

The Construction-Integration Model: A Framework for Studying Context Effects in Sentence Processing

Evelyn C. Ferstl

Department of Psychology
University of Colorado at Boulder
Boulder, CO 80309-0345
eferstl@clipr.colorado.edu

Abstract

Contextual and pragmatic knowledge facilitates the eventual interpretation of a syntactically ambiguous sentence. However, psycholinguistic studies have not provided a clear answer to when and how this non-syntactic knowledge is used. One explanation for the discrepancy of the results is that the predictions for parsing processes in context cannot be specified unless they are based on a theory of text comprehension. The construction-integration model of discourse comprehension (Kintsch, 1988) is proposed as an example for such a theory. The model is parallel and weakly interactive, and its psychological validity has been shown in a variety of applications. Three simulations for syntactic ambiguity resolutions are presented. In the first, syntactic constraints are used to account for the correct interpretation of a garden-path sentence, as well as for common misparses. In the second example, pragmatic knowledge is used to disambiguate a prepositional phrase attachment. In the final example, it is shown that the model can also account for effects of discourse context in the resolution of prepositional phrase attachment ambiguities.

Introduction

It is undisputed that pragmatic knowledge and discourse context facilitate the final interpretation of a syntactically ambiguous sentence, but the issue of how and when non-syntactic information is used during comprehension has not been settled. Two classes of models have been proposed. In syntax-first models, the initial analysis of an utterance is based on syntactic principles alone. The pragmatic and contextual consistency of the resulting structure is evaluated in a second processing stage. Examples for this type of model are the garden-path model (Frazier & Rayner, 1982) and the model based on lexical-functional grammar (Bresnan & Kaplan, 1982). In contrast, interactive models (e.g., Marslen-Wilson & Tyler, 1980; Taraban & McClelland, 1990; Altmann & Steedman, 1988) assume that information from all knowledge sources is brought to bear as soon as it becomes available. Consequently, syntax-first models predict that processing differences caused by embedding a sentence in a felicitous context cannot be observed immediately, while interactive models predict contextual

override of syntactic preferences even in the early stages of processing.

Psycholinguistic experiments have not yet provided empirical data which clearly distinguish between these theoretical approaches. The few studies explicitly addressing context effects did not yield converging results (see Ferstl, 1993, for an extensive review). Some experiments provided support for syntax-first models (e.g., Ferreira & Clifton, 1986; Britt, Perfetti, Garrod & Rayner, 1992, Exp. 3; Mitchell, Corley & Garnham, 1992; Rayner, Garrod & Perfetti, 1992). Other experiments provided support for interactive models (e.g., Crain & Steedman, 1985; Altmann & Steedman, 1988; Trueswell & Tanenhaus, 1991; Britt et al., 1992; Exp 1,2).

Two explanations for this apparent discrepancy can be identified. First, it has been proposed that the results depend on the on-line measure used in the study (Rayner et al., 1990). Eye-movement monitoring usually yields effects of context only in regressions, but not in the first-pass reading times (an exception is Britt et al., 1992). Data from subject-paced reading tasks, in contrast, often provides support for immediate use of context (but see Mitchell et al., 1992). Further studies are needed which explicitly evaluate task-specific strategies using identical materials in different paradigms.

The second factor which has been extensively discussed concerns the materials, in particular the context paragraphs. On-line context effects can only be observed if the intended bias is effective. Empirically, the bias can be confirmed by global comprehension measures, such as paraphrasing, or rating tasks. Several properties of texts have been proposed to have an impact on parsing decisions, for instance, discourse focus (Rayner et al., 1992) and tense information (Trueswell & Tanenhaus, 1990). The most general proposal is the principle of parsimony (Altmann & Steedman, 1988), stating that an interpretation is preferred which requires only the addition of a minimum number of unreferenced entities into the discourse model (or situation model, vanDijk & Kintsch, 1983). However, for a given set of context paragraphs, it is not easy to evaluate if they possess these properties, and, more importantly, if they differ only on this dimension. For example, following their principle of referential support, Altmann and Steedman (1988) wrote context paragraphs in which discourse focus and syntactic

interpretation were confounded (Clifton & Ferreira, 1989). Explicit predictions about the circumstances in which context effects are expected are only valid if they are based on a careful analysis of the materials. However, this analysis requires the description of how a given context paragraph is understood and remembered, that is, a description of how people understand text.

The study of context effects in parsing can therefore not be successful without a general theory of text comprehension. This theory must allow representations of text on different levels, so that it can account for the interactions between situational and syntactic factors. Examples for this type of model include the Cooperative Language Processor (Perfetti, 1990), NL-SOAR (Lehman, Lewis, & Newell, 1991), and CC-READER (Just & Carpenter, 1992).

