DISCOURSE PROCESSES, *39*(2&3), 129–164

Goal Plans of Action and Inferences During Comprehension of Narratives

Tom Trabasso
University of Chicago

Jennifer Wiley
University of Illinois at Chicago

A theory of how readers monitor concerns of characters and make causal inferences during reading is presented. The focus is on the reader's understanding of what characters do when goals either succeed or fail. Knowledge of goal processes enable coherent understanding to be achieved when characters change goal plans and pursue new courses of action. Four simulations of data from previously published studies on goal-related inferences during reading are reported. Narrative texts and data came from Experiments 1, 2, and 3 of McKoon and Ratcliff (1992) and Experiment 3 of van den Broek and Lorch (1993). In the simulations, use of goal plan knowledge in making causal inferences is illustrated in the discourse analysis of the texts. To simulate online processing, the causal relations between clauses or sentences were processed 1 clause or sentence at a time by a connectionist model. Estimates of accessibility of clauses or sentences in long-term memory by the model successfully predicted differences in response times from the 4 experiments. The successful simulations provide strong evidence that readers monitor and infer goal plans of action and make causal inferences based on them during reading.

This article examines the role that causal inferences play in understanding narratives that have complex goal plans of action and outcomes. The approach taken is to analyze the casual structure of texts on both local and global levels. To determine the global causal structure, each text is subjected to a goal process analysis based on Stein and Albro (1997); Stein and Levine (1989); and Stein, Trabasso, and Liwag (1993). Second, each text is analyzed in terms of the necessary causal

Correspondence and requests for reprints should be sent to Tom Trabasso, The University of Chicago, Department of Psychology, 5848 South University Avenue, Chicago, IL 60637. E-mail: trabasso@uchicago.edu.

relations among the story components. Using these sources of data about the text structures, the processing of each text can be simulated in a connectionist model (Langston & Trabasso, 1999). Measures of accessibility derived from the connectionist model are then used to predict recognition priming and sentence priming data from actual experiments. Using the texts and data of McKoon and Ratcliff (1992) and van den Broek and Lorch (1993), the main question in this simulation approach is whether causal inferences that are made over the course of comprehending a text are those based mainly on the goals of the characters (Trabasso & Magliano, 1996).

A widely held view of narrative understanding (Graesser, Millis, & Zwaan, 1997; Zwaan & Radvansky, 1998) is that readers track characters in space and time and monitor their concerns. If readers track characters in space and time, they require the use of knowledge of the spatial and temporal conditions that enable the states and actions of characters to occur. Further, if readers monitor concerns, they require the use of knowledge of human goals, goal plans, their related actions, and the outcomes that affect goal success or failure. One kind of evidence of the use of these kinds of knowledge during reading is in answering who, what, where, when, how, and why questions after reading (Graesser, 1981; Trabasso, van den Broek, & Liu, 1988). The question in this study is whether these goal-based inferences occur during reading without question prompts.

A goal plan originates when the concerns of a character are not being advanced, such as when a character's goal is blocked or has failed and something must be done about it (Stein & Albro, 1997). Goals and plans are maintained as long as a character's attempts are advancing the goal. On the other hand, subordinate goals and plans are generated when superordinate goals are blocked or fail. Goals and plans that have been blocked are reinstated when the character has achieved a subordinate goal that enables the attainment of a superordinate goal. Understanding the options that characters have with respect to goal attainment, maintenance, substitution, revision, abandonment, and enablement is necessary for the coherent understanding of the narrative. This knowledge of what happens to the concerns of characters over the course of a narrative enables the reader to monitor and update information as events, states, actions, and state changes occur over time.

To illustrate the role that goal processes play in story understanding, consider the following excerpt from the Farmer and Donkey story in Rumelhart (1977): "There once was a farmer who wanted to get his donkey into the barn. He pushed and he pulled but the donkey would not budge."

In this excerpt, the farmer has a goal of getting the donkey in the barn. The actions of pushing and pulling follow from the farmer's goal plan. The donkey resists and thwarts the farmer's goal. This goal plan of action and its outcome are not sufficient for understanding. The goal, actions, and outcome have to be understood in relation to one another. This can be accomplished by causal inferences (Trabasso,

van den Broek, & Suh, 1989). The goal motivates each action as an attempt to achieve it; the attempts cause resistance and the outcome of failure.

The causal inferences that connect the components of an unfolding goal plan form an elementary, episodic unit called a *goal–attempt–outcome* episode (Trabasso & Nickels, 1992). These goal–attempt–outcomes organize narrative events and serve as building blocks for encoding complex narratives when the episodes themselves are interconnected. In an analysis of narrations of picture stories by children and adults, Trabasso and Nickels found that these goal–attempt–outcome units serve as the basic building blocks of the structure of a story. Developmentally, from age 5 to adulthood, Trabasso and Nickels found that in encoding picture stories, the basic goal–attempt–outcome episodes emerge and increase in number and elaborated complexity as the narrators mature.

Returning to the Farmer and Donkey story, we witness a new, subordinate goal plan of action with an outcome: "The farmer thought that he could scare the donkey to go into the barn. He asked a dog to bark at the donkey. The dog refused."

In this excerpt, the farmer developed a subordinate goal plan of action to achieve his superordinate goal of getting the donkey into the barn. The subordinate goal plan involved one action, a request, and a failure outcome. The goal of scaring the donkey by asking a dog to bark was psychologically caused by the goal of wanting to get the donkey into the barn jointly with the failure caused by the donkey's refusal. In this excerpt then, another goal–attempt–outcome occurs as a result of a superordinate goal and its failure.

The story continues on with repeated subordinate goal generations to each subsequent subordinate goal failure, with repeated attempts to carry out the plan and achieve a subordinate goal, and with repeated subordinate goal outcomes of failure until the farmer successfully obtains milk from a cow by giving her hay. He gives the milk to a cat, which scratches the dog, which makes him howl, which scares the donkey that runs into the barn. The success of the subordinate goal of obtaining milk enables the goal plan concerning the cat to succeed and enables the goal plan concerning the dog to succeed, which in turn enables the goal plan concerning the donkey to succeed. In each case, the successful outcome enables an attempt at one level above the previous goal in the goal hierarchy. Each attempt, however, is motivated by a reinstatement of the previously failed goal plan at the next level up in a hierarchy of goals, attempts, and outcomes.

GOAL PROCESS ANALYSIS

The Farmer story illustrates the importance of understanding goal processes over the course of reading a story. The understanding of what characters do as a result of goal failure enables the reader to construct coherent interpretations of new goals and actions. The understanding of goals requires understanding of the contingen-

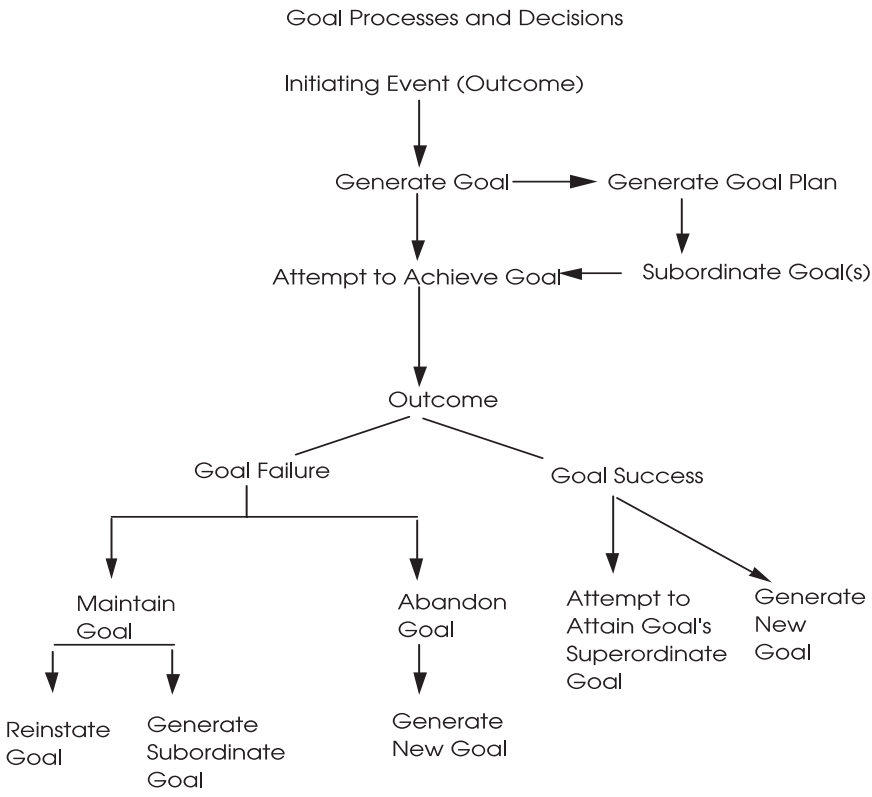

FIGURE 1 Decision tree for goal outcome processes.

cies of goal success and failure. Figure 1 depicts these contingencies in terms of an action–outcome decision tree. Figure 1 is based on the work of Stein and Albro (1997), Stein et al. (1993), and Trabasso et al. (1989).

Goals are generated when events that happen to a character impact the character's well being (Stein & Glenn, 1979; Stein & Levine, 1989; Stein et al., 1993). Initiating events that start an episode and goal outcomes, especially goal failures, are appraised in terms of which of the character's goals are affected. New goals are generated as a plan to deal with the problem. The character may generate a goal plan complete with subordinate goals and attempts. A good example is found in Stein and Glenn's (1979) Tiger Whisker story, where a woman plans to obtain a tiger's whisker to make a special medicine (main goal) by playing a song, luring the tiger from his cave and lulling him to sleep, and then cutting off the whisker. The story presents this plan and ends by saying that is what she did as a summary of the actions carried out. This kind of plan generation is indicated on the right side of Figure 1 after General Goal generation.

Once the character begins to carry out voluntary actions, attempts to achieve a goal can occur. The attempt leads to an outcome. Outcomes indicate mainly goal success or goal failure. If goal failure occurs, then the character has two main options: either to maintain or to abandon the current goal. There are two ways in which a goal is maintained. First, the character can reinstate the same goal and try again or generate a new but directly related subordinate goal, thereby modifying the a priori goal plan. For example, in the Jimmy Story of van den Broek (1988), Jimmy decides to get a job to earn money after the goal of asking his mother to buy him a bike fails by her refusal. The new goal is coherent with a goal plan to obtain a bike.

When a goal is abandoned, a new goal plan is substituted. If a character has a goal to lose weight and a goal plan of dieting does not work, then the character could decide to not lose weight. Instead, if the character's physical appearance is the superordinate goal that motivated loss of weight, the character could decide to shop for new and flattering clothes. The new goal does not represent a coherence break because it is consistent with a superordinate goal plan.

When a goal success occurs, the success may enable the pursuit of another, subordinate goal. In this case, a goal that is higher in the goal hierarchy is enabled and can motivate new attempts. For example, in Suh and Trabasso's (1993) Betty story, Betty achieves a subordinate goal of knitting a sweater. The success of this goal enables her to wrap and give the sweater as a present to her mother. The new actions are coherent with the enabled goal of giving the present.

