The Construction-Integration Model: A Framework for Studying Memory for Text.

Walter Kintsch & David M. Welsch

Institute of Cognitive Science
Box 345
University of Colorado
Boulder, Colorado 80309-0345
303/492-8663

Institute of Cognitive Science
Publication Number 90-15

To appear in Hockley, W. E. & Lewandowsky S. (Eds.) Relating theory and data: Essays on human memory. Hillsdale, NJ: Erlbaum.

The construction-integration model: A framework for studying memory for text.

by

Walter Kintsch & David M. Welsch
University of Colorado

Models of memory need not be restricted to list learning studies but can be extended to account for how people comprehend and remember text. The basic memory processes appear to be the same in both cases, but what is considerably more complex in the case of texts are the comprehension and encoding processes involved. Below we sketch a model of discourse comprehension and memory for text that has evolved over a number of years. In particular, we focus on the question how the information in a text is combined with and integrated into the reader's knowledge and personal store of experience. We use this model to account for two results that have been reported in the literature on priming. First, we analyze a study that explores the differences in the accessibility of the first- and second-mentioned actors in a sentence, as measured by priming methods. Secondly, we consider a study demonstrating that it is not the words of a sentence by themselves, but the whole discourse context that determines knowledge activation, and hence the degree of priming that is obtained experimentally. Our purpose in presenting these analyses is to show that the model presented below is able to account for these phenomena in a simple and consistent way, without complex auxiliary assumptions, although it was not specifically designed to do so.

1. The construction-integration model as a cognitive architecture

The construction-integration (CI) model is a general theory of comprehension. It has evolved from studies of story comprehension and memory for text to a point where it can be considered seriously as a possible architecture for that large area of cognition that we locate between "perception" on the one hand and "problem solving" on the other. In everyday language we name these processes "comprehension", although we do not make precise distinctions, say, between "perceiving" a scene and "comprehending" it, or "solving" a word algebra problem and "understanding" it. The CI framework allows us to construct explicit computational models of processes within that general domain.

The principle characteristics of the CI model are the following:
(1) The CI theory extends and incorporates earlier work on discourse comprehension, specifically Kintsch & van Dijk (1978) and van Dijk & Kintsch (1983). The primary feature that it adds to this earlier work is a model of knowledge activation and knowledge use in discourse comprehension (Kintsch, 1988).
(2) Comprehension in the CI framework arises from an interaction and fusion between the to-be-comprehended object, usually a text, and the general knowledge and personal experience the comprehender brings to the situation.
(3) Knowledge is represented as an associative network, the nodes of which are concepts and propositions. Global knowledge structures such as frames and schemata are generated in the context of the particular task in which they are used from the local relations of this associative net.
(4) The CI theory is a hybrid approach, combining features of symbolic systems (its first phase, construction, involves a rule-based system that constructs a network representation of the text and the knowledge that has been activated) and connectionist systems (the second phase of the model, integration, uses a constraint satisfaction mechanism that generates a consistent interpretation of the to-be-comprehended text.
(5) Purely symbolic systems need "smart" rules and sophisticated control structures such as schemata to assure that the right inferences are made at the right time, and that the relevant knowledge, and only the relevant knowledge, gets activated in each

situation. In the CI model we can do with weaker, more general, more robust rules, because the rules will not have to be precisely correct. A process of contextual integration follows the construction of a mental representation of the text, ensuring that contextually irrelevant items become deactivated in the network. In other words, the CI model does a lot of parallel inferencing with weak, simple rules, and then eliminates the inferences that are irrelevant in a given context.

Below, we briefly outline how this construction-integration model works in general, and then apply it to some interesting experimental phenomena that have been reported in the literature on priming effects.

2. The construction-integration model

2.1 Construction processes.

2.1.1 Text-based constructions: Levels of representation

There is no single representation of a text that is appropriate for all purposes. Instead, different levels of representations must be distinguished, primarily the verbal or linguistic level, the level of meaning, and the level of the situation, though at times additional levels need to be considered.

The units at the linguistic level are the words that make up the text and the (hierarchical) sentence constituents to which these words belong. The rules that are used to form these units are those of conventional phrase-structure grammars. In many applications of the CI model, this level of analysis has been neglected because surface effects play only a negligible psychological role. In some situations, however, it is crucial to include this level of analysis (e.g., to describe memory for sentences, as in Kintsch, Welsch, Schmalhofer, & Zimny, 1990).

At the semantic level, the representation used is the propositional analysis proposed in Kintsch (1974), though various related representation systems developed by other authors in linguistics, artificial intelligence, and psychology could be readily adapted for that purpose. The rules for generating text propositions can be formulated as productions. Both local meaning units and global meaning units (macropropositions) must be considered. A sketch of what a complete system would look like has been presented in Kintsch (1985). However, an explicit and complete set of rules to formally derive propositions from a text (except for restricted domains) is not available at this point, either for the Kintsch (1974) model or any of the other systems in use. Since conventions for hand-coding are well developed and highly reliable, this is not a serious problem. In the applications of the CI model that have been developed so far, the semantic/propositional level is typically central (except in poetic language use, as analyzed in Kintsch, in press, a).

While it makes some sense to postulate general rules for generating propositional representations of texts (even if we can't specify all of these rules at this point), general, comprehensive, formal rules for constructing situation models are quite implausible. Situation models are by definition domain specific and hence depend on special domain knowledge, in contrast to verbal and semantic representations which are more general and domain independent. The model of the situation described by a given text is developed from the interaction between the information provided by the text and further information derived from the reader's knowledge. Situation models emerge from this interaction. Linguistic cues in the text (Kintsch, in press, b) guide the reader in forming a situation model. The semantic relations specified by the text constrain what is situationally relevant, and activate the knowledge that is needed for building a mental model of the situation. These processes are most easily analyzed in well-defined, highly structured situations, hence much of the work done with the CI model has involved such situations (word arithmetic problems in Kintsch, 1988; computer systems in Mannes & Kintsch, in press; the UNIX operating system in Doane, Kintsch, & Polson, submitted).

