COGNITIVE PSYCHOLOGY 12, 447–472 (1980)

A Partial Resolution of the Paradox of Interference: The Role of Integrating Knowledge

LYNNE M. REDER AND JOHN R. ANDERSON
Carnegie-Mellon University

It has been noted that models of memory that posit retrieval interference imply that the more one knows about a topic, the harder it is to retrieve any one of these facts. Smith, Adams, and Schorr (*Cognitive Psychology*, 1978, 10, 438–464) regard this to be a paradox and postulate that people use world knowledge to integrate various facts about a concept and thereby avoid interference. Exploring this issue further in two experiments we discovered that integration of facts alleviates interference only when a person can perform his memory task by simply making a consistency judgment and can avoid the need to retrieve a specific fact. When foils force subjects to retrieve the specific assertion, the interference occurs among integrated facts as among unrelated facts. It appears that, when possible, subjects will judge whether they have seen a fact simply by judging if it is related to (consistent with) a theme they have studied. In other words, people judge themes rather than facts. Consistent with this interpretation, we found interference among themes; that is, the more themes were associated with a concept, the greater the interference.

The debate concerning the cause of forgetting is an old one. Some assert that erasure is necessary—we would very rapidly run out of storage space for new input if we did not throw out old information. Others assert that we have not really "lost" the old information, rather, it has merely become inaccessible for the time being. The most popular analysis of inaccessibility is interference theory. There has been a wide proliferation of ideas about the mechanisms of interference (see Crowder, 1976, for a review). A still popular view of interference is what has been called stimulus-specific interference, which asserts that the more facts one learns about a concept the harder it is to retrieve any particular fact from the concept. Traditionally evidence for this view of interference has accrued using probability of recall as the dependent measure. Our concern will be with the more recent studies of the phenomenon within the information-processing framework using reaction time as the dependent measure (e.g., Anderson, 1974, 1976; Anderson & Bower, 1973; Hayes-Roth, 1977; King & Anderson, 1976; Lewis & Anderson, 1976; Moeser, 1977, 1979; Smith, Adams, & Schorr, 1978; Thorndyke & Bower, 1974).

This research was supported by an NIMH postdoctoral fellowship to Lynne Reder and by NSF grants BNS-76-00959 and BNS-78-17463 to John Anderson. We are extremely grateful to Takashi Iwasawa for his spectacular work in programming, running, and analyzing these experiments. Correspondence concerning this paper should be sent to Lynne M. Reder, Department of Psychology, Carnegie-Mellon University, Pittsburgh, PA 15213.

0010-0285/80/040447-26$05.00/0

These reaction time studies have shown that subjects are slower to recognize a studied fact when other studied facts share some of the same concepts. For example, subjects are slower to recognize the studied fact "The doctor hated the lawyer" if they have also studied "The doctor vacationed in the Carribean" and "The lawyer loved the actress" than if no other experimental facts share concepts with the probe. This particular result has been named the *fan effect* because latencies are found to increase with the number of propositions *fanning* out of a given concept node in its propositional network representation (see Anderson, 1976). Subjects are also slower to reject unstudied probes as a function of the fan off their concepts.

In a recent paper Smith, Adams, and Schorr (1978) have discussed the apparent paradox inherent in this type of interference-theory. The paradox can be explained as follows. The more one knows about a particular topic, the more potential interference to that topic. This means the more one knows about a topic, the slower one should be to answer a specific question on that topic. This conclusion, they claimed, runs counter to the intuition that the more one knows about an area, the more questions one can answer on the topic and the faster one can answer them. Although Smith et al. did not exactly resolve the apparent paradox, they tried to show that the effect of fan can be drastically attenuated when the facts to be learned about a particular individual can be thought of as forming an integrated unit. (Examples are given below.)

We do not think that, upon careful introspection, people always agree with the intuitions expressed by Smith et al. While in some circumstances it does seem easier to remember new information the more one knows about a domain, in other circumstances just the opposite seems to be the case. A good case in point, where most people agree, involves use of computer text-editors. Initially one learns an editor slowly. Learning new editors is much easier, but if one still uses several editors within the same month, it is difficult to keep the commands straight. Thus, it is interfering to have to use several text editors at once. So, it is wrong to suppose that expertise necessarily implies the disappearance of interference. Rather it is necessary to come to an understanding of why some kinds of material gather resistance to interference with expertise and why some do not.

The present paper is intended to further investigate the nature of interference in memory and the facilitation of memory search reported by Smith et al. due to integrated or thematically related facts. Although it seems reasonable to believe that related facts should somehow make storage and retrieval easier, the exact mechanisms that enable this to occur need to be worked out. Hopefully, this set of studies will shed light on these mechanisms. Before describing these experiments, it will be useful to review in more detail the conception of fact retrieval developed by Anderson and Bower (1973) and Anderson (1976). The paradigm used by Smith et al. and the one that will be used here are both based on their work. By analyzing the modifications Smith et al. made to the standard design, we can see how the old conception of fact retrieval relates to their findings.

AN INFORMATION PROCESSING ANALYSIS OF FACT RETRIEVAL

In a prototypical fan study, subjects must learn a number of unrelated facts about a number of individuals. For example, *Marty broke the bottle, Marty did not delay the trip, Steve cooked spaghetti, Steve liked to bowl, Steve rescued the kitten*. After committing all facts to memory, subjects are presented, one at a time, with studied facts and new facts like "Steve broke the bottle." Reaction times are taken to correctly accept or reject test probes. A schematic network representation for the first studied sentence is given in Fig. 1. According to the HAM theory of Anderson and Bower and the ACT theory of Anderson (1976), a representation analogous to this one is stored in memory by the subject. In these theories, concepts are represented as terminal nodes in the semantic network. A proposition consists of a set of concepts connected together by arcs. When a subject is asked to verify if he has studied a particular sentence, the probe is converted into a propositional structure and memory is searched trying to find a match to the structure of the probe. Figure 2a illustrates how multiple propositions might be stored in memory. The processes used in searching are assumed to be of limited capacity. Search is controlled by a spreading activation process according to the ACT theory. This spreading activation conception is similar to, but not identical with, the ideas of Quillian (1968), Collins and Quillian (1972), and Collins and Loftus (1975). Concepts in memory are activated by search energy reaching those concepts. Search begins at each concept mentioned

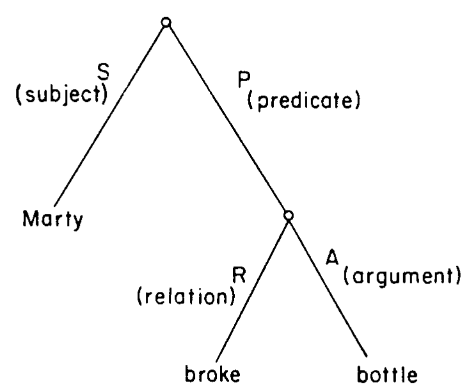

FIG. 1. Schematic memory representation of the sentence "Marty broke the bottle."

in the probe. Search travels out from these concepts along arcs to other concepts in memory. The rate at which other concepts are activated is proportional to the number of arcs emanating from a given concept. That is, a fixed amount of search energy is assumed to be present at each node corresponding to a concept in the probe. If five arcs are attached to a concept in memory, the search energy going down any one of these arcs will be one-fifth what it would have been had there been only one fact attached to the concept. (Actually the assumptions are slightly more complicated, involving the notion of arc strength—consult Anderson, 1976 for more details).

