COGNITIVE SCIENCE 8, 59–85 (1984)

Empirical Evidence for Narrative Structure*

JAMES PAUL GEE
Boston University

FRANCOIS GROSJEAN
Northeastern University

Three experimental tasks—spontaneous telling of a story, reading, and parsing the story—were used to determine whether empirical data reflect the narrative structure of stories and can be predicted by a plot unit analysis of the stories (Lehnert, 1981). It was found that spontaneous pause durations at sentence breaks were highly correlated with the importance of these breaks as predicted theoretically. Only low correlations were obtained, however, when reading pause durations were correlated with the model. As for parsing values, the value of the correlation coefficients depended on whether stories had sufficient superficial linguistic cues to help the subjects in parsing. It was concluded that spontaneous pausing not only reflects the narrative structure of stories, but can be used as a guide to constructing theories of narrative structure as well as for deciding between competing theories

Narrative has become a topic of interest in a number of areas, including literary criticism, linguistics, psycholinguistics, and anthropology. While many researchers have an avid interest in the internal structure of narratives, we still know very little about the way narratives are actually structured, be it in their telling, in their perception, or in their storage in memory (see Glenn, 1978; Hymes, 1982; Kintsch, 1974; Kintsch & Greene, 1978; Labov, 1972; Labov & Waletzky, 1967; Mandler & Johnson, 1977; Mandler, et. al., 1980; Perfetti, 1982; Rumelhart, 1975, 1977; Stein & Glenn, 1979; Thorndyke, 1977).

*The authors are coequal and their names are simply listed in alphabetical order. We gratefully acknowledge that this research was funded in part by grants from the Department of Health and Human Services: RR 07143 and NS 14923 to the Department of Psychology, Northeastern University. We would also like to thank the members of the Language and Cognition Group in the Department of Psychology at Northeastern for their encouragement and help throughout this project. Correspondence and requests for reprints should be addressed to James Paul Gee, School of Education, 605 Commonwealth Avenue, Boston, MA 02215.

One simple approach to the structure of narratives is to diagram them in terms of tree structures similar to phrase structure trees for sentences. For example, Gee and Kegl (in press) diagram a particular story as follows: the Story as a whole splits into two constituents, an Initiation and a Conclusion. The Initiation splits into an Action and a Result, while the Conclusion splits into a Problem and its Resolution. The Resolution, in turn, splits into a False Resolution and a True Resolution, each of which are themselves made up of an Action and a Result. Such simple parsings of a story have the virtue of clearly raising two questions: First, can such simple tree structures realistically represent the structure of complex stories (even given the fact that they can represent the structures of sentences) and second, how should one choose the labels for the units in the story (the nodes in the tree)? Both questions ask for a theory of the units of story structure.

Recent research on story grammars shows that not all stories can be assigned hierarchical branching structures of the sort familiar from work on English phrase structure. Such hierarchical trees may be too simple to capture the many complex relations in a narrative (for discussion, see Black & Wilensky, 1979; Gee & Kegl, in press; Kintsch, 1974; Kintsch & van Dijk; Lehnert, 1981; Rumelhart, 1975, 1977; Thorndyke, 1977). For example, one unit in a narrative may be simultaneously part of two different higher-order plot units. Units may interrupt other units, occurring temporally in the middle of higher-order units to which they do not conceptually or thematically belong, as in the case of anticipations, flashbacks or flashforwards.

The issue of what to label the units in a story creates even more vexing problems. There is a long tradition in literary criticism of analyzing narratives in terms of basic plot motifs or small thematically-named units centered around the actions of characters. For example, in what is perhaps the most famous of these sorts of analyses, V. Propp (1969, original work published 1928) argues that all Russian folktales can be characterized by the combination and organization of 31 basic motifs (e.g., "interdiction," "violation," "reconnaissance," "trickery," "lack," "receipt of a magical agent," and so forth). Such analyses are satisfying in that they specify the thematic content of the units of the narrative—their intrinsic meaning in the story as a whole. Of course, such analyses raise issues about the cultural and genre specificity of certain narrative units, and the extent to which some units will occur more widely given the human mind's universal attempt to structure experience through narrative. Lehnert (1981) has recently developed a computationally and psycholinguistically based approach to narrative structure which is articulated in terms of basic plot motifs. In addition, Lehnert offers a "grammar" that structures these basic motifs into higher-order units, eventually assigning a structure to the story as a whole that goes well beyond the power of a standard phrase structure representation.

The question we will address in this study is not whether a nonhierarchical analysis of narrative structure is necessarily superior to any conceivable hierarchical analysis (although this is an important question in its own right), but, rather, whether we can obtain empirical evidence for the narrative structure of relatively complex narratives. Whatever the theoretical model, can a task be developed so that data obtained from subjects reflects the organization of a narrative as predicted by a theory? If this is the case, can one go a step further and propose that researchers use empirical data to help obtain and analyze the structure of narratives? Since a great deal of research in a variety of fields has yet to yield agreement on how narratives are structured, and how best to represent narrative structures, these questions indeed appear to be necessary preliminaries. We will, in fact, suggest that pausing can give us crucial insight into the structure of stories and, thus, serve as a data base for the construction and validation of theories and competing hypotheses in this domain.

These types of questions have been asked at a level below the narrative; namely, at the level of the sentence. The answers achieved are instructive because they show that pausing can be used, at the level of the sentence, to test various theories of sentence structure. We will argue below that pausing is important to the understanding of narrative structure as well. Grosjean, Grosjean, and Lane (1979) asked subjects to read sentences at various rates and measured the pauses that were produced (especially at slow rates) between every two words. Then they constructed hierarchical structures based on the pause durations by first linking the words that were separated by the shortest pause, then linking the words or group of words separated by the next level pause, and so on. The structures constructed in this manner were then compared to constituent structure trees derived from a standard phrase structure grammar of English. Grosjean, Grosjean, and Lane showed that although a surface structure representation of the sentence was a fairly good predictor of their data structures (they called these "performance structures"), a number of mismatches occurred between the two. For example, the NP-VP break of a sentence often received a very short pause, whereas in constituent structure it was the most important break of the sentence.

Gee and Grosjean (1983) reanalyzed these "performance structures" for sentences and showed that they reflect prosodic structures and not, as was first thought, surface constituent structures. Prosodic structures are hierarchical structures generated by the phonological component of the grammar and based on the stress and rhythmical pattern of the sentence. These prosodic structures, however, also reflect certain aspects of the syntactic and discourse structure of the sentence. When Gee and Grosjean compared prosodic structures and performance structures, they found that the

former accounted for 92% of the variance of the latter, thereby showing that empirical data could reflect, in a clear manner, the structure of a sentence.

This work on the performance structures of sentences naturally raises the question of whether performance structures exist at the level of narratives. That is, would the structures obtained from empirical data (pausing in this case) have hierarchical structure in the sense that different idea groups would be separated from each other by longer or shorter pauses? Or, would the structures be flat in the sense that different idea groups or thematic units would be linked to one another by pause durations of roughly equal size? If the data produce some kind of structure, what does it reflect? Is it a hierarchy that is due to the experimental technique used (an uninteresting result) or a hierarchy that reflects the structure of the narrative as proposed by models of narrative structure? If the latter alternative is the correct one, then we may be in a position to propose an empirical method of validating hypotheses about narrative structure.

The only previous study that has asked these questions in this way (but, see Lehiste, 1975, for discussion of the phonetic structure of paragraphs) dealt with (in addition to another story) the narrative structure of "Goldilocks" as signed in American Sign Language (Gee & Kegl, 1983). The "pauses" (actually holds at the ends of signs) that were produced while the story was signed at various rates were used to construct a performance structure tree of the narrative. This tree was then compared to a formal analysis of the story, also represented in tree form (the story was simple enough to have a straightfoward simple hierarchical representation). The performance tree and the narrative structure tree were found to be correlated .86, indicating that the empirical tree was indeed a good reflection of the formal tree. To our knowledge, this approach has not been used with narratives in English.

