How Readers Construct Situation Models for Stories: The Role of Syntactic Cues and Causal Inferences

by

Walter Kintsch

University of Colorado
Campus Box 345
Boulder, CO 80309-0345

ICS Technical Report #90-9

To appear in A. F. Healy, S. Kosslyn, & R. M. Shiffrin (Eds.) (in press). *Essays in honor of William K. Estes* Hillsdale, NJ: Erlbaum.

To appear in the Festschrift for W. K. Estes,
Healy, Kosslyn, & Shiffrin, Eds.

How readers construct situation models for stories: The role of syntactic cues and causal inferences.

Walter Kintsch
University of Colorado

It may appear unusual to write a Festschrift article in which the person honored is not cited, but discourse comprehension is not one of Bill's many research interests. Nevertheless, I hope that his influence on the work reported here is still discernible to the attentive reader. There are two rather direct lines that can be traced from Bill's work to this research. First, the theory described below is a computational model of a psychological process. Such models are (one of) the contemporary successors to the mathematical psychology which Bill helped to found when I was a graduate student. Secondly, the emphasis on associative processes and structures which characterizes this work, obviously owes something to older traditions. Bill's work has had an important bridging function in this respect: it demonstrated the strengths of associationist theories, when precisely and explicitly formulated. Thus, I would claim Bill as a model for both the computational approach and the associationist orientation.

The present chapter is part of my attempt to develop a general theory of comprehension, focusing on story understanding and the role that syntax plays therein. Partly as a reaction against the almost exclusive concern of some linguists and psycholinguists with syntax, syntactic questions have played only a minor role in the current work on discourse processing, text comprehension, and conversation. Maybe this is too strong a statement as far as the field as a whole is concerned, but it certainly characterizes my own work in this area. Syntax plays a role in that it directs the reader how to parse a text into meaning units (propositions), but all the most interesting processes psychologically seem to occur at the level of meaning and language use.

Such an essential feature of language as grammar surely must be more important than we have given it credit for, and a recent paper by Givón (in press) in which he proposes that grammatical cues function as mental processing instructions, may give us some clues as to the manner in which syntactic cues are used in discourse processing. I shall show how Givón's ideas can be incorporated into a model of discourse comprehension, how such an extended model accounts for some salient empirical results that have been reported previously by others as well as myself, and I shall argue that by paying attention to the role of syntactic cues in processing, we can greatly expand the power of our model and clarify some currently controversial issues. Specifically, I shall show how only a single verbal signal can determine whether a text representation will be organized linearly or in terms of a topic-subtopic structure (Section 2), and how syntactic cues may indicate the presence of important causal relations in a text (Sections 3 and 4), thus providing an alternative way to construct situation models: either by means of powerful, knowledge-based, but domain-limited strategies, or by relying on weaker but more general syntactic cues.

1. A theoretical framework for discourse comprehension.

I must first give some background concerning the view of text comprehension that underlies this discussion ((Kintsch, & van Dijk, 1978; van Dijk, & Kintsch, 1983; Kintsch, 1988). At a very general level, the task of the comprehender can be considered as constructing a mental representation of the information provided by the text that is integrated with knowledge, beliefs, and goals. We take this representation to consist of concepts and propositions (complex terms which establish some sort of relation between concepts and/or other propositions) forming an interrelated network. What is related to what in this network depends in part on properties of the to-be-comprehended text, and in part on the semantic and associative relations between concepts and propositions in the comprehender's long-term memory (which includes general knowledge). In practice, we approximate these rich interrelationships by means of argument overlap: two propositions are related if they have a common argument (concept or proposition). The rationale for this approximation is that it is frequently the case that whenever two propositions are in fact related semantically, there exists a shared argument. Obviously, this is neither a sufficient nor necessary condition for semantic relatedness, but it works reasonably well in a statistical sense, and it is easily computed and does not require uncertain subjective judgments. Sometimes, however, argument overlap provides too crude an approximation, and we need a more powerful analysis (e.g., the role of causal connections in story understanding as discussed below).

Comprehension, then, consists of the construction of some sort of propositional network in our model. Below, we shall introduce a special set of construction rules that take advantage of syntactic cues signaling the importance of the various meaning elements that have been constructed.

There are, however, several further considerations which cannot be neglected. First, listening as well as reading[1] is sequential. The capacity limitations of working memory make it impossible to maintain in an active state all of the previous information in a discourse. Hence processing must be cyclical. Depending on the level of analysis desired, comprehension may be analyzed word by word

---

[1] From now on I shall speak about reading comprehension, because the experiments discussed below involve reading rather than listening.

(e.g., if we are interested in the speed-accuracy trade-off in sentence recognition, as in Kintsch, Welsch, Schmalhofer, & Zimny, 1990), or, more usually, sentence by sentence (as in Kintsch & van Dijk, 1978), or even paragraph by paragraph, if we are only concerned with the macrostructure of a long text (e.g. the campaign speeches analyzed in Kintsch & Vipond (1979). In any case we need a memory buffer to form a bridge between cycles. Therefore we assume that a few propositions from the current processing cycle, which are considered to be the most important ones, are retained in the buffer to form a connection between the previous text and the current input. If there is no connection between the contents of the buffer and the new input, extra processing is required from the reader to make the textbase cohere. This may either involve a search of previous text which is no longer actively available in working memory, resulting in the reinstatement of a connecting proposition, or it may require a knowledge based inference to form a connection when it is not explicitly expressed in the text (e.g. Kintsch & Vipond, 1979).

A second consideration concerns the way in which a text activates the reader's relevant knowledge, and how this knowledge becomes integrated with the current text. Kintsch (1988) has argued that the activation process is quite unselective and driven mostly by local associations, rather than controlled by some sort of schema that filters out anything but the contextually appropriate knowledge. In this way, a lot of irrelevant knowledge is activated, but a spreading activation process can then select those portions of the activated knowledge that actually fit into the given context. Thus, if activation is spread around the network consisting of the text propositions as well as the associatively activated knowledge elements, it will tend to collect in those portions of the network that are most strongly interconnected, while isolated elements or elements connected with inhibitory links will tend to become deactivated. Thus, the end effect is the same as in a schema theory: only contextually appropriate knowledge remains, but this effect is produced in a psychologically more plausible and computationally more flexible way.

