1 Comprehending implicit meanings in text without making inferences

Charles A. Perfetti Joseph Z. Stafura

There are two thin lines that separate descriptions of text comprehension. The first is the line between what a text says (i.e., its explicit or literal meaning) and what is inferable from the text (i.e., its implicit meaning). The second thin line is one that separates two kinds of implicit meaning processes. On one side of this line is what text researchers refer to as inferences. Although these inferences can be further sorted into various types (e.g., Graesser, Singer, and Trabasso, 1994), they have in common the idea that the meaning that is obtained is not in the words and syntax of the text but is constructed by the reader based on knowledge that is largely independent of that particular text. On the other side of the line are implicit meaning processes that are closely bound to the language of the text, the meanings of words, and the grammar of the language. These lines are "thin" because they divide meaning processes only approximately and are subject to encroachment in difficult cases. There is good company for such thin lines in closely related distinctions in linguistic descriptions, especially semantics versus pragmatics and meaning versus interpretation.

In text comprehension research, attention has been diffusely distributed across both lines, treating inferences as a set of differentiated types that contrast with literal meaning. Indeed, analyses of the variety of inferences have led to taxonomies that differentiate among nine different types (Pressley and Afflerbach, 1995) and even thirteen or fourteen types (Graesser et al., 1994). These multiple varieties cross various dimensions of mixed types, some reflecting linguistic devices in the text (thematic role assignment, referential anaphora), but most lacking linguistic sources, dependent on reader knowledge (e.g., character emotion, causal consequence) or coherence strategies (e.g., causal coherence). Taxonomies are of limited value when they are not well defined according to some classification hierarchy, however. A mature taxonomy, for example, a biological taxonomy, serves long-lasting functions (in between modification) because it is organized hierarchically around structural or functional principles (e.g., biological

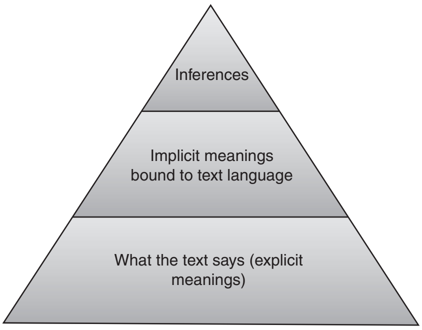

Figure 1.1 A hierarchical instrinsicalist view of explicit and implicit meaning. Texts are assumed to express explicit meanings constrained by the grammar of the language in combination with the meanings of words. Implicit meanings vary with how closely they depend on the text meanings. At the top level of the hierarchy are inferences that are not bound to the text language.

principles). Inference taxonomies have tended to reflect the earliest observational stages of classification, sorting observed cases into categories according to appearances, with emerging (ad hoc), incomplete hierarchical structures. Moreover, the processes that serve inferences appear to be highly general, functioning without specific guidance from the type of inference required (Myers and O'Brien, 1998).

A simpler approach is to capture the coarse-grain differences among the three types of thinly contrasting text meaning we referred to above: the literal; the linguistically constrained implicit meaning; and the linguistically independent, reader-constructed implicit meaning. This three-way classification assumes, not uncontroversially, that texts do contain some degree of intrinsic meaning, constrained by grammar and conventional word meanings. Figure 1.1 captures this approach in a hierarchical structure that puts the explicit meaning of the text at the foundation. Building directly on what the text says, the second level captures implicit meanings that a reader derives from the words and syntax of the currently read segment of text. The third level adds a range of implicit meanings that are heavily knowledge-dependent and cannot be made only by knowledge of words and syntax.

The triangle of Figure 1.1 is consistent with some proposals that try to capture broad, two-way distinctions in implicit meaning. Examples include text-connecting versus knowledge-based (Graesser et al., 1994); inter-sentence versus gap-filling (Cain and Oakhill, 1998); and automatic versus strategic (McKoon and Ratcliff, 1992). Such distinctions might correspond to the difference between level 2 and level 3 in Figure 1.1. However, some broad distinctions that are important in text comprehension, e.g., local vs. global (Graesser et al., 1994) seem less clearly reflected in the hierarchy without further assumptions about the role of the linguistic text in each case.

