UC Merced
Proceedings of the Annual Meeting of the Cognitive Science Society

Title
A Time-Dependent Distributed Processing Model of Stragtegy-Driven Inference Behavior

Permalink
https://escholarship.org/uc/item/1t93p7h1

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 9(0)

Authors
Eiselt, Kurt P.
Granger, Richard H.

Publication Date
1987

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

A Time-Dependent Distributed Processing Model of Strategy-Driven Inference Behavior

Kurt P. Eiselt
Richard H. Granger, Jr.
Irvine Computational Intelligence Project
Department of Information and Computer Science
University of California
Irvine, California 92717

Abstract

Experimental evidence suggests that some readers make inference decisions early on in text understanding and mold the inferences from later text to fit with the earlier inferences, while other readers postpone inference decisions until later in the text and then base their final interpretation of the text on those postponed inferences. This behavior has been called *strategy-driven inference behavior* because it was originally ascribed to different strategies used by readers to guide the course of their inference decisions. This paper presents a new theory of how this behavior comes about, attributing the observed differences in behavior not to different strategies but to very small differences in the underlying cognitive architecture. This theory is illustrated by a simple model of inference processing during text understanding. The inference processing model employs a hybrid connectionist network whose behavior is extremely sensitive to the order of activation of nodes in the network, which in turn corresponds to the order of presentation of events in the story.

1 Introduction

One of the more intriguing mysteries of natural language understanding is that of *strategy-driven inference behavior*. According to the theory of strategy-driven inference behavior, different readers consistently employ different strategies to guide their choice of an interpretation for a text. These strategies are time-dependent; that is, they are sensitive to the order of presentation of events in the text. Experimental evidence suggests that some readers make inference decisions early on in text understanding and mold the inferences from later text to fit with the earlier inferences, while other readers postpone inference decisions until later in the text and then base their final interpretation of the text on those postponed inferences. These differences in inference behavior can be elicited with the use of specially constructed texts that have two equally likely interpretations. Upon reading these texts, some readers will arrive at one interpretation while the other readers will find the alternative interpretation, and their choices appear to be entirely dependent upon the sequence of events in the story. If the order of presentation is reversed, the different sets of readers will reverse their interpretations.

Previously, we have built two computational models of text understanding in attempts to shed light on the processes underlying strategy-driven inference behavior. Though the explanations provided by these models were satisfying at the time the models were constructed, we have since found weaknesses in these explanations. This paper describes our new model of text understanding which explains observed differences in inference behavior as the result of a connectionist network in which the interpretation settled upon is determined by the order of activation of the nodes in the network.

2 Old Problems, Old Solutions

Granger and Holbrook (1983) reported the results of a psychological experiment that investigated the processes people use in making pragmatic inferences while reading text. These results provided support for Granger and Holbrook's theory that different readers employed different strategies for selecting from alternative interpretations of a single text. In this experiment, subjects read a number of short texts that had two equally plausible interpretations, such as the following text:

Text 1: Wilma began to cry.
Fred had just asked her to marry him.

Interpreting this text requires that a causal relationship between Fred's proposal and Wilma's tears be inferred. The experimental results showed that many subjects inferred that Wilma was happy about Fred's proposal and was crying "tears of joy," while approximately the same number of subjects inferred that Wilma was crying because she was saddened or upset by the proposal. However, the subjects' interpretations were not based on their predispositions toward a particular interpretation; for example, tests run on another set of subjects drawn from the same subject pool confirmed that the subjects almost unanimously associated crying with sadness and marriage with happiness. Instead, the results indicated that the subjects' interpretations were based on the *order of presentation* of the events in the story. In other words, when some subjects read Text 1, they determined that Wilma was sad based on inferences generated from the first story event, the fact that Wilma was crying, while other readers determined that she was happy based on the second story event, Fred's marriage proposal.

Granger and Holbrook theorized that some readers consistently make inference decisions as early as possible in reading and try to make inferences from later text that agree with the earlier inferences; these readers were called *perseverers*. Other readers, called *recencies*, postpone making inference decisions. When recencies eventually do make decisions, they are based on the most recently read text. A computational model of the processes suggested by this theory was developed soon thereafter (Granger, Eiselt, & Holbrook, 1983). This model, called STRATEGIST, arrives at either of two interpretations of an input text using the same component inference processes but different rules for deciding when the processes are invoked, resulting in different interpretations of the same text.

STRATEGIST was later subsumed by another computational model of inference processing. This model, ATLAST, attempts to unify lexical and pragmatic inference decision processes and offers an explanation of how readers are able to correct erroneous inference decisions (Eiselt, 1985; Granger, Eiselt, & Holbrook, 1986; Eiselt, 1987). In addition, ATLAST also provides a framework in which to further study strategy-based inference behavior.

