress fell from the fourteenth story window" may not lead to "dead" as the specific causal consequence of the fall. The word "actress" activates the context of movies where falls do not lead to death since falls are usually made by stunt persons or by dummies who do not die. "The despondent teenager jumped from the fourteenth story" might be a better sentence for a predictive inference of "dead" (p.298).
14. For McKoon & Ratcliff (1986), the data show that this inference is only partially encoded into the mental representation of a text.
15. This view suggests that if an inference is not required to give coherence to the information in a text, then the inference will not be produced. This proposal suggests also that elaborative inferences are often minimally encoded (McKoon & Ratcliff, 1992).
16. Nonetheless, the ways in which a representation can be elaborated can be indefinitely many, and in no sense it is ever complete (Garnham, 1992).
17. Keefe & McDaniel's (1993) results can be taken as supporting the view that predictive inferences are temporarily drawn and then deactivated. With more difficult material, deactivation of predictive inferences may either be delayed or prevented.
18. It is true that many mental model theorists have emphasised not just constructive processing, which is an essential part of text comprehension, but on-line elaborative inference making. However, on-line elaboration is not an essential part of a mental model's theory of text comprehension, hence McKoon & Ratcliff's (1992) attitude.
19. The general theory of inference based on mental models embraces both explicit and implicit inferences. Implicit inferences depend on constructing a single mental model while explicit inferences depend on searching for alternative models that may falsify putative conclusions. Hence, the fundamental distinction between the two types of inference is whether or not there is a deliberate search for alternative models of the discourse (Johnson-Laird, 1983).
20. Such as the lowest frequency at which nodes can sustain synchronous oscillations, the coarseness of synchronisation, and the time it takes connected nodes to synchronise.
21. These limitations concern its representation of variable bindings and concepts.
22. For instance, Feldman & Ballard (1982), Hendler (1987), Kosko (1988), Fanty (1988), Dolan & Smolensky (1988) and Derthick (1988).
23. In the selection of such a system, we will have in mind three main types of elaborative inferences: inferences about the meaning of words, predictive infer-

ences about what will happen next in a story (predictable events), and instrumental inferences. Moreover, such a system should also take into account the role played by previous context during reasoning. As we have seen above, previous context is used to determine the meaning of words, the relevance of propositions in the structure of the text, and the referents of anaphoric expressions.

24. Also, it can be argued that some of the work required to produce an inference might get done without the complete inference being calculated.
25. But, these networks are poor at reasoning using the standard semantics of classical logics, based on truth in all possible models. Recently, Derthick (1988) has showed that using alternate semantics, based on truth in a single most plausible model, there is an elegant mapping from theories expressed using the syntax of propositional logic onto connectionist networks.
26. A knowledge-explosion results, before general knowledge acts as a context that filters out unnecesary inferences. The connectionist network rapidly removes inconsistencies and irrelevancies. What is gained is flexibility and context-sensitivity (Stevenson, 1993).
27. Pronoun reference is demonstrably sensitive to focus. Anderson, Garrod & Sanford (1983) showed that pronominal reference to characters which depend upon being situated in a scene are more difficult to process when that scene is cued as completed. (An example of dependence is that *a waiter* is dependent on a restaurant scene.) However, main characters, not dependent upon a particular scene, can easily be referred to by a pronoun after a change from the scene in which they have just appeared (see, also, Garrod & Sanford, 1990).
28. Biologically motivated values upon Shastri & Ajjanagadde's (1993) reasoner system.
29. In this theory, the global component of attentional state is modelled by a set of *focus spaces*; changes in attentional state are modelled by a set of transition rules that specify the conditions for adding and deleting spaces.
30. Such a function would correspond with the multiple-constraint satisfaction in distributed systems.
31. See Ezquerro & Iza (1995, 1996).
32. For the sake of having a clear distinction between memory and context, and in order to reduce the vagueness of the term, "context" is used here to refer to the relevance and salience of information which restricts memory retrieval. Focus would be a particular part of it.
33. Sun (1994) tries to consider the attentional module as containing contextual information in order to account for more psychological phenomena during text processing.
34. A remarkable exception is ROBIN (Lange & Dyer, 1989). This is a system designed to address the problem of ambiguity resolution using evidential knowledge. However, an important shortcoming of this system is that requires a distributed connectionist system in order to get more expressive power.
35. See Clancey (1998) concerning the need of the use of explicit rules in cognitive activities, and its psychological support.