SMITH ET AL.'S EXPERIMENTS

The experiments of Smith et al. did not find much interference when the facts associated with a character could be thought of as belonging to the same theme or script. In one of their experiments, subjects first learned sets of unrelated facts about several individuals such as "Marty broke the bottle" and "Marty did not delay the trip." Later they might learn an integrating fact such as "Marty christened the ship." Smith et al. contrasted subjects who learned the third fact that integrated the preceeding pair with subjects who just learned the pair. Subjects were almost as fast to verify one of these three facts as they had been to verify one of the two original facts. That is, an integrated triple had essentially no more interference than an unintegrated pair.

The model Smith et al. proposed to account for the different pattern of data assumes a very different memory organization and search strategy than ACT. They assume that one stores the facts from an integrated triple in a "script" (see Schank & Abelson, 1977). People store the name of the script with the character. At test, it may take a long time to access the script name associated with the probed character, however once the script instantiation associated with the character is retrieved, the script can be very rapidly inspected for tagged propositions to see if a particular script fact was stored. Smith et al. assume that time to inspect tagged facts in the script is almost nil. These assumptions are consistent with response times being insensitive to the number of related facts. On the other hand, Smith et al. postulate a totally different system to explain the fan effect with unintegrated material. As an alternative, we offer one model that involves only a slight extension of the ACT theory and accounts for data from both types of material.

Figure 2 sketches the two representations we consider viable to account for Smith et al.'s data. Figure 2a is essentially the same representation that Anderson has proposed in the past. Figure 2b suggests that subjects organize facts into subnodes if they form a coherent theme (this idea has been suggested in Anderson, 1976, and in Anderson & Paulson, 1978).

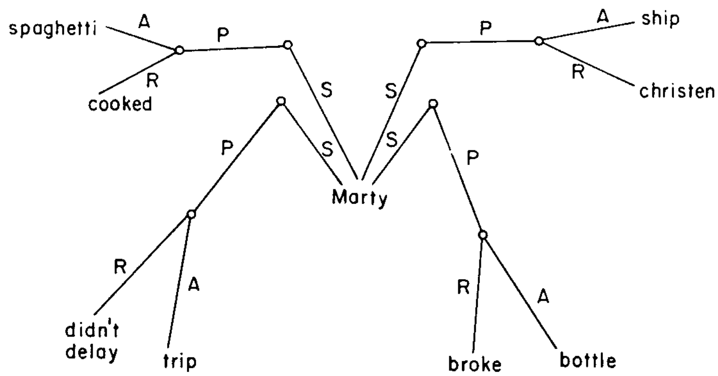

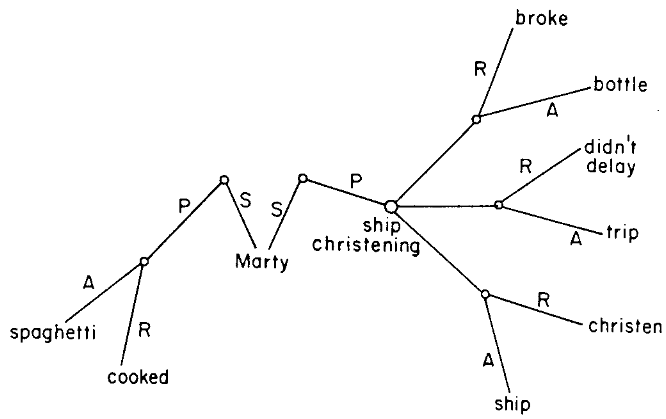

FIG. 2. Two possible memory representations for related facts about Marty. (a) ACT representation; (b) modified ACT representation using subnodes.

The retrieval scheme thought to operate on the representation in Fig. 2a is as follows: Subjects search memory from each concept node in the probe. Activation travels down connected arcs. As soon as a stored proposition is activated that is *consistent* with the topic associated with the probe, the subject responds "old." In this case, it does not matter how many facts fan off a "character node" if the facts are all related. Since any related fact will do, it should not take any longer to respond "old" with three relevant facts than with one.[1] If the test probe were *Marty christened the ship,* but *Marty broke the bottle* were activated first, the subject still responds positively. For the most part, this strategy would have been effective in Smith et al.'s experiments. Evidence for such a strategy in a

[1] Assume a horse race analogy for times for completion: Average activation rates will be slower, but the chances for a very early completion time go up.

somewhat different paradigm was provided in Reder 1976, 1979). In those experiments subjects were found to search for consistent information, not specific propositions. We will call this explanation the *consistency model.*

The retrieval and judgment scheme assumed to operate on the representation in Fig. 2b is somewhat different. Search again spreads out along all arcs emanating from a concept node. In this case however, the amount of fan would be determined by the number of different *sets* of thematically related facts, not the number of facts per se. Once search has reached a theme or script node consistent with the probe, search stops and the subject responds "old." For example, suppose "bowling" and "dining out" are themes associated with Steve, and the test probe is "Steve got a spare on a split." When search activation reaches Steve bowling, search stops. One would not expect an effect of number of thematically related facts in this case either, because the amount of fan off the subnode should not affect time to reach the subnode from the character node (consult Fig. 2b). One does not bother to ascertain whether one of the facts attached to the subnode is the correct one. Again this strategy would suffice in Smith et al.'s experiments. We will call this explanation the *subnode model.*

Note that both of these representation–search strategies are merely modifications of the original process assumed to account for the fan effect. The principal difference is a change in the subject's *criterion* for responding "old," not in the basic mechanisms underlying retrieval. That is, subjects do not bother retrieving the exact proposition as they had in other experiments, since any consistent proposition or the consistent subnode will satisfy the constraints. These models do assume, however, that the subject has some way to assess consistency. We will propose possible mechanisms for this in the discussion section.

The two models presented here imply that if subjects cannot use the more lenient criterion for deciding whether a proposition was studied, then the fan effect will re-emerge. We can force subjects to use a strict criterion by using foils (nonstudied sentences) which come from the same theme as the studied sentences. Consider the first strategy we proposed, the consistency model. If subjects were to search until they activate a proposition consistent with the probe and then respond "old," they would respond erroneously to all thematically related foils. To avoid making errors, the subject would have to resort once more to inspecting all propositions. The same shift in criterion would be required for the second model proposed. To avoid making errors to thematically related foils, subjects would have to inspect which propositions were attached to the theme node. In that case, the fan off the subnode would be a relevant variable.

Smith et al., in order to test this explanation, conducted an experiment which included *thematically-related foils* in addition to their *re-paired* *foils*. Re-paired foils used predicates that subjects had studied with a different character. For example, if the subject had studied "Marty broke the bottle" and "Steve drove the car," the re-paired foil might be "Marty drove the car." This was the type of foil used in the earlier Smith et al. studies. The thematically related foils were created by taking the original facts and inserting a new word that maintained the theme. Since Marty's theme was ship christening, the thematic foil was: "Marty broke the *champagne* bottle." The results from this experiment did not produce much more of a fan effect. Smith et al. concluded that a model such as those we proposed could not explain the data.

It seemed to us that Smith et al.'s control experiment was not adequate to rule out the models we suggested to explain the attenuated fan effect. The foils used in their thematically related condition had the property of including a new, low frequency word like "champagne." This insertion could allow the subject to reject the foils on the basis of lexical familiarity, even if the statement seemed very plausible.[2] We decided that the models proposed above still needed further experimental test before they could be conclusively rejected.