The goal of the remainder of this paper is to introduce the construction-integration model of discourse comprehension (Kintsch, 1988) as a theory of human sentence processing. In a wide variety of applications, this model has been shown to be psychologically valid (e.g., Kintsch, Welsch, Schmalhofer, & Zimny, 1990). For instance, context effects in the disambiguation of homonyms and situational effects on sentence verification have been explained. Moreover, the mechanisms of the model are sufficiently specified to enable simulating effects of syntactic ambiguity, while the knowledge representation is flexible enough to be useful in a multitude of contexts.

Kintsch (1988) has applied the construction-integration model to simulate the parsing of an isolated reduced complement sentence. However, the strength of the model is that it provides a natural and widely useful mechanism for integrating pragmatic and contextual information with the text input. Psycholinguistic data on how these factors influence sentence processing have only recently become available. Moreover, the model is an already existing, detailed implementation of a weakly interactive theory, an architecture which has been postulated within the domain of sentence processing (Altmann & Steedman, 1988). Thus, it is worthwhile to consider how the model accounts for context effects in sentence processing.

The Construction-Integration Model of Discourse Comprehension

The CI model is a hybrid model which combines a symbolic, local representation of the textual input with a connectionist constraint-satisfaction process. Knowledge is assumed to be represented as an associative network. The text is processed in cycles corresponding to words, phrases, clauses, or sentences. In the first stage of each comprehension cycle, a network is constructed whose nodes encode the input in propositional form. If two propositions share an argument, they are connected by a positive link. In addition, the most highly activated nodes from the previous processing cycle are carried over to the current one. To each proposition, associative elaborations are retrieved from memory. The most important feature of this construction process is that the linguistic rules guiding it are assumed to be "sloppy and general" (Kintsch, 1988), that is, the network contains nodes which might turn out to be inappropriate in the given context. In particular,

propositions encoding alternative syntactic interpretations are assumed to be activated in parallel. Mutually inconsistent nodes are connected by inhibitory links.

The second stage, the integration phase, consists of a spreading activation process. In this phase, non-syntactic knowledge (carried-over context propositions and elaborations) influences the flow of activation in the network. After the network has relaxed into a stable state, context inappropriate nodes are deactivated, while central nodes remain highly activated. The final activation pattern is assumed to be a representation of readers' text memory.

To summarize, the CI model is parallel and weakly interactive. Non-syntactic information is used to select the context-appropriate interpretation of an utterance, but not to propose a syntactic interpretation (as can be the case in strongly interactive models). Whereas a representation of the linguistic knowledge is not part of the present theory, the model specifies in detail the processes which make use of contextual constraints and pragmatic information. In the following section, I present a series of simulations which demonstrate the usability of the model in the domain of syntactic ambiguity resolution.

1. Example: Garden-Path Sentence

To illustrate how the CI model can account for syntactic phenomena, consider first the classic garden-path sentence (1) *The horse raced past the barn fell*. The disambiguation of this locally ambiguous sentence is accomplished using syntactic knowledge only. The correct interpretation as an embedded reduced relative clause is represented in propositional format using the three propositions (S1) *race[somebody, horse]*, (S2) *location[S1, past_barn]* and (S3) *fall[horse]*. The incorrect reading of *raced* as the active main verb of the sentence is represented using the propositions (P1) *race[horse]*, (P2) *location[P1, past_barn]*, and (P3) *fall???*

In the construction phase, both of these interpretations are activated in parallel, and concept nodes are added. To simplify the example, elaborations are not considered. The resulting network is displayed in Figure 1¹. The positive connections are established according to argument overlap, and competing propositions from the two alternative interpretations are connected by inhibitory links². In addition, the ungrammaticality of node P3 is indicated by inhibitory links to the previously established nodes P1 and P2. After the integration is completed, the three nodes representing the incorrect active reading are deactivated, while the propositions representing the embedded relative clause reading remain highly activated.

The often reported error of interpreting the sentence as "The horse raced past the barn *and* fell" is accounted for by assuming that the past participle reading of *raced* is not available when the last word is read. To simulate this, the nodes P1 and P2 can be deleted. After the integration, both propositions *race[horse]* and *fall[horse]* remain activated,

¹The simulations were carried out using the CI program (Mross & Roberts, 1992).

²Unless stated otherwise, the link strengths used are +1 and -1, and the self-activation +1.

while the question-node *fall???* loses its activation. If in addition the proposition *fall[horse]* is deleted (as can be assumed for readers who are aware of its ungrammaticality), the resulting representation corresponds to the active reading

of *raced*; i.e., to the representation of a reader who failed to recover from the garden path.