ATLAST uses marker-passing to search a relational network for paths which connect meanings of open-class words from the input text. A single path is a chain of nodes, representing objects or events, connected by links that correspond to relationships between the nodes. Any nodes in a path that are not explicitly mentioned in the text are events or objects that are inferred; therefore, these paths are called inference paths. A set of inference paths which joins all of the words in the text into a connected graph represents one possible interpretation of the text. In this respect ATLAST resembles a number of other models of text understanding that utilize marker-passing or spreading activation (e.g., Charniak, 1983; Cottrell & Small, 1983; Hirst, 1984; Quillian, 1969; Riesbeck & Martin, 1986; Waltz & Pollack, 1985).

For any given text, however, there may be a great number of possible interpretations, many of which are nonsensical. The problem then is determining which of the possible interpretations provides the best explanation of the text. ATLAST deals with this problem by applying inference evaluation metrics. These metrics are used to compare two competing inference paths and select the more appropriate one. Two inference paths compete when they connect the same two nodes in the relational network via different combinations of links and nodes. The path that fits better with the existing interpretation is then added to the interpretation. The choice of one inference path over another is made as soon as ATLAST discovers that the two paths compete; unlike STRATEGIST, ATLAST does not postpone inference decisions. As the marker-passing search mechanism finds more paths, ATLAST constructs an interpretation consisting of those paths that survive the evaluation process. When the marker-passing and evaluation processes end, the surviving inference paths make up the final interpretation of the text.

Most of the evaluation metrics attempt to make either a quantitative or qualitative judgment of the relative merits of competing inference paths in order to provide the most parsimonious interpretation: one metric favors the shorter of two paths, another favors the path that shares more nodes with the current interpretation, and still another metric favors the more specific path as determined by the relationships represented by the links in the path. If these metrics fail to yield a decision the last remaining metric is invoked; this metric alone determines the difference between perseverer and recency behavior. When the programmer wants ATLAST to model perseverer behavior, a rule that chooses the inference path found earlier is used; recency behavior is obtained by using a rule that selects the inference path found later. Thus the theory of strategy-driven inference behavior embodied in ATLAST differs significantly from that of its predecessor, STRATEGIST. In STRATEGIST, the different inference behaviors were caused by different orderings of invocation of the same inference decision processes. In ATLAST, the inference decision processes (i.e., the evaluation metrics) are invoked in the same order for both types of readers, and the difference in inference behavior is attributed to the use of a different "tie-breaker" metric that is applied only when the other metrics are unable to choose one inference path over another.

There are at least two problems with this explanation of strategy-based inference behavior. The first problem is with the nature of the evaluation metrics. In ATLAST, the perseverer/recency metric has the same status as the other metrics in that they are all rules to be applied after competing inference paths have been discovered. However, human understanders exhibit substantial differences in the behavior that would be determined by perseverer or recency metrics, but they do not appear to exhibit significant differences in the behavior that would be determined by the other metrics. In other words, while we see equally large proportions of perseverers and recencies in the laboratory, there is nothing that suggests that readers who favor parsimonious interpretations and those who do not exist in these same proportions. Yet, if the difference between perseverer and recency behavior is in fact rule-based, and similar rules are used to determine other inference decision behavior as well, we would expect to see differences in behavior along those lines. We simply do not see these other differences; in fact it is difficult to imagine how anyone could function in this world while always pursuing the least plausible interpretation of everything read or heard.[1] Thus, a model that could explain why readers exhibit differences in some types of inference behavior but not in others without resorting to arbitrary differences in rules would be superior to ATLAST.

The second problem is that perseverer/recency behavior in ATLAST depends upon a serial ordering of the competing inference paths as determined by the relative times at which they were discovered. Although ATLAST is intended to be a parallel model, the parallelism is only simulated and the inference paths are discovered and evaluated serially; ATLAST takes advantage of this latter fact. In a truly parallel system, two competing inference paths may be discovered at exactly the same time. In this case, ATLAST's perseverer and recency metrics would not be able to make a decision. A more desirable explanation of the difference between perseverers and recencies would be one that could function in a true parallel processing environment.

3 The New Solution

The new model, called CATLAST (for Connectionist ATLAST), is directly inspired by the work of Cottrell and Small (1983) and Waltz and Pollack (1985). These systems do not use the marker-passing style of spreading activation that ATLAST employs. Instead, these systems spread continuously variable numeric quantities representing activation energies through a network of nodes connected by excitation and inhibition links. An interpretation of input text is chosen not by application of rules, but by the iterative adjustment of activation energies at the nodes. While this method of making inference decisions is clearly an example of distributed processing, the representation scheme is just as clearly not distributed: these networks all adhere to the "one node equals one concept" principle (cf. Collins & Loftus, 1975; Quillian, 1968). Thus, these models share features of rule-based symbol-manipulation systems and parallel distributed processing systems (Rumelhart & McClelland, 1986).