2.1.2 Knowledge-based constructions

Knowledge in the CI model is represented as a long-term memory network of propositions (concepts are treated as a special kind of proposition). The links between these propositional nodes vary in strength, as determined by the associative and/or semantic relations between the concepts or propositions in question.

Propositions and concepts from this long-term memory network are added to the text representation (the "textbase") under construction in two ways. First, in the construction of text propositions long-term memory nodes are used as templates (Kintsch, 1985), and nodes in the neighborhood of these templates are randomly sampled from long-term memory and added to the textbase as knowledge elaborations, their sampling probability being determined by the strength of their association to the template. This associative knowledge activation plays a major role in some applications, such as Kintsch, 1988, or Mannes & Kintsch, in press, but not in others, such as Kintsch et al., 1990, and Doane et al. (submitted). A second way in which knowledge from the long-term store gets included in the text representation is through special reading goals and/or task demands. Thus, when reading a word arithmetic problem, the reader's complete knowledge about arithmetic is activated (Kintsch, 1988); all the knowledge about a computer system is activated in parallel when comprehending instructions to perform a computing task (Mannes & Kintsch, in press; Doane et al., submitted).

As a result of these text- and knowledge-based construction processes, a set of $n+m$ elements is obtained, n elements derived from the text (words, phrase units, concepts, propositions, or model elements), plus m_1 knowledge propositions which have been selected from the long-term memory net by the associative activation process described above, and m_2 knowledge propositions that have been selected in response to specific task demands, $m_1 + m_2 = m$.

2.2 Integration processes

The $n+m$ elements that have been constructed are linked together to form a network. The links among the n text propositions are determined by the relationships among these elements specified by the text: words are linked to the phrases they are a component of, phrases to sentences and so on (Kintsch et al., 1990); propositions are linked via argument overlap to their nearest neighbors (Kintsch, in press, a, b) or even to other propositions one or more steps away (Kintsch, 1988; Kintsch et al., 1990), and similarly for elements of the situation model. Numerical link strengths constitute a free parameter in the model, subject to the restriction that link strengths decrease the greater the distance between nodes.

The m_1 knowledge propositions that were activated associatively are linked to the text elements through which they were selected in the first place. Their interconnections are the same as in the long-term memory net.

For the m_2 task-specific knowledge propositions, a variety of links exist: they are related to the rest of the network through argument overlap, and they inherit whatever pattern of interrelationships existed in long-term memory (e.g., contradictory arithmetic hypotheses have negative connections in Kintsch, 1988; action plans may have negative or positive connections, depending upon whether the actions support or interfere with each other in Mannes & Kintsch, in press, and Doane et al., submitted), but there are further considerations. For instance, propositions describing the state of the world inhibit action plans that would produce the already existing state, and propositions describing the desired outcome of an instruction or command activate action plans that would produce these outcomes (Mannes & Kintsch, in press, and Doane et al., submitted). Once again, the numerical values of these positive and negative link strengths are free parameters in the model.

Thus, a matrix C of $n+m$ items is obtained, where the element c_{ij}, $1< i,j < n+m$, specifies the link strength between items i and j. We shall call this matrix the coherence matrix. We define a row vector $\mathbf{A}_1$ which specifies an activation value a_i for each of the $n+m$ items in the coherence matrix. Initially, all items that were constructed from the text are equally activated with activation strength $1/n$, while the m items selected from long-term memory have an initial activation value of 0, so that $\Sigma a_i = 1$.

Activation is allowed to spread in this network by taking the products $\mathbf{A}_1 * \mathbf{C} = \mathbf{A}_2$, $\mathbf{A}_2 * \mathbf{C} = \mathbf{A}_3$,.....$\mathbf{A}_{k-1} * \mathbf{C} = \mathbf{A}_k$, or $\mathbf{A}_1 * (\mathbf{C})^k = \mathbf{A}$, where the vector $\mathbf{A}_i$ is renormalized after each multiplication with $\mathbf{C}$, so that $\Sigma a_i = 1$, and k is selected so that the change in mean activation value after a multiplication is less than some criterion value, e.g. <.0001. The final activation vector **A** shows how strongly each item that was generated during the construction process, either from the text or from long-term memory, is activated after comprehension. If the integration process did what it was supposed to do, related items should have strengthened each other, while unrelated or contradictory items should have a negative or 0 activation value. In most applications it is useful to avoid negative activation values, so that all such values are set to 0 before normalization after each cycle of multiplication.

Table 1

Table 1 shows the final activation vector **A** for the n+m nodes in the network, as well as the interconnections among these nodes specified by the coherence matrix **C**. Note that, in general, some of the elements in **A** will have 0 activation, that, is, they will have become deactivated in the integration process either because they are not sufficiently strongly connected with the main part of the network (e.g., an irrelevant piece of knowledge that was generated by the sloppy knowledge elaboration rules of the model), or because they are inhibited by strongly activated nodes in the network (e.g., the contextually inappropriate meaning of a homonym).

The CI model achieves by means of this integration process what other approaches do with complex, powerful - but fragile - control process rules: it assures that the inferences and knowledge elaborations which have been generated will be contextually appropriate. In most theories of comprehension this happens because only the right inferences and elaborations are made in the first place. In the CI model it happens because the integration process has filtered out the contextually inappropriate inferences and elaborations that had been generated in the initial phase.

2.3 Long-term memory for text

So far, the CI model has constructed a set of $n+m$ items, each with a certain activation value, and a coherence matrix specifying the interrelationships among these items. Welsch (1989) has shown that these two pieces of information can be combined to yield a more convenient representation of the outcome of the comprehension process, as in pure connectionist models. In connectionist models, learning affects the connections among the nodes in the network. In the CI theory as described so far, comprehension has not affected the connections among the elements - the coherence matrix **C** - but rather the activation values of the elements - the activation vector **A** after integration. We can define a new matrix **M**, the memory strength matrix, of size $p+q$ * $p+q$, with elements m_{ij}, such that

$$m_{ij} = c_{ij}*a_i*a_j$$

where $p+q$ is the number of elements in **A** with positive activation values ($p \leq n$, $q \leq m$), c_{ij} is an element of **C**, and a_i is the final activation value of the *i-th* element. **M** is the long-term memory representation the CI model produces. The knowledge elaborations that were created during the construction process that were contextually irrelevant or contradictory have been eliminated. The diagonal values of **M** represent the strength of an item in long-term memory (the square of its final activation value). The off-diagonal elements represent the strength of the relation between any two items in memory.