Figure 1. The network constructed to process the sentence *The horse raced past the barn fell*. Final activation values, after the integration phase, are displayed on the left. Positive links are shown as solid lines, negative links as broken lines.

2. Example: Sentence Context and Pragmatic Knowledge

The second example concerns the use of sentence context to arrive at a conceptually consistent grammatical interpretation. Taraban and McClelland (1988) compared sentences such as (2) *The janitor cleaned the room with the broom*, and (3) *The janitor cleaned the room with the windows*. In the first case, the correct interpretation is to attach the prepositional phrase to the verb to indicate the instrument of cleaning. In the second case, the attachment of the phrase to the noun *room* as a modifier is appropriate. Again, the alternative interpretations are represented using different propositional representations. The instrument reading corresponds to the formation of the proposition containing three arguments (P1) *clean[janitor,room,broom]*. The noun modification is represented by the two propositions (S1) *clean[janitor,room]* and (S2) *with[room,broom]*.

Since in this example the meaning of the last noun is crucial, elaborations from general world knowledge are included in the simulation. Associations of strength 0.5 between the concepts *cleaning* and *broom*, and between the concepts *room* and *windows*, are assumed. The resulting network for sentence (2) is shown in Figure 2. The competing propositions are again connected by inhibitory links, and the positive connections are defined by argument overlap. In addition, the self activation of the node P1, and the negative connection strength to S1, are increased to indicate the lexical preference for the verb attachment.

After initial activation of the text propositions, the spreading activation process deactivates the inappropriate noun attachment reading. In contrast, integration of the analogously defined network representing sentence (3) leads to deactivation of the verb attachment proposition.

The networks differ only with respect to how the concepts *broom* and *windows* are connected to the text propositions.




Figure 2. The network constructed to process the sentence *The janitor cleaned the room with the broom*. Final activation values after the integration phase are shown on the left. Positive links are indicated by solid lines, negative links by broken lines.

Thus, the simulation finds the correct instrument interpretation for sentence (2) without explicit representation of the verbs' thematic roles or selection restrictions. Using only unlabeled word associations, the constraints in the network are sufficient to disambiguate the syntactic structure.

3. Example: Discourse Context

The final example demonstrates how prior discourse context influences the models' selection of the correct syntactic interpretation. According to the principle of referential support (Altmann & Steedman, 1988), the unpreferred noun attachment in sentence (3) *The janitor cleaned the room with the windows* should be easier to process if it is preceded by the context sentence (4) *There was a room with plants, and a room with windows*, than if it is preceded by (5) *There was a lounge with plants, and a room with windows*. In the first case, in which two rooms are introduced into the discourse, establishing a unique referent to the noun *room* in the target sentence requires a noun attachment, and renders the instrument reading infelicitous. In the second case, with only one room mentioned, the noun modification is redundant. As for processing the sentence in isolation, a verb attachment is preferred and processing difficulties should arise.

To simulate this effect, the context sentences were integrated first. The most highly activated propositions were then carried over into the network constructed for the target sentence. In this simple example, the propositions

carried over from sentence (4) are (C1) *with[room1,plants]*, and (C2) *with[room2,windows]*.

Therefore, including context sentence (4) in the network representing the target sentence involves adding proposition (C1) and linking it to the other nodes according to argument overlap. Similarly, context sentence (5) adds the proposition (C3) *with[lounge,plants]* to the network. In both cases, the final activation values after integration of sentence (3) indicate equally successful final disambiguation. However, an analysis of the time course yields context dependent differences. The activation values of the competing propositions S1 and P1 across the integration phase are displayed in Figure 3.

In the context in which only one room was mentioned, the incorrect proposition has a slightly higher activation level than the correct one for the first half of the integration phase. Its deactivation becomes apparent only after 14 cycles and is completed after 27 cycles. On the other hand, in the context in which two rooms were introduced, the deactivation of the inconsistent proposition is apparent early in the integration phase, at cycle 6, and is completed after 20 iterations.

The construction-integration theory does not claim that the number of cycles in the integration phase can be directly translated into processing times. However, the qualitative differences suggest that selection of one of the alternatives over the other is influenced by the preceding context. This observation is consistent with the predictions of the principle of referential support.




Figure 3. The time course of activation for the two propositions *clean[janitor,room]* (labeled "noun") and *clean[janitor,room,windows]* (labeled "instrument") as a function of context.