[1]There may be rare exceptions to this rule, but we are unaware of them. Though we have not run experiments to test for this sort of anomalous behavior in human subjects, we have run ATLAST with the metrics modified to favor the least parsimonious interpretation of a story. The resulting interpretations bore no similarity to any of the interpretations offered by human readers.

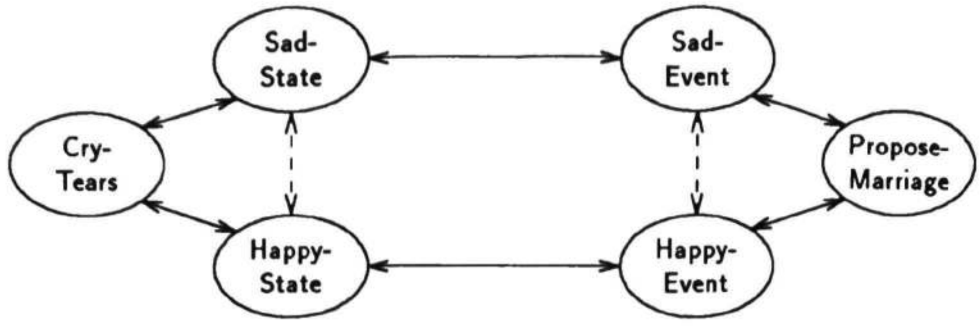

Figure 1: CATLAST's memory network.

The excitation/inhibition network used in CATLAST is simply a modified version of the relational network used in ATLAST (see Figure 1). Two uni-directional inhibitory links, represented by single dashed lines in Figure 1, have been added between pairs of incompatible nodes. The inhibition weights can take on values from 0 to −1. A solid line represents two uni-directional excitatory links. The values of excitation weights can range from 0 to 1. Some of the excitatory links are *preferred* links in that their excitation weights are higher than the norm. For the examples discussed in this paper, the preferred links are from Cry-Tears to Sad-State (but not from Sad-State to Cry-Tears) and from Propose-Marriage to Happy-Event (but not from Happy-Event to Propose-Marriage).

The activation function used to compute activation levels at the individual nodes or computing units is loosely based on the interactive activation function described by McClelland and Rumelhart (1981):

$$a_j(t+1) = min(1, max(0, C_1 a_j(t) + C_2(\sum w_{ij} a_i(t))/n_j))$$

where $a_i(t)$ is the activation energy at node i at time t, w_{ij} is the excitation or inhibition weight on the link from node i to node j, n_j is the number of inputs to node j, and C_1 and C_2 are damping factors that represent degrees of confidence in their respective terms.

With CATLAST we are trying only to find a better explanation for strategy-driven inference behavior; we are not, at this time, attempting to build a robust language understanding system. Therefore, we have taken the liberty of simplifying CATLAST's input texts down to the events explicitly stated in those texts. When CATLAST processes a text, the first event is read, the corresponding node is raised to its maximum activation level (i.e., 1), and the network is allowed to settle through iterative application of the function given above. Then the next event is read, its corresponding node is activated, and the network settles again. CATLAST's interpretation of the text is determined by comparing the final activation levels at the nodes, with high levels favored over low levels.

By adjusting the parameters, CATLAST can be tuned so that it either always selects the "sad" interpretation or it always selects the "happy" interpretation, regardless of which event is presented first. Much more interesting is the fact that CATLAST can also be tuned so that it selects the interpretation preferred by the first event presented. That is, it exhibits perseverer behavior: if Cry-Tears is activated first CATLAST chooses the "sad" interpretation, but it chooses the "happy" interpretation if Propose-Marriage is activated first.

However, equipped with only the features described so far, CATLAST cannot be tuned so that it consistently exhibits recency behavior. In order to correct this deficiency, CATLAST must be endowed with the ability to postpone making a decision until it reads the second event. This is accomplished by the addition of an inhibition threshold for the inhibitory links and an excitation threshold for the nodes.

The inhibition threshold prevents the inhibitory link from feeding inhibition energy into a destination node until the activation energy at the source node exceeds a prescribed value. Increasing this threshold has the effect of delaying the onset of inhibition, thus helping to postpone the inference decision. The excitation threshold holds the activation of a node at the minimum value (i.e., 0) until the total activation energy at that node exceeds a prescribed value. Increasing this threshold effectively slows the spread of activation energy through the network. Neither of these features alone can postpone the decision long enough to cause CATLAST to exhibit recency behavior, as far as we have been able to determine. On the other hand, the two thresholds combined do provide the necessary delay when the appropriate values have been assigned.