Whether **M** is best characterized as a surface representation of the text, a semantic textbase, a situation model, or some combination thereof, depends upon the nature of the construction processes that have taken place. If these were dominated by surface features, with little contribution from the semantics, and none from the situation, as in the children's counting rhyme analyzed in Kintsch (in press, a), **M** is primarily a surface representation. On the other hand, in Section (4) of this chapter we construct a pure textbase, arguing that the contributions of the surface structure as well as the situation model are negligible in this case. Contrast this with Mannes & Kintsch (in press), where the emphasis is entirely on the situation model: the texts are brief and simple, and we are primarily concerned with the elaborate knowledge structures that emerge in these situations. Thus, **M** represents the sum total of the construction processes that occurred during comprehension at various different levels, from the linguistic surface analysis to the building of a situation model. Surface structure, textbase, and situation model, however, are not separate mental objects, but convenient terms for us to designate the focus of text processing in different comprehension tasks.

2.4 Text Memory and Knowledge Modification

It is important to separate two effects of comprehension. On the one hand, an episodic memory trace for the text comprehended has been established - the memory matrix **M** - which supports a number of text-based behaviors, such as recognition, recall, summarization, question answering, and the like. On the other hand, the reader's long-term memory will be modified, at least temporarily, as a consequence of reading the text.

Sentence recognition (Kintsch et al., 1990) and priming (Sections 3 and 4 below) in the CI theory involve connecting the to-be-recognized sentence with the appropriate episodic memory trace as if it were itself part of the text, and noting the amount of activation that will flow into the subnet corresponding to the target item.

Recall predictions can be obtained from **M** by selecting a first-cycle element with a probability proportional to its memory strength, and then probabilistically selecting the next element, and so on, tracing a path through the whole memory network (this process needs to be elaborated with rules for backtracing and repetitions). Each such path corresponds to a single predicted recall protocol. Summarization can be handled in the same way by including a threshold, such that only sufficiently important nodes are selected (Kintsch, in press b).

We shall not describe here the modification of long-term memory effected by the establishment of an episodic memory trace of a text, because the various alternatives have not yet been explored in enough detail. It is important, however, to understand that forming an episodic text memory is a different problem, with different empirical consequences, than adding to or modifying the knowledge base itself (e.g., Mannes & Kintsch, 1987).

This brief sketch of the CI theory must suffice here. Additional detail can be found in the publications cited above. Below, we take two experimental phenomena that have been reported in the literature and show how the CI model can deal with them. If the model provides, indeed, anything like an adequate account of human comprehension processes, the phenomena observed in these studies should be explainable in terms of the CI model, without modifying its assumptions, and without the need for ad hoc mechanisms. We also want to avoid extensive parameter estimations, hence we shall restrict ourselves to qualitative simulations of the major trends in the experimental results. Our goal is simply to find out whether the CI model can account for these reasonably complex experimental data for which the model was never designed specifically. Both experiments have to do with priming effects after or during discourse comprehension.

3. First mention versus recency: Some priming results.

Gernsbacher, Hargreaves, & Beeman (1989) report some interesting results that reconcile two seemingly contradictory phenomena that have been frequently reported in the literature. On the one hand it is well known that the first-mentioned participant in a sentence has an advantage over the second-mentioned participant (it is accessed more readily, functions better as a recall cue, etc.), while on the other hand a similar advantage for the most recent clause is equally well established. Gernsbacher et al. have shown that the recency advantage holds for tests that coincide and overlap with the processing of the sentence, while for tests delayed for 1.4 sec after the end of the sentence, a first mention advantage is found.

What does this observation mean? Does it imply some new principles of comprehension, or is it an implied empirical consequence of comprehension processes as described by the CI model? If so, the CI model should be able to account for the findings of Gernsbacher et al. without requiring additional assumptions. We should simply be able to simulate the experimental procedures of Gernsbacher et al., and find their results.

This is indeed the case, as we shall show by means of an illustrative simulation of the example sentence discussed by Gernsbacher et al. In one case, we simulate comprehension of this sentence up to the last word, coincident with which the probe was presented (either the first mentioned or second-mentioned participant in the sentence). In the other case, we allow the model to complete comprehension of the sentence and then present the probe. In the first case, the more recent probe should be more activated, and the in the second case the the first-mentioned probe.

The sentence we use is

(1) *Tina gathered the kindling as Lisa set up the tent.*

In line with Gernsbacher et al.'s claim that subjects represent each clause in its own substructure, we let the CI model process (1) in two cycles, one for each clause. We simulate processing both at the level of the linguistic surface structure and that of the propositional textbase, while neglecting the level of the situation model, which probably plays no role when subjects are working with lists of unrelated sentences. Furthermore, we assume a very simple form for the coherence matrix: only neighboring nodes are linked directly, and semantic connections are twice as strong as surface connections. Thus, we obtain for the first clause of (1)

(2)

	Tina	gather	kindl	N1	V1	TINA	G(T,K)	KINDL
Tina	1	0	0	1	0	1	0	0
gather	0	1	0	0	1	0	1	0
kindlg	0	0	1	0	1	0	0	1
N1	1	0	0	1	1	0	0	0
V1	0	1	1	1	1	0	0	0
TINA	1	0	0	0	0	2	2	0
G(T,K)	0	1	0	0	0	2	2	2
KINDL	0	0	1	0	0	0	2	2

where V1 is the verb phrase *gather the kindling*, and N1 the whole clause; the first three nodes are the words "*Tina*", etc., the last three are the propositions TINA, GATHER(TINA, KINDLING), and KINDLING. Initially, the eight nodes of this network are equally activated. After three cycles of spreading activation, a reasonably stable activation vector is obtained.[1] The five surface nodes have activation values of .07, .09, .07, .07 and .05, respectively, while the final activation values for the three semantic nodes are .19, .27, and .19, respectively. Thus, GATHER(TINA, KINDLING) is the most highly activated node in the first cycle and it is carried over to the second processing cycle.