In the present study we will ask a number of subjects to tell the Fox and the Bear story (Stein & Glenn, 1979) with and without prompts. We will then average the silent pause durations that are produced at the major sentence breaks (there are 12 in all) and will examine the ensuing pattern of sentence groups. If the pausing is of the same magnitude between every sentence, we will deduce that pausing does not reflect the narrative structure of the story. If, on the other hand, some sentences are grouped more closely together by pausing than others, then we will conclude that some suprasentential pattern can be obtained from the empirical data. Whatever the outcome, we will adapt Lehnert's (1981) model of narrative structure to give a theoretical narrative structure of the Fox and the Bear story (see also, Lehnert, Black & Reiser, 1981; Reiser, Lehnert, & Black, 1981). We will then compare this prediction of the narrative structure with the data (or performance) structure of the story by correlating the values of each of the 12

breaks obtained from the theoretical model with the pause durations at these breaks. If the correlation is high (in the order of 0.7 or 0.8), we will conclude that pause data is as good an indicator of narrative story structure, as it is, at a lower level, of prosodic sentence structure. If the correlation is low, however, or at worst zero or negative, we will deduce that pausing data cannot be used to obtain the supra-sentential structure of a narrative.

EXPERIMENT 1

METHOD

Subjects

Three subjects, who had read the Fox and the Bear story aloud 3 months before in the course of another experiment, took part in the study. None reported having any speech or hearing difficulties.

Materials

Two texts were used. The first was the typed version of the Fox and the Bear story (Stein & Glenn, 1979). We numbered the sentences from 1 to 13 for convenience in later reference. The numbers were not in the text given to the subjects.

> (1) There was a fox and a bear who were friends. (2) One day they decided to catch a chicken for supper. (3) They decided to go together because neither one wanted to be left alone and they both liked fried chicken. (4) They waited until night time. (5) Then they ran very quickly to a nearby farm where they knew chickens lived. (6) The bear who felt very lazy climbed upon the roof to watch. (7) The fox then opened the door of the henhouse very carefully. (8) He grabbed a chicken and killed it. (9) As he was carrying it out of the henhouse the weight of the bear on the roof caused the roof to crack. (10) The fox heard the noise and was frightened but it was too late to run out. (11) The roof and the bear fell in killing five of the chickens. (12) The fox and the bear were trapped in the broken henhouse. (13) Soon the farmer came out to see what was the matter.

The second text, which was used as a prompt, contained the important concepts of the Fox and the Bear story. It took the following form:

> FOX BEAR CHICKEN SUPPER GO-TOGETHER NOT-LIKE-ALONE
> BOTH-LIKE-FRIED-CHICKEN WAIT NIGHT-TIME RUN FARM CHICKENS LIVE

BEAR LAZY ROOF WATCH FOX OPEN DOOR HENHOUSE GRAB CHICKEN
KILL CARRY-OUT WEIGHT BEAR ROOF CRACK FOX FRIGHTENED
TOO-LATE RUN-OUT ROOF BEAR FALL-IN-KILL CHICKENS FOX BEAR
TRAPPED FARMER COME-OUT SEE

Procedure

Subjects were run individually in sessions that lasted 20 minutes. They were seated with the experimenter in a recording studio and a microphone linked to an outside tape recorder (AKAI 4000DB) was placed in front of them. In the first subtask (–reading, +prompt), the subjects were given the prompt text of the Fox and the Bear story and were told to use the prompts to help them recall the story (they had not seen the story in 3 months and had never been told that they would be tested on it again). Once the subjects felt ready, they were asked to tell the experimenter the story, using the prompts to do so. In the second subtask (+reading 1, +prompt), the subjects were given the text of the Fox and the Bear story to read silently once. When this was done, they were asked to tell the experimenter the story again with the aid of the prompts. This subtask was repeated a second time (+reading 2, +prompt). Finally, all materials were removed and the subjects were asked to tell the story from memory(–reading, –prompt). Each subject, therefore, told the story four times.

Data Analysis

The 12 recordings—4 by each of 3 subjects—were transcribed by two judges. The transcriptions included the various slips of the tongue produced by the speakers, as well as their hesitations (false starts, filled pauses, drawls, etc.). The recordings were analyzed by means of a pen-recorder (Gould Brush 220) whose paper speed was set at 24 mm/sec. This permitted us to locate and measure the silent pauses produced by the subjects (pauses were represented by straight lines on the pen-recordings). We measured the pauses at each of the 12 sentence breaks in the text given in the previous section. This never proved to be a problem as sentences in the story represented coherent idea units and subjects always produced these. The pauses thus measured were inserted into the transcripts. One of the 12 transcripts, obtained from a +reading 2, +prompt task is given below (numbers between slash bars are silent pause durations expressed in seconds):

> Once there was a fox and a bear who were good friends /.44/ they decided to have supper togather /.48/ they decided to go together because

> neither one of them liked to go alone both of them liked fried chicken /.68/ and they decided to wait until night time /.40/ ah and they ran over to a farm where they knew chickens lived /.76/ now the bear was very lazy and he climbed up on the roof to keep watcch /.64/ the fox very carefully opened the door to the henhouse /.44/ grabbed a chicken and killed it /.76/ and he was about to carry it out when the weight of the bear cracked the roof of the henhouse /.32/ the fox was very frightened and tried to run out but it was too late /.64/ the roof and the bear both fell in killing five chickens /.64/ the fox and the bear were trapped /.56/ and later on the farmer came out to see what was the matter

The pause durations at a particular sentence break were first averaged within a subject across the four subtasks (if a subject had never paused at a sentence break we would have tabulated a 0 value, but this never happened). Then, pause durations were averaged across the three subjects. This produced 12 mean pause durations, one for each of the 12 sentences breaks in the story.

RESULTS AND DISCUSSION

Figure 1 presents both a linear and a hierarchical representation of the Fox and the Bear story based on the silent pauses produced in the spontaneous telling. The linear representation is such that a sentence (represented here by three or four key words) is separated horizontally from the next sentence by a physical distance equivalent to the mean pause duration between the two sentences. Two sentences will be close to one another if a short pause separates them and will be at a distance from one another if they are separated by a long pause. As for the hierarchical representation (which in no way presupposes that the story is indeed best represented hierarchically), the sentences have been grouped together according to the following iterative procedure: find the shortest pause (in this case 0.47 sec) and cluster the two sentences on either side (in this case, sentences 7 and 8) by linking them to a common node. Then cluster the two elements (sentences or clusters) that are separated by the next shortest pause (in this case, 0.50 between sentences 11 and 12) and continue doing this until every sentence (or sentence cluster) is part of the structure. To indicate the strength of a cluster, height has been given to the nodes: a short pause duration is represented by a low node, whereas a long pause duration is represented by a high node.

An examination of Figure 1 clearly shows that a narrative can be structured by means of spontaneous pause durations. The durations obtained from the subjects range from 0.47 seconds all the way to 0.91 and produce a series of sentence clusters: sentences 7 and 8 (linked by the smallest pause); sentences 11, 12, and 13 (linked by the next smallest pause); sentences 9 and 10; sentences 4 and 5; and, finally, sentences 1, 2, and 3. These clusters are