EXPERIMENT 1

Experiment 1 had two goals: First, we wished to determine whether an ACT-like model with a new response criterion could explain Smith et al.'s data. We attempted to do this by manipulating whether the foils were thematically related or unrelated while holding familiarity of the predicates constant. If either of the two models above is correct, we should obtain the traditional fan effect in the blocks with thematic foils and replicate the Smith et al. result with unrelated foils.

The second goal was to try to discriminate between the two proposed models, assuming they are still viable. This was done by associating more than one theme or script with a given character. By orthogonally manipulating the number of facts associated with each of two themes learned about a given character, we can look at the effect of fan in irrelevant themes and the effect of fan in relevant themes. *Irrelevant fan* is defined as the number of facts learned about the theme that is not queried by the particular probe. Imagine that a subject studied three facts about John skiing and two facts about John going to a wedding. If the test probe was about John and related to wedding, then the relevant fan would be two and the irrelevant fan would be three. The prediction for the effect of *relevant* fan is the same for the two models. However, the predictions vary for the two models with regard to the irrelevant fan manipulations.

[2] Subjects undoubtedly use more than one procedure to reject propositions. At a short delay, lexical familiarity could easily be one of these procedures. At a longer delay, Smith et al.'s related foils would have provided a fairer control.

The subnode model would predict that subjects will be slower to verify a fact about a character if there are other facts about another theme associated with that character, but that the number of facts about the other theme has no effect. This is because these facts would be attached to the theme subnode and the interference is caused by the number of theme nodes, not the number of facts studied.

The consistency model, on the other hand, has all facts fanning out from the character node. The more irrelevant facts the subject knows about a character, the longer it will take to retrieve any fact. Therefore, the consistency model predicts that three thematically related facts, though irrelevant to the probe, cause more interference than just one irrelevant fact.

Method

Procedural Overview

There were three phases to the experiment. In the first phase, subjects learned sets of facts about various characters. All facts about a given character were presented at the same time for the subject to study. In the second phase, subjects had to demonstrate that they knew which facts were associated with each character. This was accomplished by a procedure in which subjects alternated between trying to recall the sentences they had studied and giving further study to sentences they could not yet recall.

Phase three was the critical phase in which reaction times were collected to recognize whether they had studied various test sentences. The foils were of two types, separated into different test blocks. One type of foil was the unrelated foil. This foil corresponds to Smith et al.'s "re-paired" foils in that both the individual and the predicate have been presented for study, but not together. For example, the subject might have studied "Jim ordered a steak" and "Bill asked the bus driver for a transfer." An unrelated foil might be "Jim asked the bus driver for a transfer." The other type of foil is called the thematically related foil. Unlike Smith et al.'s thematically related foils, these foils were also re-paired. That is, if the subject had studied "Susan waited at the bus-stop," as well as the two facts above, a related foil would be "Bill waited at the bus-stop." (This design requires that the same theme be paired with a number of individuals).

Design and Materials

Table 1 illustrates some of the material that a subject might see. There are three orthogonal factors that define the condition within which a particular probe is tested: number of relevant facts (1 or 3) studied with the probed character that are related to the probed fact; number of facts irrelevant to the probe, but also studied with the probed character, (0, 1, or 3); whether the probe was tested in the related foil block or the unrelated foil block.

Consider the sample material in Table 1. The test probe "Alan fell while skiing down the steepest stretch" is in the 1–3 condition, where the first digit corresponds to the number of relevant facts and the second digit corresponds to the number of facts associated with Alan that are irrelevant to the probe. The subject studied only one fact about Alan relevant to skiing, and studied three other facts about Alan irrelevant to skiing. "Alan arrived on time at Grand Central Station" is an example of a probe in the 3–1 condition because the subject learned three facts about Alan relevant to the probe and only one fact about Alan that was not relevant.

The third factor is the nature of foils used in the test block, whether the foils are themat-

TABLE 1
Examples of Facts to Learn and Test Sentences Based on these Facts

Condition abbreviation	
3−1	Alan bought a ticket for the 10:00 train. Alan heard the conductor call, "All aboard." Alan arrived on time at Grand Central Station.
1−3	Alan fell while skiing down the steepest stretch.
1−1	Brian watched the freaks in the side show. Brian wanted to major in psychology.
3−0	Steven called to have a phone installed. Steven unpacked all of his boxes. Steven mailed out change of address cards.
3−3	James compared 5 different model cars. James paid the car dealer in cash. James put the license plates on his car. James checked the Amtrak schedule. James arrived at the train station early. James watched the approaching train from the platform.
Test-target:	Alan bought a ticket for the 10:00 train.
Test-foil unrelated:	Alan called to have a phone installed.
Test-foil related:	Alan watched the approaching train from the platform.

ically related or not; Table 1 gives an example of each kind. The thematically related foils can be classified according to number of relevant and number of irrelevant facts studied with the probed character. "Alan went up in the chair lift" is an example of a 1−3 related foil. "Alan watched the approaching train from the platform" is a 3−1 related foil. For the unrelated foils, however, the determination of number of "relevant" facts and "irrelevant" facts associated with the probe character is arbitrary. "Alan called to have a phone installed" could be either a 1−3 unrelated foil or a 3−1 unrelated foil. So in these two conditions (3−1 and 1−3), the assignment was made arbitrarily.

The Appendix lists all the predicates used in the experiment. The first eight "themes" listed were used in Experiment 1.[3] The themes are classified as scripts and nonscripts. This distinction will be discussed later. Not all predicates were actually studied by a particular subject since some "themes" had only two examples presented. The selection of sentences within themes and the assignment of themes to conditions and characters to themes was random for each subject. In this way, in the analysis of variance, effects due to materials would be part of the effects due to subjects (see Clark, 1973). Half the character names were male. If a predicate was paired with a male character, all references to 'her' were changed to 'his'.

Study and foil materials were constructed for each subject randomly within the following constraints: Each predicate was studied with two characters. The search process in the ACT theory involves activation spreading from all terminal nodes in the probe simultaneously.

[3] There are 11 themes in the index because the design of Experiment 2 required more themes.

This means that even if the search is slowed by multiple predicates attached to the subject node, the search from other terminal nodes will be faster if they are unencumbered by multiple predicates. This would lead to rapid recognition of the probe. Therefore, to enhance the effects of fan off the subject node, all predicate nodes had a fan of two.

A second constraint was that for a given predicate studied about *person 1* there be a different predicate from the same theme studied about *person 2* such that the predicate for the first could be used as a related foil for the second. For example, in Table 1 related but not identical facts from the train theme were associated with both Alan and James. The constraints of our design required that facts from each theme be studied with four individuals. A different *subset* of facts from a given theme was paired with each of the four characters. A given fact was studied with two of the four. Altogether there were 20 individuals. There were four characters in each of the following five conditions: one fact studied from one theme (1−0); three facts from one theme (3−0); one fact from each of two themes (1−1); one fact from one theme and three facts from another (1−3) or (3−1); and three facts from each of two themes (3−3). (Consult Table 1 again.) There were four test sentences contributing to the 1−0 cell for both targets and foils in each block, eight contributing to the 1−1 cell, four contributing to the 1−3 cell, twelve contributing to the 3−0 cell, twelve contributing to the 3−1 cell, and twenty-four contributing to the 3−3 cell. Thus, there were 64 target sentences in all. A different set of 64 foil sentences were created for the related and unrelated foil blocks. Each target was repeated in each of four blocks.

Procedural Details

Thirty-two subjects participated in the experiment conducted on a PDP 11/40 system. This system allowed up to six subjects to be run in parallel with real time monitoring of response times. The subjects read instructions that gave them a brief overview of the experiment, and then read explicit instructions about each phase prior to its initiation.