Van Dijk & Kintsch (1983) have distinguished between the construction of a mental representation of the text itself (termed a textbase) and a representation of the situation described by the text and integrated into the reader's previous knowledge (which we have termed the situation model). The textbase includes a representation of both the microstructure as well as the macrostructure of the text, reflecting its local and global organization, respectively. The structure of the situation model is not necessarily the same as the macrostructure of the text. It is independent of the rhetorical organization of the text, and reflects primarily the structure of the situation described by the text. It is a domain structure rather than a text structure.

Knowledge about the world, about the situation in question, is therefore needed to form situation models. But situation models are not mere knowledge structures. They are the product of the combination of information provided by the text and already existing world knowledge. To "really understand" a text, say a story, a great deal of specific world knowledge is often required, as well as a great deal of analysis: exactly what leads to what, and why - inferences about goals, motivations, psychological states, causal relations and implications. On the other hand, we can understand, though maybe not "really understand" a story from a different culture where we are not sure about people's goals and reactions, or we can understand a story without thinking very deeply about it, without fully analyzing its implications. There are strong, powerful methods for building situation models which require a certain amount of effort and resource commitment, but there also seem to be weaker, more general but easier methods, which we can use when we lack precise knowledge or are unwilling to take the trouble. Syntax, together with some weak, general semantic principles, allows readers to generate situation models which, while not as accurate and complete as models based upon full knowledge and deep analyses, can be quite adequate for many purposes, or may even serve as stepping stones towards a more sophisticated model.

Syntactic cues in the text signal to the reader what is likely to be important for the construction of a situation model, and some rather general semantic rules allow the reader to put these elements together into a weak or sloppy situation model. Givón (in press) has shown how the grammar provides processing instructions for the formation of a referentially coherent mental representation of a text. The grammar tells the reader where to look for what in a text. I cannot summarize here Givón's important work, and shall use only two of his examples in the analyses to be reported below: cues for topicality, such as the indefinite "this", and the signalling of importance via subject-verb-object constructions. In the latter case the syntactic signals must be filtered semantically because it makes a difference for the construction of a situation model in story understanding whether the verb is an action verb, or a verb of saying or believing.

Picking only a few examples like this from the wealth of material described by Givón may seem arbitrary. But I am concerned here with a principle. If we can show that in the cases analyzed here, syntactic cues can indeed function as mental processing instructions, we would have established an important point, and a wider range of syntactic phenomena could be examined later. First, we shall examine how syntactic signalling via an indefinite "this" can be built into the construction-integration model.

2. Grammar as mental processing instructions: Givón (in press) and a fragment completion experiment by Gernsbacher & Shroyer (1989).

Gernsbacher & Shroyer (1989) performed an experiment in which subjects received discourse fragments with or without an indefinite "this". Subjects were asked to complete these fragments. Following Givón, Gernsbacher & Shroyer hypothesized that more completions would involve nouns emphasized by an indefinite "this" than nouns without this emphasis. This prediction was confirmed in their experiment. My goal here is to explore how these phenomena can be dealt with in the construction-integration model.

The following is a sample fragment from Gernsbacher & Shroyer:

(1) I went to the coast last weekend with Sally.
We had checked the tide schedules
and planned to arrive at low tide,
because I just love beachcombing.
Right off, I found three whole sand dollars.
Sally found <u>an/this</u> egg............

We shall assume that when subjects complete a discourse fragment, they tend to talk about topics that are highly activated in the memory representation they have created. Thus, the model needs to produce a noticeably higher activation for "egg" in the "find this egg" text than in the "find an egg" text.

Suppose the text (1) is processed in six cycles, one for each line in (1). For present purposes I shall completely neglect one major aspect of the model, namely knowledge activation and integration. In many situations that is a crucial aspect of comprehension, but here it would be merely a needless complication. Thus, we obtain for the first processing cycle (the first line of text in 1) the following network of concepts and propositions (columns are labelled by the first letter of each proposition):

(2)

	I	S	A	G	C	L	T	activation
I	1	0	1	0	0	0	0	.08
SALLY	0	2	1	0	0	0	0	.14
AND[I, SALLY]	1	1	1	1	0	0	0	.20
GO[AND[..], COAST]	0	0	1	2	1	0	1	.28
COAST	0	0	0	1	1	0	0	.11
LAST-WEEKEND	0	0	0	0	0	1	1	.06
TIME[GO[..],L-W]	0	0	0	1	0	1	1	.14

The default rule of argument repetition has been used in (2) to construct the propositional network: whenever two propositions share an argument they are connected by a 1, and 0 otherwise.[2] Each proposition is assumed to be connected to itself, hence the diagonals are 1. In two cases, however, we have used Givón's syntactic signals to emphasize certain propositions by increasing their self-connection strengths from a 1 to 2: the GO[..] proposition because it is derived from the syntactic core of the sentence (subject-verb-object), and SALLY because it is marked by end focus (the unmarked form of the sentence would be "Sally and I went..."; by transposing Sally to the end of the sentence, the writer has - weakly - emphasized Sally). What this means is that in the first sentence of (1) the writer has signalled, by purely syntactic means, that these two elements should be considered as potential discourse topics.

The seven propositions constructed from the first sentence are initially equally activated, but once we start spreading activation around the network defined by (2), activation will collect in the more central nodes in the network. The reader can check this quite easily by typing the matrix (2) into a spreadsheet and postmultiplying it by a column vector of seven 1/7's. The resulting column vector is normalized so that the total activation always sums to 1, and is then used again as a multiplier for (2), until a stable vector is obtained. The last column of (2) shows this asymptotic activation vector.

The model's memory for the sentence, at this point, has three components: a list of propositions, a connectivity matrix that connects these propositions, and their activation values. It is often convenient to combine the last two components in a way that parallels more closely connectionist practice, so that the process of encoding a text in memory modifies not merely the activation value of the nodes, but the actual interconnections among the nodes. It has been shown

---

[2]Since we are only interested in the qualitative patterns, there is no need to estimate precise parameter values. Simple illustrative numbers are used instead.