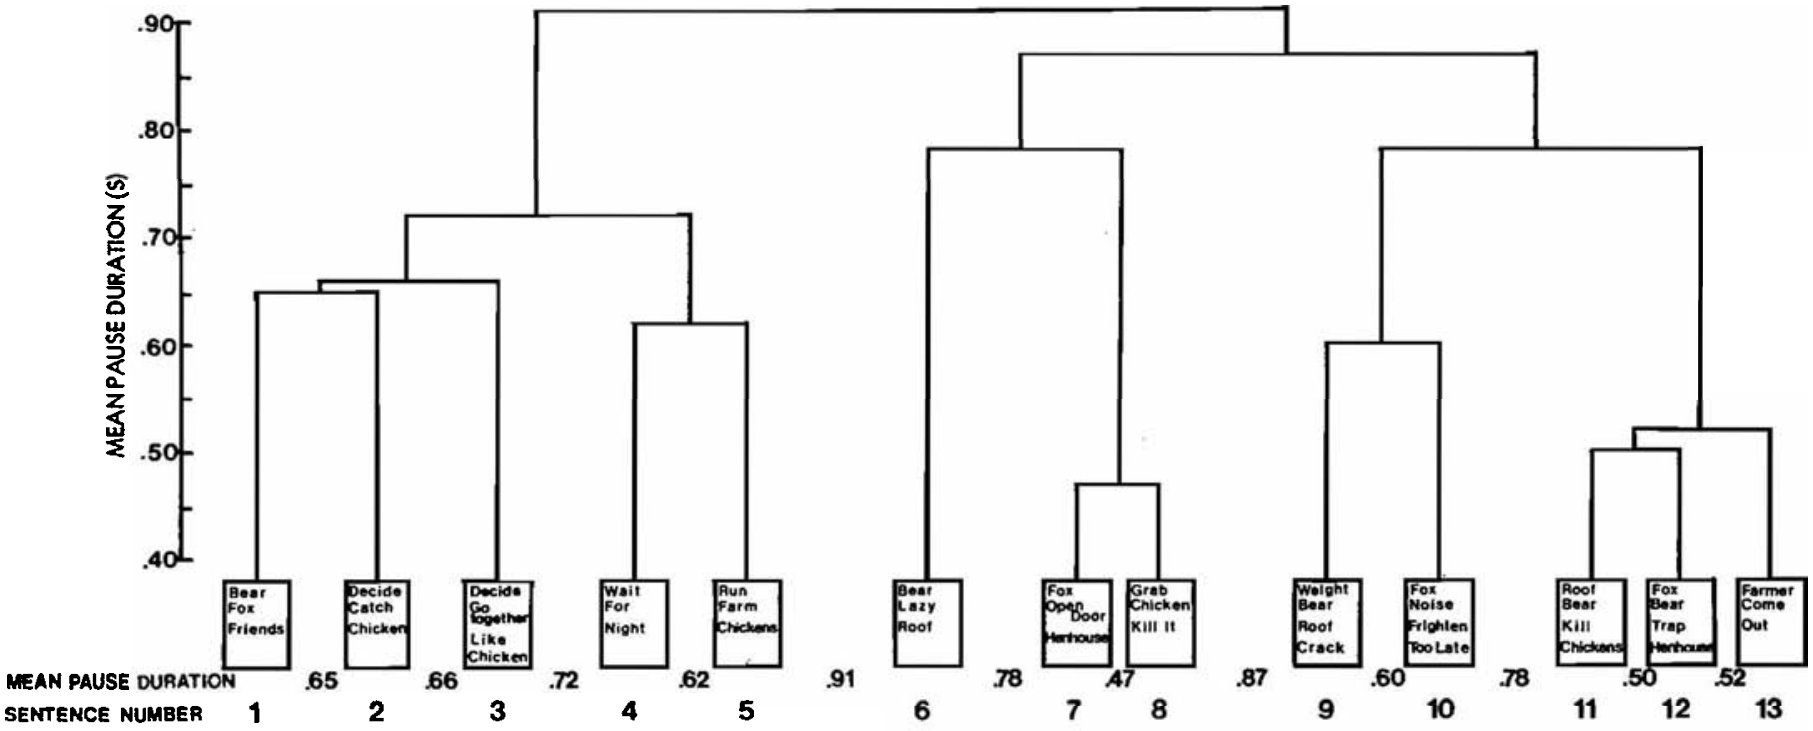

Figure 1. A linear and hierarchical representation of the Fox and the Bear story based on the silent pause durations produced in spontaneous telling. Linear distance between sentences (represented by three or four key words) and height of nodes are both proportional to mean pause duration.

themselves clustered together into higher level groupings, so that the story as a whole is a hierarchical cluster of clusters. Thus, is it certainly not the case that pause durations are identical at each sentence break and that, therefore, the importance of a break in a narrative cannot be reflected by pause durations.

Not only does sentence grouping emerge from the story-telling data, but this grouping is intuitively logical as well. The sentences most closely linked to one another (i.e., separated by the smallest pauses) are also conceptually related. For example, sentence 7 represents (in its content) the cause of sentence 8; sentence 9 the cause of 10; 11 the cause of 12; and 11 and 12 are together the cause of 13. At the higher levels, note that sentences 1–5 represent the coming to a decision and the formulation of a plan and take place at the home site of the fox and the bear, whereas sentences 6–13 represent the action carrying out the plan and take place at a different location, i.e., the farm. The longest pause in the story (.91) takes place at this change of scene/location. The intermediate groupings also seem plausible: sentences 1–3 are mental events and take place at a different time than 4–5 which contain a time elapse and an action that carries us to the location of the second part of the narrative; sentences 6–8 contain the initial actions of the fox and the bear in carrying out their plan, 9–10 contain the crisis, and 11–12 represent the denouement (thus, the pausing actually segments out the classical rising pyramid structure of Aristotlean tragedy—rise to crisis, crisis, fall to conclusion—argued by many to be *the* basis of plot).

The Prediction of the Performance Structure of the Fox and the Bear Story

In order to confirm that pause durations do indeed reflect the narrative structure of the Fox and the Bear story, we decided to take a theoretical model of narrative structure and adapt it so as to predict the importance (the complexity) of the sentence breaks in the story. We chose to work with Lehnert's (1981) basic plot unit system of story analysis. Her model was developed to study the processing strategies used in narrative summarization in the context of psychological experiments and computer simulation. We should note that our interest is not in validating Lehnert's model; we do not believe that at this stage of our knowledge of narrative structure, any analytical model can be anything more than roughly approximate. Our major interest is in using a sufficiently plausible model of story structure to render convincing the claim that pausing may well reflect narrative structure. If this is the case, we will propose that pause structure be used to help develop more accurate models of narrative/story structure. Below, we will first discuss Lehnert's model and then explain our adaptation of it to obtain complexity indices at each sentence boundary.

For Lehnert, a narrative text is a number of *simple plot units* combined and connected in various ways to make up complex plot configurations. The simple plot units that go to make up larger plot configurations are themselves made up of what Lehnert calls "affect states," of which there are basically three sorts: events that please a character or further a character's goals ("positive event state"), events which displease a character or frustrate a character's goals ("negative event state"), and mental events which represent a character's needs, desires, decisions, or "speech acts" like requests, promises, or threats ("mental event affect states"). The way in which affect states can combine with each other to make up a simple plot unit puts heavy constraints on the number and nature of the simple units. In fact, Lehnert allows four links between affect states: one state can *motivate* another, one can be *actualized* (fulfilled, realized) in or by another, one can *terminate* (conclude) another, and, finally, one can be *equivalent* to another. Our interest is in the fact that Lehnert names and describes 15 allowable simple plot units that serve as the basic building blocks of larger plot configurations. Some examples follow:

Simple Plot Units, Examples (based on Lehnert, 1981)

1. Motivation
   a) a "Mental Affect State": you need advice
   motivates
   b) a "Mental Affect State": you decide to ask a friend
2. Success
   a) a "Mental Affect State": you need a car
   is actualized in
   b) a "Positive Affect State": you win one in a raffle
3. Loss
   a) a "Positive Affect State": you buy a new car
   is terminated by
   b) a "Negative Affect State": you total your new car
4. Hidden Blessing
   a) a "Negative Affect State": your uncle dies
   turns out to be equivalent to
   b) a "Positive Affect State": you inherit a million dollars
5. Resolution
   a) a "Negative Affect State": your wallet gets stolen
   is terminated by
   b) a "Positive Affect State" for you: they catch the thief
6. Problem
   a) a "Negative Affect State": you get fired
   motivates
   b) a "Mental Affect State": you need and want another job

7. Complex Negative Event
   a) a "Negative Affect State": your wallet gets stolen
   is equivalent to
   b) a "Negative Affect State": you lose $100.

These simple plot units serve as blocks to build more complicated plot configurations. For example, we can put together "Problem," "Success," and "Resolution," according to the rules Lehnert gives, to get a larger plot configuration called "Intentional Problem Resolution":

INTENTIONAL PROBLEM RESOLUTION

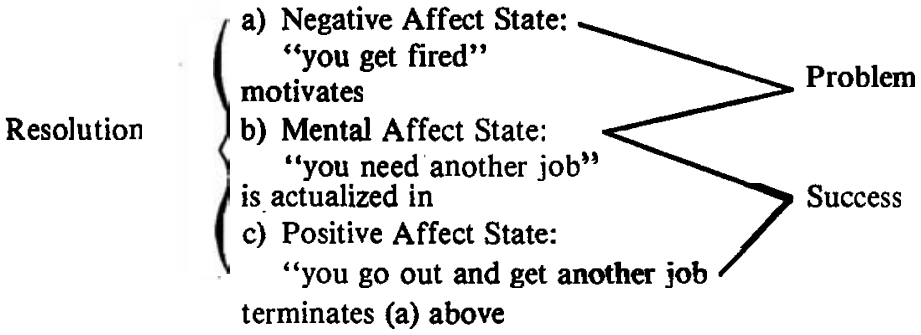

Lehnert shows how to get more and more complicated plot configurations, and also how to deal with plots with more than one character. In this latter case, the affect state of one character is linked to the affect state of another character by a cross-character link.

