DISCOURSE PROCESSING
A. Flammer and W. Kintsch (eds.)

PLANS AND GOALS IN UNDERSTANDING EPISODES

Gordon H. Bower

Department of Psychology
Stanford University
Stanford, California

People's judgments about important elements in narrative episodes were studied. Subjects chose the protagonist's goal as most important, then actions, outcomes, complications, and lastly background and modifiers. A statement's importance correlated with its likelihood of recall and inclusion in a summary. Relating episodes to a Plan schema, further experiments found that readers take longer to comprehend an action in light of a goal the greater the "distance" between them in a goal hierarchy. Furthermore, the time to comprehend a character's action increases the more independent goals the reader is monitoring for that character.

INTRODUCTION

This symposium is concerned with how people understand texts, recall them, paraphrase them, summarize them, and answer questions about them. The fact that this symposium is being sponsored by a Psychology Department attests to the progress being made in cognitive psychology. Psychologists' concern with text processing is relatively recent. Eight years ago a symposium like this could not have taken place because there simply was not enough research on the topic. But since 1974 there's been an increasing stream of research on text processing, attested to by conferences and specialized research journals.

Several years ago when I first began studying text comprehension with my student, Perry Thorndyke (see Bower, 1976; Thorndyke, 1977), we adopted the story-grammar approach then proposed by Dave Rumelhart (1975) and Tuen van Dijk (1972). I have come to realize over the ensuing years that my interest is not so much in stories as in how people understand episodes and action sequences. Stories have episodes, of course, and the telling of the episodes is arranged so as to arouse suspense, surprise, mystery, humor, or irony, thus to entertain and hold the reader's attention. But I have not been studying these affective, entertaining features of stories; rather, I've studied only how people understand and remember episodes and event sequences.

It turns out that the central part of story grammars is the way they analyze simple episodes. Nearly all the story grammars assume that an interesting episode must have at least four parts: a goal for the protagonist, some obstacle or complication to attainment of that goal, some actions designed to overcome those obstacles, and some outcome of these

actions. Another way to say this is that an episode consists of a problem and its resolution. The problem can be characterized as stemming from the protagonist's goal plus a complication or obstacle; the resolution is comprised of the protagonist's action plan plus its outcome. To describe these elements more fully, the complications typically arise either from physical obstacles, or from the conflict of several goals within the same individual, or the conflict of goals between two competing individuals or teams. The resolution of an episode describes either the winning, losing, or compromising of a goal, abandoning it, or regaining a lost state of bliss. In order for the episode to be interesting, the problem must be significant and the resolution must be novel or unexpected. Dull episodes deal either with small problems or ones which have routine, familiar solutions.

The story grammars assume that people have acquired an implicit schema or prototype about episodes. This schema has various uses. One function of the schema is as a source of questions for readers. In his theory of question-asking, August Flammer (1980) suggests that people ask questions about gaps or critical slots in the episode schema that are not filled in by, or inferable from, the text. It is further assumed that the schema helps readers identify the critical elements of a text. If episode schemas are used in analyzing texts and in parsing episodes into significant constituents, then subjects should be able to reliably identify these allegedly important elements from a mass of text. Certainly, if naive readers do not agree with the story grammars about what are the essential, important elements in an episode description, then we all have surely been following the wrong leads. After a brief survey of the relevant literature, however, I was unable to find much direct empirical study of which parts of episodes readers consider to be important and necessary.

IDENTIFYING EPISODE CONSTITUENTS

The question I asked is whether college readers will identify as important those elements of narrative episodes which story grammars claim to be critical. Furthermore, I wondered whether people would summarize the episode and recall it largely in terms of these same critical elements. As I noted, the elements are the problem (with constituents of goal and complication) and the resolution (with constituents of actions and outcome). In order to study reader's intuitions, we wrote two six-episode narratives and had people read them, judge them, and recall them.

The six distinct episodes were printed one per page in a booklet. Each episode was written to set forth a distinct goal, complication, action-plan, and outcome, these comprising four sentences. Among these we mixed four further statements which set forth descriptive information, giving background or elaborating on the properties of the other constituents. From the viewpoint of story grammars these descriptive elaborations were inessential fillers, although they tended to make the prose somewhat more readable and natural.

One of the stories was about a male university student, Paul, and his problems in paying for his schooling, getting good grades, holding down a part-time job, and having an active social life. Here, for example, is the first episode in the Paul story.