Phase 1. Each subject sat in front of a TV monitor. The screen would display one to six facts about a particular fictitious individual and the subject was told to study those facts. The facts stayed on the screen for $(n + 1)10$ sec, where n corresponds to the number of facts on the screen. After the appropriate amount of time, the screen would be erased and the next fact set displayed. After the computer had cycled through all 20 individuals and all 64 facts associated with them, there was a rest break before Phase 2.

Phase 2. In this phase, subjects had to type in all predicates they had learned to a given character's name. The screen would display both the name and the number of predicates studied. Subjects were instructed that they need not type in the exact predicate. A paraphrase that omitted function words or a phrase that abbreviated words in an unambiguous fashion would be acceptable. Subjects typed in one line for each predicate studied. If they could not remember a fact they were told to simply hit the return key once for each missing fact. After the requisite carriage returns, the screen would display the correct predicates to allow subjects to score their accuracy. They indicated their score by typing in the number corresponding to the correct number of facts recalled. Subjects were informed that they should be honest with themselves and not give themselves credit for a similar, related fact which was the wrong fact. They were told that they would never finish the last phase of the Experiment if they did not learn the material correctly in the current phase. They were also told that they would cycle through each name until they could correctly recall twice all the facts associated with that name.

After a subject finished the dropout phase of testing, a written test was given that listed the name of all the characters studied and space under each name to write all the predicates that had been studied with him or her. Subjects were given another rest break after completing this form.

Phase 3. This was the last phase of the experiment and consisted of four test blocks.

For half the subjects the four test blocks were in the order related foil, unrelated, unrelated, and related and for the other half of the subjects, the test blocks were in other unrelated foil, related, related, and unrelated. Each of the four test blocks were divided in half, with rest breaks in between. Each half-block consisted of 64 trials. Each trial proceeded as follows: One statement was presented on the screen at a time and the subject was to hit the letter "k" or "d" to indicate if the probe was studied or not. Subjects were told to keep their index fingers resting on the two appropriate keys. Reaction time was measured from presentation of the stimulus on the screen to the key press. Subjects were given feedback after each response as to their accuracy and their response time. One sec intervened between the disappearance of the feedback from the screen and the next trial. Accuracy was stressed more than response time. Subjects were told to respond "studied" to any statement that they had studied in the previous phases and "not studied" to all others. It was explained that two of the four test blocks would have foil statements that would seem more difficult than those in the other blocks. Examples of the differences between the two types of foils, illustrating the thematic relatedness of the difficult ones, were given. Each target was tested four times.

Subjects

Thirty-two subjects were recruited from the general Yale undergraduate population. They were paid $2.50 an hour for an experiment that lasted from 4–7 hours. Two of the subjects were excluded from the analysis because they failed to satisfy the criterion of recalling 90% or more correctly on the written test prior to Phase 3.

Results

There were 24 conditions in this experiment, defined by the factorial combinations of probe type, (studied vs. foil) × relevant fan (1 vs. 3) × irrelevant fan (0 vs. 1 vs. 3) × type of foil block (related vs. unrelated). An analysis of variance was performed using subjects' mean correct response times in each condition. The error term used was always the interaction of subjects with the effect(s) of interest. The standard error of the means (based on the overall interaction of subjects with conditions) was less than 35 msec. The degrees of freedom (based on the sum of *df* for each interaction of subjects with conditions) was 667. To simplify the presentation of the data, first we will describe the data collapsed over the irrelevant fan conditions.

Figure 3a presents the data from Experiment 1, collapsed over irrelevant fan, for targets and foils, and for the two types of trial blocks—the related foil blocks and the unrelated foil blocks. There was an effect of response type, $F(1,29) = 18.17$; $p < .001$ such that foils take longer to answer. There was a main effect of trial block such that subjects took longer to respond in those blocks where the foils were thematically related, $F(1,29) = 50.41$; $p < .001$. Consider the data displayed as solid lines, which represent trials from blocks with unrelated foils. These trials closely approximate the conditions used in Smith et al.'s experiments. The fan effect (difference between 1 vs. 3 relevant facts) for the target statements is 60 msec which is significant, $t(667) = 2.09$; $p < .05$. The fan

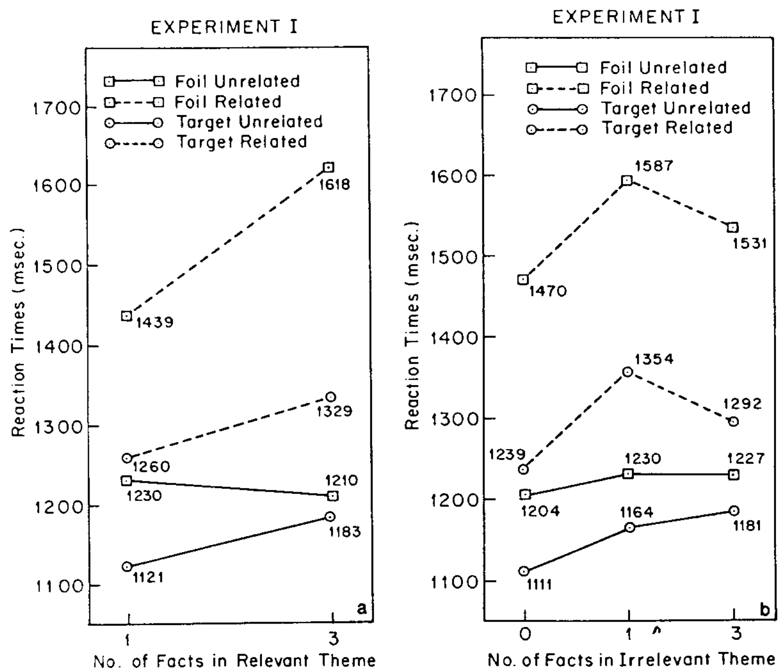

FIG. 3. Mean RT for correct recognition judgments in Experiment 1. Data are plotted as a function of the number of predicates associated with the probed character (a) that are also consistent with the probe; (b) that are unrelated to the probed predicate.

effect for the foil statements is −20 msec, which means subjects were faster to reject statements about a character for whom they learned three facts rather than for a character for whom they had learned only one fact on that topic.[4] This negative fan effect, however, is not reliable. Averaging across target and foil, the fan effect with unrelated foils is 20 msec, which is also not significant. So it seems that when unrelated foils are used, we are able to replicate the Smith et al. results that there is only a small fan effect overall.

The dashed lines of Fig. 3a display the data from the trial blocks with the related foils. In this case, there is a 69 msec fan effect for targets which is significant, $t(667) = 2.4; p < .05$, and a 178 msec fan effect for foils, also significant, $t(667) = 6.2; p < .01$. This gives an average effect of 123 msec which is highly significant, $t(667) = 4.3; p < .001$. If one collapses over type of foil block, there is an overall effect of relevant fan of 70 msec which is significant, $F(1,29) = 22.14; p < .001$. There is also a highly significant interaction of fan and type of foil block, $F(1,29) = 7.98; p <$

[4] This data also contains the arbitrary assignment for the (1−3) vs. (3−1) conditions. Consult the discussion in the design and materials section.

.001. The significant fan effect in the thematic foil block, the attenuated, nonsignificant one in the unrelated foil block, and the significant interaction of fan effect with type of foil block argue that whether one gets an interference effect or not depends on whether the subject can use related information to his advantage or not.