In sum, the construction of a coherent understanding of a story is thus achieved by inferring relations between the goals, attempts, and outcomes that occur over the course of the narrative. All goals that occur can be understood in terms of their status where goals are threatened, advanced, succeed, or fail. Any narrative text can be analyzed in relation to goal plans and the goal-related global structure of a text can thus be made explicit.

CAUSAL DISCOURSE ANALYSIS AND CONNECTIONIST MODELING

In addition to a goal process analysis, narratives can also be analyzed in terms of their causal networks based on the procedures of Trabasso et al. (1989). In this procedure, each clause is evaluated in terms of its causal relation to all other clauses (i.e., did *x* lead to or enable *y*? If not for *x,* could *y* have occurred?). Logical necessity is determined through counterfactual reasoning. Counterfactuals illustrate the logic of necessity proposed for causes by Mackie (1980) and used by Trabasso et al. as criteria for an identified causal relation. Why and how questions based on Graesser (1981) are another means for identifying causal relations: Why did the crowd cheer? Answer: Because the president arrived. Therefore, the president's arrival caused the

crowd to cheer. Using this method of discourse analysis, a causal network can be created where the causal relations for each clause in a narrative are depicted.

The results of the discourse analysis are then submitted to a connectionist model (Langston & Trabasso, 1999). The model has a long-term storage component called the text representation that contains nodes, connections between nodes, and quantitative values that change over time as each new node is processed. The nodes correspond to clauses from the discourse being modeled. Each node has an associated activation value that changes over time as the text representation is constructed. Each connection between nodes has a connection strength. Connection strengths change as new nodes and their connections are integrated into the existing text representation. The activation values and connection strengths reflect, through processing, both the current memory representation of the reader and the history of comprehension.

The network is processed by the model one node at a time to simulate reading comprehension. The model incorporates each new node and its connections into an existing text representation and spreads activation among the nodes. Once activation has settled, the activation values are used to adjust the strengths of the connections between the nodes. At any time during or after this processing, connection strength values may be obtained. Although activation values, settling rates, or connection strengths index accessibility in memory and each could be used in a simulation, Langston, Trabasso, and Magliano (1998) found empirically that connection strength was the most reliable and psychologically meaningful predictor of accessibility. The average connection strength of a node is used when the sources of activation of a given node (clause or sentence) are multiple and the connection strength between two nodes is used when one node is the source of activation of a second node (see Langston & Trabasso, 1999; Trabasso & Bartolone, 2003, for applications of both measures in simulation). Using the results of the goal and discourse analysis as input, this model will be used to predict the accessibility of goal-related probes in the following simulations of the data of McKoon and Ratcliff (1992).

SIMULATIONS

McKoon and Ratcliff (1992) reported three experiments that vary the goal structures of stories. The accessibility of goal-related information was assessed by priming. We explicitly modeled the casual structure of their stories, submitted the networks of these discourse analyses to a connectionist model, and tested the predictions based on the model's estimates of accessibility against the McKoon–Ratcliff data for Experiments 1 through 3. We now present the simulations of each of their experiments in turn.

EXPERIMENT 1

In Experiment 1, the accessibility of goal information was measured by the recognition time for each of two test words. One recognition test was on a superordinate goal word and the other was on a subordinate goal word, both of which were taken from an eight-sentence story. Participants read a story, sentence by sentence, and the recognition test word was presented after the final sentence of the story. The participant decided whether or not the test word had been in the story. The amount of time required to respond was recorded. It should be noted that the test word occurred after the story was read so that its speed of accessibility should be affected by the inferences that were made during reading.

McKoon and Ratcliff (1992) compared word recognition times in three conditions. The texts across the conditions had a common core narrative. We simulate the data based on an analysis of the example story given in the article and presumed representative. The example, core story is about an assassin who attempts to kill the president using a rifle. Each condition had a different ending to the core story. Table 1 displays the common core, experimental conditions with different story endings, and test words for the Assassin stories in Experiment 1.

In Table 1, the stories are parsed into clauses where each clause contains one main predicate. The superordinate goal (to kill) was explicit in the text. The subordinate goals are inferred and marked by brackets. The categories of the clauses are based on those of Stein and Glenn (1979) that were used by Trabasso et al. (1989) to identify episodic functions.

Goal Process Analysis of Experiment 1 Stories

A goal process analysis of the three versions in Table 1 is summarized in Figure 2. The following analysis uses this figure as a reference.

In all three conditions in Table 1, the character's superordinate goal is to kill the president. The subordinate goal is to shoot the president with a rifle. This goal is inferred from three actions of using the rifle and scope. The three conditions introduce variations on goal success or failure and goal processes of substitution, maintenance, or abandonment. In the control condition, the assassin carries out a subordinate goal plan of shooting the president with the rifle. This subordinate goal apparently succeeds in that the president is hit by the shot. The superordinate goal success of killing the president must be inferred. The success of hitting, and presumably killing the president, leads to a new goal of the assassin, namely to escape with his life. However, this goal is blocked by the sun shining in his eyes. On the basis of this analysis, we renamed the control condition: goal success, new goal blockage.

Experiment 1 manipulated two kinds of goal failure. In the try again condition, diagramed on the left side of Figure 2, the subordinate goal of shooting the presi-

TABLE 1
Common Core and Condition Story Variations (Assassin Narrative, Experiment 1, McKoon & Ratcliff, 1992)

	Common Core	*Category*
1	the president's arrival	Event
2	the crowd cheers alerted onlookers (to)	Event
3	The assassin wanted to kill the president	Goal
4	[He wanted to shoot him with a rifle]	Goal
5	He reached for his high-powered rifle	Attempt
6	to peer through its scope	Goal
7	He lifted the gun to his shoulder	Attempt
Control conditions		
(goal success, new goal blockage)		
8	[He fired] the first shot from his rifle	Attempt
9	The assassin [shot] hit the president	Outcome
10	[He wanted to escape]	Goal
11	Then he started to run toward the west	Attempt
12	The searing sun blinded his eyes	Outcome
Substitution		
(goal failure, goal substitution, goal blockage)		
8	As he lifted his rifle	Attempt
9	the scope fell off	Outcome
10	[He wanted to blow up the president]	Goal
11	So he reached for his hand grenades	Attempt
12	The searing sun blinded his eyes	Outcome
Try again		
(goal failure, goal reinstatement, goal blockage)		
8	As he lifted the rifle	Attempt
9	the scope fell off	Outcome
10	to draw a sight without the scope	Goal
11	He lay prone	Attempt
12	The searing sun blinded his eyes	Outcome
Test word		
Subordinate goal: "Rifle" or		
Superordinate goal: "Kill"		

dent with a rifle fails. The superordinate goal to kill the president is maintained as is the subordinate goal of using the rifle. The subordinate goal of the rifle goal, the use of the telescope, is abandoned, and is replaced by a new subordinate goal to use the rifle's sight instead. This goal is blocked as the sun shines in the assassin's eyes.

In the substitution condition, diagramed in the middle of Figure 2, the subordinate goal of shooting the president with a rifle fails. This subordinate goal failure leads to the abandonment of the subordinate goal of shooting the president with the rifle. The superordinate goal of killing the president is maintained. With the abandonment of the rifle as a means to kill the president, a new subordinate

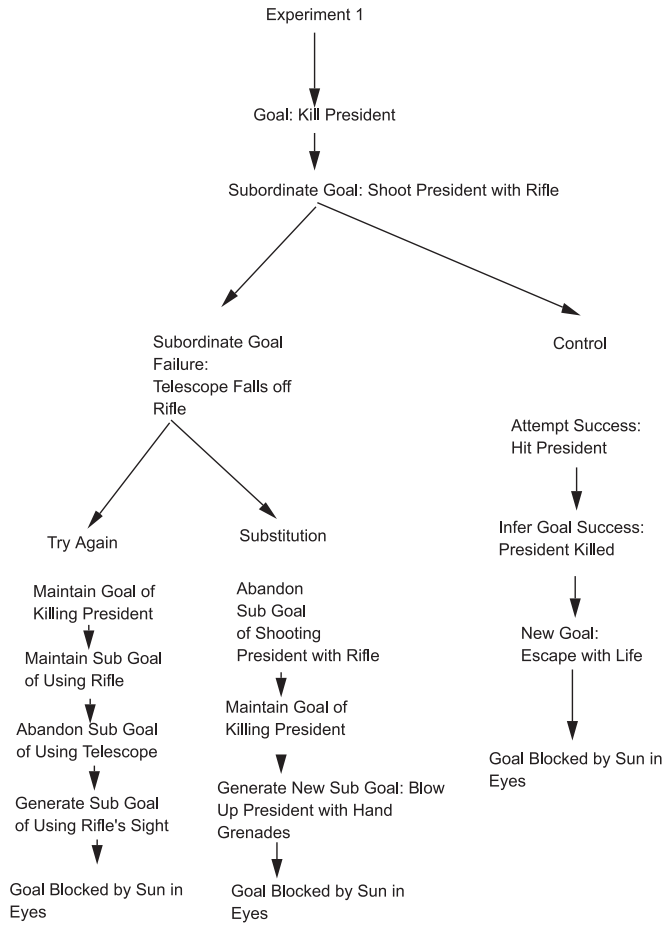

FIGURE 2 Goal process analysis of Assassin stories in Table 1.

goal is "substituted." The new means to kill the president is with a hand grenade. This goal is blocked when the assassin reaches for the hand grenade and the sun shines in his eyes.

Causal Discourse Analysis of Experiment 1 Versions

Figure 3 shows the causal networks for each of the three conditions derived from a causal discourse analysis of the story versions in Table 1 based on the procedures of Trabasso et al. (1989). The clauses for each story were first classified according to the episodic categories of Stein and Glenn (1979; setting, goal, event, attempt,

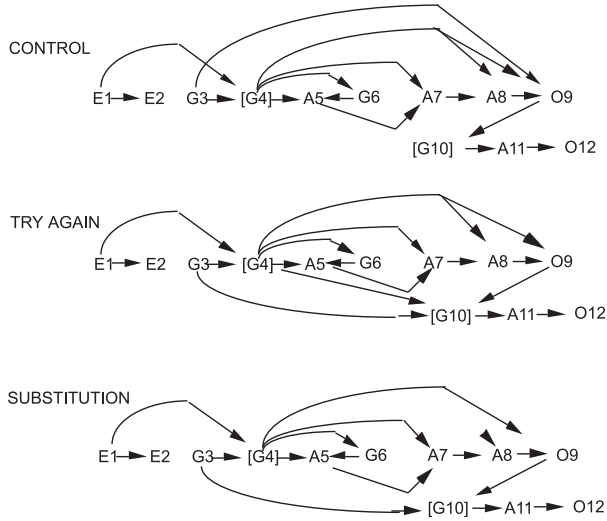

FIGURE 3 Causal networks for conditions of Experiment 1.

and outcome). The causal relations were then identified using the constraints of the episodic categories, how and why questions, and counterfactual tests.