(Goal)	Paul wanted to go to college.
(Filler)	He decided on a university in California.
(Complication)	But he didn't have enough money for expenses.
(Filler)	He had only $535 in a savings bank.
(Action)	He applied for a football scholarship.
(Filler)	He had played halfback in high school.
(Outcome)	After reviewing his case, the coaches granted him an award enabling him to go to school.
(Filler)	Paul hoped he could play first string.

The second story was about a female university student, Gail, who had a conventional set of problems--making friends, losing wieght, getting more exercise, breaking up with a boyfriend.

We had two groups of 30 college students read these stories. Some subjects simply read the stories, at 45 seconds per episode, then 15 minutes later recalled both stories when cued with the characters' names. Other subjects rank-ordered the eight statements in each episode according to their importance or significance within the episode. After they'd so ranked all statements, they re-read them and wrote a summary of each episode in two or three sentences, using less than 15 content words. They were instructed to imagine composing a telegram to relay the essential gist of the episode in as few informative phrases as possible while remaining faithful to the literal events. (This instruction prevented people from composing abstract morals as summaries.)

These two groups of subjects thus assessed each statement for its importance ranking within the episode, its likelihood of inclusion in a summary for that episode, and its likelihood of recall within the entire story. The text grammar hypothesis claims that the statements within each episode can be divided into two sets, those that are irrelevant or not essential versus those that are essential parts of any episode that has a point-namely, the goal, complication, planned action, and outcome. The hypothesis does not predict whether elements within the "essential" set will vary in importance.

The main results of this study are shown in Table 1 giving the average importance ranking, probability of being included in a summary, and probability of recall for each type of statement, averaged over the six episodes within each story. The four background fillers were combined in these statistics. Table 1 contains several interesting findings. First, the descriptive elaborations were indeed judged as irrelevant and unimportant, were least likely to be recalled, and least likely to be included in summaries of the episodes. Thus, subjects' intuitions about what are essential elements in an episode agree with our theory of the episode schema.

A second consistent finding is that subjects usually rate the goal statement as the most important statement in the episode. This average ordering arose for eleven of the twelve episodes (two stories each with six episodes). This is not simply a "first sentence" effect: half the episodes had some background fillers before the goal, yet even in those cases

subjects still rated the goal highest in importance.

Table 1

Importance ranking, probability of inclusion in a summary,and probability of recall by statement type, averaged over the six episodes in each story. Importance scaled from 1 (most important) to 8 (least)

		STATEMENT TYPE				
		Goal	Complication	Action	Outcome	Fillers
PAUL STORY	Importance Rank	1.75	2.82	3.83	3.62	5.57
	Summary Inclusion	.57	.51	.50	.83	.10
	Free Recall	.65	.66	.79	.76	.53
GAIL STORY	Importance Rank	1.71	4.33	3.85	3.81	5.57
	Summary Inclusion	.74	.18	.67	.63	.16
	Free Recall	.82	.63	.79	.59	.54

A third finding is that the importance ordering of the complication, action, and outcome varied across the two stories. Thus, the data do not support theories which assign importance to these elements simply on the basis of their role in the episode structure.

In reviewing our two stories, I noticed another factor that seemed critica in determining the importance rating subjects gave to the non-goal element of the episode. This other factor was how informative, nonredundant, or unusual a given statement was in the context of the character's goal. Som Complications or Actions were very routine and expected; stating them conveyed little new information beyond what one could already infer from the context. Consider a few of our Complications: an example of an infor mative Complication is that Paul's playing football frustrates his goal o doing well in his classes; an example of a routine,redundant Complication is that Gail lacked motivation to get more exercise; another is that Gail didn't know what to do to become less shy, so she asked a friend who suggested an assertiveness class. Among American college students Gail's "Complications" are so standard and routine that they are hardly worth mentioning.

I thought that this redundancy factor would influence the importance people assigned to the essential episode elements. So I had some new sub jects rate the elements within each episode on a scale of informativeness or unpredictability in context. This enabled us to separate the episode elements somewhat more. Thus, Complications judged to be "informative" h earlier received mean importance rankings of 2.93 (recall, 1 is the most important) whereas Complications judged as more predictable and redundant had received average importance rankings of 5.08, which is significantly lower. A similar difference in importance rankings was found for Actions rated as informative (3.60) versus those rated as redundant (4.17) with t Goal or Complication. Outcomes did not differ often enough in redundancy ratings for us to compare the importance assigned to high vs. low redunda outcomes.