The ACT-like models presented in Fig. 2 predicted that we would obtain results like those of Smith et al. with the unrelated foils and typical fan effects with thematic foils. In contrast, it does not seem that this interaction would be readily explained by the Smith et al. script model. It was necessary to first demonstrate the credibility of our two explanations before trying to discriminate between them. The data concerned with effects of irrelevant fan should help in deciding between the consistency and subnode models.

Effects of Irrelevant Fan

In Fig. 3b, the data is displayed as a function of irrelevant fan and is collapsed across relevant fan. That is, here we ignore the number of relevant predicates that were studied with the probed character; instead, the variable of interest is the number of facts studied with the probed character that are irrelevant to the theme of the probe. The dashed lines represent trials from the related foil blocks, and the solid lines from the unrelated foil blocks. There is a clear effect of irrelevant fan such that subjects take longer to judge a target or foil about a person if they had also learned something else about the person that was totally unrelated. The appropriate contrast is the 0 irrelevant fan condition versus an average of the 1 and 3 irrelevant fan conditions. This difference is 65 msec and is highly significant, $t(667) = 3.85$; $p < .001$.

Both models presented at the beginning of the paper (Fig. 2) would predict that subjects would take longer if they knew more facts about Marty, relevant or not. The important contrast to discriminate between those two models involves comparing the 1 vs. 3 irrelevant fan conditions. The consistency model, (which posits that subjects search memory and respond positively when they find the first related fact) predicts that subjects will be slower with three irrelevant facts than with one irrelevant fact. The subnode model, (in which facts are grouped into related sets attached to subnodes) predicts no difference between one and three irrelevant facts. The data indicate that subjects are no slower with three irrelevant facts than with one. If anything, subjects are actually faster with three irrelevant facts than with one. There is a −30-msec fan effect between 1 and 3 irrelevant fan, however this effect is not significant $t(667) = 1.21$. Thus, the data seem to support the subnode representation/explanation given earlier.

Differential Error Rates

Most of the fan effect in the related foil block can be attributed to the foils. There is not much more of a fan effect for the targets in the related blocks than in the unrelated blocks. The reaction time difference between targets and foils is also larger in the related blocks. The difference between targets and foils in the related block is magnified if one takes into account error rates. Subjects show much higher error rates on foils and particularly high error rates for foils with three relevant fan. (Consult Table 2.) The error rate for foils in the three fan condition is greater than 25%. The high error rates and long reaction times suggest subjects alternate between two criteria in the related foil block. Subjects do not always want to use the stringent criterion that is necessary for optimal performance in this block; rather, they sometimes behave as if they were still in the easy block. Consistent with this hypothesis is the significant triple interaction for error rates of *fan* × *foil-type* × *target-foil*, $F(1,29) = 18.56$; $p < .001$, such that the greatest error rate is in the three-fan, thematic, foil condition. The tendency to respond with a lax criterion may be greater in the three fan case because the three facts create a stronger thematic impression.

The proposal, then, is that the reaction times to target probes in the related foil block are a mixture of two processes: a careful search that the task demands and a more cursory search that is sufficient in the unrelated-foil test block. Since subjects are much more likely to make false alarms by accepting a thematic foil than incorrectly rejecting a presented statement, most errors show up in the foils. In the related foil block, when a subject responds quickly, saying that a statement is true, he will be correct half the time. These fast guesses only affect the latencies for target responses—driving the reaction times down.[5]

Since the error rates were higher for three fan related foils than for one fan related foils we expect that if accuracy could be equated, the targets will be slower and will show a much larger fan effect in the thematic foil block. This was the motivation for Experiment 2.

[5] This notion, that the target responses are a mixture of careful searches and fast guesses, implies that the latencies for errors to foils should be faster than latencies for errors to targets in the thematic 3-fan conditions. In fact, in these conditions, the mean RT for errors to 3-fan targets was 1604 msec, while it was 1499 msec for errors to foils. In contrast, the mean RT for errors in the *unrelated* block was 1152 msec for 3-fan targets, while it was 1148 msec for errors to 3-fan foils. Another prediction from our fast guess analysis is that the difference between RTs for targets and foils should be greater in the 3-fan condition than in the 1-fan condition in the related block. This is because in the related block, correct responses to targets will be a mixture of fast guesses and careful searches, while correct responses to foils will all be of the latter type. The amount of fast guesses made to targets in the 3-fan condition should be greater than in the 1-fan condition due to differences in degree of perceived thematicity. This prediction was also confirmed.

TABLE 2
Mean Reaction Times (in msec) and Accuracy for the Conditions of Experiment 1

	Related foils				Unrelated foils			
	Target		Foil		Target		Foil	
Fan	Percentage correct	RT	Percentage correct	RT	Percentage correct	RT	Percentage correct	RT
1–0	.954	1185	.829	1367	.942	1069	.896	1198
1–1	.915	1343	.842	1516	.860	1121	.912	1262
1–3	.875	1253	.883	1436	.896	1172	.954	1230
3–0	.928	1293	.729	1572	.915	1153	.922	1209
3–1	.890	1364	.726	1657	.899	1207	.931	1197
3–3	.909	1331	.754	1625	.876	1189	.934	1224

EXPERIMENT 2

In Experiment 2, we monitored subjects' accuracy after each trial. They were informed when they were responding too carelessly. In Experiment 1, accuracy was considerably higher in the 1-fan thematic-foil condition than in the 3-fan condition. For this reason we tested subjects memory for three fan probes in different blocks from one fan probes. (During the acquisition phase the 3-fan- and 1-fan facts were presented together.) By asking questions from different fan conditions in different blocks, we could achieve the same level of accuracy for both fan conditions. In both test blocks, we forced subjects to perform within one standard deviation of 90% accuracy. It required greater vigilance of course, to get accuracy up to 90% in the "3-fan" block of trials. This was the only way however, to achieve approximately equal levels of accuracy in different conditions. The division of 3-fan blocks from 1-fan blocks was not obvious to the subjects. In interviews after the experiment, subjects reported that they were unaware of this division manipulation.

Method

Procedural Overview

The general design and procedure was quite similar to that of Experiment 1 with a few modifications. The most striking difference in procedure from Experiment 1 was that we computed the subject's accuracy in a given block after each trial. The subject would be informed to modify his or her behavior if accuracy was above or below one standard deviation of 90%. The formula was as follows where N is the number of trials: If $.9N - .3\sqrt{N}$ was greater than the number of correct trials in the block, the subject was told to be more careful each time an error occurred. If the number correct was greater than $.9N + .3\sqrt{N}$, the subject was told to go faster after each correct trial on which his RT was more than a running average. This calculation was done only after the number of trials, N, was greater than 9 in each block. The exact messages were: "you are making too many errors—be more careful" and "speed up—you are going too slow."

Design and Materials

The number of 1-fan facts studied was increased from 16 to 28 while the number of 3-fan facts stayed at 48. In this way, the ratio of 1-fan probes to 3-fan probes would not be so low and the 1-fan blocks would be closer in length to the 3-fan blocks.

The same three orthogonal factors as in Experiment 1 defined a condition: number of facts associated with the probed character related to the probe, number of irrelevant facts associated with the probed character, and whether the probe was tested in a related-foil block or an unrelated block. The same constraints were used in constructing study and foil materials: each predicate was studied with two characters. Rather than four characters in each of five conditions, in the second experiment there were ten characters in the 1−0 condition, six in the 1−1 condition, six in the 3−0, six in the 1−3 condition, and two in the 3−3 condition. There were 24 characters and 11 themes or scripts (listed in the Appendix). For each block, there were 28 targets and 28 foils with relevant fan of one. There were 48 targets and 48 foils with relevant fan of three. Again each target fact was tested in both types of foil block.