(Kintsch & Welsch, in press) that if c(i,j) is the original connection strength of nodes i and j, and a(i) is the activation value of node i and b(j) is the activation value of node j, a modified long-term memory network can be obtained with connection strength of element l(i,j) given by $l(i,j) = a(i)*b(j)*c(i,j)$. Figure 1 shows the network after the first processing cycle, and illustrates the computation of the connection strengths in this network for the connections between the GO[..] proposition and its neighbors. Strength values are shown only for the links emanating from GO: the first number is the original value in the coherence matrix; it is multiplied by the product of the activation values of the two nodes being linked. For example, GO and AND are linked by a 1 in (2), and their activation values are .28 and .20, respectively, yielding a memory strength of .056.

We can now continue processing the next three lines of (1) in the same way, except that we have to decide how many and which propositions from the first processing cycle to keep in the buffer in order to connect this cycle with what is yet to come. Buffer estimates in previous work have ranged from 1 - 4 propositions, averaging about 2, which I shall use for these calculations. This means that the two most highly activated propositions are carried over into the second processing cycle where they are re-processed together with the new text.

For Cycle 2, another coherence matrix is constructed with AND[..] and GO[..]. as buffer contents, and CHECK[AND[..],TIDESCHED] and TIDESCHEDULE as the new input. The CHECK[..] proposition is derived from the main verb of the sentence, hence it is is emphasized.

- Figure 1 -

In the third cycle the syntax suggests PLAN[..] as another candidate for the situation model, but it is rejected because main verbs with a sentence complement ("plan" has the complement "to arrive..") are not considered emphasized. Similarly, we do not emphasize verbs of saying, believing, or feeling - for convenience, one could label this class as "psychological" verbs - because the situation model of a story should be built around actions rather than psychological states. Thus, in the fourth cycle, "love" receives no emphasis in our analysis. In the fifth and sixth cycles we have syntactic emphases on FIND[I,SANDDOLLAR] and FIND[SALLY,EGG]. In addition, for the "this-egg" version of the text, EGG is emphasized. As a result, the activation values for this cycle differ for the focused and non-focused version of the text:

	Focused	Non-Focused
Buffer:		
AND[I,SALLY]	.24	.27
GO[AND[..],COAST]	.13	.15
Input:		
FIND[SALLY,EGG]	.37	.39
EGG	.26	.18

EGG is more highly activated when it is emphasized by an indefinite "this" than when it is not. At the same time, the other propositions in the sentence are somewhat less strongly activated in the focused condition than in the non-focused condition (a consequence of normalizing the activation vector). This observation agrees with a recent report by Gernsbacher that concepts marked by indefinite "this" and spoken stress *suppress* the activation of other (potentially competing) topics (Gernsbacher & Jescheniak, 1990).

Thus, we have achieved our first goal: we have shown how the construction-integration model can account for the findings of Gernsbacher & Shroyer. One of Gernsbacher & Shroyer's subjects continued the version with the focus on egg in the following way:

(3) It looked like it came from a lizard.
We couldn't tell whether it had hatched,
so we put it back where we found it in case it was still alive.

We have simulated the way the model would process the combined texts (1) and (3), both the version with "an egg" and "this egg". Note that the texts are exactly the same, the only difference is the little word "this" that replaces "an" in the non-focused version. Nevertheless, the resulting memory representation is quite different, especially for the portion of the network corresponding to (3), for which the correlation between the memory strengths calculated for the two versions is only r=.62.

However, the memory trace for this story consists not only of the strength values of its elements, but also of their pattern of interrelations. This is strikingly different for the focused and non-focused versions for the portion of the text corresponding to (3). In the focused version, EGG is the central node for the second half of the story; EGG clearly has subtopic status, in that everything else is related to it, which is not the case in the non-focused version of the text. The grammar - a single verbal signal in this case - has determined what sort of situation model was constructed - a linear sequence of Going-to-the-Coast episodes in one version, versus a topic/subtopic structure in the other.

Other syntactic means of indicating the discourse function of particular concepts or propositions (see Givón, in press) can be handled within the construction-integration theory in much the same way as the indefinite "this" in Gernsbacher & Shroyer (1989). For instance, particle movement provides another example of syntactic emphasis: In He picked up a friend, the friend is likely to be persist longer as a discourse object than in He picked a friend up (Chen, 1986), and is more strongly activated psychologically (Clifford, 1990) - findings which are readily incorporated within the framework proposed above.

3. Causal reasoning in story understanding: Trabasso & van den Broek (1985)

Conventional stories are about people, about some problem in which the hero gets involved, and how this problem gets resolved. The domain knowledge the reader needs to understand these stories is about people, their motivations, goals, social relations, and so on. This is the same knowledge we need to interact with others in our daily lives. Even very young children already know a lot about such matters (and learn more from the stories they are told). Much of this knowledge takes the form of causal relations: a goal causes an action, a mental state causes a goal, etc. (The term "causal" is used here in the broad sense of an enabling relation.) A situation model for a story, therefore, consists primarily of the specification of the causal relations among the various events and actions in the story (in addition there is information about the setting). Trabasso and his colleagues (e.g. Trabasso & van den Broek, 1985) have extensively explored the role of these causal relations in story understanding.

There are at least three different aspects of a story that could, in principle, contribute to the formation of a coherent memory representation: causal relations, as claimed by Trabasso & van den Broek, but also the pattern of argument overlap among the story propositions as well as the syntactic signals, as in the example analyzed above.[3]

Below, a story used by Trabasso & van den Broek (1985, also Omanson, 1982) is analyzed in the same way as the "Going-to-the-Coast" story above. How different will the resulting structure be from the causal structure of the story, as determined by Trabasso & van

---

[3] While Trabasso & van den Broek (1985) have shown the importance of causal inferences for story understanding, they did not actually explore whether argument repetition is also a factor: they merely demonstrated that a coherence measure based on the identity of words in the text is irrelevant.

den Broek? Furthermore, what will happen if, in addition to the default structure generated by argument repetition, we build into our network the causal relations identified by Trabasso & van den Broek?

The story is as follows:

(4) One day Mark and Sally were sailing their toy sailboat in the pond. / Suddenly, the sailboat began to sink. Mark was surprised. He lifted the boat up with a stick / and found a turtle on top of it. The turtle became frightened and tried to crawl off the boat. The turtle put Mark in a playful mood. / Mark thought the turtle was hurt. Mark had always wanted Sally to see a turtle, so he waded out to the turtle and brought it back to her. / Sally thought Mark was going to hurt the turtle. Sally felt sorry for Mark. Sally tried to touch the turtle, but the turtle bit her. / Sally didn't like this and threw the turtle into the pond. The turtle crashed into the sailboat. Sally knew she had made a mistake.