4 The New Solution in Action

Figure 2 shows the time course of the activation of the individual nodes while CATLAST exhibits perseverer behavior on an abbreviated version of Text 1. In this example, the first event is read and Cry-Tears is raised to its maximum activation level.[2] As activation energy feeds into Sad-State and Happy-State, the activation level at Cry-Tears dips briefly and then recovers. Sad-State and Happy-State follow almost the same course of activation for awhile, though the level at Sad-State is always slightly higher because the weight on the excitation link from Cry-Tears to Sad-State is greater than the weight on the link from Cry-Tears to Happy-State.

Activation energy then spreads to neighboring nodes. Because Sad-State is always higher than Happy-State, the excitation threshold at Sad-Event is exceeded before that threshold is reached at Happy-Event. The activation level at Sad-Event begins to rise, which further encourages the increase at Sad-State. As the level at Sad-State approaches the maximum, the inhibition threshold on the inhibitory link from Sad-State to Happy-State is exceeded and the activation level at Happy-State begins to plummet. Meanwhile, activation energy from Sad-Event spreads to Propose-Marriage which in turn spreads to Happy-Event. Receiving support from both Propose-Marriage and the still active Happy-State, Happy-Event exhibits a brief increase in activation which quickly dies as Happy-State drags Happy-Event down. The network settles into a stable state in which all nodes are at maximum activation levels except for Happy-State and Happy-Event, which are at minimum levels; this occurs before the second event is read. Activating Propose-Marriage does not change the distribution of energy in the network as Propose-Marriage is already at the maximum level and the network is stable, so CATLAST maintains its "sad" interpretation of Text 1. Had CATLAST been presented with the same events but

[2]For those who wish to try this at home, the parameters are set as follows: C_1 is 0.6, C_2 is 0.9, the inhibition weight is -1.0, the standard excitation weight is 0.9, the preferred excitation weight is 0.95, the inhibition threshold is 0.999, and the excitation threshold is 0.192.