The conclusion from this post-hoc analysis is that the importance assigne to a Complication or Action will usually be higher the more unexpected an informative it is in light of the goal and the other elements.

Consider now the likelihood that different episode elements are included in
the telegraphic summaries (see lines 2 and 4 of Table 1). Irrelevant
fillers hardly appear at all in summaries; Actions, Complications, and
Outcomes are likely to appear, but their exact ordering varies. For
example, Complications appear in summaries of the Paul story but hardly
ever in summaries of the Gail story. This difference probably reflects
again the predictable versus unpredictable nature of the complications
in the two stories. Across the two stories, Complications rated as highly
informative were included in episode summaries 74 percent of the time,
whereas Complications judged to be redundant and predictable were included
only 15 percent of the time. Thus, deletion of predictable Complications
in summaries seems to follow Grice's Conversational Postulate--that is, one
should be brief and not say what your audience can readily infer.

Table 1 also shows that free recall percentages were related to the impor-
tance ranking of the elements of the episode. We computed the Spearman
correlation coefficient among the three variables--recall, summary, and
importance--across the five categories and two stories (so N=10). The
results show moderately strong correlations: importance correlates .72
with likelihood of inclusion in a summary and .62 with free recall; and the
likelihood of inclusion in a summary correlates .70 with likelihood that
the statement will be recalled. While much common variance is being cap-
tured by these measures, the variance unaccounted for still always exceeds
50 percent. Some of this is due to uncontrolled differences in content,
in redundancy of the structural elements of the several episodes, and so
on. However, perhaps we should be satisfied with the conclusion that to a
first approximation, readers may be viewed as identifying and assigning
greatest importance to statements stipulating the goal, the complications
encountered, the actions undertaken, and the outcome, whereas they devalue
and skip over background statements, descriptive elaborations, and details.
Readers then use these structurally critical elements they've identified
in order to reconstruct the text in recall.

EPISODES ELABORATE UPON PLAN SCHEMATA

One may notice that the constituents of episodes which we have identified
are almost the same as the elements of a Plan schema underlying intentional
actions. Plans have goals, actions, outcomes, and may encounter complica-
tions. Thus, people's knowledge of narrative episodes certainly includes
their knowledge about Plans. In this view, readers use their general Plan
schema to understand intentional action sequences, and the Plan organizes
behaviors according to their goals. Studies by John Black and I (1980) and
Edward Lichtenstein and William Brewer (1980) have found that action Plans
have a hierarchical structure, that goal-directed actions at higher-levels
of the goal-tree are remembered better than non-goal-directed actions at
lower, more detailed levels. Also, people do best at processing and
recalling a text when it mentions the Plan elements in their stereotypic
order.

I want to examine more closely now how plans and goals are used by readers
in processing narratives. This topic is discussed in detail in the book by
Roger Schank and Bob Abelson, (1977) and their student, Bob Wilensky,
(1978), wrote a computer simulation program which understands plan-based
stories. The program was called PAM, the initials standing for Plan
Applying Mechanism. The basic assumption is that people understand events
or statements in narratives by trying to explain them. Thus, actions are

to be understood by reference to the actor's plan; plans are understood by
reference to the goal they serve; goals are understood by reference either
to a superordinate goal, or a state or theme that gives rise to the goal.

Wilensky's PAM program follows a specific algorithm in understanding each
event as it occurs. First, it checks whether the action satisfies an
on-going expectation--for example, whether it fits into a known plan for
the actor. If so, then that's the explanation of the event and it is thus
incorporated into the reader's developing representation of the story.
Second, if the immediate predictions fail for this action, then the reader
tries to infer a plan which includes this action, then checks to see
whether this plan serves a known goal. Third, if a goal is stated or in-
ferred, the reader supposedly checks whether it is consistent with a higher
goal or theme the actor has.

THE DISTANCE EFFECT IN GOAL-ACTION PAIRS

You might have noticed that some actions will be psychologically close to a
given goal but farther away from other goals. That is, a given action may
relate to its goal either directly or indirectly through several inter-
mediate steps or sub-goals. This intuitive notion of the logical distance
between a goal and an action can be explained using the idea of a goal-
subgoal hierarchy or a goal-reduction tree. A goal reduction tree decom-
poses a top-level goal into subgoals, and those into further subgoals or
actions that can be performed. Figure 1 illustrates part of a goal reduc-
tion tree for someone's knowledge about how to steal money, which can be
done, let's say, through embezzlement, armed robbery, or stealthy burglary.
To carry out armed robbery, one should have a gun, a get-away plan, and
select a suitable target like a bank. To rob a bank requires that you get
information about the bank's cash reserves, what kind of security systems
they have, and so on.