Comprehension of this story has been simulated by means of the construction-integration model with the following parameters:

<u>Input cycles</u>: sentence or phrase boundaries, so that cycles consist of at least 7 propositions; this scheme results in 6 processing cycles, which are indicated by slashes in (4).

<u>Buffer</u>: the two most strongly activated propositions from the previous cycle.

<u>Coherence matrix</u>: a 1 if two propositions are related by argument overlap, and 0 otherwise; self-connections are 1, except for propositions that are emphasized syntactically, which are assigned a value of 2.

<u>Syntactic emphasis</u>: derives from the main verb and subject in a sentence, except for main verbs with a sentence complement, and verbs of saying, believing, and feeling.

Note that these parameters are in no way "best" estimates, but rather a-priori guesses, based on previous experience with the model and/or general expectations of what would be reasonable values, as in the case of buffer size and input size, or just simplicity, as in choosing values of 0,1, and 2 for the coherence matrix. I have chosen this informal approach to parameter estimation because I am not primarily interested in fitting data quantitatively, but in demonstrating some robust principles about comprehension.

- Figure 2 -

Figure 2 illustrates the construction of the coherence matrix for the second processing cycle. The propositions held over in the buffer from the first processing cycle, as well as the input propositions of the current cycle are shown. The buffer consists of AND[...] and SAIL[...]. Propositions that are syntactically emphasized are marked with stars. Links among propositions indicate argument overlap. (The bold arrows are causal links and will be discussed later). Among the current input propositions, SINK[..] and LIFT[..] are emphasized since they are derived from the main verbs of sentences, but SURPRISE[..] is not, because it is a "psychological" verb. Figure 3 shows the pattern of activation after the integration is completed.

- Figure 3 -

After completing the simulation through all six cycles, a long-term memory matrix can be calculated as in Figure 1. This long-term memory matrix can be used to predict recall data by assuming that the likelihood of recall of an item is monotonically related to its memory strength as calculated. Indeed, recall data collected by (Omanson, 1982) on 20 subjects correlate reasonably well with these predictions, Spearman rho = .76. A non-parametric measure of correlation had to be used, because the relationship between recall likelihood and predicted memory strength is highly non-linear (regression on a 4th order polynomial gives equally high correlations). Since Omanson's data were not scored in terms of propositional units but in terms of slightly more global idea units, appropriate propositions were combined to match his analysis.

Omanson (1982) was interested in a story grammar analysis, for which purpose he categorized all statements in (4) as belonging to one of the categories Setting, Initiating Event, Internal Response, Attempt, Consequence, and Reaction (after Stein & Glenn, 1979). His main point was that statements belonging to a certain story grammar category are recalled/judged differently than statements in some other category. This was taken as evidence for the psychological validity of the story grammar analysis. Since we have predicted memory strength values for all statements, the average memory strength of statements in each story grammar category can be calculated. These memory strength values derived from the construction-integration model correlate (Spearman's rho) .80 with the averages of the measures Omanson had collected (importance judgments, summaries, immediate recall, delayed recall). It is interesting that we can thus account for Omanson's finding that story grammar category affects the way statements in a story are processed, although our analysis makes no use whatever of the concepts of story grammar. The story grammar results simply fall out as consequences of normal discourse processing; story grammar appears to be an analytic scheme imposed by the theorist, not something that actually guides the reader's processing of a story. This question deserves further scrutiny, but, if the results reported above are confirmed, we may conclude that general principles of discourse comprehension may be sufficient to account for the phenomena described by story grammars.

But let us return to the role of causal relations in story understanding. So far, our network does not take into account causal relations; it is based entirely on the default-semantics of argument overlap, together with the emphasis on situation-model relevant propositions cued by the syntax. So far, the situation model used has been a very weak one: the syntax signals mostly topics and subtopics as well as the main actors and action components. Now suppose we assume that subjects develop a "strong" situation model by figuring out all the causal networks determined by Trabasso and van den Broek (1985; Figure 1). It is a simple matter to add these causal relations to our coherence matrices. What is involved is shown in Figure 2 by the bold arrows: there is a causal path from SAIL to SINK, and from SINK to LIFT. Note that these relations are asymmetric. We can therefore obtain a new coherence matrix which is identical with the previous one, except that for each causal relation among propositions a value of 1 is added in the appropriate cell of the matrix. Recalculating activation values based on the new coherence matrices containing causal links yields a new pattern of memory strengths. Surprisingly, this new pattern correlates with the one obtained without the causal links r= .98. Whether or not we include the causal links makes almost no difference at all!

What is going on can be glimpsed from Figure 2: the syntactic cues we have used overlap to such an extent with the causal links that they make them redundant. In other words, a "weak" situation model, as described above, will do just as well for present purposes as a "strong," causal one. What does this mean psychologically? It means that whether or not subjects actually infer causal relations during reading, more or less the same recall pattern is predicted. At least for recall, readers do not *have to* infer causal links in a story, they can do just as well with weak, general strategies, based on syntactic and semantic relations as reflected by argument overlap. We shall show in the next section, however, that this is not so when it comes to forming higher-order macrostructures: the model must have access to causal information in this case.

We now have two models, a complete model that defines coherence in terms of causal links (C) plus syntactic emphasis (S) plus argument overlap (A), and a weaker model that defines coherence in terms of the latter two relations only. Consider two more models, one that is obtained by defining coherence in terms of argument overlap only (this is the model of (Kintsch & van Dijk, 1978)), and one that relies only on the causal relations of Trabasso & van den Broek (Trabasso &

van den Broek, 1985, Figure 1). For each of these models a memory strength matrix was calculated as described above. The interrelations among these models, as well as the correlations between the various models and the Omanson data are as follows:

	A+S+C	A+S[1]	A[1]	C[1]	DATA[2]
A+S+C	-	.98	.90	.40	.79
A+S		-	.84	.49	.76
A			-	.16	.55
C				-	.61

1 Pearson r, n=36

2 Spearman rho, n=20

Note that the correlations between the various versions of the model and the data merely present a lower bound for the theory. These correlations were obtained not by adjusting the parameters of a model to maximize goodness of fit, but by simply postulating a certain set of (plausible) parameter values, i.e. by giving equal weight to argument overlap, causal relations, and syntactic cues. (The other parameters of the model, input and buffer size, are at least historically motivated through estimates obtained in other studies). Thus, while the model has a potentially large number of parameters, it does appear to be reasonably robust.