The hierarchical model is descriptive, referring to the levels of meaning from a text perspective rather than from the comprehension processes that produce meanings. Its value is to organize observations about text processes that can then be linked to comprehension processes. For example, one can see that much of the research interest in inferential comprehension is at the top of the pyramid, especially when questions concern comprehension skill. Thus, much of the work that locates poor inference making as a source of comprehension failure is at this level, where the application of knowledge or strategies beyond routine text comprehension is critical (e.g., Cain, Oakhill, Barnes, and Bryant, 2001; Oakhill, Cain, and Bryant, 2003).

In what follows, we take a closer look at the intermediate level, illustrating a kind of implicit meaning process that is often considered a lower-level automatic process. This intermediate level is below the level of which inference processing is typically examined.

The intermediate level: close-to-the-text inferences

One class of close-to-the-text inferences that has attracted some research attention is the syntactically triggered inference that secures coreference (e.g., Gordon and Hendrick, 1997). Simple pronoun binding is a nearly invisible process that ordinarily must be automatic. For example, in "John bought himself a present," "himself" is automatically linked to "John" because of the morphological trigger of the reflexive pronoun. This same fact of English morphosyntax prevents "him" from referring to John in "John bought him a present."

Other cases show complexities that demonstrate an implicit knowledge of syntax is what is responsible. For example, in the following sentences, readers bind a pronoun with a coreferential noun in accordance with a syntactic-binding principle. Thus, readers interpret "Bill" and "he" as coreferential in both (1a) and (1b), even though the main clause contains the proper noun in one case and a pronoun in the other.

1. a. If **Bill** does well on the exam, **he** will pass the course.
   b. **Bill** will pass the course, if **he** does well on the exam.

However, in 2a and 2b readers show a sharp difference in making a coreferential interpretation.

2. a. If **he** does well on the exam, **Bill** will pass the course.
   b. **He** will pass the course, if **Bill** does well on the exam.

Gordon and Hendrick (1997) confirmed that readers find only (2a) as an acceptable sentence for understanding "he" as referring to Bill. The general point, which covers a wide range of linguistic structures, is that an implicit knowledge of the language itself guides coreferential meaning processes so readily that one might be reluctant to consider them inferences at all.

However, not all text-dependent coreferential processes are triggered automatically by implicit syntactic knowledge. Below we consider coreferential processes that depend on word knowledge, which is less deterministic than syntactic knowledge and thus perhaps more variable in its use.

Word-to-text integration as an intermediate-level inference

Word-to-text integration is a second kind of close-to-the-text "inference" process that we have examined in research on reading comprehension. The quotes indicate that we have not heretofore referred to these processes as inferences and, indeed, one can question whether they are. However, what we have in previous work referred to as "word-to-text" integration (Perfetti and Stafura, 2014; Perfetti, Yang, and Schmalhofer, 2008) involves close-to-the-text meaning processes that seem to have some of the features of inferences. Moreover, part of our general claim in this chapter is that the borders that separate kinds of meaning processes may be permeable and it seems useful to think about a borderline case.

Word-to-text integration (WTI) is a recurring process in language comprehension. In its broadest sense it can include parsing processes, in which a syntactically coded word is fit to a partially constructed syntactic structure (phrase, clause, or sentence) in memory. Such a process must occur on virtually every word as it is encountered. However, in its more typical sense within text comprehension, WTI is the fitting of a word into a meaning representation that has been constructed, based on the reading of the text to that point.

Exactly what this meaning representation is could be a matter of some uncertainty. We assume, for now, that the relevant meaning representation is a mental model – a representation of what the text says – its basic linguistic meaning that is updated by WTI processes. However, the relevant representation could be simply a propositional level meaning representation. In fact, the two are so closely related during immediate memory-limited processing that separating them goes beyond what we can address in our research so far. For now, we assume that at the point a word is read, a motivated reader is constructing a mental model (Johnson-Laird, 1980) of what the text is about (a "situation model," (van Dijk and Kintsch, 1983) that is based primarily on the linguistic information in the text with some limited influences of extra-text knowledge. Considerations of working memory are relevant here. During text reading – at the moment a word is encountered – processing is occupied at the local text level by reading this single word and configuring it into syntactic structures and meaning relations that are active in working memory. Although in the long run, both syntactic and propositional information may fade, such information is prominent temporarily within a resource-limited processing system.