References

Alba, J.W. & Hasher, L. (1983). Is memory schematic?. *Psychological Bulletin, 93*, 203-231.

Anderson, A., Garrod, S. & Sanford, A.J. (1983). The accessibility of pronominal antecedents as a function of episode shifts in narrative texts. *Quaterly Journal of Experimental Psychology, 35*A, 427-440.

Aretoulaki, M. & Tsujii, J. (1994). A symbolic-connectionist hybrid abstract generation system. *Workshop on Combining Symbolic and Connectionist Processing, ECAI'94*, 114-122.

Baddeley, A. (1986). *Working Memory*. London: Oxford University Press, Oxford Psychology Series, No. 11.

Bakunas, B. (1996). Promoting idea production by novice writers through the use of discourse-related prompts. *Applied Psycholinguistics, 17*, 385-400.

Balota, D.A., Flores D'Arcais, G.B. & Rayner, K. (Eds.) (1990). *Comprehension Processes and Reading*. Hillsdale, NJ: LEA.

Boden, M. (1993). Is Cognitive Science a fraud?. In *Proceedings ICCS-93*. Donostia, University of Basque Country.

Carpenter, P.A. & Just, M.A. (1977). Integrative processes in comprehension. In D. LaBerge & S. J. Samuels (Eds.), *Basic processes in reading: Perception and comprehension*. Hillsdale, NJ: Erlbaum.

Chafe, W.L. (1972). Discourse structure and human knowledge. In J.B. Carrol & R.O. Freedle (Eds.), *Language Comprehension and the Acquisition of Knowledge*. Washington: Winston.

Clancey, W.J. (1998). Interactive coordination processes: How the brain accomplishes what we take for granted in computer languages. In Z. Pylyshyn (ed.), *Constraining cognitive theories: Issues and options* (pp. 165-190). Greenwich: Ablex Publishing Corporation.

Clark, H.H. (1977). Bridging. In P.N. Johnson-Laird & P.C. Wason (Eds.), *Thinking: Readings in Cognitive Science*. Cambridge: Cambridge University Press.

Collins, A. & Michalsky (1989). The logic of plausible reasoning: A core theory, *Cognitive Science, 13*, 1-49.

Corbett, A.T. (1984). Pronominal adjectives and the disambiguation of anaphoric nouns. *Journal of Verbal Learning and Verbal Behavior, 17*, 683-695.

Corbett, A.T. & Chang, F.R. (1983). Pronoun disambiguation: Accessing potential antecedents, *Memory and Cognition, 11*, 283-294.

Corbett, A.T. & Dosher, B.A. (1978). Instrument inferences in sentence encoding. *Journal of Verbal Learning and Verbal Behavior, 17*, 479-491.

Dell, G.S., McKoon, G. & Ratcliff, R. (1983). The activation of antecedent information in the processing of anaphoric reference in reading. *Journal of Verbal Learning and Verbal Behavior, 22*, 121-132.

Derthick, M. (1988). Mundane reasoning by parallel constraint satisfaction. *Technical Report CMU-CS-88-182*. Carnegie-Mellon University, Pittsburgh, PA.

Dolan, C.P. & Smolensky, P. (1988). Implementing a connectionist production system using tensor products.*CSD-880088*. Computer Science Department, UCLA.

Duffy, S.A. (1986). Role of expectations in sentence integration. *Journal of Experimental Psychology: Learning, Memory and Cognition, 12*, 208-219.

Ezquerro, J. & Iza, M. (1995). "A Hybrid Architecture for Text Comprehension: Elaborative Inferences and Attentional Focus". *Pragmatics and Cognition*, *3*(2), 245-277.

Ezquerro, J. & Iza, M. (1996). "Knowledge Representation, Reflexive Resoning and Discourse Processing". *Theoria. 26*(11), 125-145.

Fanty, M.A. (1988). *Learning in Structured Connectionist Networks*. Ph.D. Thesis, Department of Computer Science, University of Rochester.

Feldman, J.A. & Ballard, D.H. (1982). Connectionist models and their properties. *Cognitive Science, 6* (3), 205-254.

Garnham, A. (1979). Instantiation of verbs. *Quaterly Journal of Experimental Psychology, 31*, 207-214.

Garnham, A. (1981a). Mental models as representations of text. *Memory & Cognition, 9*, 560-565.