In our analysis of the Fox and the Bear, we first isolated the simple plot units and larger plot configurations in the story, using Lehnert's model, and then we developed an algorithm to determine the importance or complexity of each sentence break. Applying Lehnert's model to our story was basically a method of applying a set of templates. Each of the applications to our story could be checked in two ways: did a piece of our story fit a name Lehnert gave to a plot unit or larger plot configuration and did that piece fit the internal structure Lehnert assigned to that label (in terms of affect states and links for plot units, or in terms of plot units for larger plot configurations)?

In Figure 2 we present a visual diagram of our theoretical analysis of the Fox and the Bear story in terms of our adaptation of Lehnert's model. The story is made up of two larger intersecting plot configurations, enclosed in large squares, which Lehnert calls "Fleeting Success" and "Regrettable Mistake/Sabotage." The Fleeting Success configuration is composed, for Lehnert, of two simple plot units, an initial "Success" (sentences 2, 3, 5, 7, 8) and a subsequent "Loss" (sentences 7, 8, 11, 12, 13). The simple plot unit

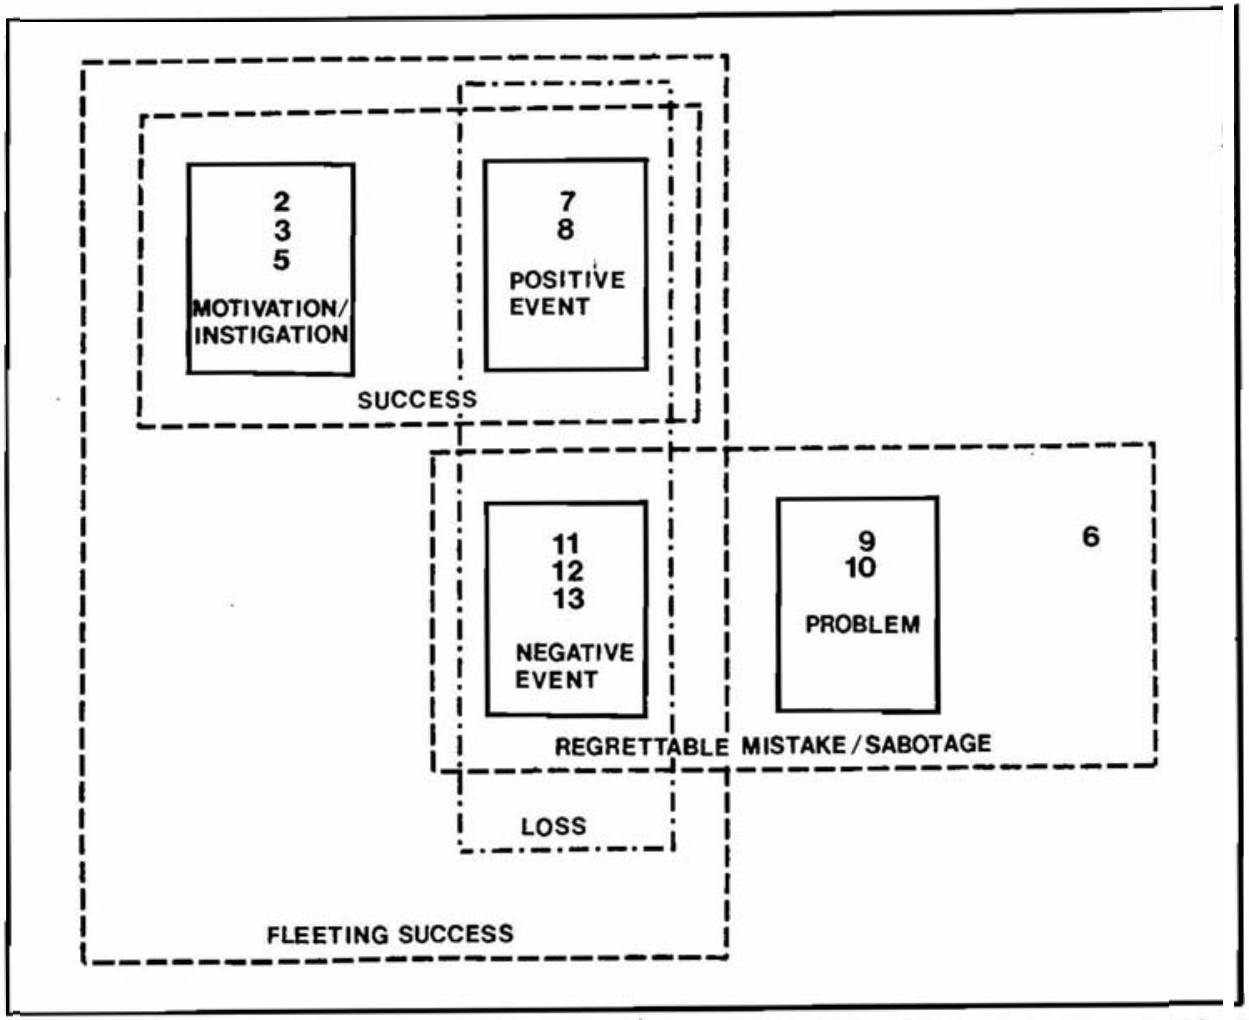

Figure 2. Diagram of theoretical narrative analysis of the Fox and the Bear story in terms of our adaptation of Lehnert (1981). Numbers represent the sentence numbers in the text of the story. Constituents are enclosed within boxes.

of "Success" is made up of a mental affect state (the desire and decision in 2, 3, 5) that motivates a positive or successful outcome (getting hold of the chicken in 7 and 8). The mental events in 2–5 that motivate what is to follow are a bit more complicated than a simple mental event. Sentences 2–5 actually make up what Lehnert calls a simple plot unit of Motivation (2, 3) and another simple plot unit of a successful act (3, 5—getting off to the farm) that instigates the action proper. Thus, we refer to 2, 3, and 5 as "Motivation/Instigation." The simple plot unit of "Loss," the other component of the Fleeting Success, is made up of a positive event (7, 8) that is terminated by or issues in a negative event (11, 12, 13). What makes this whole unit a Fleeting Success is the fact that the positive event that is the outcome of the Success unit is the same one that is turned around by being terminated by the negative outcome of the Loss unit.

The "Regrettable Mistake/Sabotage" configuration, the other large plot configuration in the story, is made up of an action (which Lehnert allows to be positive, negative, or neutral) by one character (the bear) that causes a simple plot unit called a "Problem" (a problem for another character—in this case, the fox/bear team). This Problem, in turn, causes a negative event (here a series of negative events) in 11, 12, and 13.

This leaves only sentences 1 and 4. Sentence 1 is a piece of Exposition information and sentence 4 is a piece of Setting information (it announces a time change). Such categories are not discussed by Lehnert (because they are not really part of the plot proper). We assume, along with the standard procedure in literary criticism, that such pieces of exposition/setting information are outside the main line of the narrative proper and are connected to the narrative as a whole and not to any particular plot unit (see Scholes & Kellogg, 1966; Sternberg, 1978). We represent this by placing them in Figure 2, outside any particular plot unit or configuration.

The Fox and the Bear is an interesting story for narrative analysis. Though it is a simple folktale, it has a number of features that go beyond any simple hierarchical approach to story structure. The story has discontinuous constituents (6 is temporally interrupted from its constituent (9–13) by 7 and 8; 4 interrupts the Motivation/Instigation unit by coming between 3 and 5). The story also has constituents that are simultaneously part of two or more constituents (e.g., 11, 12, and 13 are simultaneously part of "Loss," "Fleeting Success," and "Regrettable Mistake/Sabotage"), and it has setting and exposition information (1 and 4) that lie outside any particular plot unit. Because of this, we have represented the structure of the story essentially in terms of an old-style "immediate constituent" diagram of boxes within boxes (see Hockett, 1958; Wells, 1947).