Obviously, very similar predictions are obtained whether or not causal relations are included in the coherence matrix, and these predictions fit the data about equally well. The argument-overlap-only model correlates fairly well with the complete model, but does not predict the data as well. The cause-only model is more different from the complete model, and has a very low correlation with the argument-overlap-only model. Argument overlap and causal analysis make quite different contributions to the coherence matrix, while syntactic emphasis and causal analysis appear to be redundant.

4. Macrostructures and summaries: Kintsch (1976)

In general, the macrostructure of a text cannot be obtained merely by deleting all the unimportant propositions of the microstructure. Constructive processes involving generalization and construction rules are usually involved (van Dijk, 1980). However, there are special cases where the deletion of macro-irrelevant proposition suffices to obtain a good macrostructure. One such case that has been reported before is the (first part) of the Boccaccio story analyzed in Kintsch (1976). Below, I shall show how the construction-integration model predicts summarization data for this text. To create a macrostructure, some *ad hoc* rules had to be used in Kintsch (1976). Will the extension of the construction-integration model introduced above generate the same macrostructure as had been obtained earlier, without the ad hoc arguments?

I am going to analyze the Boccaccio story discussed in Kintsch (1976) in the same way as the Going-to-the-Coast and the Mark-Sally-&-Turtle stories above, indeed with the same parameters. The Boccaccio story is about a merchant named Landolfo, who is involved in a risky venture, goes bankrupt, becomes a successful pirate, is shipwrecked and loses it all again, is saved, finds a chest of jewels, and finally returns home happily. The first part - the one analyzed here - is analyzed in detail in Table 5.4 in the original paper. We are concerned here with the following text fragment:

(5) This Landolfo, then, having made the sort of preliminary calculations merchants normally make, purchased a very large ship, loaded it with a mixed cargo of goods paid for out of his own pocket, and sailed with them to Cyprus. But on his arrival, he discovered that several other ships had docked there, carrying the same kind of goods as those he had brought over himself. And for this reason, not only did he have to sell his goods at bargain prices, but he was forced to give away his goods, thus being brought to the verge of ruin.

This text breaks down into four processing cycles. The coherence matrix for each cycle is obtained by adding three separate components: one based on argument overlap, one based on syntactic cues, and one based on the causal relations in the story. A long-term memory matrix is obtained for this story in exactly the same way as has been described above.

Given the long-term memory strengths predicted by the model, what sort of retrieval processes should we assume to generate a macrostructure from this memory representation? Consider how we could obtain a modal summary. Suppose the subject picks the most highly activated proposition from the first processing cycle. From that proposition, the strongest link will be followed to another proposition, which will be included in the summary if its strength is above some threshold value. This link-following-plus-editing process is continued until the end of the story is reached. The propositions traversed in this process will be the modal summary. If at any point the strongest link leads to a below-threshold proposition, the next strongest link is tried. If no link can be found to an above-threshold proposition, the process backtracks and tries to find another path from the previously selected proposition. There are two comments to be made about this retrieval model. First, the restriction to following the strongest link is merely a simplification. In general, choices could be probabilistic, so that the model would generate a whole family of summaries, differing in their likelihood, rather than just a single modal summary. Second, the same process, minus the response threshold, could be used to derive recall predictions from the model, instead of the simplified assumptions used above. Thus, each prediction would involve a whole recall path. Averaging many such recall paths, we would probably generate predictions close to those obtained from the mere response strengths of each separate item, as was done above for the Sally-Mark-&-Turtle story.

Figure 4 shows what happens for the Landolfo story: We start out with Landolfo is a merchant, which leads to He purchases a ship, which leads to He loads the ship with cargo, which leads to He sails to Cyprus, which leads to He discovers other ships, which leads to Which are docked in Cyprus, which does not link to a strong enough proposition; backing up one step, we find another link to He sells his goods at bargain prices, which leads to Landolfo is ruined.

- Figure 4 -

The Kintsch (1976) summary includes all the statements listed above, plus three other, redundant ones: He pays for the goods he bought, He arrives in Cyprus, and He gives away his goods. Clearly, we have produced, if anything, a better summary - and without ad hoc assumptions. If we model the construction of a situation model, based upon general syntactic cues and some general semantic constraints, the macrostructure will be implied (excepting construction and generalization processes).

As before, the memory strength values for the propositions of the Landolfo story are almost the same whether causal links are or are not included in the coherence matrices. Specifically, exactly the same summary results as shown in Figure 4 if causal links are deleted. Hence the conclusion reached above is supported that if we take into account syntactic signals, precise causal inferences are not necessary to predict either recall or summarization.

However, Kintsch (1976) reports not only a first-order summary, as described above, but also a second-order summary, where the whole first part of the Landolfo story gets summarized as The merchant Landolfo was ruined. This second-order summary is well motivated, both empirically and theoretically (macrostructures have always been thought of as hierarchical). Can the construction-integration model produce this second-order summary, too? It can, but only if it is given the causal link information. Figure 5 shows the pattern of activation that is obtained when the eight summary propositions of Figure 4 are integrated. The same values are entered into the coherence matrix for the summary as were used for the microstructure analysis. The strongest proposition turns out to be Landolfo is a merchant, and it is linked most strongly to Landolfo is ruined - exactly what is needed for the second-order summary. Landolfo is ruined wins out over its competitors, because it is at the end of a long causal chain: one event enables the next, and ruin is the end result. Activation flows from cause to effect, and hence favors the end result. Without the causal links this does not happen, and for the given summary matrix, all events would be equally activated.