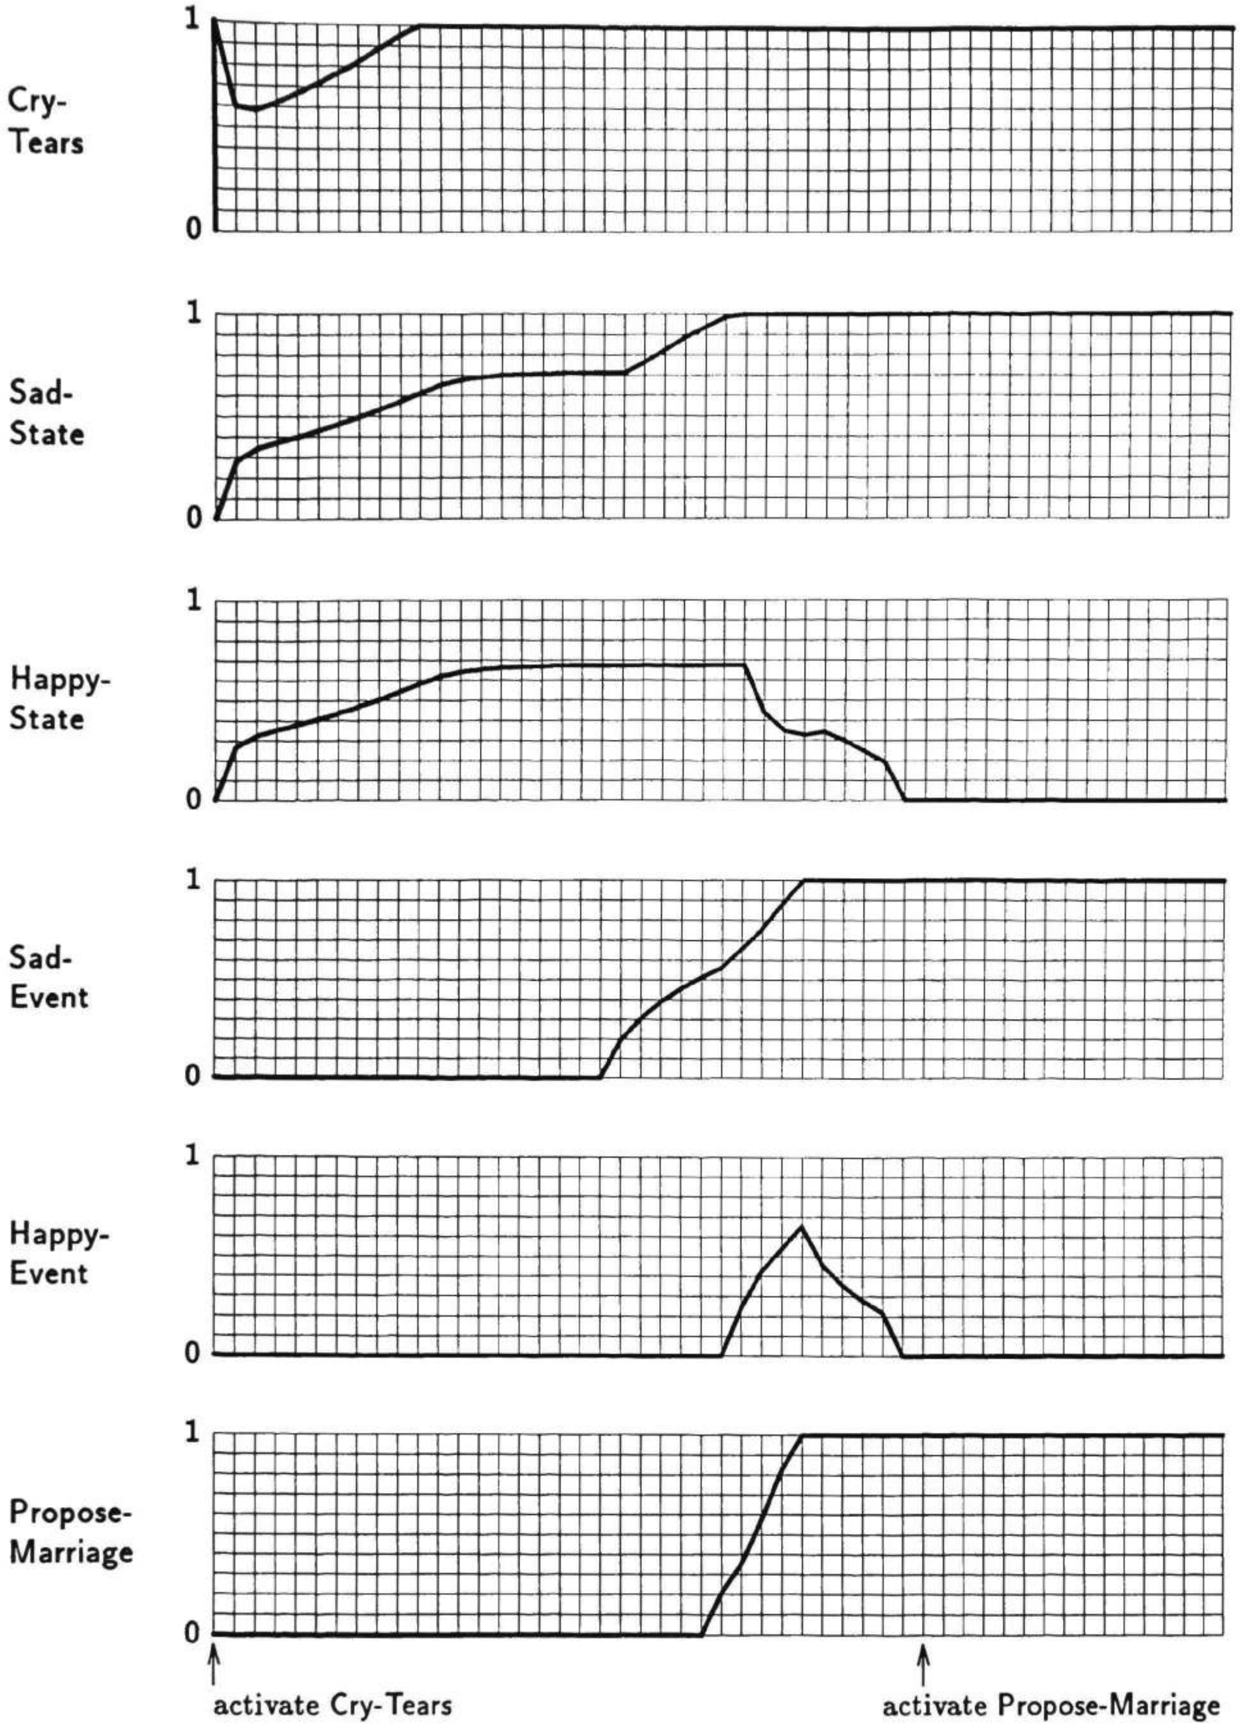

Figure 2: The time course of activation as CATLAST exhibits perseverer behavior.

in the reverse order, it would have settled upon the "happy" interpretation. Because of the symmetry of the network, the resulting time course of activation would have been the same as that shown in Figure 2, with the exception that the order of the labels assigned to the individual graphs would be inverted. Thus, Cry-Tears would exhibit the time course that Figure 2 now shows for Propose-Marriage, Sad-State would display the time course now shown for Happy-Event, and so on.