To obtain a measure of complexity at each sentence break, we used a very simple notion of connectivity defined on the immediate constituent diagram in Figure 2. The measure we used was as follows:

Complexity Index[1]

The complexity of the boundary following any sentence is the number of boxes (lines) in the diagram that occur between that sentence and the sentence immediately following it in the same constituent (either because the sentences are members of the same base constituent, e.g., sentences 3 and 5, or because they are members of the next level constituent, e.g, sentences 5 and 7), to which is added the number of sentences that actually intervene in the text between that sentence and the sentence immediately following it in the same constituent (e.g., sentence 5 is immediately followed by sentence 7 in the diagram, but sentence 6 temporally intervenes in the text).

This method of measuring complexity sees two factors as contributing to the complexity of the boundary following a sentence in the text. First, this boundary is made complex if it is a site where one or more constituent boundaries of the narrative structure fall, and second, this boundary is made complex if the sentence following it temporally is not immediately part of the same thematic constituent as the sentence preceding it. Table 1 below shows how the value for each boundary is computed. We should note that, for the sake of simplicity, we assume that the exposition/setting sentences (sentences 1 and 4) only have to cross the narrative line which surrounds the whole story.

The coefficient of correlation between the 12 complexity indices and the corresponding pause durations is 0.88 ($t<0.01$). This means that as the narrative complexity of a break between two sentences increases, the pause produced by a speaker also increases—and in a very systematic way. This very encouraging result leads to two conclusions. First, the pauses that are produced during the spontaneous telling of the story are well predicted by a theoretical analysis of the story structure (in this sense it also validates Lehnert's analysis). Second, there is now evidence that experimental data can be used to help analyze the structure of a narrative and to choose among competing analyses. Thus, pausing is sensitive not only to prosodic structure at the level of the sentence (Gee & Grosjean, 1983), but also to narrative structure. Used circumspectly and in conjunction with other methods, the analysis of pausing may prove to be a very valuable tool in building theories of narrative structure.

[1]A slightly more formal statement of the Complexity Index is as follows: The complexity of the boundary following any sentence *n* is the number of boxes (lines) in the diagram between *n* and the sentence immediately following it in the same *minimal constituent,* where two sentences *n* and *m* are in the same minimal constituent X just in case *n* and *m* are within the box labeled X and there is no box Y such that *n* and *m* are within Y and Y is within X. To this number is added the number of sentences in the text that temporally intervene between *n* and the sentence immediately following it in the same minimal constituent.

TABLE I

Complexity values at the boundary following each sentence in the Fox and the Bear story and the way they were computed from our complexity Index.

Sentence	*Value at End of Sentence*	*Computation*
1	1	Exposition—crósses into the main line of the narrative—one box
2	0	2 is followed by 3 in its constituent, there are no boxes between 2 and 3 in the diagram and no sentence in the text interrupts them
3	1	3 is followed bv 5 in its constituent—there are no boxes between them, but one sentence (4) temporally intervenes between 3 and 5 in the text
4	1	Setting—crosses into the main line of the narrative—one box
5	4	The sentence immediately following 5 in the next level constituent is 7. There are three boxes between 5 and 7 (Motivation/Instigation, Loss, and Positive Event) and one sentence temporally intervenes between 5 and 7 in the text (6)
6	3	The sentence immediately following 6 and in the next level constituent with it is 9—there is one box between 6 and 9 (Problem) and two sentences temporally intervene in the text between 6 and 9 (7 and 8)
7	0	There are no boxes in the diagram or sentences in the text between 7 and the sentence immediately following it in the same constituent (8).
8	6	The sentence following 8 in the next level constituent is 11. Between 8 and 11 lie four boxes (Positive Event, Success, Regrettable Mistake/Sabotage, Negative Event) and two sentences temporally intervene in the text between 8 and 11 (9 and 10)
9	0	There are no boxes in the diagram or sentences in the text between sentence 9 and the sentence following it in the same constituent (10).
10	4	The next sentence following 10 in the next level constituent is 11 (both are in the Regrettable Mistake/Sabotage constituent). There are four boxes between 10 and 11 (Problem, Fleeting Success, Loss, Negative Event), but no sentences temporally intervene in the text.
11	0	No sentences or boxes intervene
12	0	No sentences or boxes intervene

EXPERIMENT 2

The results obtained in Experiment 1 are extremely encouraging in that they show a close relationship between the structure of a narrative obtained from experimental data and its theoretical structure. There are two aims to the second experiment. The first, is to obtain converging evidence for the rela-

tionship found in the first experiment. To do this, we will analyze the pause durations obtained in the telling of a second story—the Tiger's Whiskers (Stein & Glenn, 1979). We will develop a theoretical structure for the story following Lehnert's model and will then compute a correlation coefficient between the pause durations obtained from the speakers and the complexity indices based on the model. The second aim is to determine whether other experimental tasks can reflect the narrative structure of a story. We will therefore ask our subjects to read the two stories at varying rates and we will measure the pauses at the sentence breaks. We will then correlate these with the complexity indices obtained from our theoretical analysis of the story. In addition, we will ask our subjects to parse the two stories and we will again correlate the data with the complexity indices. If, in each case, the coefficient correlation is quite high, we will conclude that reading and parsing are two more tasks that can be used to study narrative structure. It may well be, however, that these tasks are fundamentally different from spontaneous story telling (as they appear to be at first sight) and that they do not reflect narrative story structure. We will then have to conclude that researchers should concentrate their efforts on spontaneous story telling in order to obtain evidence for narrative story structure.

METHOD

Subjects

The same three subjects took part in Experiment 2. They were run individually on all tasks.

Materials

The Fox and the Bear story used in the reading and parsing tasks is given in the Materials section of Experiment 1. The Tiger's Whiskers text is also taken from Stein and Glenn (1979) and reads as follows (sentences have been numbered in the text for later discussion):

> (1) Once there was a woman who needed a tiger's whisker. (2) She was afraid of tigers but she needed a whisker to make a medicine for her husband who had gotten very sick. (3) She thought and thought about how to get a tiger's whisker. (4) She decided to use a trick. (5) She knew that tigers loved food and music. (6) She thought that if she brought food to a lonely tiger and played soft music the tiger would be nice to her and she could get the whisker. (7) So she did just that. (8) She went to a tiger's cave where a lonely tiger lived. (9) She put a bowl of food in front of the opening to the cave. (10) Then she sang soft music. (11) The tiger

came out and ate the food. (12) He then walked over to the lady and thanked her for the delicious food and lovely music. (13) The lady then cut off one of his whiskers and ran down the hill very quickly. (14) The tiger felt lonely and sad again.

The prompt text for this story was the following:

WOMAN NEED TIGER'S WHISKERS AFRAID NEED MEDICINE HUSBAND
THINK HOW GET TIGER'S WHISKERS DECIDE USE TRICK KNOW
TIGERS LOVE FOOD MUSIC BRING FOOD TIGER MUSIC TIGER NICE
GET WHISKER DO JUST THAT GO TIGER'S CAVE LONELY TIGER
LIVE BOWL FOOD FRONT CAVE SING SOFT MUSIC TIGER EAT FOOD
THANK FOOD MUSIC LADY CUT-OFF WHISKER RUN-DOWN HILL TIGER
LONELY SAD AGAIN

Procedure

The spontaneous story telling data for the Tiger's Whiskers were obtained in exactly the same manner as those of the Fox and the Bear. Subjects were run through the four subtasks individually (-reading +prompt, +reading 1 +prompt, +reading 2 +prompt, -reading -prompt) and their productions were recorded on an AKAI 4000DB taperecorder.

The reading data were obtained by reanalyzing the results from an earlier experiment dealing with the prosodic structures of sentences. In this study, which took place some 3 months prior to the telling and parsing tasks, the subjects were seated in a studio and asked to read both stories to themselves. Then they read each story aloud at three different rates: a normal reading rate, a rate that they estimated to be half their normal reading rate, and a rate that they estimated to be a quarter their normal rate. They gave two readings at each rate before moving down to the next rate. The readings were recorded on an AKAI 4000DB taperecorder.

Finally, the parsing data were obtained by asking the same subjects to estimate the importance of the sentence breaks in the two stories. After reading each story to themselves a number of times, the subjects were asked to find the most important break in the story and to give it a number from a 1–10 scale in which 1 represented an unimportant break and 10 an important break. They were then asked to take each subpart and locate within it the most important break and give it a rating. They were to continue this until every sentence break had a rating.