Figure 5

This is an interesting result because it indicates that causal inferences are necessary for story understanding. They are not really needed for comprehension, recall, or the generation of a first-order summary, in the sense that we get essentially the same memory trace with or without causal inferences, because the syntax emphasizes the same propositions that play a role in the causal network. Thus, the causal inference does not add much to other information already gleaned from the text. Generating a second-order summary, however, is another matter: all the model has to work with now is a list of propositions, the syntactic information from the text has already been used up in selecting this particular list from the original input. Now the only way to structure the macrostructure further hierarchically is to rely on the causal network.

5. Conclusions

It has often been said that "comprehension is problem solving". The construction-integration model sets the accent differently. It emphasizes the bottom-up, perception-like aspects of comprehension, as distinct from the controlled, conscious problem solving processes. Our model describes normal comprehension as a highly data-driven process based upon constraint satisfaction, more like perception, but it does not deny that as comprehension difficulties arise, more controlled, problem-solving type behavior plays a role. Indeed, we have just seen an example of how comprehension shades into higher-level processes. Just to read and retell a story, we may or may not infer causal links; weaker, more general operations will suffice. But to reach a higher understanding of a story, we can't do without a causal analysis.

Indeed, syntax guided processing may play a significant role in the actual process of generating causal inferences by selecting from the multitude of propositions in a text those propositions that are most important and that, in a story, are most likely to be relevant causally. Thus, it is not necessary to test for the presence or absence of causal relations among all the numerous pairs of propositions in a story, but only for those propositions that, on other grounds, have already been determined to be important.

There are, therefore, more than one ways to obtain a satisfactory situation model. If the reader has the requisite knowledge, a knowledge-based situation model can be constructed. But, if a text is well written, syntactic cues can tell the reader what is important in the text, even in the absence of specific domain knowledge, allowing the reader to construct an adequate situation model. Perfetti (1989) has made a similar point: without denying the importance of knowledge based strategies in comprehension, he finds compelling evidence that general reading strategies and abilities also play a significant role in comprehension.

This essay is an exploration of the linguistic richness of texts that so far has not been considered sufficiently in psychological processing models of text comprehension. Of course we must simplify, even radically, if we want to understand language processes. As a first approximation, it makes sense to study comprehension as the integration of a network of propositions, but we also need to look more closely at the language itself to refine our model of language comprehension. Givón (in press) has suggested that syntax be viewed as a set of processing instructions. Realizing this suggestion within the construction-integration model of text comprehension has yielded a rich pay-off: if we attend to the processing instructions the syntax gives us, our model of language comprehension will be much more powerful and capable of dealing with a range of phenomena that were previously out of its reach.

We have finally found a use for syntax in a psychological processing model. It provides the comprehender with "weak" but general methods for comprehension, to be complemented by "strong" knowledge-based and domain-specific methods. As in problem solving, weak and strong methods have their respective advantages and uses, and the complete comprehender would not forego either.

Footnote

This research was supported by Grant MH 15872 from the National Institute of Mental Health. I thank Paul van den Broek for making his data available to me, and Paul van den Broek, Stefanie Doane, Evelyn Ferstl, Alice Healy, Eileen Kintsch, and Valerie Reyna for their comments on this work.

References

Chen, P. (1986) Discourse and particle movement in English. Studies in Language, 10, 79-96.

Clifford, J. (1990) Grammar and the activation of referents in working memory: The case of particle movement. Unpublished manuscript, University of Colorado.

Gernsbacher, M. A., & Jescheniak, J. (1990). Cataphoric devices in spoken discourse. Unpublished manuscript, University of Oregon.

Gernsbacher, M. A., & Shroyer, S. (1989). The Cacioppo use of the indefinite this in spoken narratives. Memory & Cognition, 17, 536-540.

Givón, T. (in press). The grammar of referential coherence as mental processing instructions. Cognitive Science, ,

Kintsch, W. (1976). Memory for prose. In C. N. Cofer (Ed.), The Structure of Human Memory (pp. 90-113). San Francisco: Freeman.

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

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

Kintsch, W., & Vipond, D. (1979). Reading comprehension and readability in educational practice and psychological theory. In L. G. Nilsson (Ed.), Perspectives of memory research (pp. 325-366). Hillsdale, N. J.: Erlbaum.

Kintsch, W., & Welsch, D. (in press). The construction-integration model: A framework for studying memory for text. In W. E. Hockley,

& S. Lewandowsky (Ed.), Relating theory and data: Essays on human memory. Hillsdalle, N. J.: Erlbaum.

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

Omanson, R. C. (1982). The relation between centrality and story category variation. Journal of Verbal Learning and Verbal Behavior, 21, 326-337.

Perfetti, C. A. (1989). There are generalized abilities and one of them is reading . In L. B. Resnick (Ed.), Knowing, learning, and instruction (pp. 307-336.). Hillsdale, N. J.: Erlbaum.

Stein, N. L., & Glenn, C. G. (1979). An analysis of story comprehension in elementary school children. In R. O. Freedle (Ed.), New directions in discourse processing (pp. 53-120). Hillsdale, N. J.: Erlbaum.

Trabasso, T., & van den Broek, P. (1985). Causal thinking and the representation of narrative events. Journal of Memory and Language, 24, 612-630.

van Dijk, T. A. (1980). Macrostructures . The Hague: Mouton.

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

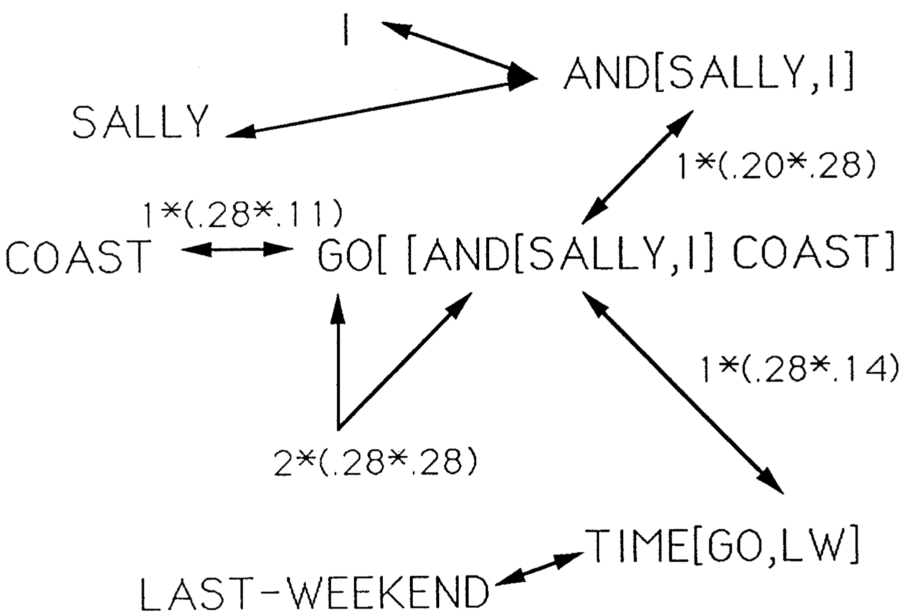

Figure 1. Long-term memory strength values for GO and its neighbors.