4. Example: Discourse Context

The previous example demonstrated how the principle of referential support can be accounted for within the CI framework. As noted above, though, there are empirical results which suggest that, under certain circumstances, discourse context does not influence initial parsing decisions. For instance, Ferreira and Clifton (1986) did not find faster reading times for noun attachment sentences when they were preceded by a felicitous context. Although the

contexts were constructed using the principle of referential support, there was an important difference to the Altmann and Steedman (1988) materials. The referential information was not in the discourse focus, and it was not mentioned immediately before the target sentence. In terms of the CI model, the proposition modifying the object noun is not carried over from the previous comprehension cycle, and can thus not facilitate the selection process immediately.

Rayner et al. (1992) used contexts in which the referential information either was in the discourse focus, i.e., was mentioned in the sentence preceding the prepositional phrase, or in which it was not, i.e., in which the referential information was separated from the target by one intervening sentence. For the prepositional phrase, differences were only found for the total reading times, but not for the first pass reading times. Assuming a small buffer size in the CI model, the proposition representing the referential information is not directly carried over into the target sentence network. Thus, the effects illustrated in Example 3 cannot be expected for these materials. If we further assume that the initial comprehension cycle is reflected in first pass reading times, then the model's predictions are consistent with the empirical findings. The function of rereading, or backtracking, on the other hand, is at present not implemented in the model. However, it seems safe to assume that rereading initiates a second comprehension cycle, using the same textual input. If this re-construction is based on the previously constructed network, additional information can be gained by retrieving more elaborations from long-term memory, including propositions representing prior discourse information. In Rayner et al.'s focus condition, the crucial referential proposition is more easily accessible (through more and stronger connections to the target sentence), and thus it is more likely to be added to the comprehension net.

This account, of course, is still speculative. However, it serves to show how the CI model can be used to analyze the interactions between syntactic and contextual information. In particular, it can be used as a tool to determine which type of context information is likely to influence parsing processes. Further empirical research has to be carried out to see whether the model can be used to make specific predictions which go beyond recency or accessibility arguments.

Conclusions

The results presented here establish the construction-integration model of discourse comprehension as a candidate for a viable theory of sentence processing. The weakly interactive architecture was used to model garden-path effects, as well as to illustrate effects of pragmatic knowledge and discourse context.

Two extensions of the model seem warranted. First, the parser which constructs propositions and establishes their connections is not yet part of the model. For a full description of how syntactic knowledge aids comprehension, a specification of the mechanisms used in the construction phase would be desirable. In particular, it is necessary to assume that the subcategorization structure of lexical items is part of the linguistic representation (cf. Jurafsky, 1992).

Second, the time course of comprehension has not yet been taken into account. In all of the presented examples, the disambiguating information was presented at the sentence boundary, and the simulations were carried out under the simplifying assumption that the entire sentence is integrated at once. To be able to describe on-line effects, processing costs associated with constructing the propositional networks have to be quantified. As shown in

the last example, the function of reading strategies such as backtracking have to be formalized in the model. Moreover, a more detailed description of the interplay between the construction and integration phases is needed, and the locations at which at least partial integration takes place have to be identified (e.g., phrase boundaries). Simulations are currently being carried out using several processing cycles per sentence, in order to explore this issue.

Despite these limitations, the construction-integration model provides a promising approach in sentence processing research. Although there is a multitude of parsing models, many of them are mainly concerned with the processing of single sentences (e.g., Marcus, 1980; McRoy & Hirst, 1990). Other researchers recognize the necessity to take the discourse context into account, but have not included specific mechanisms in their models (e.g., Jurafsky, 1992; Johnson-Laird, 1983). Thus, even the ability to account for global context effects within a psychologically well-founded theory is invaluable. Moreover, the demonstration of the general applicability of the construction-integration model in the domain of syntactic ambiguity resolution has extended the generality of the theory.

Acknowledgments

I would like to thank Walter Kintsch for his encouragement and advice. Paul Smolensky and three anonymous reviewers made valuable comments on an earlier draft of this paper.