Garnham, A. (1981b). Anaphoric reference to instances, instantiated and non-instantiated categories: A reading-time study. *British Journal of Psychology, 72*, 377-384.

Garnham, A. (1982). Testing psychological theories about inference making. *Memory and Cognition, 10*, 341-349.

Garnham, A. (1985). *Psycholinguistics. Central topics*. Cambridge: University Press.

Garnham, A. (1987). *Mental Models as Representations of Discourse and Text*. Chichester: Ellis Horwood Limited.

Garnham, A. (1992). Minimalism versus constructionism: a false dichotomy in theories of inference during reading. In *Psycholoquy, 3*(63), Reading Inference 1 (1).

Garnham, A. & Oakhill, J. (1985). On-line resolution of anaphoric pronouns: Effects of inference making and verb semantics. *British Journal of Psychology, 76*, 385-393.

Garnham, A. & Oakhill, J. (1992). Discourse processing and text representation from a "Mental Models" perspective, *Language and Cognitive Processes, 7*(3/4), 193-204.

Garrod, S., Freudenthal, D. & Boyle, E. (1994). The role of different types of anaphor in the on-line resolution of sentences in a discourse. *Journal of Memory and Language, 33*(1), 39-68.

Garrod, S., O'Brien, E.J., Morris, R.K. & Rayner, K. (1990). Elaborative inferencing as an active or passive process. *Journal of Experimental Psychology: Learning, Memory and Cognition, 16*(2), 250-257.

Garrod, S. & Sanford, A.J. (1985). On the real-time character of interpretation during reading. *Language and Cognitive Processes, 1*(1), 43-59.

Garrod, S. & Sanford, A.J. (1990). Referential processes in reading: Focusing on roles and individuals. In D.A. Balota, G.B. Flores d'Arcais & K. Rayner (Eds.), *Comprehension Processes and Reading* (pp. 465-486). Hillsdale, NJ: LEA.

Gernsbacher, M.A. & Robertson, R.W. (1992). Knowledge activation versus sentence mapping when representing fictional character's emotional states. *Language and Cognitive Processes, 7*(3/4), 353-371.

Grosz, B. (1977). The representation and use of focus in dialogue understanding. *Technical Note 15*. SRI International Artificial Intelligence Center.

Grosz, B. (1986). A theory of discourse structure. In *Proceedings of the Eighth Annual Conference of the Cognitive Science Society* (pp. 432-441).

Haviland, S.E. & Clark, H.H. (1974). What's new? Acquiring new information as a process in comprehension. *Journal of Verbal Learning and Verbal Behaviour, 13*, 512-521.

Hendler, J.A. (1987). Marker-passing and microfeatures. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence* (pp. 151-154). Milan, Italy.

Hudson, S.B., Tanenhaus, M.K. & Dell, G.S. (1986). The effect of the discourse center on the local coherence of a discourse. In *Proceedings Eighth Annual Conference of the Cognitive Science Society* (pp. 96-101). Amherst, MA: LEA.

Johnson-Laird, P.N. (1983). *Mental models*. Cambridge: Cambridge University Press.

Just, M.A. & Carpenter, P.A. (1980). Inference processes during reading: From eye fixations to comprehension. *Psychological Review, 87*(4), 329-354.

Just, M.A. & Carpenter, P.A. (1992). A capacity theory of comprehension: Individual differences in working memory. *Psychological Review, 99*, 122-149.

Keefe, D.E. & McDaniel, M.A. (1993). The time course and durability of predictive inferences. *Journal of Memory and Language, 32*, 446-463.

Kintsch, W. (1974). *The Representation of Meaning in Memory*. Hillsdale, NJ: Erlbaum.

Kintsch, W. (1988). The role of knowledge in discourse processing: A construction-integration model. *Psychological Review, 95*, 163-182.

Kintsch, W. (1992). A cognitive architecture for comprehension. In H.L. Pick, P. van den Broek & D.C. Knill (Eds.), *Cognition: Conceptual and Methodological Issues* (pp. 143-164). Washington: American Psychological Association.

Kintsch, W. (1994). Text comprehension, memory, and learning. *American Psychologist, 49*, 294-303.

Kintsch, W. & van Dijk (1978). Towards a model of text comprehension and production. *Psychological Review, 85*, 363-394.

Kintsch, W. & Vipond, D. (1979). Reading comprehension and readability in educational and psychological theory. In G. Nilsson (Ed.), *Perspectives on Memory Research*. Hillsdale, NJ: LEA.