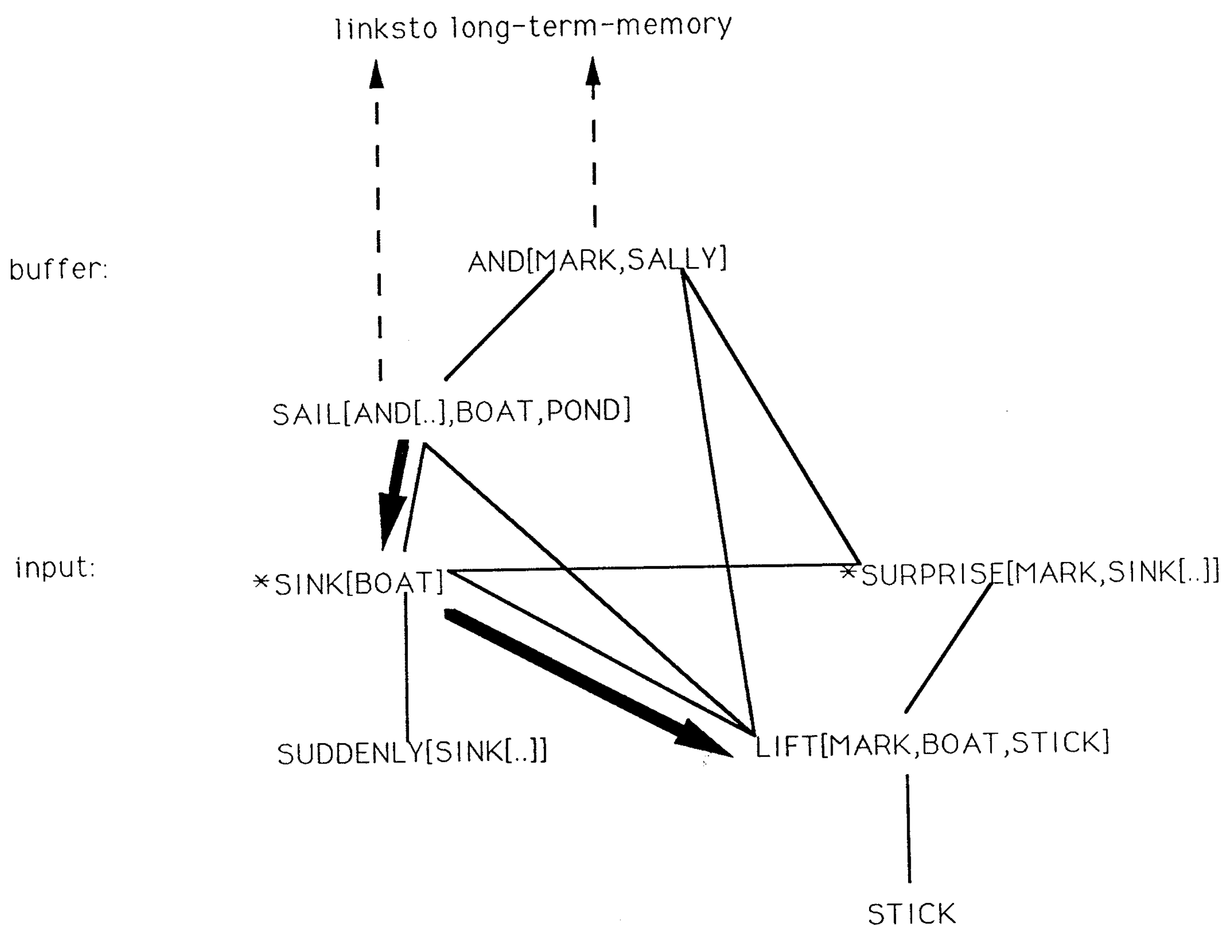

Figure 2: Buffer and input propositions for the second processing cycle. Lines denote argument overlap, bold arrows causal relations. Stars mark syntactically emphasized propositions.

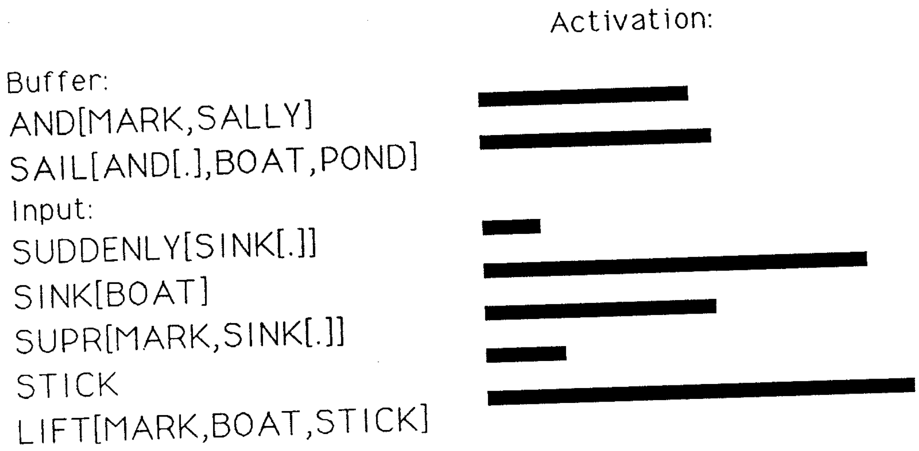

Figure 3: Activation after the second processing cycle.