The first seven clauses define the core story of each network. In the core story, the first event, E1, is that of the president's arrival. Event E1 psychologically causes Event E2, a second event of a crowd cheering and onlookers being alerted. In the third clause, G3, the assassin's superordinate goal, namely to kill the president, is psychologically caused by the president's arrival. This goal, G3, motivates an inferred, subordinate goal, G4, to shoot the president with a rifle. This goal is inferred from the action in Attempt A5, reaching for a rifle, and follows the action-plus-goal assumption of Graesser (1981). The inferred subordinate goal, G4, is psychologically caused by the president's arrival, Event E1. Event E1 and Goal G3 are jointly necessary for Goal G4. That is, if the president had not arrived or if the assassin had not wanted to kill the president, the assassin would not have wanted to shoot him, in the circumstances of the story. Further, Goal G4 explains why the assassin reached for a rifle to satisfy the superordinate goal, G3. The subordinate goal of using the rifle, G4, motivates a subordinate goal, G6, to peer through the rifle's scope (to shoot), and G4 motivates an attempt, A5, to reach for the rifle. Reaching for the rifle, A5, enables the assassin to lift the gun to his shoulder, A7. If he had not reached for the rifle, he would not have been able to lift the gun to his shoulder. The goal, G4, is jointly necessary with A5 for A7 because if he had not wanted to shoot the president, he would not have lifted the gun to his shoulder. This completes the causal analysis of the core story.

Control (goal success, new goal blockage). In the top network of Figure 3, the attempt, A8, the firing of the first shot by the assassin, is enabled by Attempt 7, the lifting the gun to his shoulder, and is motivated by the subordinate goal, G4, wanting to shoot the president. The attempt/outcome, O9, of the shot hitting the president, is motivated by two goals, G3 and G4, the desires to kill the president and to shoot him. The outcome, O9, is also an event in a new episode that psychologically initiates a new goal, G10, to escape, and that, in turn motivates an attempt, A11, to run. This attempt, A11, of running toward the west, enables the sun to blind him, O12. If the assassin had not shot the president, he would have ran toward the west and if he had not run toward the west, then the sun would not have blinded him, in the circumstances of the story.

Substitution (goal failure, goal substitution, goal blockage). The main difference between the substitution and control stories lies in the fact that the Outcome, O9, is changed so that it is a goal failure only for the subordinate goal, G4, the using of the rifle but not of G3, the killing of the president. The failure of the subordinate goal, G4, leads to its abandonment and a goal substitution, a subordinate goal, G10, to blow up the president, inferred from A11, the action of reaching for the hand grenades. Goal G10 is motivated by Goal G3, the goal to kill the president, and by the failed subordinate goal outcome, O9, of not being able to use the rifle. The new goal, G10, motivated an attempt, A11, to reach for the hand grenades. This attempt, All, enables an outcome, O12, the sun's blinding him. The analysis indicates that the main difference between the causal networks of the control and substitution conditions is the distal connection of the superordinate Goal G3. The shift is one from motivating a successful outcome to being reinstated to motivate a new subordinate goal. Otherwise, the causal networks are identical.

Try again (goal failure, goal reinstatement, goal blockage). The Try Again version differs from the control in the same way as did the substitution version. The outcome, O9, was changed so that it is now a failure of only G4, the subordinate goal to use the rifle. The other difference between try again and the other two versions is that the superordinate goal, G3, and the subordinate goal, G4, jointly motivate a new, subordinate goal, G10, to draw a sight without the scope. This subordinate goal maintains both goals (hence, the name try again). The subordinate goal, G10, motivates attempt, A11, to lay prone, and this attempt enables the outcome, O12, of the sun shining in his eyes.

Comparison of networks. The three networks reveal that in all three versions, the subordinate goal, G4, of using the rifle is more highly connected than the superordinate goal, G3, of killing the president. The word "rifle" or anaphoric references to it appear in five clauses (G4, A5, G6, A7, and A8) whereas the word "kill" occurs only in one clause, G3, confounding frequency with the level of the

goal. The difference in frequency between these words was not noted in McKoon and Ratcliff (1992) but it possibly confounds their interpretation of their results in terms of goal accessibility because they made the superordinate goal word less frequent and therefore less available.

The superordinate goal clause that contains the target word, kill, has two causal connections in each condition whereas the subordinate goal clause that contains the target word, rifle, has either 5, 6, and 5 connections, respectively, in the control, try again, and substitution versions. This difference in connections would make "rifle" more accessible than "kill" (Langston & Trabasso, 1999; Trabasso & van den Broek, 1985; van den Broek, 1988; van den Broek & Trabasso, 1986) because "rifle" would have many more means of being activated and accessed than the word "kill." The difference in connections for the subordinate goal across versions would make "rifle" slightly more accessible in the try again condition than in the other two conditions. Finally, the similarity of connections for the superordinate goal across versions would make it equally accessible across conditions. There is a strong relation between connections in the network and connection strength, so the connectionist model measures of accessibility and response times should reflect these expectations.

Connectionist Modeling

The three networks from Figure 3 were each submitted to the Langston and Trabasso (1999) model, one pair of connected nodes at a time from the first connected pair to the last in the network. Because each target word appeared in more than one clause, we assumed the accessibility of the word would be determined by the node with the highest activation from multiple sources in the story. Accessibility was therefore measured by the strongest average causal connection strength of a node that contained the target word. These were goal nodes G3 for "kill" and G4 for "rifle." The respective measures of average connection strength for each of the three texts are shown as predictors from the model for the observed response times of Experiment 1 in Figure 4.

Variation in average connection strength across conditions of the clauses that contained "kill" and "rifle" accounted for 85% of the variation in mean response times across conditions, $F(1, 2) = 23.17$, $p < .01$. It should be noted that the connection strength predictions are "free" and independent. The connection strengths were generated a priori and independently of response times so that the latter were free to vary. Regression analysis is used here as a goodness-of-fit measure of the Predictions of the Recognition Times × Connection Strength measures. This analysis indicates that the connection strength measures were highly accurate in predicting response times that, in turn, measure accessibility of information in long-term memory. We note that the simulation, based on a causal analysis of the text, accurately predicted the greater accessibility of the subordinate goal word, rifle, over the superordinate goal word, kill. Second, the simulation also predicted

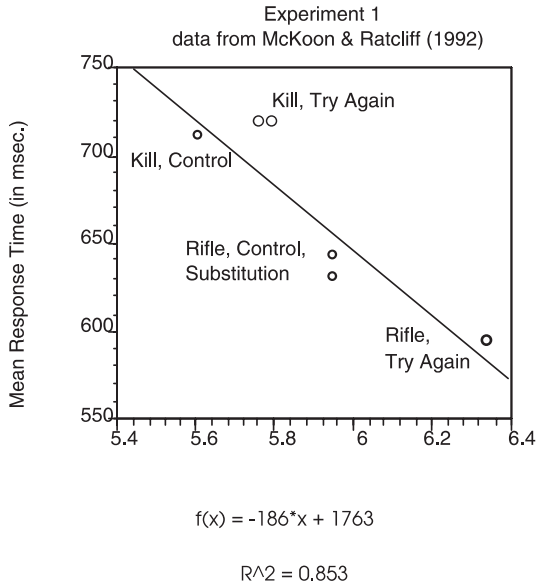

FIGURE 4 Accessibility of words: Prediction of Observed Response Times for Conditions of Experiment 1 × Average Connection Strength.

the greater availability of "rifle" in the try-again condition versus the other two rifle conditions, and little difference among the "kill" conditions.

The connection strength measures assume that readers made the inferences revealed in the analysis and their high degree of accuracy supports the validity of this assumption.

EXPERIMENT 2

In Experiment 2, McKoon and Ratcliff (1992) varied what they termed "global and local inconsistency" across four versions of a story. Participants read one of two story versions within each of the inconsistency conditions. After reading a story, a single target word was presented and readers had to judge whether the word had appeared in the story. Table 2 gives their example stories for each of the four conditions.

Goal Process Analysis of Experiment 2 Versions

A goal process analysis of the Experiment 2 versions in Table 2 is summarized in Figure 5. The following analysis of the stories in Table 2 uses Figure 5 as a reference. In the globally inconsistent stories, the main superordinate goal is to prepare for an upcoming tennis match. The first, related subordinate goal is to obtain a workout. This is followed by a second, related subordinate goal to find an oppo-

TABLE 2
Locally Inconsistent Versions for Experiment 2

Common Core	*Category*
1. Diane wanted to lose weight	Goal
2. She thought she should lose at least 20 pounds	Goal
3. Diane thought cycling might help her lose weight	Goal
4. to find her bike	Goal
5. She went to the garage	Attempt
Control (goal success, new goal success)	
6. Diane peddled 5 miles each day for 3 months and	Attempt
7. became very thin	Outcome
8. to complete her degree	Goal
9. She decided to go back to school	Goal
10. It took her several years but	Outcome
11. Diane reached her goal	Outcome
Problem (goal failure, substitution, and success)	
6. Diane's bike was broken and	Outcome
7. she could not afford a new one	Setting
8. to buy grapefruit and yogurt	Goal
9. So she went to the grocery store	Attempt
10. It took several years but	Outcome
11. Diane reached her goal	Outcome
Test word: weight	

nent for the workout. In the control Continuation, the finding of an opponent, the subordinate goal meets with success, which enables the success of the subordinate goal of obtaining of a workout. A new goal is generated following goal success, to go home. This goal is enabled by the completion of the match outdoors. In the problem continuation, the second, subordinate goal of finding an opponent fails; this failure also leads to the failure of the first subordinate goal of a workout. However, the superordinate goal of preparing for a match is maintained and it motivates a new subordinate goal of viewing tapes. The subordinate goal of a workout is abandoned. The attempt is to go home to view the tapes to prepare for the match, the superordinate goal.