References

- Altmann, G. T. M., & Steedman, M. (1988). Interaction with context during human sentence processing. *Cognition*, *30*, 191-238.
- Bresnan, J., & Kaplan, R. M. (1982). Introduction: Grammars as mental representations of language. In J. Bresnan (Ed.), *The mental representation of grammatical relations*. Cambridge, MA: MIT Press.
- Britt, M. A., Perfetti C. A., Garrod, S., & Rayner, K. (1992). Parsing in discourse: Context effects and their limits. *Journal of Memory and Language*, *31*, 293-314.
- Clifton, C., Jr., & Ferreira, F. (1989). Ambiguity in context. *Language and Cognitive Processes*, *4*, S177-104.
- Crain, S., & Steedman, M. (1985). On not being led up the garden path: The use of context by the psychological syntax parser. In D. R. Dowty, L. Karttunen, & A. M. Zwicky (Eds.), *Natural language parsing: Psychological, computational, and theoretical perspectives* (pp. 320-358). Cambridge: Cambridge University Press.
- Ferreira, F., & Clifton, C. (1986). The independence of syntactic processing. *Journal of Memory and Language*, *25*, 348-368.
- Ferstl, E. C. (1993). *The role of lexical information and discourse context in syntactic processing: A review of psycholinguistic studies*. (Tech. Report No. 93-03), Boulder: University of Colorado, Institute of Cognitive Science.
- Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.
- Frazier, L., & Rayner, K. (1982). Making and correcting errors during sentence comprehension: Eye movements in

- the analysis of structurally ambiguous sentences. Cognitive Psychology, 14, 178-210.
- Johnson-Laird, P. N. (1983). Mental models. Cambridge, MA: Harvard University Press.
- Jurafsky, D. (1992). An on-line computational model of human sentence interpretation: A theory of the representation and use of linguistic knowledge. (Tech. Report No. UCB/CSD 92/676). Berkeley: University of California, Computer Science Division.
- Just, M. A., & Carpenter, P. A. (1992). A capacity theory of comprehension: Individual differences in working memory. Psychological Review, 99, 122-149.
- Kintsch, W. (1988). The role of knowledge in discourse comprehension: A construction-integration model. Psychological Review, 95, 163-182.
- Kintsch, W., Welsch, D. Schmalhofer, F., & Zimny, S. (1990). Sentence memory: A theoretical analysis. Journal of Memory and Language, 29, 133-159.
- Lehman, J. F., Lewis, R. L., & Newell, A. (1991). Natural language comprehension in SOAR (Tech. Rep. No. CMU-CS-91-117). Pittsburgh, PA: Carnegie Mellon University, School of Computer Science.
- Marcus, M. P. (1980). A theory of syntactic recognition for natural language. Cambridge, MA: MIT Press.
- Marslen-Wilson, W. D., & Tyler, L. K. (1980). The temporal structure of spoken language understanding. Cognition, 8, 1-71.
- McRoy, S. W., & Hirst, G. (1990). Race-based parsing and syntactic disambiguation. Cognitive Science, 14, 313-353.
- Mitchell, D. C., Corley, M. M. B., & Garnham, A. (1992). Effects of context in human sentence parsing: Evidence against a discourse-based proposal mechanism. Journal of Experimental Psychology: Learning, Memory, and Cognition, 18, 69-88.
- Mross, E. F., & Roberts, J. O. (1992). The construction-integration model: A program and manual. (Tech. Report No. 92-14). Boulder: University of Colorado, Institute of Cognitive Science.
- Perfetti, C. A. (1990). The cooperative language processor: Semantic influences in an autonomous syntax. In D. A. Balota, G. B. Flores d'Arcais, & K. Rayner (Eds.), Comprehension processes in reading (pp. 205-230). Hillsdale, NJ.: Lawrence Erlbaum Associates.
- Rayner, K., Carlson, M., & Frazier, L. (1983). The interaction of syntax and semantics during sentence processing: Eye movements in the analysis of semantically biased sentences. Journal of Verbal Learning and Verbal Behavior, 22, 358-374.
- Rayner, K., Flores d'Arcais, G. B., & Balota, D. A. (1990). Comprehension processes in reading: Final thoughts. In D. A. Balota, G. B. Flores d'Arcais, & K. Rayner (Eds.), Comprehension processes in reading (pp. 631-638). Hillsdale, NJ.: Lawrence Erlbaum Associates.
- Rayner, K., Garrod, S., & Perfetti, C. A. (1992). During parsing discourse influences are delayed. Cognition, 45, 109-139.
- Taraban, R., & McClelland, J. L. (1988). Constituent attachment and thematic role assignment in sentence processing: Influences of content-based expectations. Journal of Memory and Language, 27, 597-632.
- Taraban, R., & McClelland, J. L. (1990). Parsing and comprehension: A multiple-constraint view. In D. A. Balota, G. B. Flores d'Arcais, & K. Rayner (Eds.), Comprehension processes in reading (pp. 231-263). Hillsdale, NJ.: Lawrence Erlbaum Associates.
- Trueswell, J. C., & Tanenhaus, M. K. (1991). Tense, temporal context and syntactic ambiguity resolution. Language and Cognitive Processes, 6, 303-338.
- vanDijk, T. A., & Kintsch, W. (1983). Strategies of discourse comprehension. New York: Academic Press.