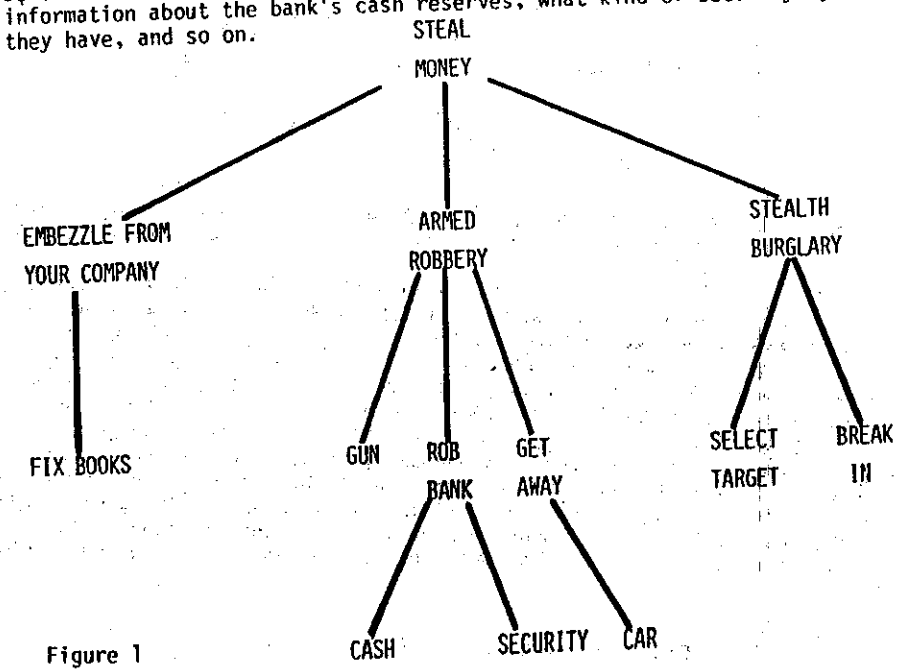

Figure 1

In such a goal-reduction tree, "understanding an action in light of a goal" would mean finding or computing a connecting link of the correct kind between the two elements in the tree. Notice that certain goal-to-action pairs are quite close in the tree whereas others are farther away. When someone thinks about this plan or goal-reduction tree, it would not be available in active memory all at once; rather, it would be retrieved piecemeal from long-term memory. We may think of the links in this figure as a set of one-step productions in memory which encode rules of the form "IF you want to achieve goal G, THEN do subgoals or actions A, B, and C." Thus, if one wants to find a connection between a goal and some action, the productions starting from that goal will be fired, entering its subgoals into active memory, and these in turn will fire their productions, entering their subgoals into active memory. If in this expanding activation process the specified action is encountered, then a connection has been found, so we can say that the reader has understood the action in terms of the plan.

If the retrieval and activation of each link in the goal-tree takes time, then comprehension or interpretation time should take longer for those action-goal pairs that are farther apart in the network. For instance, an action like "John checked out the security guards at the bank" would be understood quickly when preceded by a Near goal like "He wanted to rob a bank" but more slowly when preceded by a Far goal like "He wanted to steal some money".

DISTANCE EFFECT WITH NATURAL GOAL TREES

A Stanford student, Carolyn Foss, and I performed an experiment to see whether this analysis was worthwhile. First, we had to make up many pairs of goals and actions which were psychologically Near or Far from one another according to a plausible goal-reduction tree. Unfortunately these materials could only be chosen informally, by guessing about prototypical goal trees for many standard plans of our subjects. A principled way to select Near versus Far goal-action pairs is to choose three elements along a goal-reduction chain, as in the example above of a top-goal ("stealing"), a subgoal ("armed robbery"), and a lower action ("check out security guards"). In our experiment, we fixed the action and then preceded it either with a Near subgoal or with a Far, higher goal. Thus, the Near subgoal-action pair was nested within the Far goal-action pair, with a shorter distance.

Subjects were timed as they read each statement within a number of four-line episodes. They read for comprehension and had to answer a question after reading each episode. The subject pressed a button to present himself with each successive statement of the text on a CRT. Subjects were not aware that they were being timed for line-by-line reading. The time between button-presses presumably measures the time the subject required to read the statement, comprehend it, and integrate it into his interpretation of the text. We expect that actions will be understood faster when they follow Near rather than Far goals.