Kosko, B. (1988). Bidirectional associative memories. *IEEE Transactions on Systems, Man, and Cybernetics, 18*(1).

Lange, T. & Dyer, M.G. (1989). Frame selection in a connectionist model of high-level inferencing. *Technical Report UCLA-AI-89-05*. Computer Science Department, UCLA.

Mannes, S.M. (1994). Strategic processing of text. *Journal of Educational Psychology, 86*, 377-388.

Mannes, S.M. & Kintsch, W. (1991). Routine computing tasks: planning as understanding. *Cognitive Science, 15*, 305-342.

McKoon, G. & Ratcliff, R. (1980). The comprehension processes and memory structures involved in anaphoric reference. *Journal of Verbal Learning and Verbal Behaviour, 19*, 668-682.

McKoon, G. & Ratcliff, R. (1981). The comprehension processes and memory structures involved in instrumental inference. *Journal of Verbal Learning and Verbal Behaviour, 20*, 671-682.

McKoon, G. & Ratcliff, R. (1986). Inferences about predictable events. *Journal of Experimental Psychology: Learning, Memory and Cognition, 12*, 82-91.

McKoon, G. & Ratcliff, R. (1989a). Semantic associations and elaborative inference. *Journal of Experimental Psychology: Learning, Memory and Cognition, 15*(2), 326-338.

McKoon, G. & Ratcliff, R. (1989b). Inferences about contextually defined categories. *Journal of Experimental Psychology: Learning, Memory and Cognition, 15*, 1134-1146.

McKoon, G. & Ratcliff, R. (1992). Inference during reading. *Psychological Review, 99*, 440-466.

Mehl, S. (1994). Forward inferences in text generation. In A.G. Cohn (Ed.), *Proceedings ECAI'94* (525-529). Wiley.

Murray, J.D., Klin, C.M. & Myers, J.L. (1993). Forward inferences in narrative text. *Journal of Memory and Language, 32*, 464-473.

Myers, J.L., Shinjo, M. & Duffy, S.A. (1987). Degree of causal relatedness and memory. *Journal of Memory and Language, 26*, 453-465

Noordman, L.G.M. & Vonk, W. (1992). Reader's knowledge and the control of inferences in reading, *Language and Cognitive Processes, 7*(3/4), 373-391.

O'Brien, E.J., Duffy, S.A. & Myers, J.L. (1986). Anaphoric inference during reading. *Journal of Experimental Psychology: Learning, Memory and Cognition, 12*, 346-352.

O'Brien, E.J., Shank, D.M., Myers, J.L. & Rayner, K. (1988). Elaborative inferences during reading: do they occur online?. *Journal of Expeerimental Psychology: Learning, Memory and Cognition, 14*, 410-420.

Potts, G.R., Keenan, J.M. & Golding, J.M. (1988). Assessing the occurrence of elaborative inferences: Lexical decision versus naming. *Journal of Memory and Language, 27*, 339-415.

Ratcliff, R. & McKoon, G. (1978). Priming in item recognition: Evidence for the propositional structure of sentences. *Journal of Verbal Learning and Verbal Behavior, 17*, 403-417.

Rohwer, R. (1992). A representation of representation applied to a discussion of variable binding. In M. Oaksford & G. Brown (Eds.), *Proceedings Neurodynamics and Psychology Workshop*. Bangor, Wales.

Rohwer, R. (1993). A representation of representation applied to a discussion of variable binding. In M. Oaksford & G. Brown (Eds.), *Neurodynamics and psychology*. Academic Press.

Roth, E.M. & Shoben, E.J. (1983). The effect of context on the structure of categories. *Cognitive Psychology, 15*, 346-378.

Sanford, A.J. (1990). On the nature of text driven inference. In D.A. Balota, G.B. Flores D'Arcais & K. Rayner (Eds.), *Comprehension processes and reading* (515-533). Hillsdale, NJ: LEA.

Sanford, A.J. & Garrod, S.C. (1981). *Understanding Written Language: Explorations in Comprehension Beyond the Sentence*. Chichester: John Wiley.

Schank, R.C. (1982). *Dynamic memory*. Cambridge: Cambridge University Press.

Schank, R.C. (1986). *Explanation Patterns: Understanding Mechanically and Creatively*. LEA.