In the locally inconsistent stories, the main superordinate goal is to lose weight. The first subordinate goal is to go biking and this motivates a second subordinate goal to find a bike. In the control Continuation condition, the second subordinate goal succeeds and this enables the attempts at attaining the first subordinate goal of biking. The attempts are motivated by and result in the success of the superordinate goal of losing weight. In the problem continuation, the subordinate goals of finding and using a bike to go biking fail. The superordinate goal of losing weight, however, is maintained. The failed, subordinate goal of biking is abandoned. The superordinate goal of losing weight motivates a new subordinate goal to diet. This

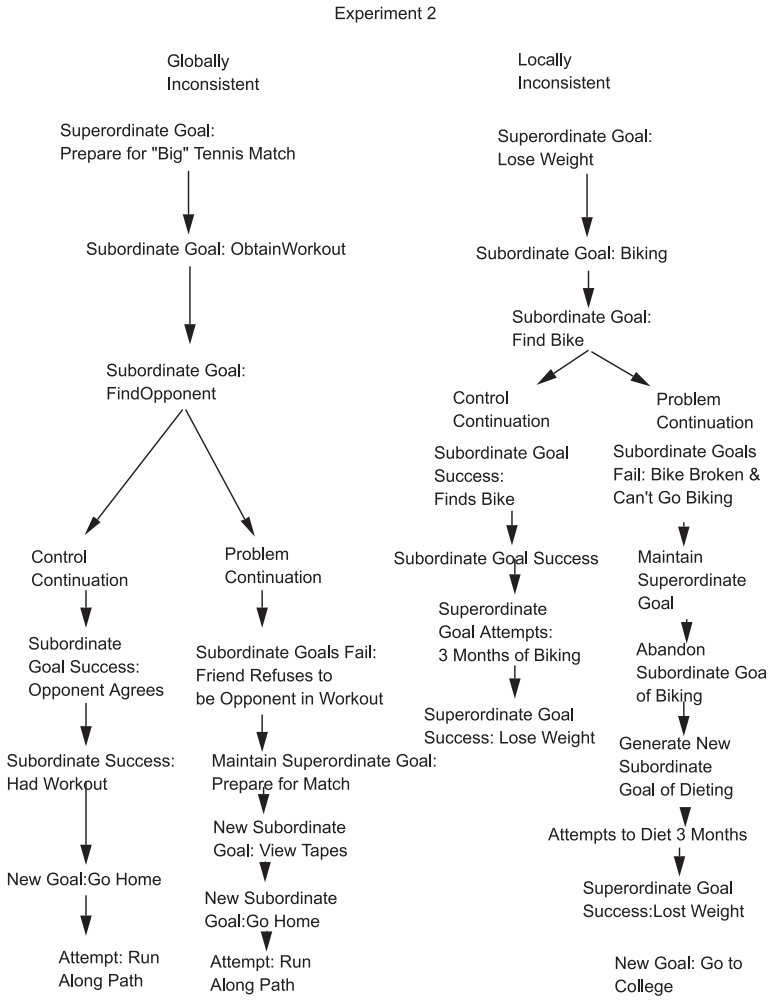

FIGURE 5 Goal processes for conditions of Experiment 2.

goal motivates attempts over a period of time and the superordinate goal succeeds in that weight is achieved. The new goal of going to college is not motivated in the story.

Discourse analysis and causal networks of Experiment 2 versions. Tables 2 and 3 display the parsed and episodically categorized clauses for the McKoon–Ratcliff published story versions for their locally and globally inconsistent conditions in Experiment 2. Figure 6 displays the causal networks that were

TABLE 3
Globally Inconsistent Versions for Experiment 2

Common Core	*Category*
1. Curtis spied a tennis court in the park	Event
2. His arm was healing from an injury	Setting
3. And (he) needed a workout	Goal
4. Before the big match	Event/[Goal]
5. So he needed an opponent	Goal
6. To join him	Goal
7. Curtis waved to a friend	Attempt
Control (goal success, new goal)	
8. The friend came over and	Outcome
9. Was an exhausting opponent	Outcome
10. For a drink	Goal
11. Curtis decided to go home	Goal
12. Curtis ran happily along the path	Attempt
Problem (goal failure, goal substitution)	
8. Curtis's friend did not want to be Curtis's opponent	Outcome
9. So Curtis decided to study videotapes of his serve instead	Goal
10. [Curtis decided] to go home	Goal
11. Curtis ran happily along the path	Attempt
Test word: workout	

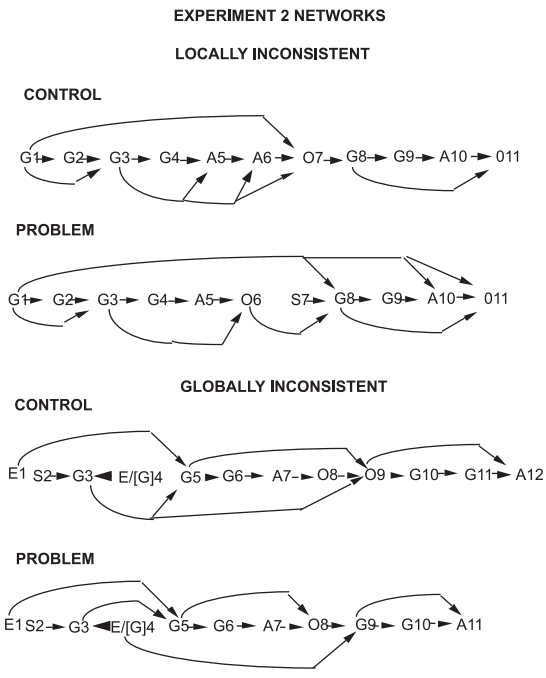

FIGURE 6 Causal networks for conditions of Experiment 2.

derived from application of the Trabasso et al. (1989) analysis of the versions in Tables 2 and 3.

As in Experiment 1, the Experiment 2 stories had a core version and variations in endings for each condition. In the locally inconsistent core (the first six clauses of Table 2 and in the top two causal networks of Figure 3), the superordinate goal, G1, is Diane's desire to lose weight. This goal motivates a subordinate goal, G2, of losing 20 pounds. Both goals, G1 and G2, motivate a subordinate goal, G3, to lose the weight by cycling. This subordinate goal dominates the action over the next four clauses. The goal, G3, motivates a subordinate goal, G4, to find a bike. The attempt, A5, of going to the garage is motivated by the Goal G4, to find a bike, and by the Goal G3, to cycle. Each of these relations satisfies criteria of necessity in the circumstances; for example, if Diane had not wanted to lose weight by cycling, then she would not have wanted to find a bike; if Diane had not wanted to find a bike to lose weight by cycling, then she would not have looked for a bike in the garage, in the circumstances of the story.

The causal network for the control (goal success–new goal success) version of the locally inconsistent condition is shown at the top of Figure 2. In the control network, the attempt, A6, of peddling 5 miles per day for 3 months, is motivated by the goal to cycle, G3, and is enabled by her attempt to go to the garage to find the bike, A5. The Attempt A6 enables a goal success, O7. This attempt and the successful outcome were each motivated by the goals of G1, losing weight, and G3, cycling to do so. A new goal, G8, to complete a degree, is generated. It may have been enabled by the outcome of losing weight, O7. The goal, G8, of completing a degree motivates a subordinate goal to return to school, G9. The subordinate G9 motivates a number of attempts, A10, and a successful outcome, O11, that is also enabled by the attempts in A10.

The "local inconsistency" of the control condition is the new goal that is generated at G8 after success of a prior goal at O7. There seems to be no clear relation between O7 and G8. It is possible to infer that going to college was successful upon looking fit so the success of losing weight could have been an enabling condition for the new goal of going to college. Why else would a narrator juxtapose these two episodes? In the interests of communicative felicity, an enabling relation is inferred, though its existence has little bearing on the accessibility of G1. Further, O7 and G8 are referentially locally coherent because the same character, Diane, is involved in both clauses.

In the core of the globally inconsistent versions (Table 3 and the two networks in the lower half of Figure 6), there are initiating events. The first is E1, the discovery of a tennis court in the park by Curtis. The second is an initiating event with its inferred goal, E/[G]4, which informs the reader that Curtis has a goal to compete in a forthcoming tennis match. A setting, S2, indicates that Curtis is healing from an injury and this condition psychologically causes a subordinate goal, G3, to workout. The superordinate goal to compete in the tennis match, E/[G4], is a jointly necessary condition for G3, to workout. The goal of a workout, G3, motivates a

subordinate goal of finding an opponent, G5. This goal, G5, is also enabled by the discovery of the tennis court, E1, so that both the presence of the tennis court and the desire to work out are jointly necessary conditions for finding an opponent. Once established, G5 motivates a subordinate goal, G6, to have someone join him. The subordinate goal, G6, motivates an attempt, A7, to wave to a friend.

The main contrast between the control and problem versions of the globally inconsistent condition lies in the Outcome O9. In the control (goal success, new goal) version, attempt, A7, to wave to obtain an opponent, psychologically causes the friend to come over in Outcome O8. Outcome 08 enables Outcome O9, the friend's being an opponent. Outcome O9 enables the success of Goal G5 where an opponent is found and another subordinate goal success outcome, O9, attaining a workout goal, G3. In the problem (goal failure, goal substitution) version, the same attempt, A7, enables a goal failure Outcome O8, not finding an opponent for Goal G5.

The different outcomes have different causal consequences. In the control version, the successful, strenuous workout in O9 physically causes a goal, G10, of thirst. This goal, G9, motivates an attempt to obtain water, G10, which, in turn, motivates another goal, G11, to go home to obtain the drink in Attempt A12.

In the problem conditions, the failure to obtain an opponent, O8, does not lead to reinstatement of the subordinate goal, G5, to find an opponent and have a workout. Rather, its failure psychologically causes its abandonment and the generation of a new subordinate goal, G9, to watch video tapes. This goal, G9, is consistent with preparation for the upcoming match in the event/goal E/[G]4, and this subordinate Goal G9 is psychologically caused by that event and the subordinate Goal G5 failure and abandonment. In the problem version, the last subordinate goal, G10, of going home, and the last attempt, A10, of running along the path are motivated by the new subordinate Goal G9, to view tapes. The Goal G10 motivates Attempt A11.

The aforementioned goal process analysis indicates that McKoon and Ratcliff (1992) erroneously labeled the problem version as globally inconsistent. The theme, as pointed out by Singer, Graesser, and Trabasso (1994), is one of an upcoming match, E/G4. The subordinate goals of a workout in G3, an opponent in G5, and viewing tapes in G9 are consistent with a superordinate goal, G4, to prepare for the match.

Connectionist modeling. The four networks of Figure 6 were each submitted to the Langston and Trabasso (1999) connectionist model for processing. Accessibility of the respective target words, "weight" and "workout," was measured by the average causal connection strength of Goal G3 in each story network after the entire story was processed. Goal G3 contained the target words in the story.

The accessibility predictions using average connection strength measures are shown for the observed response times of Experiment 2 in Figure 7. The average

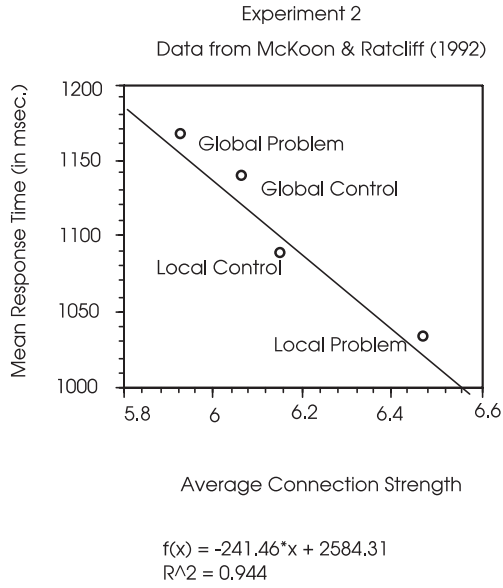

FIGURE 7 Accessibility of words: Prediction of Observed Response Times for Conditions of Experiment 2 × Average Connection Strength.

connection strengths of the clauses that contained "weight" and "workout" predicted 94.4% of the variance in response times, $F(1, 2) = 33.40$, $p = .028$.