Figure 3 shows the time course of node activation as CATLAST processes Text 1 as a recency.[3] After reading the first event, CATLAST's behavior as a recency parallels its behavior as a perseverer, but raising the excitation threshold has prevented activation energy from spreading beyond Sad-State and Happy-State; the network stabilizes with a slightly higher activation level at Sad-State than at Happy-State, again because of the increased weight on the excitatory link from Cry-Tears to Sad-State.

CATLAST reads the second story event and raises Propose-Marriage to the maximum activation level. Because both Sad-Event and Happy-Event are now receiving activation energy from two sources, their activation levels rise sharply, though Happy-Event rises a bit more quickly because it is preferred by Propose-Marriage. At the exact time that Happy-Event reaches maximum activation, Sad-Event is just far enough behind to make all the difference in CATLAST's behavior because the energy at Happy-Event has exceeded the inhibition threshold but the energy at Sad-Event has not. Happy-Event begins to inhibit Sad-Event but not vice-versa, and Sad-Event begins to fall. At the same time, Sad-State and Happy-State have reached maximum activation and inhibit each other. They then fall below the inhibition threshold, their activation levels rise to maximum, and they inhibit each other again, thus displaying oscillatory behavior. Sad-State, though, loses reinforcement as Sad-Event declines, activation levels at both nodes decrease to the minimum, and CATLAST settles upon the "happy" interpretation. Again, if the story events had been presented in the reverse order, CATLAST would have settled on the alternate interpretation. The resulting time course of activation can be constructed by inverting the order of the labels of Figure 3 as described previously.

5 Conclusion

The explanation of strategy-driven inference behavior offered by CATLAST is that such behavior is not strategy-driven after all. CATLAST explains differences in inference behavior as resulting from exactly the same architecture with only very slight differences in the computing units' sensitivity to activation energy. (In the examples given above, the abrupt change from perseverer to recency behavior was caused by an increase in the excitation threshold of only 0.001.) Whether CATLAST behaves as a recency or a perseverer, the model's preference for the most parsimonious interpretation is preserved, as such preferences are inherent in connectionist architectures (e.g., given two competing inference paths in a connectionist network, the nodes in the shorter path will have higher levels of activation than the nodes in the longer path and the shorter path will be favored). In addition, the model does not rely on any serial ordering which arises through simulation

[3]The values of the parameters for this example differ from the previous example only in that the excitation threshold is now 0.193 instead of 0.192.