I've discussed only the case where the goal precedes the action, and where we measure the time required to understand the action. However, if understanding simply requires connecting up a goal with an action, then one might expect a similar distance effect when the goal follows the action and we measure the time required to comprehend the goal and its connections to that prior action. Thus, the subject would be timed on the second sentence

as he read the Near sequence "John decided to rob a bank. He wanted to steal some money" versus the Far sequence "John checked out the security guards at the bank. He wanted to steal some money." Presumably, when the action is stated first, the person infers a plan and goal for it; then when the target goal is read, it will produce either a relatively direct match to the predicted goal in the Near case or will require several steps of inference in order to link up through subordinate goals in the Far case. Therefore, we predicted that the effect of distance on comprehension would be about the same whether the subject were comprehending the action in light of the goal, or vice versa. To test this, we had our subjects read four-line episodes where the middle lines were equally often in the action-goal order and in the goal-action order.

The results of this experiment are shown in Figure 2, which depicts the average time required to comprehend a target sentence. The top line depict

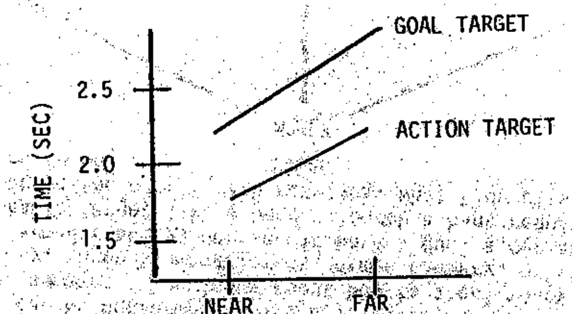

Figure 2

the time to understand a goal following a Near or Far action; the bottom line is the time to understand an action following a goal.

The first conspicuous result is that a target sentence that is Near to its preceding context sentence is comprehended about one-third second faster than are targets that are Far from their preceding context. So, this is the distance effect we were seeking.

A second result in Figure 2 is that readers are about one half second faster in understanding an action following a goal than in understanding a goal following an action. This goal-then-action sequence is, of course, the prototypical as well as causal order of these elements in the Plan schema. Thus, we may conclude that people more quickly understand statements when they occur in the same order as the slots in the schema used t encode the sequence.

Figure 2 shows no interaction between the order of the goal and action, a the distance between them. The two factors have additive effects on comprehension time.

DISTANCE EFFECT WITH SPECIALLY TRAINED GOAL TREES

Although this experiment succeeded in demonstrating distance effects, Carolyn Foss and I were bothered that we had no measure of the distance between a goal and action except our intuitions, which at best provide on an ordering of more or less distance within a given goal-subgoal chain. Our intuitive guesses might be wrong about the goal tree of many of our subjects. Also, it's not clear to what extent the intuitive sense of goal-action distance we were using was just associative strength of connection between the two predicates. Thus, to take just one example, "steal" and "rob bank" are more closely associated than are "steal" and "gun". To counter-argue this point, if one accepts the idea that people store plans, then "associations" are just the consequence of the causal order of these events in the Plan.

For such reasons, we decided to stop using naturalistic materials of unknown organization and instead have the subject learn a novel goal hierarchy which we could specify precisely. Therefore, Carolyn Foss and I ran a second experiment in which we first had subjects read a text describing a novel procedure; then, after they had thoroughly learned the goal-tree of that procedure, subjects made a number of timed judgments using their knowledge of this tree.

The text the subjects studied described the procedure for joining a fictitious Top Secret Club. The goal-hierarchy implicit in the text is shown in Figure 3 below.