The simulation of McKoon and Ratcliff's (1992) Experiment 2 data was highly successful and replicates the success of the simulation of the Experiment 1 data. A theory that assumes that local and distal causal inferences are made during reading provides an excellent, quantitative account of the accessibility of the target words across the variations in the stories. The discourse analysis and causal networks show that what McKoon and Ratcliff (1992) labeled as *inconsistent* can be regarded as coherent when the goal structure of a narrative is considered. What they labeled as *globally inconsistent* is not at all inconsistent with an inferred, superordinate goal of preparing for a forthcoming tennis match.

EXPERIMENT 3

In Experiment 3, McKoon and Ratcliff (1992) had readers read a pair stories in a self-paced manner, one paragraph at a time. After the reader read a pair of stories, the reader judged whether each sentence in a test series of 24 sentences occurred in the story. A target sentence was embedded in the series and was primed by a preceding goal, an action near the goal, or an action near the target sentence, all of which had actually been presented in the story. The recognition time of the target

sentence for each of these three primes served as the dependent variable. The logic was that if a reader makes global goal-related inferences, then processing the target information should be fastest in the preceding goal condition.

McKoon and Ratcliff (1992) performed a form of discourse analysis that led to their conditions and expectations. Their analysis was a description of the goal and outcome structure for one of their 12 stories for Experiment 3. McKoon and Ratcliff's story structures appear to have a hierarchical goal structure. A general hierarchical goal structure in which their goal–outcome structure is embedded is shown in Figure 8. Figure 8 is limited to a four-goal hierarchy for purposes of analysis and discussion. The actual number of goals in the McKoon–Ratcliff stories was found in the discourse analysis reported later to be greater than four in each of their 12 stories.

The Farmer and his Donkey story, cited in the beginning of the article, conforms to the hierarchy depicted in Figure 8. The goal–attempt–outcome structures that

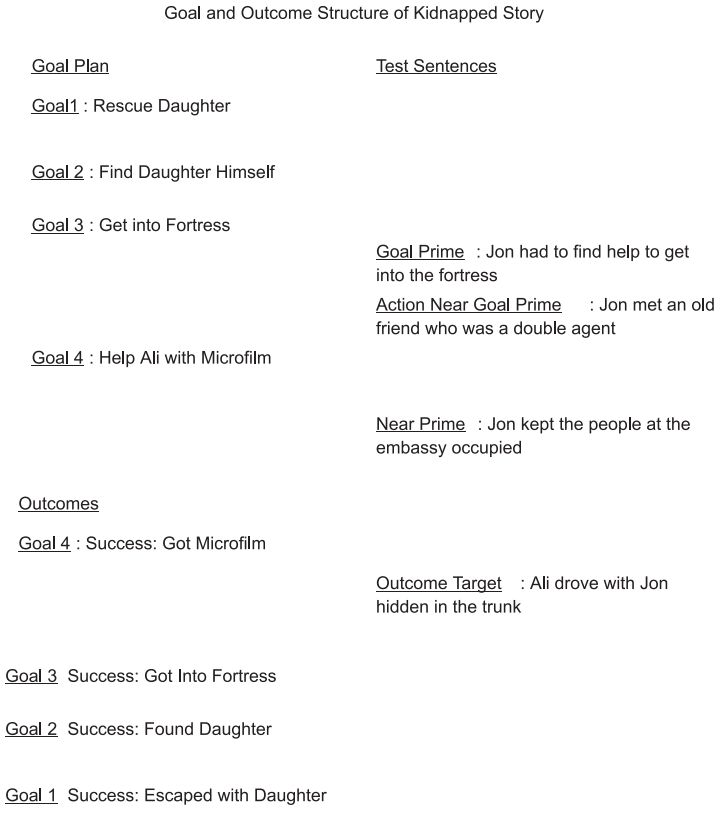

FIGURE 8 General hierarchical goal structure.

are repeated at each level represent a goal followed by attempts that lead to an outcome failure or obstacle for the first three goals. For the fourth goal, a goal success occurs that enables the attainment of the third goal. Attainment of the third goal enables attainment of the second goal and attainment of the second goal enables the attainment of the first goal.

For example, the Farmer wants to scare the donkey into the barn so he asks the dog to bark at the donkey but the dog refuses. The failed goal and outcome jointly cause a new, subordinate goal plan to be generated. The farmer then attempts a new goal, namely to have the cat scratch the dog to get him to bark at the donkey. New subordinate goals keep failing until finally one succeeds. The cow giving the farmer milk enables a successful attempt and outcome of getting the cat to scratch the dog. This success enables the success in attainment of goals successively higher in the hierarchy.

The important goal processes are that each new subordinate goal is motivated by its immediate superordinate goal and is psychologically caused by that goal's outcome failure, and that the successful attainment of a subordinate goal enables the carrying out of attempts to attain its immediate superordinate goal. Understanding of subordinate goal generation in the face of failure and superordinate goal attainment in the face of subordinate goal success enables the reader to achieve coherence.

In the Kidnapped story, Goal 1 in Figure 8 represents the top goal in the goal hierarchy. This goal motivates attempts that lead to a goal obstacle. For example, the desire to rescue his daughter leads Jon to find out that "The fortress appeared to be completely impenetrable," an obstacle to the success of the rescue goal . In Figure 8, the attempt and outcome that are related to a goal are shown on the left side of the network. In the actual McKoon–Ratcliff stories, this structure summarizes a number of goals, actions, and outcomes that occur.

The blocked goal and outcome jointly cause new subordinate goals to be generated. In the Kidnapped version, Jon finds out that his daughter is being held as a prisoner in a fortress. This obstacle and his desire to rescue her lead him to want to get into the fortress.

Figure 9 displays the structure of the Kidnapped Story reported by McKoon and Ratcliff (1992). In addition, Figure 9 shows the prime and target sentences, also reported by McKoon and Ratcliff, that occurred as recognition targets after the reading of the Kidnapped story. Their placement in the hierarchical structure is proximal.

In Figure 9, the top goal, labeled Goal 1, of the main character, Jon, is to rescue his daughter (who has been kidnapped). The superordinate goal was to have her back safely, so this goal is actually a subordinate goal. The second goal in their reported structure, Goal 2, is to find his daughter himself. The third goal, Goal 3, is to get into the fortress (where she is held captive). The fourth goal, Goal 4, is to help Ali (a confederate) with the microfilm to gain his assistance to get into the fortress. The fourth goal's success outcome, Outcome 4, is obtaining the microfilm for Ali.

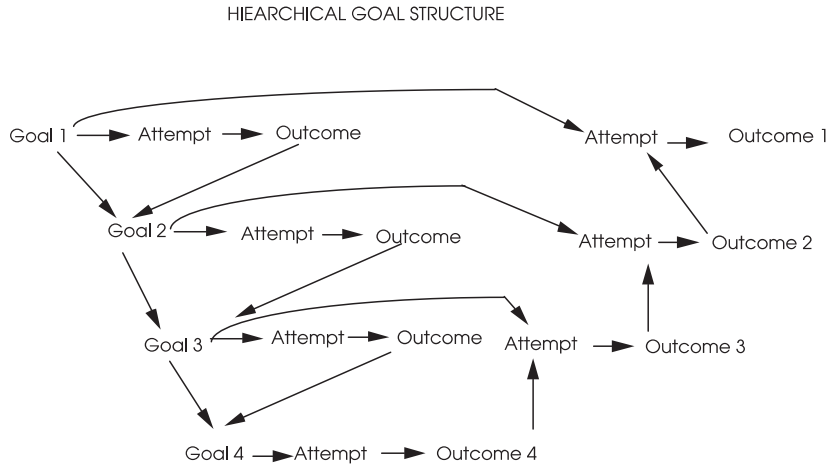

FIGURE 9 General hierarchical goal structure of Kidnapped Story (Experiment 3, McKoon & Ratcliff, 1992).

This success enables another goal outcome success, namely that Ali helps him. The third goal's success outcome, Outcome 3, is getting into the fortress. The second goal's success outcome, Outcome 2, is finding his daughter. The first goal's success outcome, Outcome 1, is escaping with his daughter. The superordinate goal success of having his daughter back home occurs in the last sentence of the story and is not reported as part of the structure by McKoon and Ratcliff (1992).

In the Kidnapping structure and recognition tests depicted Figure 9, McKoon and Ratcliff (1992) did not test the accessibility of the superordinate or either of the two highest goals in the hierarchy. Rather, they tested accessibility of a target subordinate goal attempt by priming of other subordinate goals and attempts at the third and fourth levels in their four-goal hierarchy. Further, the goal primes are especially distal from the target in the surface text and do not seem to be directly related to the target sentence. In Figure 9, the target recognition sentence, "Ali drove with Jon hidden in the truck," is an attempt to get into the fortress (and not an outcome as claimed). This attempt is related to Goal 3 (to get into the fortress), but it is not directly related to the goal prime. The goal prime was "Jon had to find help to get into the fortress." This goal is not reported as part of the structure and is a subordinate goal, motivated by Goal 3. The action near the goal prime was an initiating event, "Jon met an old friend who was a double agent." This sentence is not directly related to Goal 3, but is related to the goal prime. The near prime was an attempt motivated by a subordinate goal of Goal 4, to get the microfilm: "Jon kept the people at the embassy occupied." This attempt is only indirectly related to the goal prime, as it enabled success of Goal 4, which in turn allowed Jon to get help from Ali.

The research question is how accessible are target sentences following each of the prime sentences? However, instead of only varying in their relation to a preceding goal, these primes vary in both their content and in their surface and causal distance from the target sentence. The recognition priming data, obtained after reading, and the actual primes and targets used do not constitute a test of global goal inferences being made during reading. Each of the primes has some more or less indirect connection to the target that can be revealed only by a sentence-by-sentence causal discourse analysis.

The simulation of Experiment 3 provides such a discourse analysis. The causal network permits a quantitative assessment of causal distance in the causal network representation. Causal distance indicates how far away from one another two related sentences are. When processed by the connectionist model, the simulation permits quantitative predictions of accessibility. In priming (Langston & Trabasso, 1999), the most valid measure is the connection strength between the prime and target sentence. This measure assumes that the processing of the prime spreads activation and that the distance between the prime and the target affects the speed of activation. Langston and Trabasso reported several simulations where the connection strength is related to causal distance in the network and is a valid predictor of accessibility of one unit of information by activation of another. Hence, the causal distance between a prime and target sentence should also be a good predictor of priming or accessibility of the target from the priming sentence.