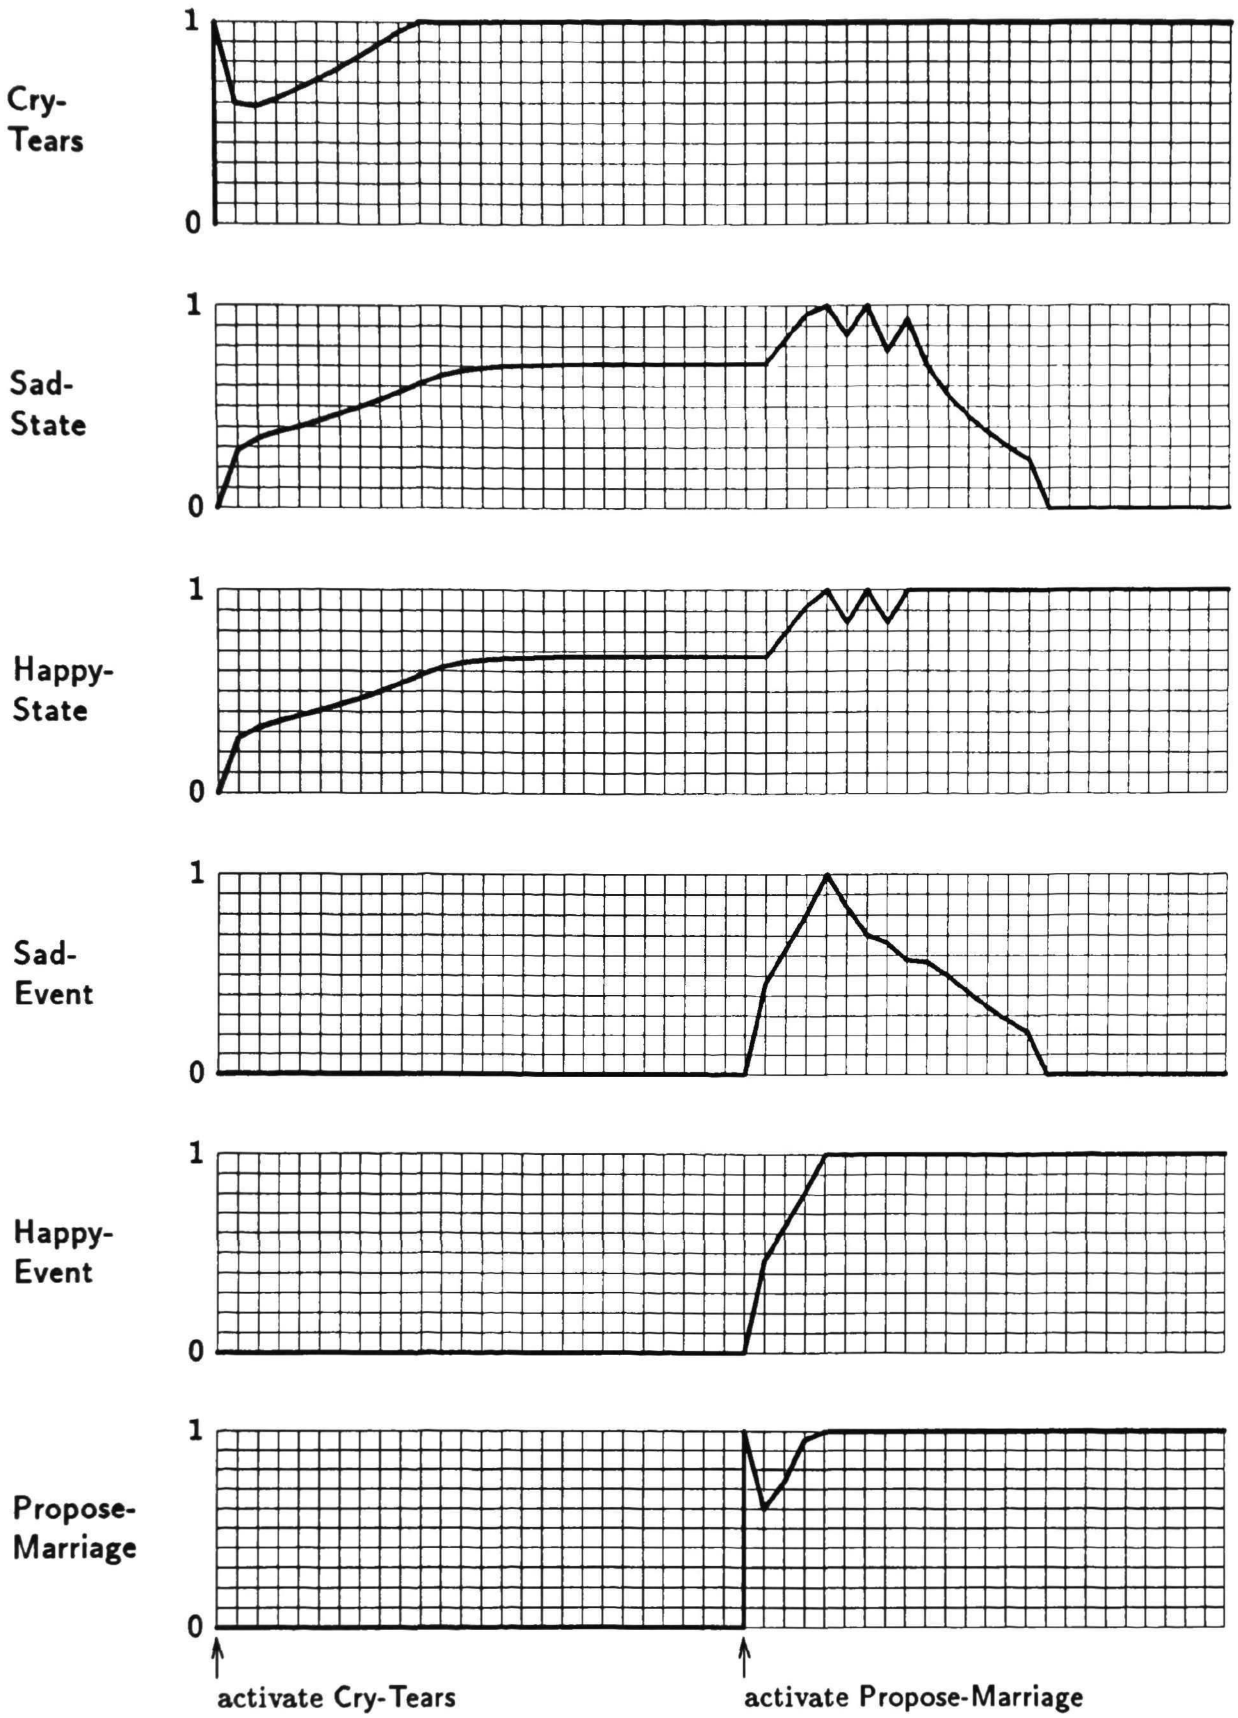

Figure 3: The time course of activation as CATLAST exhibits recency behavior.

of a parallel process on a serial processor. Thus, CATLAST provides solutions to the two problems discussed earlier that were left unsolved by ATLAST.

