Journal of Experimental Psychology: General
1999, Vol. 128, No. 4, 563–579

0096-3445/99/$3.00

Memory Retrieval and Suppression: The Inhibition of Situation Models

Gabriel A. Radvansky
University of Notre Dame

When people retrieve newly learned facts on a recognition test, they are often increasingly slowed by the number of other newly learned facts that have a concept in common with the probed fact. This is called the fan effect. Assuming that people are using situation models of the learned information, the author considers whether the inhibition of competing representations is one of the processes involved in the fan effect. Evidence was found for negative priming of related but irrelevant situation models, thus supporting the idea that the inhibition of highly related memory traces is used in long-term memory retrieval. As such, this is a form of retrieval-based inhibition.

The inhibition of information by means of active suppression mechanisms has been a topic of increasing interest over the past few years in cognitive psychology. The operation of inhibitory processes has been implicated in a number of domains of human cognition, including perceptual attention (Tipper, 1985), lexical access (Simpson & Kang, 1994), language comprehension (Gernsbacher & Faust, 1991), and aging (Hasher & Zacks, 1988). The aim of this article is to further explore whether inhibitory processes are involved in long-term memory retrieval.

In this article, the term *suppression* is used to refer to the mechanism that reduces the activation of a mental representation, and the term *inhibition* is used to refer to the action and result of reducing activation. In general, suppression is an attentional mechanism typically thought to be invoked when a cognitive process needs to select one item from a set of interfering distractors. People often need to segregate a target item from a collection of similar items. One part of this process involves the activation of a representation of the target item itself, and another part of this process involves the active inhibition of related sources of competing information that may have also been initially activated. Suppression keeps these inappropriate representations from entering the mainstream of processing so that errors may be minimized.

There are two sources of evidence to suggest that inhibition can be involved in long-term memory. The first is studies of *directed forgetting* (Bjork, 1970). In directed-forgetting studies, people are given a set of information. At some point, they are cued to remember part of the set (the to-be-remembered information) and to forget the rest (the to-be-forgotten or TBF information). It has been suggested that TBF information is actively inhibited in memory (Bjork, 1989). For example, it has been found that people take longer to reject TBF items relative to new (control) items on a speeded recognition test (Bjork, Abramowitz, & Krantz, 1970, as reported in Bjork, 1989; Neumann & DeSchepper, 1992; Zacks, Radvansky & Hasher, 1996). There is some agreement that this slowdown is a result of suppression. However, it is unclear whether the increase in response time is due to the greater effort needed to bring an inhibited TBF item to threshold (Bjork, 1989) or to response competition as a result of incomplete inhibition of the TBF information (Zacks et al., 1996). The nature of the directed-forgetting effect aside, these sorts of data are limited in breadth with regard to the role of suppression in long-term memory retrieval because they do not speak to whether suppression is involved in long-term memory retrieval when information has not been marked as TBF.

Other evidence comes from a set of experiments reported by M. C. Anderson and Spellman (1995) in an experimental design dubbed the *retrieval-practice* paradigm (see also M. C. Anderson, Bjork, & Bjork, 1994; M. C. Anderson, de Kok, & Childs, 1997; M. C. Anderson & McCulloch, 1999; May, Einstein, & Knight, 1997). In this paradigm, people were first given a list of categorized words in which all of the experimental items were presented. They then engaged in repeated practice on a subset of that list by recalling items using a category-word stem cue, such as "Green—Em_____" for the word "Emerald." Other green things in the list, such as "Lettuce," were not practiced. Later, people were given a recall test for the entire list. Obviously, performance was enhanced for practiced items. The important finding was that memory for the unpracticed items from the practiced category was worse relative to unpracticed, unrelated control items. The explanation was that the retrieval of one item from a category involves the inhibition of other members of the category. Practiced items presumably received enough activation through practiced retrieval to overcome much of the effect of this inhibition. Unpracticed items were only suppressed, thus making it more difficult to later retrieve these items.

The results from the retrieval-practice paradigm suggest

This research was supported in part by National Institutes of Health B/START Grant 1RO3 MH56315-01.

I would like to thank John Anderson, Mike Anderson, Laura Carlson-Radvansky, Jackie Curiel, Nancy Franklin, Brad Gibson, and Yuhong Jiang for their helpful comments on an earlier version of this article. I would also like to thank Amee Appel, David Copeland, Simon Dunn, Jeremy Heckman, Adam Singleton, and Dan Wolters for their assistance in gathering the data.