Discourse Analysis of Experiment 3 Stories

Each of the 12 stories of Experiment 3 was first separated into sentences. In this analysis, sentences rather than clauses served as the unit because sentences served both as the target and primes in the postreading verification task. The sentences for each story were each classified according to the episodic categories of Stein and Glenn's (1979) story grammar (setting, event, goal or internal response, attempt, and outcome). The classification was done by Tom Trabasso and Mark Langston and was highly reliable. The average Kappa was .89 (range .78 to .96 where .75 is taken to be an index of high reliability). Differences between the authors were resolved by discussion.

Causal networks. The causal relations were independently identified using the constraints of the episodic categories, how and why questions, and counterfactual tests. Trabasso and Langston independently identified causal relations among the clauses for each version. The average Kappa was .925 (range .82 to .95), indicating high reliability on identification of causes. Differences between the authors were again resolved by discussion.

Because there were 12 stories, we focus on the causal analysis of a representative story, called Monster. Table 4 displays the parsed and categorized clauses for the

TABLE 4
Monster Story Clauses With Episodic Categories and Surface Number of Sentences

Category and Surface Number	*Sentence*
Setting 1	Jimbo was an addict who lived in New York City.
Setting 2	He used drugs a lot, and over the years he had developed a large heroin habit to support.
Setting 3	He was also deeply in debt to several loan sharks.
Global	
Goal 4	Jimbo needed to find a lot of money so that he could pay off his debts (and keep his health), and buy more heroin to support his habit.
Attempt 5	After a breakfast of day old pizza, Jimbo went out into the streets to look for opportunities to make some big bucks.
Attempt 6	Jimbo tried his hand at pick-pocketing and followed an elderly gentleman into one of the new buildings downtown.
Attempt 7	When Jimbo got into the elevator to follow the man, he pressed the wrong button and wound up at the top of the unfinished building.
Outcome 8	While wandering around up there he stumbled on the hideout of a large flying creature.
Internal Response 9	Jimbo was sure it was the creature many believed responsible for several recent murders in the city.
Internal Response 10	The police thought the eye witnesses were crazy, but it was getting hard to ignore so many people.
Internal Response 11	They reluctantly agreed that something strange was going on in their city.
Internal Response 12	Everyone was scared.
Goal 13	Jimbo thought that there must be a way to make some money on the information he had just discovered.
Attempt 14	He quickly left the building and went to the coffee shop around the corner.
Attempt 15	He sat down, shaking with fear and excitement at the same time.
Outcome16	As he poured lots of sugar into his coffee, Jimbo came up with a plan.
Primed	
Goal 17	He planned to ransom the location of the monster's hideout for a large sum of money.
Primed action near goal	
Attempt 18	Jimbo called city hall and claimed that he had information that could save the city.
Outcome 19	The mayor, who was being criticized for not doing enough to find the monster, decided to take a chance and talk to Jimbo.
Outcome 20	When Jimbo saw the mayor the next day, he told him he had found the monster's hideout and would disclose the location for 1 million dollars.
Goal 21	Of course, the mayor didn't want to take any chances with a wild-eyed drug addict.

(*continued*)

TABLE 4 (*Continued*)

Category and Surface Number	*Sentence*
Outcome/Goal22	The mayor agreed to sign a contract giving Jimbo 1 million dollars after he killed the monster.
Internal	
Response 23	Jimbo knew he couldn't do the whole job himself.
Setting 24	He didn't even have any weapons.
Goal 25	He would have to get help so Jimbo decided to go to the police.
Outcome 26	The police chief was only too glad to help.
Setting 27	One of his relatives had mysteriously vanished and he was beginning to believe that the monster everyone was talking about was responsible.
Goal 28	The police chief insisted Jimbo himself lead the way to the monster's hideout.
Internal Response 29	Jimbo was not too thrilled about having to go back to the monster's hideout, but he figured once they got there he could stay well behind the officers.
Primed action near target	
Attempt 30	He climbed into one of the police cars and led the police to the building where the monster had its nest.
Attempt 31	Jimbo pointed the way to the top floor, where they waited for the monster to return.
Attempt 32	Jimbo hid behind as many of the officers as possible.
Outcome 33	After several hours, the monster flew into its lair.
Outcome 34	As soon as the monster was in its nest, the officers opened fire with automatic rifles, stopping when they were sure the creature was dead.
Target	
Attempt 35	Jimbo went back to city hall and reminded the mayor of the matter of 1 million dollars, his payment for the information on the creature.
Outcome 36	The mayor signed a check for a million and presented it to Jimbo, who ran straight to the bank and cashed his check.
Global goal	
Outcome 37	Now he had enough money to pay off his debts and stay high indefinitely.

Note. The primes and target are bolded and the global goal and outcome are bolded and underlined.

Monster story. The corresponding causal network for the Monster story in Table 4 is displayed in Figure 10. The Monster story was parsed into 37 sentences so that a detailed presentation of the analysis is unwieldy. Instead, the focus is on clauses and connections that relate the goal and action sentence primes and the target sentence. The 12 parsed stories and their causal networks are available from Tom Trabasso.

In Table 4, the superordinate goal is Goal 4 (G4 in the network in Figure 10). This sentence is actually a complex set of higher order goals—to pay off debts, to

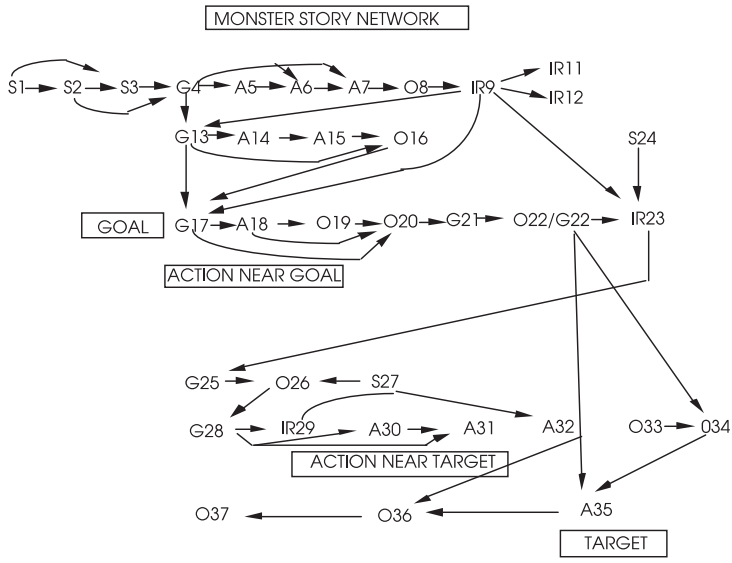

FIGURE 10 Causal network for the Monster story.

maintain health, and to buy heroin to support his habit. These higher order goals motivate the goal that dominates the story, namely, Jimbo wants to obtain a large amount of money. He obtains this goal, which is related to the higher order goals, explicitly in the last sentence of the story in Outcome 37 (O37 in Figure 10). "Jimbo finally has *the money to pay off his debts and stay high on drugs*" (italics added). Because the goal success and higher order goals are explicitly stated, there would be no need for the reader to access the superordinate goals at this point. No causal inference was therefore assumed between the global goal, G4, and Outcome O37 in the Kidnapped story. Other conditions such as O38, the mayor's signing of the check and giving it to Jimbo, and Jimbo's cashing the check, are jointly necessary in the circumstances for his obtaining the money in O37.

The goal prime for the Monster story, marked in Table 4, is a subordinate goal, G17, a plan to exchange information on the location of the monster's hideout for money. This goal is at level 3 as was the case in the Kidnapped story. The target was an attempt, A35, where Jimbo goes to city hall to ask for his reward. The target is not directly connected to G17, the goal prime, but is connected to an outcome and goal, O/G22, a promise by the mayor to pay the reward to Jimbo. If the mayor had not agreed to sign a contract giving Jimbo 1 million dollars, Jimbo would not have gone to city hall to ask for his reward after he had killed the monster. The subordinate goal prime, G17, motivates a sequence of goals, attempts, and outcomes in an episode before Jimbo goes to the mayor and receives the promised reward. The action near the goal was an attempt, A18, directly connected to G17. The at-

tempt is to carry out a goal plan by calling city hall and telling them that he has the information. The fact that G17 is directly connected to A18 indicates that their accessibility as measured by connection strength will be close in value. The action near the target is Attempt A31, where Jimbo climbs into a police car and leads the police to the monster, and is indirectly connected to the target.

An analysis of the 12 stories as to the locations of the primes and targets shows some interesting surface characteristics. In Table 5, the average surface sentence number of the global goal, the goal prime, the action near the goal prime, the action near the target, the target, and the global goal outcome are listed. Global goal information was always explicitly stated as part of an outcome in the last sentence of the story. The global goal occurred, on average, in the fourth sentence. The subordinate goal prime and its action were near each other and occurred on average at the 30th and 33rd sentences, respectively. The target was mainly an attempt (11 out of 12 cases) rather than a goal outcome, and occurred at the 33rd sentence, on average. Thus there was considerable surface distance between the goal and action near the goal primes and the target, approximately 17 and 15 sentences, respectively. The action near the target prime had a surface distance of only 3 to 4 sentences. If the reader relied exclusively on local coherence, this difference could contribute to the higher accessibility of the action near the target prime. This is confirmed by the

TABLE 5
Surface Sentence Number of Sentences that Contain Global Goal and Outcome, Primed Goal and Actions, and Target for McKoon and Ratcliff (1992) Stories

Story	*Global Goal*	*Goal Prime*	*Action Near Goal Prime*	*Action Near Target Prime*	*Target*	*Global Goal Outcome*
1. Monster	4	17	18	30	35	37
2. Kidnapped	4	18	19	20	24	30
3. Writer	4	16	20	31	34	40
4. Gambler	5	17	18	30	33	36
5. Witness	5	15	16	30	32	34
6. Baby	3	18	23	29	31	44
7. Scientist	4	12	18	26	30	33
8. Matchmaker	3	16	18	30	36	42
9. Drugs	4	14	17	32	36	40
10. Transferred	9	14	18	29	34	48
11. Western	3	16	17	22	27	36
12. Hired gun	5	27	29	43	46	50
Mdn	4.00	16.00	18.00	30.00	33.50	38.50
M	4.42	16.67	19.25	29.33	33.17	39.17
Causal distance to target		2.99	3.18	2.92		

finding that the amount of variance in the response times accounted for by the linear order of primes is 82.5%.

In the causal network, however, the accessibility of the target for the different prime sentences tells a different story. Accessibility is related to the number of nodes that intervene between the target and the prime. The number of intervening nodes is called "causal distance." Causal distance is highly correlated with the connection strength between nodes (Langston & Trabasso, 1999). This measure has been used to successfully predict accessibility of one node from another in priming studies such as the data of Suh and Trabasso (1993). The number of intervening nodes between each prime and the target was found for each story. The overall average number of sentences intervening between the primes was about 3 sentences. In terms of order, the highest number was for the action near the goal prime at 3.18 sentences, the next highest was for the goal prime at 2.99 sentences, and the lowest was for the action near the target at 2.92 sentences. These data, as indexes of accessibility, indicate that the primes were, on average, closer to the target causally than in the surface structure. The mean causal distances of the primes accounted for 90% of the variance in the amount of priming found by McKoon and Ratcliff (1992) for Experiment 3. The connectionist model's connection strength measures should account for about the same degree of variance and presents an alternative explanation for priming to that of linear surface order and local processing without global inferences.