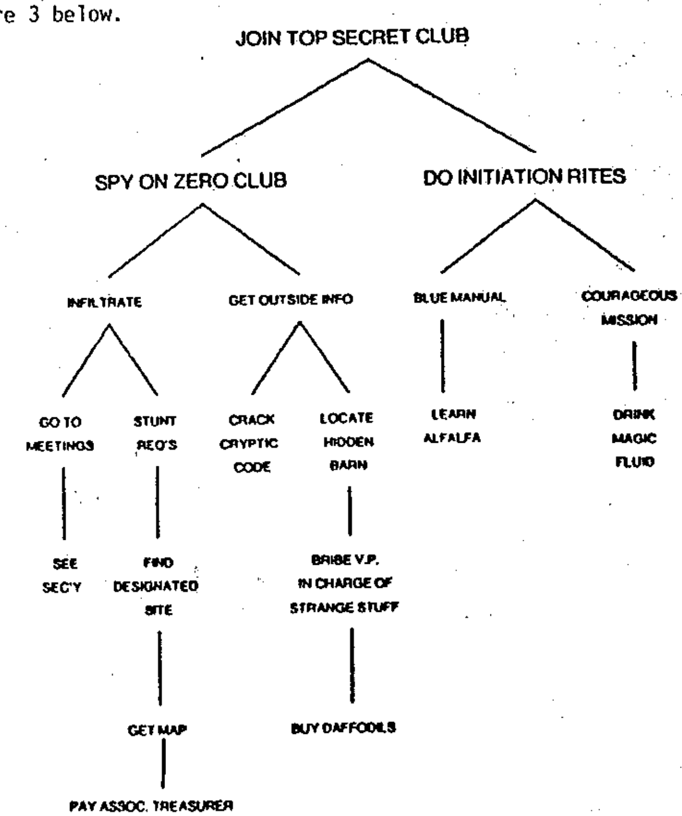

FIGURE 3
GOAL-STRUCTURE FOR EXPERIMENT 2

Thus, in order to join the Top Secret Club, the candidate must spy on its enemy, the Zero Club, and perform several initiation rites. To spy on the Zero Club, the candidate has to infiltrate the club and also get some outside information about it. To get that required that he crack a cryptic code and locate their treasury in a hidden barn, and so on. This goal-tree consists of 16 subgoals nested along six branches. It was rendered into prose resembling instructions for imaginary games like "Dungeons and Dragons" or "Startrek" with which most of our subjects were familiar. The subjects never saw the goal tree as set forth in Figure 3. Rather, they studied the text until they learned it well before the testing phase began.

For the test phase, subjects were told that some CIA agents had found burned and shredded copies of the procedure for joining the club, and they were trying to piece together the original complete procedure. These agents would formulate a plan and submit it for evaluation to the subject, since he was the expert; he was to decide quickly whether or not the proposed plan was well formed. The proposed plans were formatted as two separate clauses: first, a clause would appear on the CRT such as "In order to (infiltrate the Zero Club)"; after the person read that, he pushed a button which showed the second clause, something like "John had to see the secretary". The subject had to decide whether the action in the second clause was a subordinate or descendant of the goal mentioned in the first clause. Thus, it is proper to say that "In order to infiltrate the Zero club, John had to see the secretary"--that is a correct plan because the action in the second clause falls below the goal in the first clause. An Incorrect plan is one where the second clause refers to an action that is either above it in the tree or on a side branch from the first clause. Thus, it is incorrect to say that "In order to see the secretary, John infiltrated the Zero Club" or "In order to infiltrate the Zero Club, John had to carry out a courageous mission".

Training our subjects on this novel goal hierarchy provides several theoretical advantages. Importantly, we know what the goal structure is, and know that it's roughly the same for each subject. Also we know that the degree of learning of the various links in the tree is about the same, so associative strengths won't be varying randomly. Finally, we now have a simple measure of distance between any two nodes in the network as well as the amount of branching or fanning that occurs between two nodes. This measure enables us to plot parametric functions.

In the experiment subjects judged 56 plans once, then repeated the test series. There were slightly more Correct than Incorrect plans, and half of each type were Near or Far goal-action pairs. The basic result is shown in Figure 4.

Importantly, the time to decide that a goal-action plan is correct increased nearly linearly with the number of steps between the elements in the hierarchy. Each step increased reaction time by about half a second. Such a function would be expected if the person searched links downward from the goal at about half a second per step.

Second, Figure 4 shows that subjects answer quicker the second time through the tests. This would occur either if the subject is strengthening and facilitating the same links he'd used before, or if he is learning distant goal-action dependencies, accessing them directly, and by-passing derivation of their relationship the second time.