The area labeled "Working Memory" in Figure 1 shows what happens when the second clause of (1) is processed. Words are presented one at a time, and Figure 1 shows the network that has been constructed when all but the last word of the sentence have been read. Once again, we assume links of strength 1 among the surface elements of the sentence, and strength 2 among the meaning elements. The last word, "tent", is presented simultaneously with either the test word "Tina" or "Lisa", the first- or second-mentioned participant, respectively. Only the proposition (concept) TINA ties directly into working memory at all when the target item is "Tina", while both the word "Lisa" and the corresponding semantic element LISA connect to fairly strongly activated nodes in the working-memory network. Hence the results in Figure 2: the more recent participant is more strongly activated than the first-mentioned participant. Under these conditions, the mean reaction time in Gernsbacher et al. (Exp. 1) was 1,118 msec for the first-mentioned participant, but only 1,065 msec for the more recently mentioned one.

Figures 1 & 2

Now consider what happens when the same probes are presented well after the sentence has been read. Both clauses of the sentence are now represented as a long-term memory trace in which the original links have been modified by the activation values of the nodes connected. For (1), we obtain the network shown in the area labelled "Long-term Memory" in Figure 3. Both the probes "Tina" and

---

[1] Changes at this point are no more than .01; this loose criterion will be used throughout in this example.

"Lisa" connect to this network in three places, as shown in Figure 3, but the first-mentioned participant "Tina" is linked to stronger nodes, and hence receives more activation than the more recent "Lisa". The activation values on a delayed test for these two probes are shown in Figure 4. In Gernsbacher et al. (Exp. 3), we find a reaction time of 726 msec for the more highly activated first-mentioned participant, versus 788 msec for the second-mentioned participant.

Figures 3 & 4

The CI model thus implies the Gernsbacher et al. results without any additional assumptions. It is true that the model has a large number of parameters, but these predictions do not depend so much on particular parameter choices: if we choose a buffer of size 2, so that the proposition TINA also gets carried over into the second processing cycle, the difference in the activation values between the first- and second mentioned participants is somewhat reduced on the concurrent test, but the pattern of predictions remains the same. Likewise, if we weight the original link strengths less heavily in favor of semantic links a similar pattern is nevertheless obtained..

The CI model implies that the Gernsbacher et al. results occur for structural reasons and not because we made some lucky guesses about the parameters of the model.These reasons are quite in agreement with the explanation given by Gernsbacher et al. themselves, as well as with many other results in the memory literature. The model gives the kind of results that it does as long as we let it process the two clauses of (1) separately, in agreement with Gernsbacher et al.'s separate structure assumption. While it is still working on the second clause, the partial representation it has constructed is active in working memory, and information in the second clause must be more accessible than in the first clause. On the other hand, if we are dealing with the long-term memory representation of the sentence, a primacy effect will be obtained for much the same reason as in list learning studies. More resources are available for processing the first part of the sentence, resulting in a stronger memory representation. One should be able to reduce this primacy effect by making the first clause longer than the second (and thereby reducing the activation available per element in the representation), or by providing more processing opportunities for the relevant material in the second clause, e.g., by adding a suitable third clause. Nevertheless, it is hard to see how one could completely avoid a primacy effect by such means.

Figures 2 and 4 also illustrate a limitation of the CI model as it is presently formulated. In the actual experiment, concurrent targets took much longer overall than delayed targets, a result which is intuitively quite plausible. However, the average activation values for the two targets that were predicted by the model are not directly comparable at all. This is an undesirable consequence of the normalization procedure used in the calculations: if activation vectors are always made to sum to 1, we cannot meaningfully compare vectors of different length. Hence we have no way of accounting for the different overall levels of performance on the two tests.

4. Contextually relevant aspects of meaning: More priming.

In a well-known study McKoon & Ratcliff (1988) demonstrated that priming effects were highly dependent upon the meaning of a text as a whole, not just on the words used in the text. Which feature of a word is important depends on what is stressed in the context. McKoon & Ratcliff point out that a model in which activation spreads from words whose meanings are fixed regardless of context to related words is counterindicated by these data. We shall show, however, that a spreading activation process such as that posited by the CI model provides a good account of their data. In the CI model, words "mean" different things in different contexts because their pattern of associations with other concepts and propositions is different.

Consider the three texts (3), (4) and (5), and the priming pair (6a and b):

(3) *The still life would require great accuracy. The painter searched many days to find the color most suited to use in the painting of the ripe tomato.*

(4) *The child psychologist watched the infant play with her toys. The little girl found a tomato to roll across the floor with her nose.*

(5) *While eating his bacon and tomato sandwich, the painter searched for the colors most suited to the still life painting of the rustic English countryside.*

(6) *(a)* *Tomatoes are red.*
(b) *Tomatoes are round.*

Tomatoes are red is a semantic feature emphasized in (3), but not in (4) and (5). On the other hand, (4) emphasizes *Tomatoes are round* , while neither of these aspects is particularly relevant in the case of (5). Correspondingly, McKoon & Ratcliff (1988; Exp.1) observed a matching effect in the response times for sentence targets such as (6a) and (6b): when the text emphasized color, *red* was more primed, but when the text emphasized rolling, *round* was more primed. Response times for matching targets averaged 1,270 msec (with 4.7% errors), versus 1,390 msec (with 9.7% errors) for targets that did not match. For control sentences such as (5), both targets were responded to equally, on the other hand. Intuitively, this context dependency of meaning makes a great deal of sense. We seem to be dealing here with an important characteristic of the semantic processing of text. Is it implied by the basic architecture of the construction-integration model, or do we need to introduce some ad hoc mechanism to account for this phenomenon?

We proceed by simulating comprehension and test in the same way as in the example above. The CI model "reads" the three texts in a succession of iterative cycles, and is then given either one of the target sentences (6a or b). More activation should flow into (6a) after reading (3) than after reading (4), and conversely for (6b), but there should be no difference after reading (5). Once again, we do not attempt a quantitative fit and estimate no parameters. We are simply working with illustrative numbers, attempting to demonstrate the right kind of qualitative trends.