Not all inference decision behavior is automatic or unconscious, but we believe the perseverer/recency behavior discussed in this paper certainly is. Anecdotal evidence from the experiment described in Section 2 provides some support for this claim: several experimental subjects not only were consistent in the interpretations they assigned to test stories, they protested that other interpretations were entirely implausible (Granger et al., 1983). If attentional or conscious processes had been involved in interpreting a given story, one would expect that the subjects would have had some memory of different possible interpretations they must have entertained while reading the text; in at least these cases they did not. Thus it seems reasonable to try to explain this aspect of the human language understanding mechanism as the result of subtle differences in the underlying cognitive architecture, though it is not necessarily the case that all aspects of language understanding can be modeled adequately at this level.

Finally, some might conclude that the question of how perseverer/recency behavior comes about does not warrant the attention we have given it. After all, this behavior is usually only elicited by specially constructed texts; everyday texts are seldom if ever constructed in such a way as to leave the reader with two equally likely interpretations and no clue as to which interpretation is correct. However, the fact that people do consistently exhibit these differences, even if it occurs only in the laboratory, informs us about the human language understanding mechanism—it provides constraints which must guide the development of a comprehensive model of language understanding, however far in the future that may be.

6 Acknowledgments

This research was supported by the National Science Foundation under grant IST-85-12419. We wish to thank Jeff Schlimmer for the insights he provided during the development of CATLAST. We also wish to thank Kathleen Delling for helping to beat a large body of data into submission.

7 References

Charniak, E. (1983). Passing markers: A theory of contextual influence in language comprehension. *Cognitive Science*, 7(3), 171-190.

Collins, A.M., & Loftus, E.F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, *82*(6), 407-428.

Cottrell, G.W., & Small, S.L. (1983). A connectionist scheme for modelling word sense disambiguation. *Cognition and Brain Theory*, *6*, 89-120.

Eiselt, K.P. (1985). A parallel-process model of on-line inference processing. *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, Los Angeles, CA.

Eiselt, K.P. (1987). Recovering from erroneous inferences. *Proceedings of the National Conference on Artificial Intelligence*, Seattle, WA (to appear).

Granger, R.H., Eiselt, K.P., & Holbrook, J.K. (1983). STRATEGIST: A program that models strategy-driven and content-driven inference behavior. *Proceedings of the National Conference on Artificial Intelligence*, Washington, DC.

Granger, R.H., Eiselt, K.P., & Holbrook, J.K. (1986). Parsing with parallelism: A spreading-activation model of inference processing during text understanding. In J.L. Kolodner & C.K. Riesbeck (Eds.), *Experience, memory, and reasoning.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Granger, R.H., & Holbrook, J.K. (1983). Perseverers, recencies, and deferrers: New experimental evidence for multiple inference strategies in understanding. *Proceedings of the Fifth Annual Conference of the Cognitive Science Society*, Rochester, NY.

Hirst, G. (1984). Jumping to conclusions: Psychological reality and unreality in a word disambiguation program. *Proceedings of the Sixth Annual Conference of the Cognitive Science Society*, Boulder, CO.

McClelland, J.L., & Rumelhart, D.E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review, 88*, 375-407.

Quillian, M.R. (1968). Semantic memory. In M. Minsky (Ed.), *Semantic information processing.* Cambridge, MA: MIT Press.

Quillian, M.R. (1969). The teachable language comprehender: A simulation program and theory of language. *Communications of the ACM, 12*(8), 459-476.

Riesbeck, C.K., & Martin, C.E. (1986). Direct memory access parsing. In J.L. Kolodner & C.K. Riesbeck (Eds.), *Experience, memory, and reasoning.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Rumelhart, D.E., & McClelland, J.L. (Eds). (1986). *Parallel distributed processing: Explorations in the microstructure of cognition. Volume 1: Foundations.* Cambridge, MA: MIT Press.

Waltz, D.L., & Pollack, J.B. (1985). Massively parallel parsing: A strongly interactive model of natural language interpretation. *Cognitive Science, 9*(1), 51-74.