Connectionist modeling. Each of the 12 causal networks was submitted independently to the Langston and Trabasso (1999) model. Accessibility of the respective target sentences was measured by the connection strength of the sentence prime to the target sentence. Table 6 shows the respective connection strengths obtained between the goal, action near goal, and action near target primes to the target for each story.

The measures of connection strength between nodes for each of the prime–target pairs are compared with the observed recognition times of Experiment 3 in Figure 11. The average connection strength between the primes and the target accounted for 99% of the variance in mean recognition times, $F(1, 2) = 200.33$, $p = .06$. Thus, the accessibility as measured by recognition priming was nearly perfectly accounted for by the connection strength between the prime and the target. In a step-wise regression, the connection strength measure accounted for all of the variance. The linear location of the sentences in the story was not entered as a variable, indicating that it did not account for any unique variance in the response times.

The fact that connection strength between the primes and the target predicted the observed recognition times very accurately supports the conclusion that readers made the causal inferences revealed by the discourse analysis and did not rely exclusively on local coherence to represent the story.

TABLE 6
Connection Strength Between Primes and Target Sentences

	Goal-Target CS	*Action Near Goal-Target CS*	*Action Near Target-Target CS*
Monster	4.00	5.00	4.00
Kidnapped	3.00	2.00	4.00
Writer	5.01	4.01	4.01
Gambler	3.00	2.00	4.00
Witness	1.00	0.00	5.00
Baby	0.00	4.03	5.02
Scientist	2.00	3.00	5.01
Matchmaker	5.01	4.00	4.00
Drugs	5.01	3.00	3.00
Transferred	3.00	3.00	4.00
Western	4.01	3.00	5.01
Hired gun	5.02	4.01	4.01
M CS	3.34	3.09	4.26

Note. CS = connection strength.

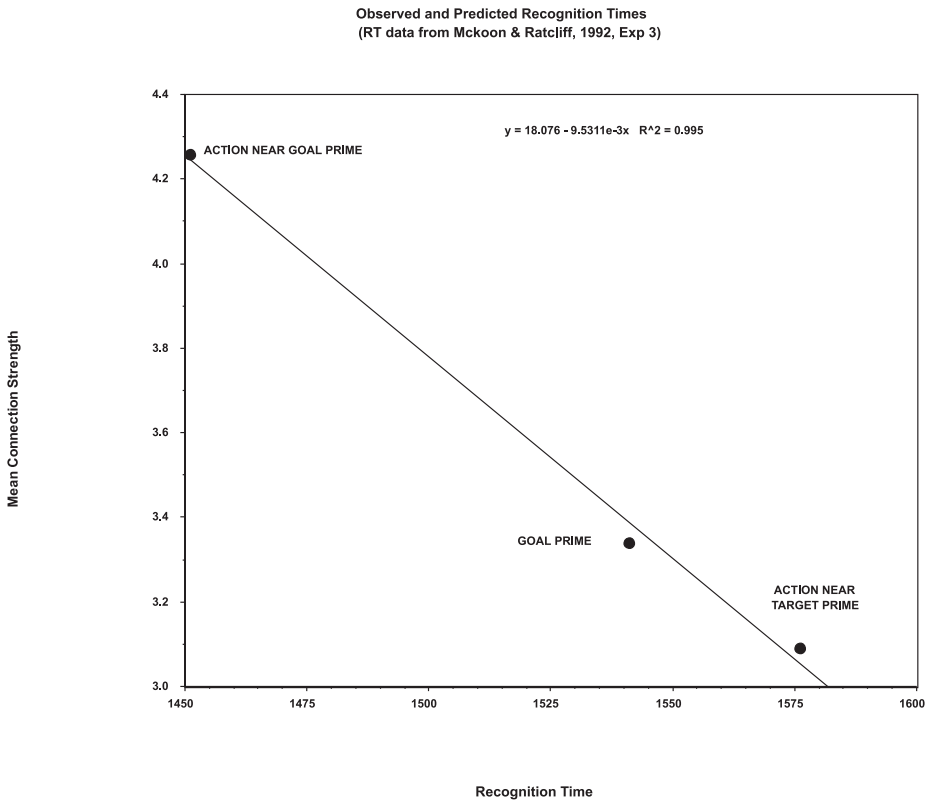

FIGURE 11 Accessibility of sentences: Prediction of Observed Response Times for Conditions of Experiment 3 × Prime-Target Connection Strength.

Van den Broek and Lorch (1993, Experiment 3) performed a similar experiment to McKoon and Ratcliff's Experiment 3 (1992). The van den Broek and Lorch study yielded an opposite result to that of McKoon and Ratcliff. Van den Broek and Lorch (1993) found significant priming of the goals but negligible priming by actions near the goal prime. The difference in findings between the experiments of van den Broek et al. (1993) and McKoon and Ratcliff (1993) may lie in the difference in their respective story structures. In the stories of van den Broek et al. (1993), the control action near the goal was not connected to the primed goal or target, whereas in the McKoon and Ratcliff (1992) stories, the action near the goal were often connected directly or had a short causal distance from the goal. If so, the accessibility of the goal and the action would be similar in value. If the action was not connected or if the causal distance was high, then it would be less accessible than the goal. To test these implications, we simulated the van den Broek and Lorch (1993) study by using their causal network (p. 78) and their data (Table 4, p. 92).

Figure 12 shows the results of the relation between the amount of priming (in milliseconds) and the prime–target connection strengths. The ordering of the priming data corresponds to the ordering of the connection strength between the prime and target sentences. The amount of variance in the priming data accounted for by the model was 81.4%.

The apparent contradictory findings of van den Broek and Lorch (1993) and McKoon and Ratcliff (1992) can be reconciled by an analysis of the causal relations between the primes and the targets in the respective studies. In van den Broek

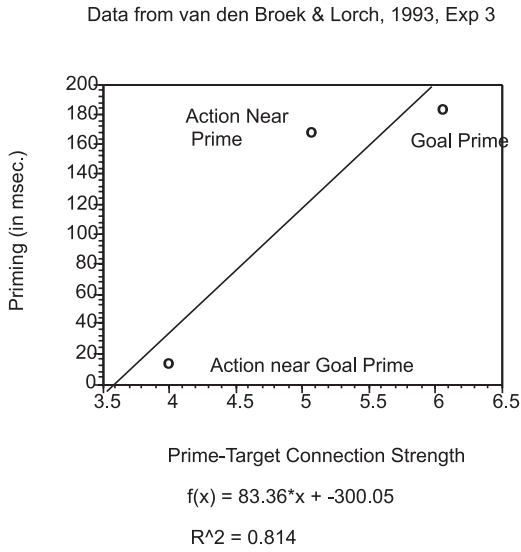

FIGURE 12 Simulation of van den Broek and Lorch (1993) Experiment 3 priming data.

and Lorch, the goals were more directly connected to the target than the action near the target prime, thereby yielding a higher degree of accessibility for the goals. In contrast, in the McKoon–Ratcliff study, the prime near the target is more directly connected to the target. A simulation using the same theory accounts for the differences in accessibility within and between studies. When the structural differences of the stories are taken into account by a causal discourse analysis, the data are not contradictory. The differences lie in the story structures.

GENERAL DISCUSSION

The successful simulations of accessibility of text information, as measured by recognition times, provide an alternative account for the data of four experiments, three by McKoon and Ratcliff (1992) and one by van den Broek and Lorch (1993). The success of the simulations supports the assumption that readers use knowledge of goal plans of actions and outcomes to make causal inferences that explain why and how events, actions, and states occur in a narrative.

The present simulations provide a strongly supported alternative interpretation to the minimalist hypothesis offered by McKoon and Ratcliff (1992). The simulations made no distinction between local or global inferences. Causal inferences may be "local" (i.e., between adjacent sentences or clauses) or "distal" (i.e., separated by one or more clauses or sentences). The simulations also do not distinguish between "local" or "global coherence." Rather, coherence is a result of processing and occurs when the reader can find meaningful relations between what is known and what is being read. The simulations assume that processing occurs in an activated portion of long-term memory with the current contents of working memory. The distinctions in the minimalist position on local versus global inferences and coherence come from assumptions about the limitations of working memory and no access to long-term memory unless there is a local coherence break. These distinctions in processing are not necessary if access to long-term memory is allowed as in the Ericsson and Kintsch (1995) long-term working memory model. The activation of long-term memory during reading and understanding text enables the direct access of what is known and what is represented in long-term memory with the current contents of working memory. This activation of long-term memory overcomes the restrictions of assuming a separate, short-term, working memory store and requiring that access to long-term memory be deliberate or strategic and dependent on failures to achieve local coherence.

Knowledge of goal processes played an important role in the present discourse analysis and connectionist simulations. We relied heavily on knowledge on goals and plans by Stein and Albro (1997) and Stein et al. (1993). The knowledge that readers have of goals, plans, actions, and outcomes enables coherent understanding of what characters do when things go well or when things badly. New actions

and goals following outcomes are to be understood in the context of what prior goals exist and what happened to advance or fail them. These goal–outcome contingencies enable understanding that prior goals are being maintained or abandoned when new goals and actions occur.

Connection strength was assumed to index the degree of activation of information in a clause or sentence by another clause or sentence or by several clauses and sentences. This assumption was supported empirically by the success of the four simulations. This access and connection strength interpretation depends on an implicit memory retrieval process. The memory retrieval process is one of spreading activation from a word, clause, or sentence to related information in long-term memory. The connection strength is a measure of how much activation there is between two or more nodes in the long-term memory representation. The greater the connection strength, the greater the amount of activation from one clause or sentence to another or from one clause or sentence to several clauses or sentences. The degree of activation between two clauses or sentences or between a clause or sentence and all other clauses and sentences that have been processed changes over the course of the reading of a text. The continuous updating of connection strengths during online understanding of the story and the access by spreading activation are similar to the mechanisms of the resonance theory of O'Brien and Myers (1999).

The important theoretical assumption is that activation and connection strengths are highly related. The access of text information is affected by where the reader is in processing the text and by the number of connections that clauses or sentences have to other clauses or sentences in the network at that point in time. These connections and their strengths change dynamically during online processing but are fixed after processing is completed. Accessibility depends on what is activating what. In the cases where access is made from several clauses or words, as in Experiments 1 and 2, the number of connections or average connection strength is an excellent predictor of the speed of access. The greater the number of connections, the greater the amount of activation of the clause of sentence. In the case where access is from one clause or sentence to another, as in McKoon and Ratcliff's (1992) Experiment 3 and Lorch and van den Broek's (1993) Experiment 3, the greater the connection strength between the clauses or sentences then the greater the amount of activation from one clause or sentence to another clause or sentence. The choice of which measures are used depends on a task analysis and assumptions about what is contributing to the activation of a clause or sentence of interest.