In this experiment, there were four trial blocks with 152 trials per block. (Of these 56 trials involved 1-fan probes and 96 trials involved 3-fan probes.) The first portion of each block was devoted to testing questions of just one-relevant or just three-relevant fan and the second portion to questions of the other fan. There was a rest break in the middle of each 1-fan portion and two in the 3-fan portion and one between the 1-fan and 3-fan blocks. Half the subjects received related foil blocks first and half received unrelated foil blocks first. The related and unrelated blocks were tested twice and they alternated in presentation. Within each of these groups half were tested on 1-fan first, then 3-fan, for the first two blocks and then this was reversed for the last two blocks. The other half of each group was presented the 1-fan and 3-fan questions in opposite order.

Subjects

The counterbalancing aspects of the design required that the number of subjects be a factor of four. Nineteen subjects had to be run to get 16 who could pass the 90% accuracy on the written test prior to the RT phase. An additional requirement was that subjects have their average error rate in the 3-fan, thematic foil block (this was the most difficult condition in Experiment 1) be less than 20%. The 16 subjects who met the 90% learning criterion also met this criterion.

Results

As in Experiment 1, an analysis of variance was performed on subjects' mean correct response times in each condition. The error term used was always the interaction of subjects with the effect(s) of interest. The standard error of the means, based on the overall interaction of subjects with conditions, was less than 43 msec; $df = 345$.

The reaction time data for correct responses are displayed in Fig. 4a as a function of relevant fan. Once again, the data are collapsed over irrelevant fan. There is a significant effect of foil block $F(1,15) = 43.19; p < .001$ such that questions in the related foil block take longer than questions in the unrelated foil block. There is also an effect of response type such that foils are more difficult than target responses, $F(1,15) = 99.35; p < .001$. There is an interaction of foil block and response type, $F(1,15) = 28.48$; $p < .001$, in that the difference between targets and foils is greater in the related block.

As before, there is an effect of relevant fan, $F(1,15) = 28.60; p < .001$, and an interaction of fan and foil block, $F(1,15) = 28.51; p < .001$. Happily, unlike in Experiment 1, there is no triple interaction of response type by foil block by fan, i.e., the target and foil slopes are parallel in the thematic block $F(1,15) = 2.53; p < .132$. More specifically, there is now a much bigger fan effect for targets in the difficult block than there had been before. Apparently, forcing subjects to be accurate had the intended effect. Now subjects display the same interference in both the target and foil conditions. The RTs and accuracy for each condition are displayed in Table 3. By contrasting Table 3 with Table 2, it is clear that subjects were much more balanced in accuracy among conditions in Experiment 2. (However, they are still least accurate in the 3-fan, related foil conditions.)

TABLE 3
Mean reaction Times (in msec) and Accuracy for the Conditions of Experiment 2

	Related foils				Unrelated foils			
	Target		Foil		Target		Foil	
Fan	Percentage correct	RT	Percentage correct	RT	Percentage correct	RT	Percentage correct	RT
1−0	.939	1128	.909	1344	.934	1023	.944	1110
1−1	.889	1220	.927	1352	.905	1110	.952	1133
1−3	.898	1299	.906	1388	.833	1146	.966	1129
3−0	.907	1478	.816	1728	.910	1129	.929	1216
3−1	.892	1528	.828	1748	.898	1180	.939	1187
3−3	.889	1571	.842	1706	.882	1222	.952	1239

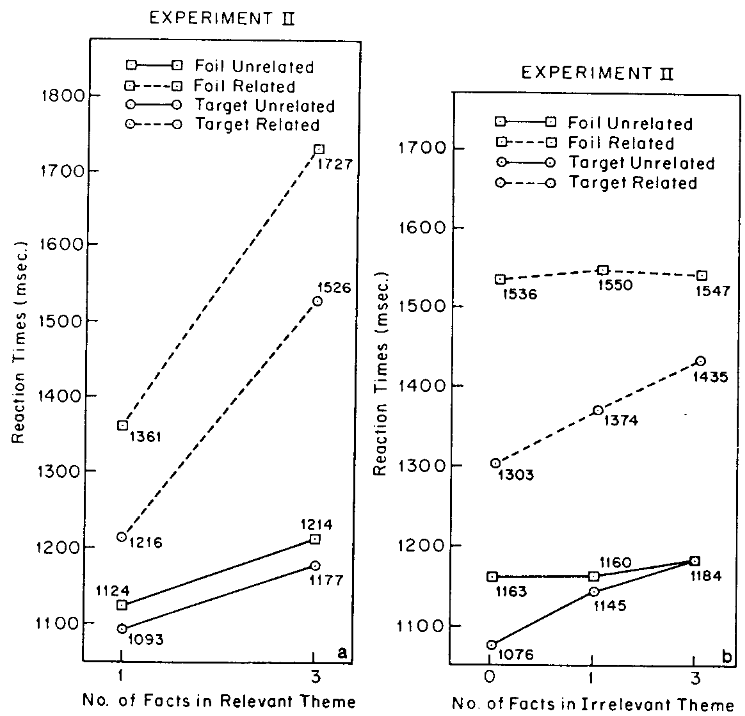

FIG. 4. Mean RT for correct recognition judgments in Experiment 2. Data are plotted as a function of the number of predicates associated with the probed character (a) that are also consistent with the probe; (b) that are unrelated to the probe.

Figure 4b displays the data collapsed over relevant fan as a function of the number of facts associated with a character about a theme other than the one being probed. The difference between no irrelevant theme and an extra theme (average of 1 and 3 fan) is 53 msec, $t(345) = 2.85$; $p < .01$ which replicates the results of Experiment 1. The difference between 1 and 3 fan is 24 msec which is not significant, $t(345) = 1.29$ and is in the opposite direction to the insignificant effect in Experiment 1. So it seems safe to conclude that there is an effect of an additional theme (0 vs. 1 and 3) but no effect of number of facts studied in that irrelevant theme (1 vs. 3). Figure 5 displays the effects of irrelevant fan collapsed over the two experiments. There appears to be little difference between 1 and 3 fan in either the related or unrelated foil block.

Using Table 3, we can also make another comparison with the results of Smith et al. In their data, the integrated triple was verified only 20 msec slower than the unintegrated pair, which was not significant. To make the analogous contrast, we can compare the 1−1 (unintegrated pair) and 3−0 (integrated triple) conditions. In the unrelated blocks, subjects are 50

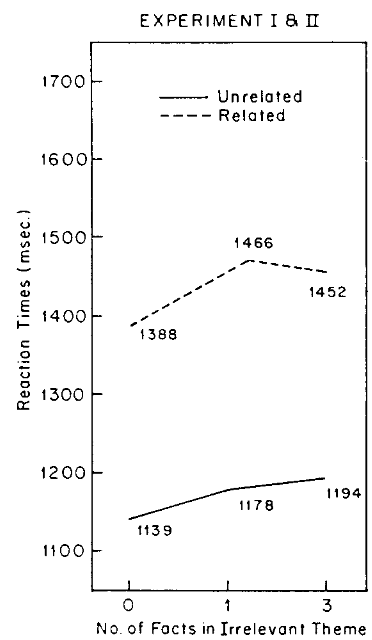

FIG. 5. Mean RT for correct recognition averaged over Experiments 1 and 2 and plotted as a function of the number of facts associated with the probed character but unrelated to the predicate in question.

msec faster to the pair. This difference is not reliable and is essentially the same result (a 20 msec effect) that Smith et al. found. In the related blocks subjects were 317 msec (consult Table 3) faster to verify facts with a fan of two (1−1) than facts with a fan of three (3−0). This fan effect is significant, $t(345) = 7.37$; $p < .001$. The difference holds for targets and foils. Hence, by this measure too we have shown that there can be a fan effect for integrated material.