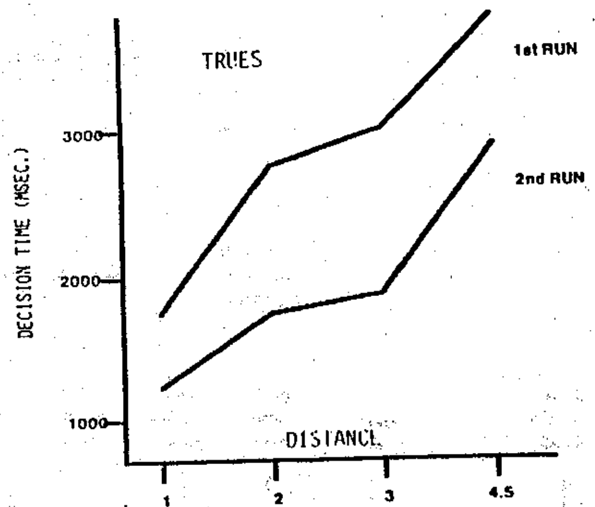

Figure 4

A third effect not shown in Figure 4 was an interference or fan effect of slower search due to branching: for a given step-distance between the goal and action, the decision time was longer the greater the branching along the path connecting them. This would result if the link-searching process is slowed by dividing its resources at branching points.

Turning to the False judgments, we were surprised to find no difference whatsoever between Near, Far, or Lateral False pairs. Mean RT for the Near, Far, and Lateral falses were 2212, 2162, and 2132, respectively. These do not differ significantly. A downward search algorithm for the goal-tree implies no difference for these cases. Downward search means to start from the goal in the first clause and retrieve its descendants below in successive generations; if any of them matches the action in the second clause, respond "Correct"; if none of the descendants match, respond "False". This downward search algorithm explains the lack of differences among the Near, Far, and Lateral falses because in this experiment they all had the same average number of descendants. However, I should point out that this was not a planned or controlled comparison in this experiment, so we are not certain about our conclusions regarding False decisions.

As noted, we can represent the procedural hierarchy in this experiment as a set of one-step productions in memory that link goals to subgoals. The process of searching through the graph structure would then be simulated by the firing of productions, whereby a goal activates its immediate descendants, which fire their productions, activating their descendants. Thus will activation spread across generations. This is one way to implement the node-search procedure that is so familiar in semantic networks.

To summarize, we've found that the time to decide that an action is plausible in light of a goal increases almost linearly with the derivational distance between the two in the goal tree. Branching slows down the search, and repetition of particular pathways strengthens them and speeds up the search.

FURTHER EXPERIMENTS ON GOAL-HIERARCHIES

This experiment with a novel goal hierarchy has yielded orderly results on the time people take to perform memory search and verification. The technique can be exploited to examine a number of questions, some of which we plan to pursue. First, we plan to look at reading time for the second clause rather than decisions regarding proper plans; reading comprehension should be quicker for shorter goal-action pathways. Second, the goal-tree itself can be varied structurally so that one can study more systematically the effects of branching in the goal tree. Third, the test could list a conjunction of actions and ask the subject to decide whether all of them were necessary and sufficient to achieve some superordinate goal. Fourth, in the Foss experiment, the subject learned one large goal-tree and the tests checked the time required to retrieve different segments of the tree; no novel compositions or arrangements were required. The request for novelty in planning suggests further experiments in which we first teach the subject several pieces of disjoint plan hierarchies, and then measure how long he takes on tasks that require him to retrieve and assemble the plan-pieces in a particular order.

CURRENT EXTENSIONS ON GOAL MONITORING

I will briefly describe two extensions of our goal-action research. One project concerns how the reader monitors several goals for the actor. Imagine that the opening of a story describes several separate, independent goals that the actor wants to achieve as the opportunities arise. We conceive of this as the reader setting up a goal-list for that character in short-term memory, and then monitoring for an action relevant to any of these active goals. Later when the text describes an action, we imagine that to understand it, the reader tries to connect it up to some one of the active goals for this character. We may liken this process to Sternberg's memory-scanning task in which the subject searches for a probe digit amongst a memory list. Therefore, one predicts a set-size effect: that is, the more independent goals one has to keep in mind for a character, the longer it should take to decide that an action fits into a plan for some one of these goals.

The materials of this experiment are illustrated in Figure 5. The experimental subject reads many small vignettes in which a list of 1, 3, or 5 goals is introduced, then 0 or 3 irrelevant interpolated sentences occur to produce differing amounts of de-activation, then an action statement occurs. The subject decides as quickly as he can whether the test action is plausibily consistent with some one of the goals. Figure 5 illustrates a trial with 3 goals, with 3 interpolated sentences, and shows an example of a True action as well as a False action (only one would be presented per trial).