Data Analysis

The spontaneous productions of the Tiger's Whiskers were analyzed in exactly the same way as those of the Fox and the Bear story. Transcriptions and pen-recordings were made and pause durations were calculated at each of the 13 sentence breaks. As one of the subjects systematically left out a sentence and rearranged several others, the spontaneous pause durations were pooled across eight spontaneous productions (four by each of two subjects).

The readings at various rates were analyzed, as were the spontaneous productions, by means of a pen-recorder (Gould Brush 220). Reading pauses at each of the sentence breaks were located and measured in the readings and then tabulated. In the Fox and the Bear story, a particular sentence break could recieve as many as 18 pauses, as each of the three subjects read the story six times in all. Thus, pauses at each break were summed and divided by 18. (If a subject did not pause at that break in one of the readings, a zero pause duration was tabulated). As for the Tiger's Whiskers, means were computed in a similar manner, but on 12 readings only (the subject who had had problems with the spontaneous productions was omitted from the reading analysis).

Finally, the parsing data was averaged over three subjects for the Fox and the Bear story and over only two subjects for the Tiger's Whiskers. The mean parsing values at each break were therefore based on three values in the first story and two values in the second story.

RESULTS AND DISCUSSION

Table 2 presents the data obtained from the three tasks—telling, reading, and parsing—and the two stories. (The mean pause durations for the telling of the Fox and the Bear story come from Experiment 1). Complexity indices obtained from the theoretical analysis of the stories have been added to the table.

In our discussion of these results, we will first study the telling of the Tiger's Whiskers and then the reading and the parsing of the two stories.

Story Telling

In Figure 3, we present our analysis of the Tiger's Whiskers based on Lehnert's model.

In Lehnert's terms, the story is an example of the large plot configuration "Intentional Problem Resolution." This plot configuration is made up of three components: a Problem, a Success, and a Resolution. Let us con-

TABLE II

Average pausing values obtained from spontaneous telling and reading, and average parsing values obtained from parsing. These values were calculated at each of the sentence breaks in two different stories: The Fox and the Bear and The Tiger's Whiskers. Complexity indices calculated from the theoretical analysis of the stories were also added.

The Fox and the Bear

Sentence Number	*Complexity Indices*	*Telling Mean Pause Duration*	*Reading Mean Pause Duration*	*Parsing Mean Parsing Values*
1	1	.65	.80	4
2	0	.66	.93	5.3
3	1	.72	1.02	5.3
4	1	.62	.87	2
5	4	.91	1.06	6.7
6	3	.78	1.06	4
7	0	.47	.88	3.3
8	6	.87	1.01	6.3
9	0	.60	1.04	4.7
10	4	.78	.94	6.3
11	0	.50	1.00	7.3
12	0	.52	1.01	2

The Tiger's Whiskers

Sentence Number	*Complexity Indices*	*Telling Mean Pause Duration*	*Reading Mean Pause Duration*	*Parsing Mean Parsing Values*
1	0	.51	.94	3
2	4	.94	1.03	8
3	0	.37	.98	4
4	1	.68	1.09	7
5	1	.68	.88	3
6	2	.68	1.00	9
7	3	.74	.94	9
8	0	.39	.92	5
9	0	.60	.86	2
10	2	.70	.94	8
11	0	.49	.84	5
12	3	.71	.98	8
13	0	.36	.71	3

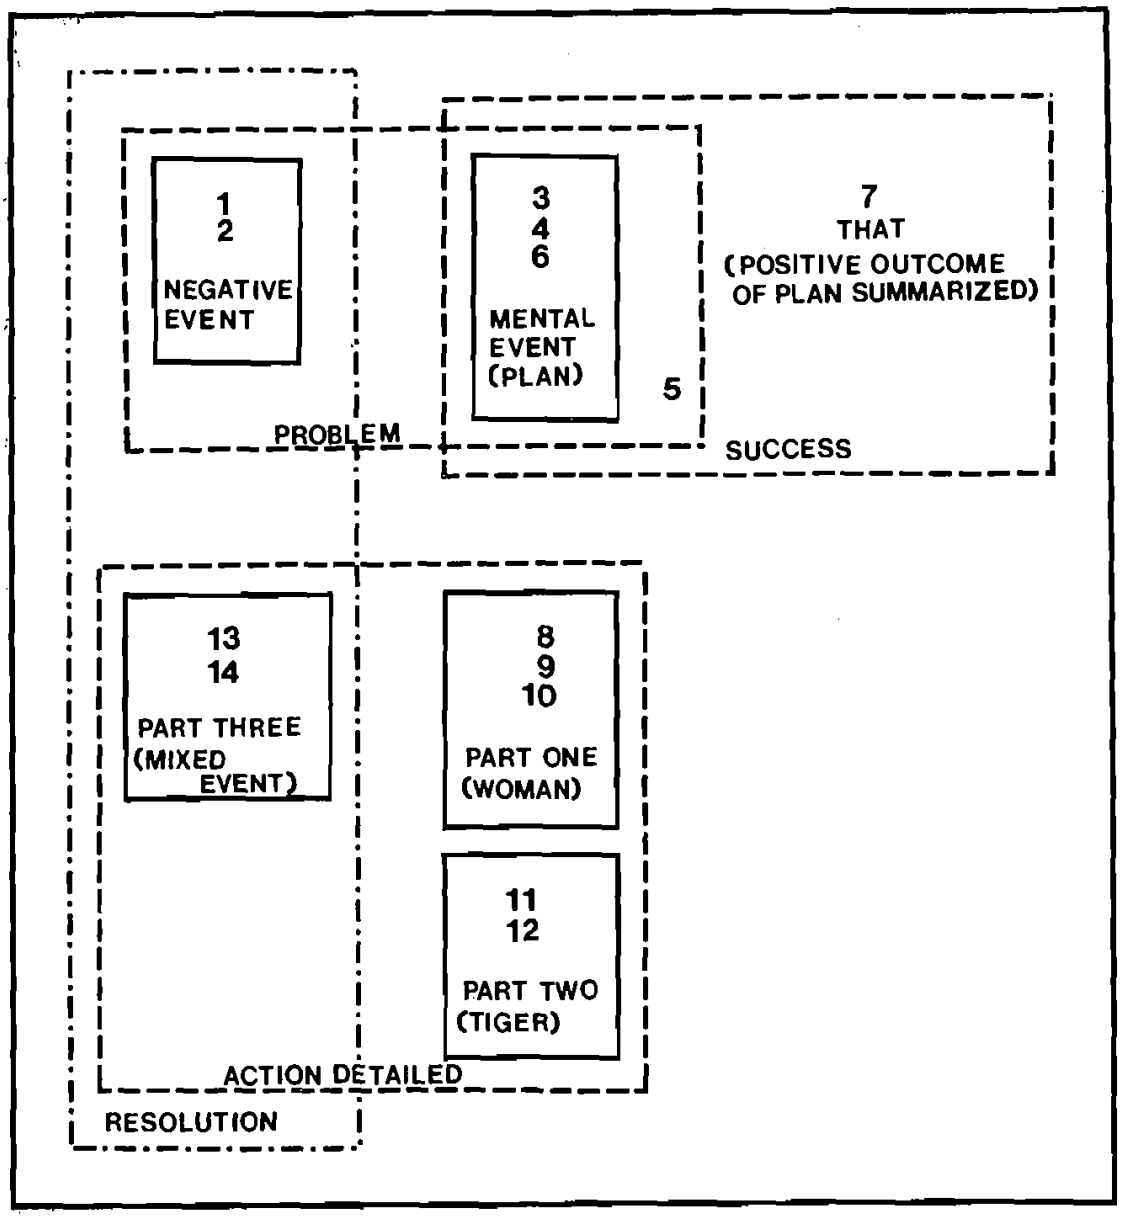

Figure 3. Diagram of theoretical narrative analysis of the Tiger's Whiskers in terms of our adaptation of Lehnert (1981). Numbers represent the sentence numbers in the text of the story. Constituents are enclosed within boxes.

sider each of these three plot units in turn. The Problem plot unit is made up of a negative affect state (the husband being sick and needing a tiger's whisker in sentences 1 and 2) that motivates a mental event directed toward the solution of the negative state (the plan to remove the negative state of affairs, represented in sentences 3, 4, and 6). In Lehnert's system there is also another way to look at this Problem unit: Sentences 1 and 2 could also be taken to be, and in fact also represent, not only a negative event, but a mental event (need) that motivates another mental event, the plan in 3, 4, and 6. In this case, 1–6 make up a plot unit which Lehnert calls Motivation, i.e., the Motivation for the rest of the action that follows. We label 1–6

"Problem," but it could also bear the label "Motivation". In addition, we should also point out that sentences 1 and 2 represent and encode exposition information—we know this by the durative main verbs and the formula, "once there was..."