Since we are not dealing with on-line priming effects as in the previous example (the target sentences are presented after the whole text has been read), we can neglect surface relations and simply construct a propositional network in which neighboring propositions are related by a value of 1, all other relations being 0. Thus, among the propositions derived from the first sentence of (3), (REQUIRE, STILLIFE, ACCURACY) turns out to be the most strongly activated one once the integration process has settled. It is therefore carried over to the next cycle which is comprised of the first part of the second sentence (seven propositions, up to the word *color*.) Up to now we have not included any knowledge elaborations in our network, for the simple reason that including them would not have made much difference for our purposes. Now, however, we encounter a knowledge association that will make a difference, namely, the association *red* to *color*. We add that as a 9th proposition to our matrix, but since textual relations should probably be weighted more strongly than knowledge elaborations, we assign it a link strength of only .5 (this differential weighting merely decreases quantitatively the priming effects predicted by the model). Performing the integration now yields (FIND, PAINTER, COLOR) as the most highly activated proposition of this cycle. The third cycle then consists of this buffer proposition, 5 more text-derived propositions from the latter part of the sentence, and two knowledge elaborations, the associations *red* and *round* to *tomato*.

The long-term memory trace which is the product of these three cycles of processing is illustrated graphically in Figure 5. We also show in this figure the two target sentences *Tomatoes are red* and *Tomatoes are round* (although in actual fact they are never presented simultaneously and the calculations are made separately, of course). It is obvious from Figure 5 that more activation will flow into *Tomatoes are red* than into *Tomatoes are round.* Indeed, this is what happens. We add (RED,TOMATO) to the LTM matrix graphically represented in Figure 5, connecting it in the same way as the matrix element RED. This new coherence matrix is then used to integrate a starting vector in which the activation of each element is equal to the corresponding diagonal element in the matrix, except for the target sentence, which starts out with an activation of 0. The integration process yields an activation value of 7 (normalized value x 10,000). Using (ROUND,TOMATO) in the same way produces an activation value of 4, on the other hand.

Figures 5 & 6

However, if we process (4) in exactly the same way, once with RED and once with ROUND, as shown in Figure 6, we find an advantage for ROUND over RED, with activation values of 31 and 1, respectively. Thus, the matching effect reported by McKoon & Ratcliff turns out to be a byproduct of comprehension in the CI model. Understanding a sentence implies contextual semantic elaboration. We should note that, once again in agreement with the experimental results, (6a) and (6b) are equally activated if they derive their activation from (5): both have activation values of 1 in that case.

While the CI model thus correctly predicted greater activation of (6a) than (6b) after reading (3), and the opposite effect after reading (4), the model also predicted a much greater differential in the activation values in the second case than in the first. We don't know whether that is a correct prediction or a flaw in the model. The experimental data are averaged over many sentences like (3) and (4). Even though the experimenters will have tried to make all these sentences as comparable as possible, there certainly remain large inter-sentence differences - the matching effect will be bigger for some sentences than for others. The CI model should be able to account for all of these differences - if we simulate all the experimental materials separately, and determine the optimal parameter values. Thus, at some point we shall no longer have to maintain the fiction that all sentences are alike, with normally distributed random error. We can fully deal with the idiosyncratic properties of each sentence. The CI model also has the potential for dealing with differences between readers, who of course are not all alike either, in terms of differences in the buffer capacity and knowledge base.

5. Conclusion

We have taken two results from priming experiments in the literature and presented calculations that show that the CI model can generate predictions that are qualitatively in accord with these data. What is the point of such an exercise? We have not predicted any surprising new phenomena - clearly, these are postdictions. We have not even proposed any surprising new explanations for these experimental phenomena. In the case of Gernsbacher et al. (1989), all the CI model does is provide an explicit mechanism for processes that the authors claimed were involved. ("Separate substructures" is mapped into the processing cycles of the model, "primacy" appears as a consequence of resource availability). The CI model explanation in terms of spreading activation for the McKoon & Ratcliff (1988) data is hardly novel either, though these authors had rejected a simple, word based spreading activation theory. Nor have we presented elegant fits of model predictions to the data.

If not novel data, nor novel explanations, and not even precise fits, then what does the CI analysis offer? If the CI theory is supposed to serve as an architecture for cognition, or at least that part of cognition that we term "comprehension", and if it is supposed to be a viable alternative to architectures such as SOAR or ACT*, then demonstrations that the CI framework offers a flexible and powerful framework for the analysis of a wide variety of phenomena, such as the ones offered here and in Kintsch (in press), play a very important role. They show that without having to introduce all sorts of ad hoc mechanisms and without the need for tricky parameter estimations, the CI model can indeed account for experimental results for which it was never specifically designed. (Note that to do so, we need to have reliable data: nothing but confusion would have resulted had we tried to apply the model to the contradictory, apparently inconsistent data on first- and second-mention before Gernsbacher et al. established just what the experimental facts were.)

We also need to find out the limits of the CI approach . How far does the realm of "comprehension" extend, and is the model really adequate throughout that range? Where and how does it have to be complemented by more controlled, goal-directed problem-solving processes on the one hand, and perceptual mechanisms on the other? In terms of Newell's analysis (Newell, 1990), comprehension is at the mid-level of cognitive processes. The CI theory is not well developed at the millisecond level (though our analyses of priming phenomena reach down to that level), nor at the level of minute- and hour-long deliberations (though Doane, Kintsch, & Polson, submitted, certainly makes some inroads there). Informal, qualitative theoretical analysis of existing experimental results, like those presented above, are a quick and simple means for exploring and testing the limits of the CI theory of comprehension proposed here. By demonstrating that the explanations that have been proposed in the literature for well-established experimental phenomena already have corresponding, explicit mechanisms in the CI model, as we have done in the case of the the Gernsbacher et al. data, or that there exist plausible alternative mechanisms within the model to explain the phenomena in question which have not been considered fully by the original authors, as was the case with McKoon & Ratcliff, we are making a strong case for the generality of the CI theory of comprehension.

References

Doane, S. M., Kintsch, W., & Polson, P. G. (submitted). Modelling UNIX command production: What experts must know.