We illustrate WTI, on the assumption that it involves a connection to a working mental model of the text, with two slightly different texts, which are based on experiments by Stafura and Perfetti (in revision).

3. a. While Cathy was riding her bike in the park, dark clouds began to gather, and it started to storm.

By the end of this sentence, we assume that the reader's mental model represents meaning information that includes referents and events, rather than the literal text. At the top level, the model is decomposable into two substructures: *Situation* + *Event:*

<SITUATION: Cathy on bike, in the park, dark clouds>
<EVENT: Storm>

The reader's text memory provides access to the situation, which can be updated through events described in the text. In this case, the situation is updated immediately by the storm event, as illustrated below with the updated element in italics.

<SITUATION: *Storm*, Cathy on bike in park>

Suppose the next sentence adds to text (3a) as repeated below in (3b):

3. b. While Cathy was riding her bike in the park, dark clouds began to gather, and it started to storm. The rain ruined her beautiful sweater.

The new information from the second sentence – that the rain ruined the sweater – can be added to the situation model, introducing the result of the recently comprehended event, the ruination of a sweater.

<SITUATION: *rain-ruined sweater*, Cathy, storm>

But comprehension proceeds not only sentence by sentence but also word by word. The noun phrase that begins the new sentence – *The rain* – is understood immediately in relation to the situation model. The immediacy of this connection of word reading to text understanding is enabled by encoding and memory processes that are sensitive to conceptual/featural overlap (Kintsch, 1988; Myers and O'Brien, 1998). The reader understands "the rain" to refer to (and augment) the storm event and integrates it into the situation model.

Studying WTI through ERPs

These word-level integration processes can be observed through various reading measures that are sensitive to the word-by-word time course of processing, including event-related potential (ERP) methods and eye tracking. ERP measures allow for the observation of this integration as it occurs – a comprehension process rather than a comprehension product. Our experiments, using short texts such as (3) above, measure ERPs initiated by a target word that appears at the beginning of a new sentence, across a sentence boundary from the text that allows integration of "rain" in the current example. When the target word appears, the N400 component, an indicator of the fit between the word and its context (Kutas and Federmeier, 2000), is reduced in amplitude, indicating word-to-text integration (Yang, Perfetti, and Schmalhofer, 2007). This reduction is relative to a baseline condition, illustrated in text (4) below.

4. When Cathy saw there were no dark clouds in the sky, she took her bike for a ride in the park. The rain that was predicted never occurred.

In text (4) there is no antecedent event for "rain." Reflecting the difficulty the reader has in integrating "rain" into the mental model, the N400 on the word "rain" has a more pronounced negative deflection. To describe this in more traditional psycholinguistic terms, coreferencing is not achieved because the prior sentence has no clear antecedent for "the rain."

Thus, text (4) is a baseline condition against which N400 reductions produced by text manipulations (paraphrase, inference, and explicit mention) are compared. Important is the fact that the second sentence of text (4) – and thus, the whole of text (4) – is perfectly comprehensible. However, the reader's situation model is different from that in (3) – there is no storm. So, when the word "rain" is encountered, there is no rapid integration process that adds "rain" to the situation model. Instead a new mental structure is built around "the rain" and (finally) its nonoccurrence.