The Success plot unit is made up of the mental events in 3–6 (actually these sentences elaborate the time stages and development of a mental event) which is actualized in the positive affect state summarized in Sentence 7 (the accomplishment of the woman's plan). This positive outcome to the mental planning in 3–6 is only summarized in 7; it is then actually spelled out in detail in the rest of the narrative (8–14).

In Lehnert's system, the rest of the narrative is made up of the events involving cross-character links. Sentences 8–10 (involving the woman) and sentences 11–12 (involving the tiger) constitute what Lehnert calls a "Shared Event" (an event which two characters are affected by in the same way), while sentence 13 (involving the woman) and sentence 14 (involving the tiger) together constitute what she calls a "Mixed Event" (the same event is experienced differently by both characters). These three units, 8–10 (woman), 11–12 (tiger), and 13–14 (both) also constitute the three parts of the plan which is specified in the Mental Event at 3–6 and whose positive outcome is summarized in sentence 7. Thus, we label these units "Part One" (bringing food and music to the tiger), "Part Two" (the tiger being nice to the woman in return), and "Part Three" (the woman getting the tiger's whisker). Thus, all of the sentences 8–14 are the actional realization of the plan at 3–6 (a plan whose positive outcome is summarized at sentence 7), but spelled out in detail only in 8–14. In Lehnert's terms, sentences 13–14 terminate the Negative Event at 1–2 (i.e., 13–14 "supplant" or "displace" the affective impact of the event at 1–2). The reader should note that 13 gives a positive outcome (for the woman) and 14 a negative effect of that outcome (for the tiger).

We treat Sentence 5 as a piece of setting information for the mental event in 3–6—it is a way of letting the reader know a piece of knowledge necessary to make sense of the woman's reasoning. Thus, we put it outside this unit, as we did for the setting information in the Fox and the Bear (with respect to the narrative as a whole).

Sentence 7 shows an interesting device: "did...that" refers anaphorically back to the plan in 3–6 and cataphorically forward to the realization of the plan in 8–14. We represent this in the diagram by letting 7 serve as the statement that the plan is ultimately actualized (thus making up a Success unit with 3–6), but letting the spelling out of the plan in detail in 8–14 give issue to the actual resolution of the negative state of affairs at 1–2 (thus making up a Resolution with 1–2).

We again measure the complexity of each sentence boundary using the Complexity Index formulated in our discussion of the Fox and the Bear story. For example, the value after Sentence 2 is 4 because four boxes inter-

vene between Sentence 2 and the sentence immediately following it in the next level constituent (Sentence 3), and no sentences intervene temporally in the text between Sentences 2 and 3; the value after 12 is 3 because three boxes intervene between 12 and 13 and the sentence immediately following it in the next level constituent (the whole constituent made up of 8–14), while no sentences temporally intervene in the text; finally, the value following 7 is 4 because four boxes intervene between 7 and the sentence immediately following it in the next level constituent (the whole narrative, i.e., Sentence 8), and no sentences temporally intervene in the text. The values of each boundary are listed in Table 2 above.

The correlation between the 13 mean pause values obtained at the 13 sentence breaks and the corresponding complexity indices computed from the theoretical analysis of the story is 0.88 ($t<0.01$). This result brings converging evidence to the fact that spontaneous pausing in the telling of stories reflects the narrative structure of these stories. In addition, this result confirms that narrative analysis, in terms of basic plot motifs, is a fruitful way to study narrative/story structure.

Story Reading and Parsing

A second aim of this experiment was to determine whether other experimental tasks would also reflect the narrative structure of stories. Two such tasks—reading and parsing—were used and the data obtained from these tasks (mean pausing and parsing values) were presented in Table 2 above.

Mean pausing values are correlated 0.36 (ns) with the Fox and the Bear complexity indices and 0.51 ($t<0.05$) with the Tiger's Whiskers indices. From this we can conclude that pausing values obtained from reading (as opposed to values obtained from spontaneous telling) do not reflect the narrative structure of stories very well, although the coefficients are positive. This may be because subjects do not actually structure the story as a whole during oral reading, but go from sentence to sentence. Or it may well be that sentence length alone, or in interaction with story structure, determines the value of the pause duration following each sentence.

A more interesting picture emerges from the parsing data. The Fox and the Bear parsing values are correlated 0.47 (ns) with the complexity indices, but those for the Tiger's Whiskers are correlated 0.86 ($t<0.01$) with their respective complexity indices. This interesting difference (which is indicated by the oral reading correlations) may be due to the fact that the Tiger's Whiskers contains clear linguistic clues to its narrative structure whereas the Fox and the Bear does not.

The story starts with a two sentence statement of a NEED/PROBLEM. It then falls into two clear parts based on the distribution of the main verbs of the sentences. The first half has only mental state verbs: "thought,"

"decided," "knew," "thought." The second half has only physical action verbs: "went," "put," "sang," "come out," "walked," "cut off." Sentence 7 stands right in the middle, anaphorically attached to the first half ("that" refers to the content of the mental states), but spelled out in detail by the attached subnarrative contained in the second half. In the second half of the narrative (sentences 8–14), we have an alternation of grammatical subjects that clearly keys the smaller units within this second half: "the woman/she" (sentences 8, 9, 10), "the tiger/he" ((11, 12), "the lady" (13), "the tiger" (14). The narrative adds an emotion verb "felt" in the last sentence, separate from both the mental verbs of the first half of the narrative and the physical verbs of the second half. Furthermore, the two halves of the narrative (after the statement of NEED/PROBLEM) are quite symmetrical. The issue of the woman's thought in Sentence 3 is the decision to use a trick in Sentence 4. The trick is spelled out in Sentence 6 and has three parts: a) bringing food and music to the tiger, b) the tiger being nice to the woman, c) the woman getting the tiger's whisker. These three parts are exactly realized as the three substructures of the second half of the narrative (8–10, 11–12, 13–14). Thus, this story is linguistically transparent in a way that the Fox and the Bear is not.

To conclude this section, it would appear that reading is not an appropriate task to use in order to obtain the performance structure of a narrative. We should note that at a lower level, however—that of sentence structure—reading data are well correlated with the prosodic structure (Gee & Grosjean, 1983). Parsing, on the other hand, may prove to be a useful task, but only if the structure is marked by superficial linguistic cues. As for spontaneous story telling, it is the most successful of the three tasks and the pauses that are obtained from it are a good reflection of the structure of narrative.

GENERAL DISCUSSION

A narrative is a sequence of sentences that flow along in time one after the other. But the sentences in a narrative contract with each other more than merely temporal relations. The fact that sentences or sequences of sentences separated in time can be conceptually and thematically related to each other, and that they can be simultaneously related in complex ways to sentences both preceding and following them, gives narratives what has been called "spatial" structure, as well as temporal order. By spatial structure we mean a complex network of relations that exist simultaneously in a mental, multidimensional space. It is aspects of this spatial structure that are represented in our narrative diagrams. Such complex relations go beyond the power of a simple hierarchical tree structure of the sort we use to represent sentential,

syntactic structure. What we have shown in this study is that pausing above the level of the sentence reflects these narrative relations. In addition, it appears to be capable of reflecting rather complex and intricate, though fundamentally basic, aspects of narrative structure.

Once we are convinced that pausing is a guide to narrative structure, it can be used to test alternative models of narrative structure, to develop new models, and to investigate particular questions in narrative theory. The correlations between our theoretical models and the pausing data are not perfect. This is, of course, due in part to "noise" in the data. But it is undoubtedly also due to the fact that the model of narrative structure we have used, as any such model, is only a rough approximation. For example, in the Fox and the Bear story, the boundary after Sentence 5 is somewhat undervalued by our complexity index. This boundary has a somewhat longer pause than we predict, probably because the unit Motivation/Instigation (2, 3, 5) also plays the role of the Introduction to and Motivation for the rest of the story as a whole. Or, to take another example, the pause after Sentence 2 in the Tiger's Whiskers is so large (we predict it to be the largest, but it could be given an even larger value) probably because the unit in Sentences 1–2 plays three different roles, i.e., a Negative Event (sickness) making up a Problem unit with 3, 4, and 6; a Mental Event (need) making up a Motivation unit with 3, 4, and 6; and Expository information for the rest of the story. This type of comparison between the data and the theoretical predictions should lead to a number of important questions concerning narrative structure and what makes for complexity in narratives.