Seifert, C.M., McKoon, G., Abelson, R.P. & Ratcliff, R. (1986). Memory connections between thematically similar episodes. *Journal of Experimental Psychology: Learning, Memory and Cognition, 12*, 220-231.

Shastri, L. y Ajjanagadde, V. (1993). From simple associations to systematic reasoning: A connectionist representation of rules, variables, and dynamic bindings using temporal synchrony. *Behavioral and Brain Sciences, 16* (4).

Shastri, L. y Grannes, D.J. (1995). Dealing with negated knowledge and inconsistency in a neurally motivated model of memory and reflexive reasoning. *TR-95-041*. International Computer Science Institute. Berkely, California.

Singer, M. (1979). The temporal locus of inference in the comprehension of brief passages: Recognising and verifying implications about instruments. *Perceptual and Motor Skills, 49*, 539-550.

Singer, M. (1980). The role of case-filling inferences in the comprehension of brief passages. *Discourse Processes, 3*, 185-201.

Singer, M. (1981). Verifying the asssertions and implications of language. *Journal of Verbal Learning and Verbal Behavior, 20*, 46-60.

Singer, M. (1996). Comprehending consistent and inconsistent causal text sequences: A construction-integration analysis. *Discourse Processes, 21*(1), 1-22.

Singer, M. & Ferreira, F. (1983). Inferring consequences in story comprehension. *Journal of Verbal Learning and Verbal Behavior, 22*, 437-448

Singer, M., Graesser, A.C. & Trabasso, T. (1994). Minimal or global inference during reading. *Journal of Memory and Language, 33*, 421-441.

Stevenson, R.J. (1993). *Language, Thought and Representation*. Chichester: Wiley.

Stevenson, R.J., Crawley, R.A. & Kleinman, D. (1994). Thematic roles, focus, and the representation of events. *Language and Cognitive Processes*, *9*(4), 519-548.

Suh, S. & Trabasso, T. (1993). Inferences during reading: Converging evidence from discourse analysis, talk-aloud protocols, and recognition priming. *Journal of Memory and Language, 32*, 279-300.

Sun, R. (1992). A connectionist model for commonsense reasoning incorporating rules and similarities. *Knowledge Acquisition, 4*, 293-321.

Sun, R. (1993a). Structuring Knowledge In Vague Domains. *Research Paper*. Department of Computer Science, The University of Alabama.

Sun, R. (1993b). Beyond associative memories: Logics and variables in connectionist models. *Information Sciences, 70*, 49-73.

Sun, R. (1993c). An efficient feature-based connectionist inheritance scheme. *IEEE Transactions on Systems, Man, and Cybernetics, 23*(2), 1-12.

Sun, R. (1994). *Integrating Rules and Connectionism for Robust Commonsense Reasoning*. New York: John Wiley.

Thorndyke, P.W. (1976). The role of inferences in discourse comprehension. *Journal of Verbal Learning and Verbal Behavior, 15*, 437-446.

Touretzky, D.S. & Hinton, G.E. (1988). A distributed connectionist production system. *Cognitive Science, 12*, 423-466.

Trabasso, T. (1991) The development of Coherence in Narratives by understanding intentional action. In G. Denhière and J.-P. Rossi (Eds.), *Text and Text Processing* (pp. 297-314). North-Holland.

Trabasso, T. & Suh, S. (1993). Understanding text: Achieving explanatory coherence through on-line inferences and mental operations in working memory. *Discourse Processes, 16*, 3-34.

van den Broek, P. (1990). Causal inferences and the comprehension of narrative text. In A.C. Graesser & G.H. Bower (Eds.), *The psychology of learning and motivation* (175-196). New York: Academic.

van den Broek, P. (1995). The role of readers' standards for coherence in the generation of inferences during reading. In R.F. Lorch & E.J. O'Brien (Eds.) *Sources of Coherence in Reading* (pp. 353-373). Hillsdale, NJ: LEA

van den Broek, P. (1996). A "landscape" view of reading: Fluctuating patterns of activation and the construction of a stable memory representation. In B.K. Britton & A.C. Graesser (eds.) *Models of Understanding Text*, (pp. 165-185), Mahwah, NJ: LEA.

Whitney, P., Ritchie, B.G. & Crane, R.S. (1992). The effect of foregrounding on readers' use of predictive inferences. *Memory and Cognition, 20*, 424-432.

(Art. recibido: 6-7-2000, aceptado: 9-4-2001)