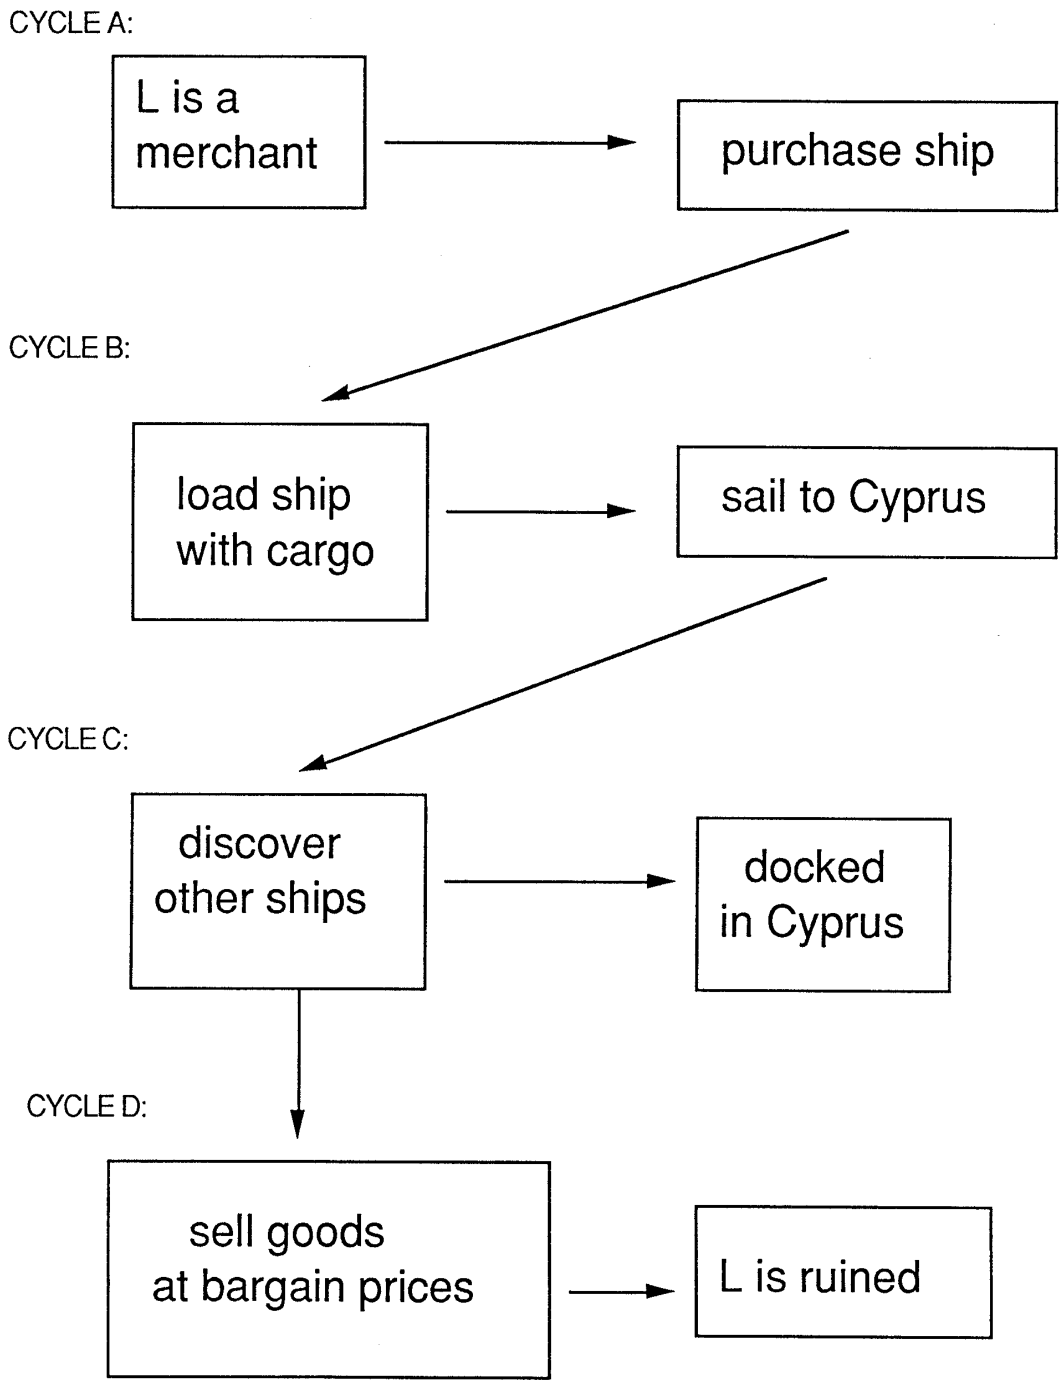

Figure 4. The construction of a summary for the Landolfo story.

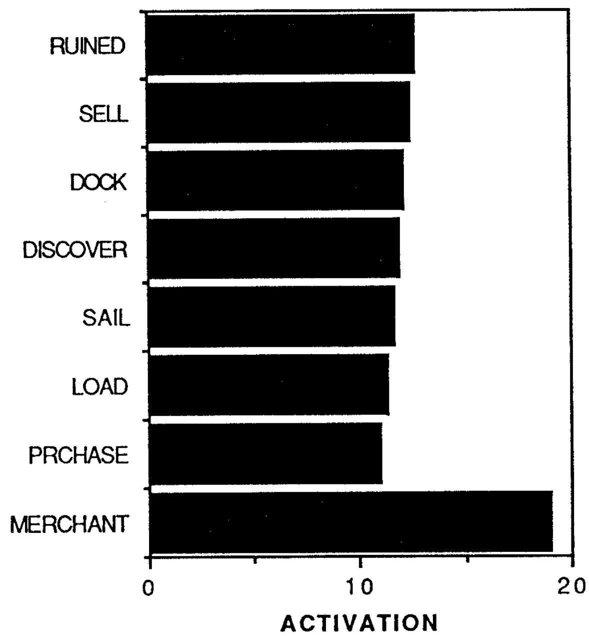

Figure 5. Activation values of the macropropositions in the Landolfo story.

University of Colorado at Boulder

Institute of Cognitive Science

Muenzinger Psychology Building
Campus Box 344
Boulder, Colorado 80309-0344
(303) 492-5063
FAX: (303) 492-7177
email: cogsci@psych.colorado.edu
URL: http://psychwww.colorado.edu/ics/