So far the results are confirming expectations. Decision time for an action increases with the number of active goals, and the slope (increase per goal) is less for True than for False action probes. The steeper slope for Falses would arise if each goal-action comparison takes much longer to decide mismatch than to decide match due to searching for ever more remote connections between mismatching elements. We are also finding that the interpolated material slows down all decision times and increases errors; this was predicted since interpolated material deactivates the goal elements, so time is needed to reactivate the goals to compare to the action probe.

JOHN WANTED TO	[CATCH A FISH, / LISTEN TO MUSIC, / TELEPHONE SALLY.]	GOAL LIST
	[JOHN IS A HARVARD MAN. / HARVARD IS IN CAMBRIDGE. / IT HAS ABOUT 6000 UNDERGRADUATES.]	INTERPOLATED FILLERS
SO, JOHN GOT OUT	[HIS PHONOGRAPH]	TRUE ACTION
(OR)	[HIS CHAINSAW]	FALSE ACTION

Figure 5

In sum, this experiment seems to be working rather well, with results as expected. We plan to extend it to study the effects of keeping track of two characters' goal lists, deleting goals from the list as they're achieved, using action probes that simultaneously satisfy several of the goals, and so on. The experimental paradigm seems useful for answering many questions about how people track multiple goals and up-date the active goal-list for particular characters in a story.

A second extension Carolyn Foss and I are currently exploring is to measure comprehension and decision time using goal-hierarchies in the context of actual stories. Unfortunately, most of the fairy tales Carolyn examined had very shallow goal-trees, which are not ideal for investigating causal distance effects. One story that is notorious for having a long chain of subgoals is "The Old Farmer and His Stubborn Donkey" (Rumelhart, 1975), so we are using that in a current experiment. The experimental subjects are being timed while deciding the truth of causal implications of the form "In order to achieve X, the farmer did Y". As before, the basic question is whether decision time will increase with the distance between the goal and the action in the hierarchy. I'm hoping that the results of this study will confirm our earlier findings but with a goal hierarchy the subject has learned within a real story.

CONCLUDING REMARKS

It is time that I bring this paper to a close. I have been discussing episodes and the plan-goal analysis of actions, and have investigated how people comprehend actions in light of goals, or goals in light of actions. The guiding theme is that readers search for explanations of narrative events, and that the difficulty of comprehension, and hence reading time, increases the greater the derivational distance between a goal and a relevant action in a planning space. Also, tracking the goals of a character

can be thought of as maintaining those goals in active memory as explanatory sources for later events, with the time to find a given goal-to-actio linkage depending on how many goals are active, how long is the link up, how activated are the correct versus incorrect goals, and so on.

These findings are not especially surprising given the theoretical analysi of the comprehension tasks in terms of goal hierarchies and memory search through activated elements in short-term memory. But the power of such ideas from cognitive psychology is their ability to explain different phenomena. The value of a theoretical framework is sometimes just to enable us to think systematically about certain phenomena and to frame questions about them in such manner that the answers seem almost obvious. Interestingly, researchers' feeling of understanding events in nature by substantiating their theoretical expectations runs almost exactly parallel to readers' feeling of understanding story events because they substantiat predictions they're made about the characters. This seems entirely fitting since both the scientist and the comprehender are just trying to explain events that engage their attention.

FOOTNOTE

Research reported here was supported by a grant MH-13905 to the author fro the United States N.I.M.H.

REFERENCES

Black, J. B. & Bower, G. H. Story understanding as problem-solving. Poetics, 1980, 9, 223-250.

Bower, G. H. Experiments on story understanding and recall. Quarterly Journal of Experimental Psychology, 1976, 28, 511-534.

Flammer, A. Toward a theory of question-asking. Research Bulletin, Nr. 22, University of Fribourg, Psychologisches Institut, Fribourg, CH. 1980.

Lichtenstein, E. H. & Brewer, W. F. Memory for goal-directed events. Cognitive Psychology, 1980, 12, 412-445.

Rumelhart, D. E. Notes on a schema for stories. In: D. Bobrow and A. Collins (Eds.), Representation and understanding. New York: Academic Press, 1975, Pp. 237-272.

Schank, R. C. & Abelson, R. P. Scripts, plans, goals, and understanding. Hillsdale, N.J.: Erlbaum. 1977.

Thorndyke, P. W. Cognitive structures in comprehension and memory of narrative discourse. Cognitive Psychology, 1977, 9, 77-110

van Dijk, T. Some aspects of text grammars. The Hague: Mouton. 1972.

Wilensky, R. Understanding goal-based stories. Ph.D. Dissertation, Research Report #140, Computer Science Dept., Yale University, 1978.