The fact that text (4) is sensible is important, because in most research that uses the N400 as an indicator of semantic processing, something that is sensible is compared with something that is anomalous, or nonsensible. For example, in a classic N400 study, an ERP is recorded on the final word of a sentence that makes it sensible, e.g., "The pizza was too hot to eat"; and this is compared with a version in which the final word makes the sentence anomalous, e.g., ". . . too hot to drink" (Kutas and Hillyard, 1980). In these situations, the N400 differences are dramatic and appear amenable to explanations based on expectancy violations (Federmeier, 2007; Lau, Almeida, Hines, and Poeppel, 2009). The N400 is assumed to be sensitive to poor fit (failures of expectations to be met, or failures to make sense of what occurs). However, in our case, a comparison is made across sensible texts. The texts differ only in the degree to which they invite an immediate word-to-text integration process. Moreover, the critical word is across a sentence boundary from the relevant antecedent in the first sentence. Taken together, these two features of our experiments are more compatible with a post-lexical integration process (Brown and Hagoort, 1993) than with an expectancy explanation. It is difficult to imagine what word a reader might predict as the first word (or the head of the first noun phrase) across a sentence boundary; nearly any grammatical sentence beginning can continue on with coherent ties to the preceding sentence.[1]

The paraphrase effect. We have referred to the fact that the word "rain" is better integrated (immediately) with text (3) than with text (4) as the *paraphrase effect*. The ordinary sense of paraphrase, expressing an idea in words different from its original expression, does not apply very well to our use of "paraphrase." A well-written text does not use different words in a new sentence to express the same idea as in the previous sentence, as if avoiding repetition were the goal. Instead, paraphrase is about the text moving the situation model a bit forward while maintaining coherence. Thus, in text (3), "rain" is *not* another way of saying "storm." Rather, the text moves the mental model forward in a baby step by referring to a correlate or consequence of the storm with the word "rain." Perhaps reserving the concept of "updating" a mental model for more substantial semantic changes (O'Brien, Rizzella, Albrecht, and Halleran, 1998), we can refer to these paraphrase-induced changes as

[1] Several studies have examined the influence of message-level factors on word-level processing across sentences (e.g., Kuperberg, Paczynski, and Ditman, 2011; Otten and van Berkum, 2008). These studies suggest that multiple levels of linguistic information can influence word-level processing, consistent with the idea that N400 amplitudes are sensitive to both expectancy and fit (i.e., integration).

"fine tuning" the mental model, an updating of individual referents and events that are part of the mental model that does not involve removing a semantic contradiction.

Thus, the paraphrase effect reflects online comprehension – a sensitivity to moment-by-moment tuning of the mental model. We discovered in these experiments that skilled comprehenders showed the paraphrase effect more robustly than less-skilled comprehenders did (Perfetti et al., 2008; Stafura and Perfetti, unpublished). We suggested that less-skilled comprehenders were showing *sluggish* word-to-text integration. Such sluggishness may result in difficulty maintaining coherence across sentences, considering the largely incremental processes involved in comprehension (Just and Carpenter, 1980; Tyler and Marslen-Wilson, 1977).

Explicit meaning process or inference?

If, as suggested by Figure 1.1, meaning processes can be approximately ordered along three levels of text constraint, one might ask whether the paraphrase effect is an explicit text process or a kind of inference.[2] We address this question by referring again to our *storm – rain* example. Referring to "rain" after "storm" as a paraphrase is, as we suggested, a small update, a fine-tuning in the mental model that maintains coherence. However, the text could also leap forward with a giant step, as it would if instead of (3) we had (5).

5. While Cathy was riding her bike in the park, dark clouds began to gather. The rain ruined her beautiful sweater.

Here, when the reader encounters "rain" there is no explicitly stated storm event in the mental model; so there is nothing to which "rain" gets attached, and no event or proposition to modify. Instead, the reader constructs a new event: rain. It's not a big stretch, given the dark clouds, to infer a rain event. And to keep the text coherent, this bridging inference (e.g., Graesser et al., 1994; Singer and Halldorson, 1996) is readily made, although with some detectable cost to processing efficiency. Yang et al. (2007) observed that for texts of this type, the N400 amplitude was nonsignificantly different from baseline. Reading "The rain …" in

[2] If so, one might be tempted to argue that the result that skilled comprehenders show a stronger paraphrase effect than less skilled comprehenders is another example of less skilled comprehenders being poor at inference generation relative to more skilled comprehenders. We would counter such an argument by insisting the lexical basis of this problem points to something other than some generalized deficit in inferential processes.

sentence (5) was similar to reading "The rain ..." in sentence (4). Neither case provided an immediate integration opportunity, although we assume that most readers make the bridging inference in (5), provided they have a sufficiently high standard for coherence (van den Broek, Risden, and Husebye-Hartmann, 1995).