DISCUSSION

The data from Experiments 1 and 2 argue quite strongly for the explanation put forth in the introduction that the attenuation of the fan effect found by Smith et al. can be explained by a subnode assumption without modifying many of the ACT assumptions about the representation of knowledge or search of the data base. Smith et al. had proposed that when subjects learn information that can be integrated into a script, their representation of that input is qualitatively different from the representation that might be used by subjects learning unrelated facts. Given that we were able to replicate the findings of Smith et al. using integrated facts and were able to replicate the typical fan effects with the same subjects using this same set of facts, it seems quite unlikely that subjects use a qualitatively different form of representation.[6] Rather, they are using a different process for retrieval. That is, the difference in response pattern from related foil blocks to unrelated foil blocks can best be explained assuming that subjects merely shift their criterion of acceptance. With related foils, subjects must match the probe exactly just as in the standard fan experiments. In the unrelated foil blocks, subjects respond positively as soon as they discover that the character queried is associated with the theme of the probe. The data from the manipulation of irrelevant fan indicates that subjects store thematically related facts at subnodes. During search, subjects need to determine only if the theme of the probe is a subnode attached to the node of the character referenced in the probe.

The findings of Smith et al. are of considerable importance because they represent search strategies which people can often use in everyday life. That is, normally we are not forced to decide whether a specific fact is presented and distinguish it from highly plausible foils. Rather, we typically decide if a particular assertion seems correct given what else we already know, (see Reder, 1979, for a fuller discussion). Given that in most circumstances people judge truth or plausibility rather than deciding if a particular assertion was presented, high fan out of a node may be of less consequence. There are a number of ways to take advantage of a set of facts which are thematically related. The subnode model described here is one way. The consistency model (Fig. 2a in our introductory discussion) provides an alternate means. It proves not to be the correct model in this paradigm but evidence has been uncovered for its operation in other paradigms (Reder, 1976; 1979).

The subnode model raises some interesting questions about basic memory processes. For this scheme to work the system must be able to tell when it is at an appropriate subnode and then focus all of its search capacity at that subnode. The question naturally arises as to how the system could know that the subnode was appropriate without searching the predicates attached to that node. If one has to search the predicates before he knows he is at the appropriate subnode how does the subnode provide any savings in search time? There are two types of answers to this issue. One is to explain how the system could know that it was the right subnode without examining the predicates. The other is to explain how a benefit might accrue even if the predicates were examined. We can think of explanations of both varieties and do not yet have the data to determine whether either or both explanations are correct.

The spreading activation mechanism allows one to know search is at the

[6] The results of Moeser (1977, 1979) can be handled in much the same way. She did not use thematically related foils just as Smith et al. had not. Her lack of a fan effect in certain conditions can be explained by positing that subjects formed thematic subnodes and used a low criterion for acceptance. We are confident that with the appropriate foils the fan effect would re-emerge in her task as well.

correct subnode without inspecting the predicates. To take a concrete example, consider the sentence: "John enjoyed the trapeze artists." Imagine that the subject has studied two themes about John, one involving the circus and one involving jogging. There should be a *circus* subnode. We will assume that this subnode is strongly associated to the concept *circus* as well as to any studied predicates. The concept *circus*, in turn is strongly associated with *trapeze artists*. When the subject reads "trapeze artists," activation spreads from that concept to the circus concept and from the circus concept to subnode. This means that presentation of the above probe, whether studied or foil, would cause an intersection at the subnode. That is, both the predicate (trapeze artist) and the subject (John) would spread activation to the circus subnode. Therefore, this subnode will be much more active than other subnodes attached to John. This level of activation could be used to select the correct subnode. So the general principle is that associative relatedness can be used through the agency of spreading activation to select the correct subnode. In the case of an irrelevant foil, no subnode would be very active and it would be possible to respond quickly with a correct rejection.

It is also possible to explain the advantage of a subnode structure even if use of the subnode required inspection of predicates attached to it. Because of the consistency of the predicates, it would only be necessary to inspect one to determine search was focused at the correct subnode. So, continuing with the circus example, if the following was found attached to a subnode: "bought a new pair of Adidas running shoes," the system would know that it was at the wrong subnode. On the other hand, if the following was found: "laughed at the clowns," the system could know, because of the connection between clowns and trapeze artists, that it was at the right subnode. Since any of the predicates at a subnode would do, it would only be necessary to consider the most rapidly activated. This should give the process of subnode judgment an independence of number of facts stored at that subnode. This explanation is a merging of the subnode and consistency models.

Scripts vs. Nonscripts

It was pointed out to us by Roger Schank (Note 1) that our materials were of two rather different types. The sets 1–4 in the Appendix were what he would call *scripts*[7] in that they had a strong temporal structure. Certain events should occur before others. The sets 5–11 have a more loose thematic structure. The question naturally arises as to whether the results would be different for the two sets. Therefore, we collapsed Experiments 1 and 2 together, but now looked for differences between the two types of material. Figure 6 plots the data for the scripts and nonscripts. The pattern

[7] We would like to thank Jaime Carbonell for being our "script-acceptability" judge.

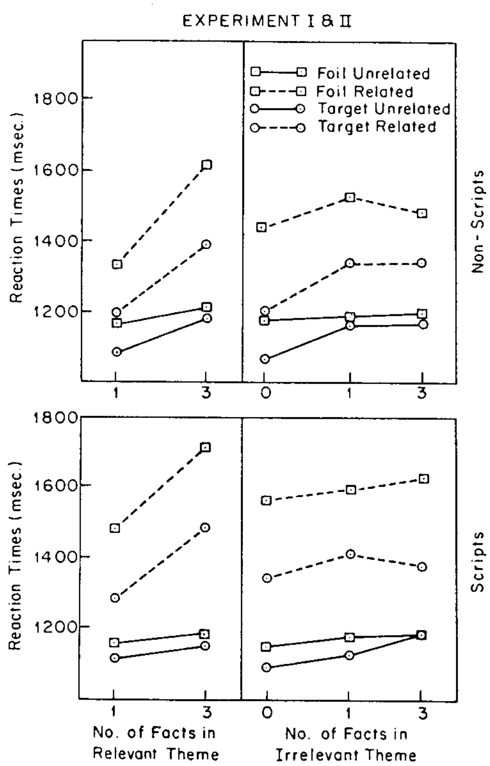

FIG. 6. Data from Experiments 1 and 2 partitioned into responses to "true themes" and "nonscripts."

of results is virtually indistinguishable. So it seems clear that the attenuation of the fan effect in the unrelated conditions is not due to scripts, per se, but is a general thematicity effect. Also, the reappearance of the fan effect on the related foil condition indicates that both scripts and other thematically related material are vulnerable to interference. It is also noteworthy that both types of material show a "subnode effect" with a clear increase in RT from 0 to 1 irrelevant propositions but no further effect of three irrelevant propositions.

We think it is important that both scripts and nonscripts show the same qualitative and quantitative patterns of data. This indicates the generality of the processes that enable thematic relatedness to speed memory search.