Accessibility of information in a long-term, text memory representation is predicable from knowledge of causal networks in a variety of contexts. Using the connection strength between nodes to predict the accessibility of one node from another, Langston and Trabasso (1999) simulated the amount of time taken to read a second sentence that varied in its causal distance from the first sentence, as well

as the ratings of causal relatedness for the pairs of sentences (data from Myers, Shinjo, & Duffy, 1987), the number of goal references made during reading in thinking-out-loud protocols (data from Lutz & Radvansky, 1997; Suh & Trabasso, 1993), and the priming of online goal inferences (data from Lutz & Radvansky, 1997; Suh & Trabasso, 1993).

Langston and Trabasso (1999) also reported simulations where access may be made from multiple sources in the text. In these cases, they used the average connection strength of a node to index its activation and accessibility from all other nodes in the mental representation. Using average connection strength, Langston and Trabasso (1999) successfully simulated importance judgments (data from Trabasso & Sperry, 1985), immediate and delayed recall of sentences (data from Trabasso et al., 1995; Trabasso et al., 1994), and the recall and coherence of entire stories (data from Trabasso, Suh, & Payton, 1994; Trabasso, Suh, Payton, & Jain, 1995). In addition, Trabasso and Bartolone (2003) reported successful simulations of decision making and counterfactual reasoning (data of Kahneman & Tversky, 1982), whereas Trabasso and Wiley (2002) successfully simulated hindsight bias (Fischhoff, 1975; Hawkins & Hastie, 1990) using before and after prediction data from Wasserman, Lempert, and Hastie (1991).

The important advance of the simulations is their ability to mimic comprehension of narratives as a dynamic understanding, integration, and learning system as the reader reads the text. Inferences are central to connecting and integrating ideas online and to updating the long-term memory representation of the text. As information is understood and integrated, the accessibility of what is in the long-term memory representation changes. Information that is connected receives more activation and gains in connection strength or accessibility. Information that is not connected loses activation and decreases in connection strength or accessibility.

The processes of dynamic understanding need not be conscious nor strategic; nor need they be limited by working memory. The Langston and Trabasso (1999) model is implemented in the software of the Construction–Integration model of Kintsch and others (Goldman & Varma, 1995; Kintsch, 1988, 1992; Kintsch, Welsch, Schmalhofer, & Zimny, 1990; Tapiero & Denhiere, 1995) but allows access to long-term memory during the processing of new information. Kintsch and Welch (1991); Ericsson and Kintsch (1995); and Kintsch, Patel, and Ericsson (1999) also made arguments on the relaxation of working memory processing constraints. Information that is stored in long-term memory is kept accessible by means of retrieval cues in short-term memory. These cues could operate at the level of words, clauses, or sentences and send activation to stored text representations in a manner suggested by O'Brien and Myers (1999).

This approach modeled readers as having direct access to relevant information in long-term memory, particularly knowledge of a character's goals, motivations, and plans. The access to this kind of knowledge enables readers to form representations of the text in long-term memory. The extremely good fit of these simula-

tions, based on causal networks derived from considering both local and global levels of analysis, suggests that such inferences are computed, and the goal plan of characters are considered, during reading of narratives.

ACKNOWLEDGMENTS

This research was funded by Grant No. 38895 from the National Institute of Child Health and Human Development.

We thank Gail McKoon for providing copies of the 12 stories used in Experiment 3. The parsed versions and causal networks of these stories are available from Tom Trabasso. We thank Mark Langston for programming the connectionist model and for his assistance in deriving the causal networks of Experiment 3. A Windows version of the program on CD is available from Tom Trabasso. We also thank Dale Barr for creating a Web site that enabled the use of the connectionist model. Please contact Tom Trabasso for current Web address.

REFERENCES

Ericsson, K. A., & Kintsch, W. (1995). Long-term working memory. *Psychological Review, 102,* 211–245.

Fischhoff, B. (1975). Hindsight ≠ foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance, 1,* 288–299.

Goldman, S. R., & Varma, S. (1995). CAPping the construction-integration model of discourse comprehension. In C. A. Weaver, III, S. Mannes, & C. R. Fletcher (Eds.), *Discourse comprehension: Essays in honor of Walter Kintsch* (pp. 337–358). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Graesser, A. C. (1981). *Prose comprehension beyond the word.* New York: Springer-Verlag.

Graesser, A. C., Millis, K. K., & Zwaan, R. A. (1997). Discourse comprehension. *Annual Review of Psychology, 48,* 163–189.

Hawkins, S. A., & Hastie, R. (1990). Hindsight: Biased judgment of past events after the outcomes are known. *Psychological Bulletin, 107,* 311–327.

Kahneman, D., & Tversky, A. (1982). The simulation heuristic. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 201–208). New York: Cambridge University Press.

Kintsch, W. (1988). The role of knowledge in discourse comprehension: A constructive-integration model. *Psychological Review, 95,* 163–182.

Kintsch, W. (1992). How readers construct situation models for stories: The role of syntactic cues and causal inferences. In A. F. Healy, S. M. Kosslyn, & R. M. Shiffrin (Eds.), *Essays in honor of William K. Estes* (pp. 261–278). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Kintsch, W., Patel, V. L., & Ericsson, K. A. (1999). The role of long-term working memory in text comprehension. *Psychologica: An International Journal of Psychology in the Orient, 42,* 186–198.

Kintsch, W., & Welch, D. M. (1991). The construction-integration model: A framework for studying memory for text. In W. E. Hockley & S. Lewandowsky (Eds.), *Relating theory and data: Essays on human memory in honor of Bennet B. Murdock* (pp. 367–385). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

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

Langston, M., & Trabasso, T. (1999). Modeling causal integration and availability of information during comprehension of narrative texts. In H. van Oostendorp & S. Goldman (Eds.), *The construction of mental representations during reading* (pp. 29–69). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Langston, M. C., Trabasso, T., & Magliano, J. P. (1998). Modeling on-line comprehension. In A. Ram & K. Moorman (Eds.), *Computational models of reading and understanding* (pp. 181–225). Cambridge, MA: MIT Press.

Lutz, M. F., & Radvansky, G. A. (1997). The fate of completed goal information in narrative comprehension. *Journal of Memory & Language, 36,* 293–310.

Mackie, J. L. (1980). *The cement of the universe*. Oxford, England: Clarendon.

McKoon, G., & Ratcliff, R. (1992). Inference during reading. *Psychological Review, 99,* 440–466.

Myers, J. L., Shinjo, M., & Duffy, S. A. (1987). Degree of causal relatedness and memory. *Journal of Memory & Language, 26,* 453–465.

O'Brien, E. J., & Myers, J. L. (1999). Text comprehension: A view from the bottom up. In S. R. Goldman, P. van den Broek, & A. C. Graesser (Eds.), *Narrative comprehension, causality, and coherence: Essays in honor of Tom Trabasso* (pp. 35–53). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Rumelhart, D. E. (1977). Understanding and summarizing brief stories. In D. LaBerge & S. J. Samuels (Eds.), *Basic processes in reading: Perception and cognition* (pp. 265–303). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Singer, M., Graesser, A. C., & Trabasso, T. (1994). Minimal or global inference during reading. *Journal of Memory and Language, 33,* 421–441.

Stein, N. L., & Albro, E. R. (1997). Building complexity and coherence: Children's use of goal-structured knowledge in telling good stories. In M. Bamberg (Ed.), *Narrative development: Five approaches* (pp. 5–70). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

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* (Vol. 2, pp. 53–120). Norwood, NJ: Ablex.

Stein, N. L., & Levine, L. (1989) The causal organization of emotion knowledge: A developmental study. *Cognition and Emotion, 3,* 343–378.

Stein, N. L., Trabasso, T., & Liwag, M. (1993). The representation and organization of emotional experience: Unfolding the emotional episode. In M. L. Lewis & J. Haviland (Eds.), *Handbook of emotions* (pp. 279–300). New York: Guilford.

Suh, S., & Trabasso, T. (1993). Inferences during on-line processing: Converging evidence from discourse analysis, talk-aloud protocols, and recognition priming. *Journal of Memory and Language, 32,* 279–301.

Tapiero, I., & Denhiere, G. (1995). Simulating recall and recognition by using Kintsch's construction-integration model. In C. A. Weaver, III, S. Mannes, & C. R. Fletcher (Eds.), *Discourse comprehension: Essays in honor of Walter Kintsch* (pp. 211–232). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Trabasso, T., & Bartolone, J. (2003). Counterfactual thinking. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29,* 904–923.

Trabasso, T., & Magliano, J. P. (1996). Conscious understanding during comprehension. *Discourse Processes, 21,* 255–287.

Trabasso, T., & Nickels, M. (1992). The development of goal plans of action in the narration of picture stories. *Discourse Processes, 15,* 249–275.

Trabasso, T., & Sperry, L. L. (1985). Causal relatedness and importance of story events. *Journal of Memory & Language, 24,* 595–611.

Trabasso, T., Suh, S., & Payton, P. (1994). Explanatory coherence in communication about narrative understanding of events. In M. A. Gernsbacher & T. Givon (Eds.), *Text coherence as a mental entity* (pp. 189–214). Amsterdam: John Benjamins.

Trabasso, T., Suh, S., Payton, P., & Jain, R. (1995). Explanatory inferences and other strategies during comprehension and their effect on recall. In R. F. Lorch, Jr., & E. J. O'Brien (Eds.), *Sources of coherence in reading* (pp. 219–239). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

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

Trabasso, T., van den Broek, P. W., & Liu, L. (1988). A model for generating questions that assess and promote comprehension. *Question Exchange, 2,* 25–38.

Trabasso, T., van den Broek, P., & Suh, S. (1989). Logical necessity and transitivity of causal relations in stories. *Discourse Processes, 12,* 1–25.

Trabasso, T., & Wiley, J. (2002, January). *Dynamic understanding and updating of memory: Simulation of hindsight bias.* Paper presented at the Winter Text Conference, Jackson Hole, WY.

van den Broek, P. W. (1988). The effect of causal relations and goal failure position on the importance of story statements. *Journal of Memory and Language*, *27*, 1–22.

van den Broek, P. W. (1990). The causal inference maker: Towards a process model of inference generation in text comprehension. In D. A. Balota, G. B. Flores d'Arcais, & K. Rayner (Eds.), *Comprehension processes in reading* (pp. 423–446). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

van den Broek, P., & Lorch, Jr., R. (1993). Network representations of causal relations in memory for narrative texts: Evidence from primed recognition. *Discourse Processes, 16*, 75–98.

van den Broek, P., & Trabasso, T. (1986). Causal networks versus goal-hierarchies as predictors of importance of story statements. *Discourse Processes, 9,* 1–15.

Wasserman, D., Lempert, R. O., & Hastie, R. (1991). Hindsight and causality. *Personality and Social Psychology Bulletin, 17,* 30–35.

Zwaan, R. A., & Radvansky, G. A. (1998). Situation models in language comprehension and memory. *Psychological Bulletin, 123,*162–185.