Gernsbacher, M. A., Hargreaves, D. J., & Beeman, M. (1989). Building and accessing clausal representations: The advantage of first mention versus the advantage of clause recency. Journal of Memory and Language, 28, 735-755.

Kintsch, W. (1974). The representation of meaning in memory . Hillsdale: Erlbaum.

Kintsch, W. (1985). Text processing: A psychological model. In T. A. v. Dijk (Ed.), Handbook of discourse analysis (pp. 231-243). London: Academic Press.

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

Kintsch, W. (in press a). Kognitionspsychologische Modelle des Textverstehens: Literarische Texte. In K. Reusser, & M. Reusser (Ed.), Verstehen lernen - Verstehen lehren Bern: Hans Huber.

Kintsch, W . (in press b). How readers construct situation models for stories: The role of syntactic cues and causal inferences. In A. F. Healy, S. Kosslyn, & R. M. Shiffrin (Ed.), Essays in honor of William K. Estes Hillsdale, N.J.: Erlbaum.

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

Kintsch, W., Welsch, D., Schmalhofer, F., & Zimny, S. (1990). Sentence memory: A theoretical analysis. Journal of Memory and Language, 29, 133-159.

Mannes, S., & Kintsch, W. (1987). Knowledge organization and text organization. Cognition and Instruction, 4, 91-115.

Mannes, S. M., & Kintsch, W. (in press). Planning routine computing tasks: Understanding what to do. Cognitive Science, , McKoon, G., & Ratcliff, R. (1988). Contextually relevant aspects of meaning. Journal of experimental psychology: Learning, memory, and cognition, 14, 331-343.

Newell, A. (1990). The 1987 William James Lectures: Unified Theories of Cognition . Cambridge, MS: Harvard University Press.

van Dijk, T. A., & Kintsch, W. (1983). Strategies of discourse comprehension . New York: Academic Press.

Welsch, D. (1989). A connectionist approach to recognition memory for sentences. Unpublished M. A., University of Colorado, Boulder,

Footnote

This work was supported by Grant MH 15872 from the National Institute of Mental Health. We thank Mike Mozer for his suggestions regarding the long-term memory representation in the model and Paula Schwanenflugel for her comments on the manuscript.

List of Figures

Figure 1. The contents of working memory at the moment when the target items are presented. The proposition GATHER[TINA,KINDLING], shown in the highlighted area, is carried over from the first cycle in the short-term memory buffer. Surface elements are printed in small letters, propositions in capitals. The activation values of each element are indicated. The target items *Tina* or *Lisa* occur concurrently with the last word of the sentence, *tent*.

Figure 2. Activation values (x 100) for the targets *Tina* and *Lisa* when tested concurrently with the last word of the sentence. The extent to which the activation derives from the propositional and surface level is indicated.

Figure 3. The target items with the long-term memory trace of (1). S2, N2, and V2 represent word groups, S1 the whole sentence. Surface elements are printed in small letters, propositions in capitals. Strength values are shown for the memory elements directly connected with the target items.

Figure 4. Activation values (x 100) for the targets *Tina* and *Lisa* when tested against the long-term memory trace of the sentence. The extent to which the activation derives from the propositional and surface level is indicated.

Figure 5. A fragment of the long-term memory trace for (3) with the targets *Tomatoes are red* and *Tomatoes are round.* The memory strength values (x 10000) of the nodes are also indicated.

Figure 6. A fragment of the long-term memory trace for (4) with the targets *Tomatoes are red* and *Tomatoes are round.* The

memory strength values (x 10000) of the nodes are also indicated.

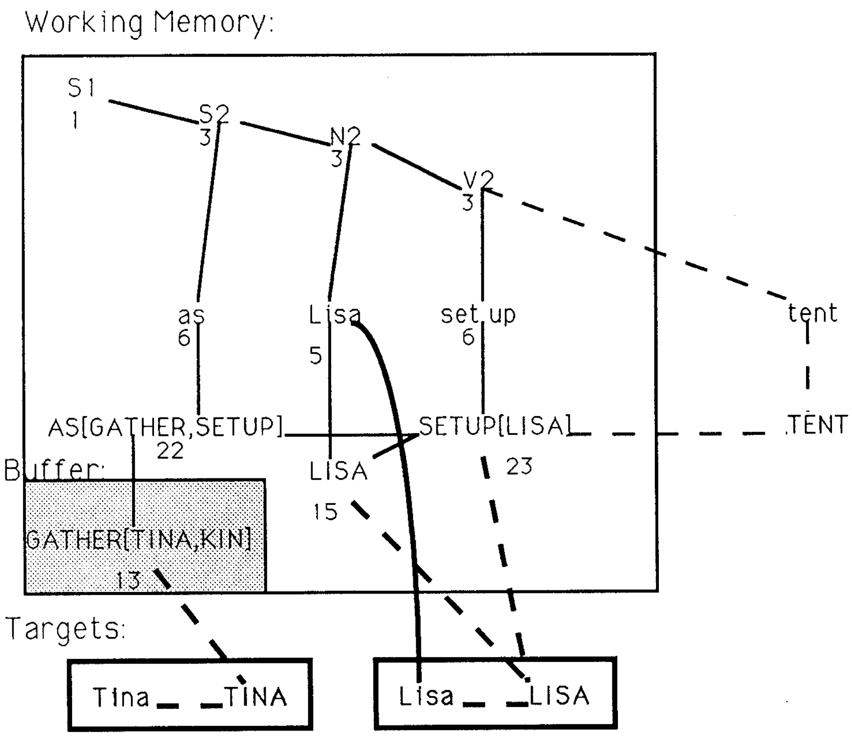
Working Memory:
S1
1
S2
3
N2
3
V2
3
as
6
Lisa
5
set up
6
tent
AS[GATHER,SETUP]
22
SETUP[LISA]
23
TENT
LISA
15
Buffer:
GATHER[TINA,KIN]
13
Targets:
Tina
TINA
Lisa
LISA