Although our result may seem merely "methodological", it, in fact, has a number of implications for research in many areas. Literary critics have for some time pointed out the importance of "spatial form" to narrative and the role it plays in much modern poetry and prose (see Frank, 1978; Jakobson, 1960; Smitten & Daghistany, 1981). We have demonstrated that aspects of such form are empirically discoverable, somewhat ironically, in the temporal flow of the narrative.

A long line of research in the ethnography of speaking has indicated the importance, in oral narratives and oral communication generally, of temporal variables (pausing and other aspects of prosody) in signalling thematic structure (such as the boundary between episodes) and keying the interpretive work of the hearer (see, among many others, Gumperz, 1982; Michaels, 1981; Scollon & Scollon, 1979, 1981; Tedlock, 1972). Many of the analyses this literature has offered are impressive, but impressionistic. We take our work here to support the claims of this research as to the importance of pausing (and prosody generally) in signalling thematic and discourse structure, and to offer a method of empirical investigation of the role of pausing in discourse.

Both linguists and researchers in artificial intelligence have sought for a theory of narrative structure as a necessary step after the construction of a

theory of the sentence. Such a theory has been difficult to develop and even harder to test, especially since we do not, for the most part, have intuitions of grammaticality to fall back on once we leave the level of the sentence. Nonetheless, theories of narrative are liable to yield substantive insight into the structure and structuring power of the human mind. We have tried to show that a rich theory of narrative structure, one that incorporates intrinsically meaningful labels for narrative units and goes beyond simple hierarchical representations, can be grounded in empirical data—data which can, in addition, suggest various revisions in the theory.

The fact that pausing indicates narrative structure both in English and in the ASL (Gee & Kegl, in press)—language modalities which physically realize "pausing" in quite different ways—shows the importance of prosody to the structure of narrative generally, and dictates the need for a theory of prosody abstract enough to accommodate all languages and language modalities. It looks as if we are tapping into a basic principle of the rhythmical organization of experience by the human mind. It would be interesting to investigate the nature of pausing in other sorts of connected discourse, for example in expository texts (where pausing should reflect the structure of argumentation) and connected conversation. We suspect that here too, pausing will reflect aspects of structure.

If pausing (and other aspects of prosody) is the key to narrative structure, then investigations of the structure of language beyond the sentence are going to require the mutual cooperation of linguists, psycholinguists, and theoreticians of narrative, whether they be in literary criticism or artificial intelligence. But the question arises as to why pausing should play this role. Scollon and Scollon (1979) suggest that pausing is not really a marker of discourse structure per se, but rather it reflects the flow of attention in discourse. As the narrator reaches larger boundaries (discontinuities) in the structure of the narrative, his attention is raised from the narrative line proper to the audience, and his degree of care in speaking increases. This shift of focus, in turn, affects the prosodic structure of the narrative, causing careful speech and pausing to occur at larger boundaries in the narrative. Thus, the Scollons raise the issue of how we can tease apart the abstract structuring of the content of the narrative and the actual performance. Obviously the structure of an oral text as recorded is highly situational, but we can still ask to what extent the cognitive structure of the text is independent of any particular performance. The method we have introduced here of pooling across several performances and across several subjects helps, we believe, to get at the underlying cognitive structure of the text, and yet does this through performance variables.

Dell Hymes (1981) has recently suggested that the "richness of syntax which linguistics finds in every normal child may be accompanied by a richness of narrative organization." This narrative organization constitutes a kind of "rhetoric of action" in that it embodies implicit cultural schemas

for the organization of experience. It is to be hoped that experimental data of the sort we have presented will give us some access into the nature of the unconscious structures by which we order our world and our lives.

REFERENCES

Black, J. B., & Wilensky, R. (1979). An evaluation of story grammars. *Cognitive Science, 3,* 213–230.

Frank, J. (1968). *The widening gyre: Crisis and mastery in modern literature.* Bloomington, IN: Indiana University Press.

Gee, J. P., & Grosjean, F. (1983). Performance structures: A linguistic and psycholinguistic appraisal. *Cognitive Psychology, 15,* 411–458.

Gee, J. P., & Kegl, J. A. (1983). Narrative/story structure, pausing, and ASL. *Discourse Processes,* 243–258.

Glenn, C. G. (1978). The role of episodic structure and of story length on children's recall of single stories. *Journal of Verbal Learning and Verbal Behavior, 17,* 229–247.

Grosjean, F., Grosjean, L., & Lane, H. (1979). The patterns of silence: Performance structures in sentence production. *Cognitive Psychology, 11,* 58–81.

Gumperz, J. J. (1982). *Discourse strategies.* Cambridge: Cambridge University Press.

Hockett, C. F. (1958). *A course in modern linguistics.* New York: Macmillan.

Hymes, D. (1982). Narrative form as a "grammar" of experience: Native Americans and a glimpse of English. *Journal of Education, 2,* 121–142.

Jakboson, R. (1960). Concluding statement: Linguistics and poetics. In T. A. Sebeak (Ed.), *Style in language.* Cambridge, MA: MIT Press.

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

Kintsch, W., & van Dijk, T. A. (1975). Comment on rappelle et on resume des histoires. *Languages, 9,* 98–116.

Kintsch, W., & Greene, E. (1978). The role of culture-specific schemata in the comprehension and recall of stories. *Discourse Processes, 1,* 1–13.

Labov, W. (1972). *Language in the inner city.* Philadelphia, PA: University of Pennsylvania Press.

Labov, W., & Waletzky, J. (1967). Narrative analysis: Oral versions of personal experience. In J. Helm (Ed.), *Essays on the verbal and visual arts.*

Lehiste, I. (1975). The phonetic structure of paragraphs. In A. Cohen & S. G. Nooteboom (Eds.), *Structure and process in speech perception.* New York: Springer-Verlag.

Lehnert, W. G. (1981). Plot units and narrative summarizations. *Cognitive Science, 4,* 293–331.

Lehnert, W. G., Black, J., & Reiser, B. (1981). Summarizing narratives. *Proceedings of the 7th International Joint Conference on Artificial Intelligence.* Vancouver, BC, Canada.

Mandler, J. M., & Johnson, N. S. (1977). Remembrance of things parsed: Story structure and recall. *Cognitive Psychology, 9,* 111–151.

Michaels, S. (1981). "Sharing time": Children's narrative styles and differential access to literacy. *Language in Society, 10,* 423–442.

Perfetti, C. A. (Ed.) (1982). Special issue. *Discourse Processes, 5,* (3–4).

Propp, V. (1968). *Morphology of the folktale.* Austin, TX: University of Texas Press. (original work published 1928)

Reiser, B. R., Lehnert, W. G., & Black, J. B. (1981). Recognizing thematic units in narratives. *Proceedings of the Third Annual Conference of the Cognitive Science Society.* Berkeley.

Rumelhart, D. E. (1975). Notes on a schema for stories. In D. G. Bobrow & A. M. Collins, (Eds.), *Representation and understanding.* New York: Academic Press.

Scholes, R., & Kellogg, R. (1966). *The nature of narrative.* New York: Oxford University Press.

Scollon, R., & Scollon, S. B. K. (1979). *Linguistic convergence.* New York: Academic Press.

Scollon, R., & Scollon, S. B. K. (1981). *Narrative, literacy and face in interethnic communication.* Norwood, NJ: Ablex.

Smitten, J. R., & Daghistany, A. (Eds.) (1981). *Spatial form in narrative.* Ithaca, NY: Cornell University Press.

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). Norwood, NJ: Ablex.

Sternberg, M. (1978). *Expository modes in narrative discourse.* Baltimore, MD: Johns Hopkins Press.

Tedlock, D. (1972). *Finding the center: Narrative poetry of the Zuni Indians.* New York: Dial Press.

Thorndyke, P. W. (1977). Cognitive structures in comprehension and memory of narrative discourse. *Cognitive Psychology, 9,* 77–110.

Wells, R. S. (1947). Immediate constituents. *Language, 23.* (Reprinted in M. Joos (Ed.), *Readings in linguistics.* New York: American Council of Learned Societies, 1958).