OTHER POSSIBLE RESOLUTIONS OF THE PARADOX

With this research in hand, we would like to make some further remarks about the issue of how an expert in a field might be able to quickly retrieve information in her or his speciality. The answers are several. First, we already discussed how integration into subnodes can facilitate search. That is, the expert's knowledge is organized into subnodes. But this, at best, puts an expert on a par with the novice.

A second factor is that the expert may have many answers stored to questions that the novice must think through. For example a considerable part of the advantage a chess expert has over the novice is that many implications of moves that the novice must work out are directly stored for the expert. Even though there may be more fan out of particular moves, the time saved by not having to compute the likely results of an action can more than compensate.

The third factor which can affect the expert's response time might be strengths of associations of the facts emanating from predicate nodes. Consider an electrician asked to verify the assertion "Electricity conducts in water" or a runner asked "A good mile time is 4 minutes." Even though the electrician knows a lot about electricity, he may not know much more about water than the average person. That would mean the water node would allow activation to spread at the normal rates. The ACT model actually specifies that search speed is a function of the *strength* of the various arcs fanning out of a node (for simplicity we have been assuming equal strength) as well as the number of arcs. Presumably the fact that water conducts well is a stronger association to water for the electrician than it is for most of us. Similarly, if asked to generate associations to "4 minutes," a runner would say "good mile time" much sooner than a nonrunner.

Another solution to the expert paradox besides the strong links from all concepts in a probe is the notion of multiple paths or redundant paths that can all be used to verify the probe. This solution, like the subnode solution, predicts no difference between novice and expert.

SUMMARY AND CONCLUSION

The findings of these experiments can be succinctly summarized: We found a large effect for relevant fan in related foil blocks, a small effect of relevant fan in unrelated blocks, and an intermediate effect for irrelevant fan regardless of the trial block. These results argue that one obtains interference effects even with integrated facts if the decision process (presented vs. not presented) does not allow consistency judgments. The extent of interference from the irrelevant facts attached to the probed character is not affected by the type of foil employed. This implies that thematically related material is organized into subnodes. That interference increased from 0 to 1 irrelevant facts, but not 1 to 3 also argues for a subnode organization.

The role of thematicity in speeding search of memory takes two forms.

First, by virtue of thematic subnodes, plausibility judgments are much easier than retrieving specific facts. Second, by having many subnodes in memory, the number of propositions that must be inspected is smaller. Even if the task does not allow plausibility judgments, thematic subnodes allow search to terminate before inspecting irrelevant information.

APPENDIX

Themes Used in Experiments 1 and 2

1-script
checked the weekend Amtrak schedule.
arrived at the train station early.
bought a ticket for the 10:00 train.
waited for the express on the platform.
heard the conductor call, "All aboard."
arrived on time at Grand Central Station.

2-script
sorted his clothes into colors and whites.
set the wash cycle and water temperature.
asked the laundry attendant for change.
cleaned the lint trap of the washer.
added liquid bleach to the rinse cycle.
had to wait to use the large dryer.

3-script
peeled the apples and chilled the dough.
rolled out the two 9-inch pie crusts.
spooned the filling into the crust.
baked the apple pie at 425 degrees.
let the steaming apple pie cool off.
cut the apple pie into six servings.

4-script
compared five different model cars.
took the car for a short test drive.
paid the down payment for the car in cash.
got the compulsory automobile insurance.
got a car loan from the local bank.
put new licence plates on his car.

5-nonscript
nervously watched the tightrope walker.
thought the clowns were very funny.
liked listening to the circus music.
was given a ticket by the ringmaster.
watched the freaks in the side show.
liked the trapeze artists the best.

6-nonscript
read and signed the one-year contract.
put down a two month security deposit.
wrote out and mailed new address cards.
hired a large moving van and movers.
unpacked all of his boxes of dishes.
called to have a telephone installed.

7-nonscript
wanted to major in Psychology at Yale.
signed up for three psychology classes.
participated in freshmman orientation.
didn't like his assigned roommate.
didn't mind living in the dormitory.
thought dorm cafeteria food was awful.

8-nonscript
jogged a mile outdoors to warm up.
preferred to run on the inside lane.
did sprints to increase his speed.
ran at least four miles every day.
wanted to make the college track team.
bought new Adidas training shoes.

9-nonscript
found a parking spot near the beach.
took the picnic basket out of the car.
covered himself thoroughly with suntan oil.
tried to keep sand out of his food.
wished there were bigger waves for surfing.
worried that his suntan would peel.

10-nonscript
bought a pair of downhill skis at the Co-op.
groaned at the price of good ski boots.
thought the lower slopes were too crowded.
rode a ski lift to the top of the mountain.
fell while skiing down the steepest stretch.
decided snowshoeing might be a better idea.

11-nonscript
decided to play squash with his friend.
reserved a squash court for the afternoon.
asked the squash team coach for pointers.
had trouble controlling the backhand shots.
broke his racquet reaching for a low ball.
felt very tired after playing squash.

REFERENCES

Anderson, J. R. Retrieval of propositional information from long-term memory. *Cognitive Psychology*, 1974, 5, 451–474.

Anderson, J. R. *Language, memory, and thought*. Hillsdale, NJ: Erlbaum, 1976.

Anderson, J. R., & Bower, G. H. *Human associative memory*. Washington, DC: Hemisphere Press, 1973.

Anderson, J. R., & Paulson, R. Interference in memory for pictorial information. *Cognitive Psychology*, 1978, 10, 178–202.

Clark, H. H. The language-as-fixed-effect-fallacy: A critique of language statistics in psychological research. *Journal of Verbal Learning and Verbal Behavior*, 1973, 12, 335–359.

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

Collins, A. M., & Quillian, M. R. Experiments on semantic memory and language comprehension. In L. Gregg (Ed.), *Cognition and learning*. New York:Wiley, 1972.

Crowder, R. G. *Principles of learning and memory*. Hillsdale, NJ: Erlbaum, 1976.

Hayes-Roth, B. Evolution of cognitive structures and processes. *Psychological Review*, 1977, 84, 260–278.

King, D. R. W., & Anderson, J. R. Long-term memory search: An intersecting activation process. *Journal of Verbal Learning and Verbal Behavior*, 1976, 15, 587–606.

Lewis, C. H., & Anderson, J. R. Interference with real world knowledge. *Cognitive Psychology*, 1976, 7, 311–335.

Moeser, S. D. Recognition processes in episodic memory. *Canadian Journal of Psychology*, 1977, 31, 41–70.

Moeser, S. D. The role of experimental design in investigations of the fan effect. *Journal of Experimental Psychology: Human Learning and Memory*, 1979, 5, 125–134.

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

Reder, L. M. *The role of elaborations in the processing of prose*. Unpublished doctoral dissertation. University of Michigan, 1976.

Reder, L. M. The role of elaborations in memory for prose. *Cognitive Psychology*, 1979, 11, 221–234.

Schank, R. C., & Abelson, R. P. *Scripts, plans, goals, and understanding: An inquiry to human knowledge structures*. Hillsdale, NJ: Erlbaum, 1977.

Smith, E. E., Adams, N., & Schorr, D. Fact retrieval and the paradox of interference. *Cognitive Psychology*, 1978, 10, 438–464.

Thorndyke, P. W., & Bower, G. H. Storage and retrieval processes in sentence memory. *Cognitive Psychology*, 1974, 5, 515–543.

REFERENCE NOTE

1. Roger Schank, personal communication.

(Accepted March 2, 1980)