Fig 1

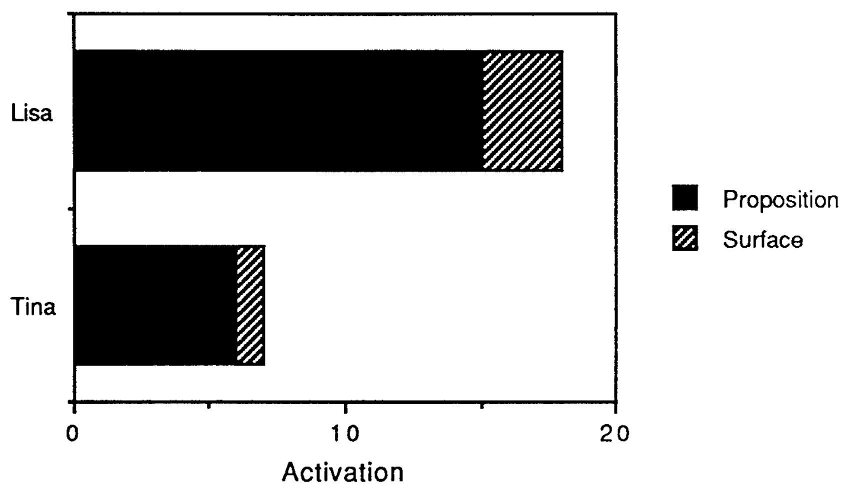

Fig 2

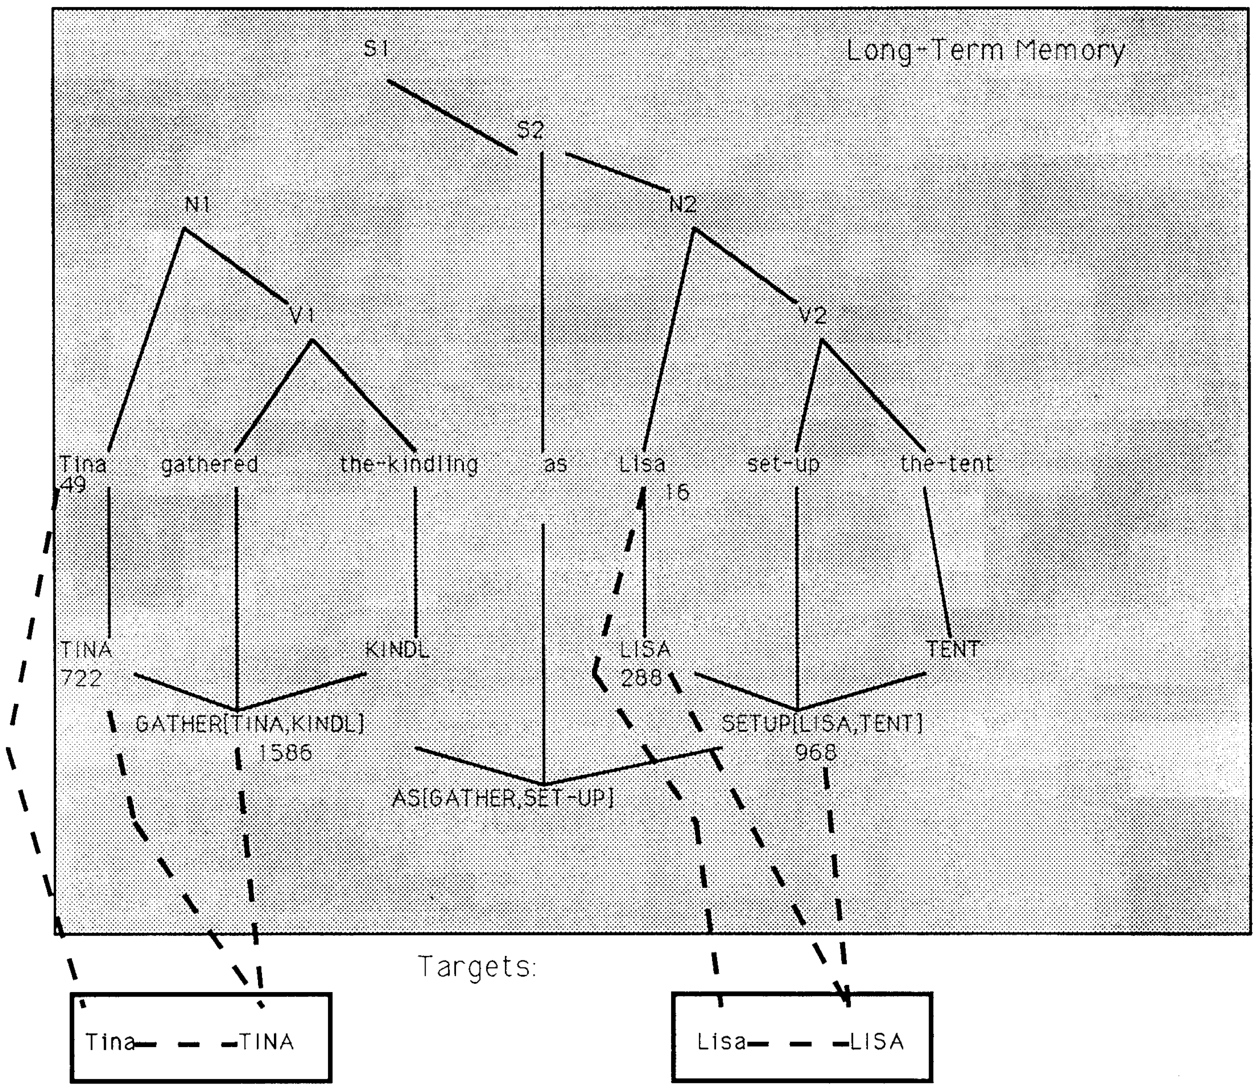
Long-Term Memory
S1
S2
N1
N2
V1
V2
Tina
49
gathered
the-kindling
as
Lisa
16
set-up
the-tent
TINA
722
KINDL
LISA
288
TENT
GATHER[TINA,KINDL]
1586
SETUP[LISA,TENT]
968
AS[GATHER,SET-UP]
Targets:
Tina TINA
Lisa LISA