Instead of being stuck with a bridging inference at the second sentence of (3), the reader might make a forward or predictive inference during the first sentence. Such an inference would be generated while reading the clause "dark clouds began to gather" (i.e., a prediction that it will rain). But such an inference has little warrant. Maybe rain would be the next event, and maybe it wouldn't. Certainly the comprehension of the dark clouds in the first sentence allows a readiness to understand "the rain" when it does appear in the next sentence (hence, the N400 to (5) is not more negative than in (4)), but this does not compel a forward inference. Predictive inferences are variable, probabilistic, and generally less compelling than inferences that are needed to support coherence (Graesser et al., 1994; McKoon and Ratcliff, 1992). In reading, there is little or no clear evidence that such inferences are made immediately so that their effects can be measured upon reading on a single word – as opposed to later, when a new text segment affirms the event. The N400 results of Yang et al. (2007) suggest that the forward inference was not made consistently on the first word where such an inference becomes possible, and thus readers had to make a more costly bridging inference when they came to the word "rain" in the second sentence.

To locate bridging and forward inferences in the hierarchical model of Figure 1.1, the predictive inference is at the top of the pyramid, dependent on the reader's knowledge and perhaps on a level of engagement high enough to motivate prediction efforts. Bridging is a bit more difficult to place. While a bridging inference depends specifically on a linguistic gap in the text – it is triggered by a coherence break in the text – it also relies heavily on the reader's knowledge and on a standard of coherence, both of which are outside the narrow scope of the text language.

Returning to the paraphrase effect, we locate the effect at the intermediate level of the meaning processes model of Figure 1.1. There are two important elements in the meaning processes captured by the paraphrase effect and by WTI in general. One is the meaning of the word currently being read. The other is the preceding text – either the text itself, including a specific linguistic unit (a phrase or a full clause), or a mental model of the text with referents and events rather than linguistic units. In either case, the process is one of referential binding, in which the meaning of the currently read word is connected to an antecedent, a linguistic phrase, or a referent based on that phrase. To choose tentatively the mental model description: The meaning process triggered by the paraphrase word is a coreferential process that updates the mental model in a modest way. Referring to text (3), the word "rain" fine-tunes the event referred to by "started to storm." (This fine-tuning allows more substantial updating with the new information about the ruined sweater.) So the two aspects of WTI are the meaning of the word and the meaning of an immediately preceding (available in memory) stretch of text. The integration process selects the referential meaning of the word that is congruent with the meaning of the text and adds new information to the text representation.

To summarize, research on the paraphrase effect exemplifies important integration processes that occur routinely during reading. These processes are quite general, continuously applied to linguistic structures of various kinds – not just single words. In fact, the general case requires noun phrases, which are needed to establish referents, and clauses (or phrasal modifiers), which are needed to establish events. Comprehension depends in part on local word-to-text integration processes that may be said to include inferencing, but are heavily dependent on the use of word meanings. It is knowledge of word meanings and the ability to integrate their context appropriate meaning into a mental model that is critical at this level, which is intermediate to mainly intrinsically text-constrained meaning processes and unconstrained knowledge-rich and strategic processes.

Mechanisms of word-to-text integration

Readers obtain the critical intermediate levels of meaning through mechanisms that retrieve word meanings and their morphosyntactic information, integrating the results of these retrievals into the representation(s) of the text. (These include syntactic representations, which we are not addressing here.) We can ask more specifically about the nature of the processes that bring about the WTI exemplified through the paraphrase effect. The general picture is of interleaved processes of word meaning retrieval and memory resonance that produce a fit between the word and the text representation.

To illustrate, we repeat the brief text of (3b), which shows a text integration process that can occur at the word "rain" in the second sentence.

3 b. While Cathy was riding her bike in the park, dark clouds began to gather, and it started to storm. The rain ruined her beautiful sweater.

The critical word is "rain," where we observe ERP indicators of word-to-text integration – an N400 reduction compared with a baseline sentence