Fig 3

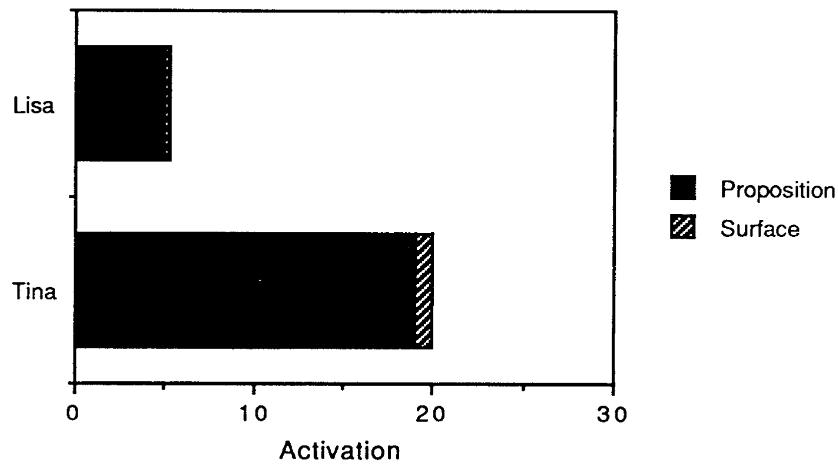

Fig 4

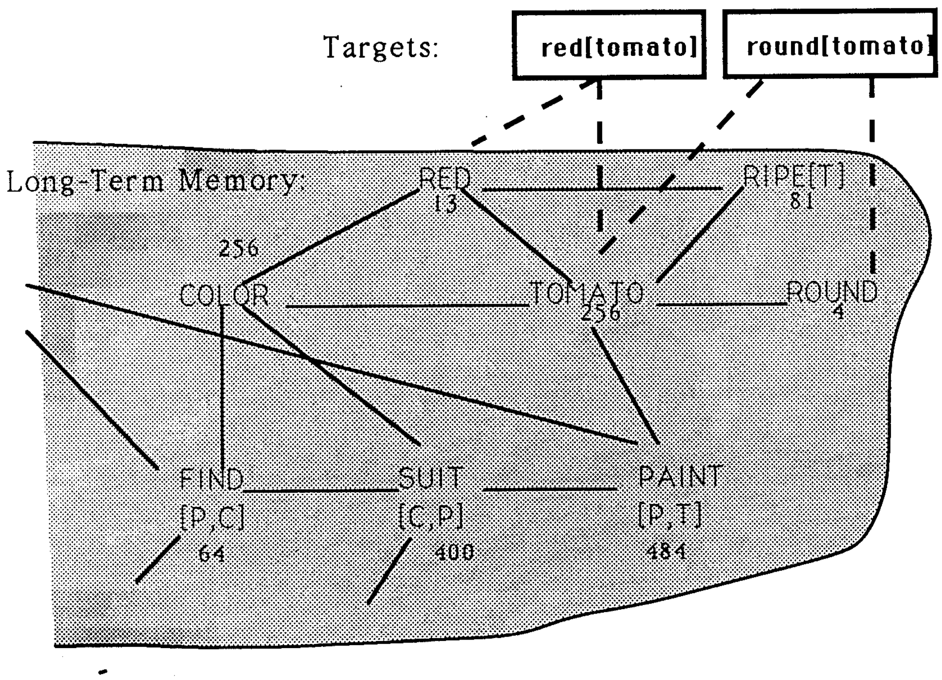

Targets:
red[tomato]
round[tomato]
Long-Term Memory:
RED
13
RIPE[T]
81
256
COLOR
TOMATO
256
ROUND
4
FIND
[P,C]
64
SUIT
[C,P]
400
PAINT
[P,T]
484

Fig 5

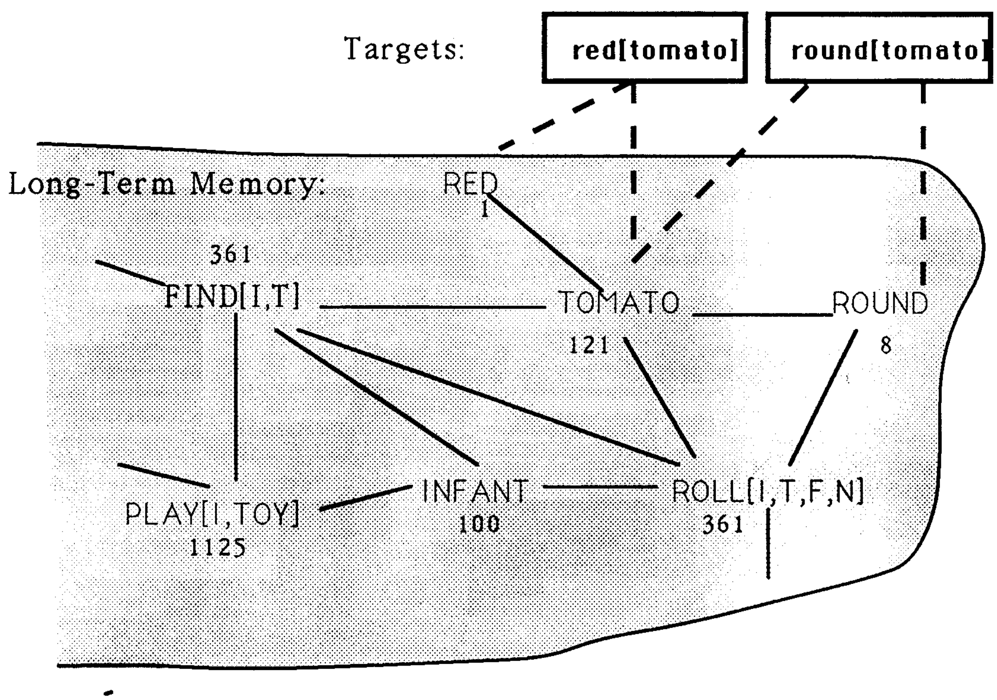
Targets:
red[tomato]
round[tomato]
Long-Term Memory:
RED
1
361
FIND[I,T]
TOMATO
121
ROUND
8
PLAY[I,TOY]
1125
INFANT
100
ROLL[I,T,F,N]
361

Fig 6