Correspondence concerning this article should be addressed to G. A. Radvansky, Department of Psychology, University of Notre Dame, Notre Dame, Indiana 46556. Email may be sent to Gabriel.A.Radvansky.1@nd.edu.

that suppression is involved in long-term memory retrieval. This is referred to more generally as *retrieval-based inhibition* because the act of retrieving one item from memory causes other items to be inhibited. One of the aims of the current experiments was to provide additional evidence for the operation of suppression in retrieval-based inhibition using a different experimental methodology. M. C. Anderson and Spellman's (1995) inhibitory effect was obtained using materials that had strong preexperimental associations, namely, semantic categories. In contrast, the current experiments used more episodic information that was learned during the experiment. Also, M. C. Anderson and Spellman's effect required the practice of some items and not others. It has yet to be established that inhibition would also be observed in the absence of such unbalanced practice. The current experiments did not rely on a differential practice. Finally, M. C. Anderson and Spellman's critical dependent measure was the proportion recalled. The current experiments used a chronometric measure of inhibition, namely, negative priming (see, e.g., Tipper, 1985), to get a more "on-line" estimation of the operation of suppression.

Suppression is expected to be involved in long-term memory retrieval when the selection of a memory trace needs to be made from a set of similar traces. Suppression would be involved when related but irrelevant traces interfere with the retrieval of a desired target trace. This set of conditions is observed with the *fan effect* (J. R. Anderson, 1974), an increase in response time on a recognition test with an increase in the number of experimentally learned associations with a concept. In these studies, people learn a set of sentences, such as "The hippie is in the park." The fan effect is the finding that the more facts that are associated with a concept (e.g., hippie), the longer it takes to later verify any one of those facts on a recognition test.

The specific issue addressed here concerns the nature of the retrieval processes involved in producing the fan effect. One possibility is that the fan effect reflects only the operation of the activation of information. Activation views, such as the Adaptive Control of Thought (ACT) family of models (see, e.g., J. R. Anderson, 1976, 1983, 1993), take the position that the fan effect arises from competition among various activated portions of memory. The more information that is activated, the more difficult it is to select out any one memory trace. A second possibility is that the fan effect reflects the operation of both activation and inhibition. Inhibition views, such as models of lateral inhibition (see, e.g., Blaxton & Neely, 1983) or attentional suppression (see, e.g., M. C. Anderson & Neely, 1996; Dagenbach & Carr, 1994), take the position that while a target trace primarily receives activation, the related but irrelevant traces need to be inhibited (for a more detailed taxonomy of inhibitory and noninhibitory retrieval models, see M. C. Anderson & Bjork, 1994). It is likely that these related but irrelevant traces are initially activated but are inhibited soon after, although the time course of such inhibition is not considered here. Note that although this is referred to as the inhibition view, I am assuming that suppression is only one of the mechanisms operating on trials that produce the fan effect. The reason behind this labeling is to highlight the additional contribution of suppression that results in inhibition, which is absent in activation views.

This article is not the first to suggest that suppression is involved in the fan effect. Cantor and Engle (1993) offered this possibility in an account of differences between high- and low-span individuals in a fan-effect study modeled after Myers, O'Brien, Balota, and Toyofuku (1984). The idea was that people with low working memory spans have decreased efficiency in inhibiting irrelevant information, leading them to produce larger fan effects. This interpretation was rejected on the grounds that the low-span people did not produce higher false-alarm rates than miss rates and that false-alarm rates were similar for high- and low-span individuals. Such differences might be expected if inappropriate information were not efficiently inhibited and became available during retrieval, leading to greater levels of familiarity, thus producing higher false-alarm rates. Furthermore, people who have less efficient suppression mechanisms should be more susceptible to this. This line of reasoning rests on the assumption that suppression necessarily prevents individual concepts from becoming inappropriately mixed during retrieval in this fan-effect paradigm. However, it is unclear why this assumption would necessarily be true. Moreover, this is an indirect assessment of the operation of suppression.

The idea that suppression is involved in the fan effect was also used in the interpretation of data showing that older adults show larger fan effects than do younger adults (Gerard, Zacks, Hasher, & Radvansky, 1991; Radvansky, Zacks, & Hasher, 1996). This interpretation stems from Hasher and Zacks' (1988) view that one of the changes that occurs in cognitive processing as part of the natural aging process is a decrease in the ability to use suppression. If older adults are less efficient at inhibiting, they should show a larger fan effect. Although this interpretation is consistent with the decreased inhibition hypothesis, there was no direct evidence that a decrease in the effectiveness of suppression caused the larger fan effects for the older adults.

The current experiment examined whether suppression is involved in memory retrieval without relying on differences between groups, such as working memory span and age.

Situation Models

The original interpretation of the fan effect was in terms of a network model of memory (J. R. Anderson, 1974). The idea was that facts were stored as propositions in which study-sentence concepts corresponded to nodes in a network and that these nodes were joined by associative links. Retrieval consisted of activation spreading from nodes matching the concepts in a memory probe along the links associated with those nodes. A probe was verified as having been studied when the spread of activation from two concepts intersected, thus forming a proposition. The fan effect arose because the number of links "fanning" off a concept divided up the activation available to search each link, and so, retrieval time lengthened accordingly. As such, the term fan effect is theoretically loaded.

More recent research has suggested that propositional representations are not the most effective vehicle for describing the data from traditional fan-effect studies. Instead, the pattern of response times is better described from a situation-model view (see Radvansky, 1999, for a review). The basic idea is that when people memorize study sentences, they often comprehend what these sentences are about, that is, the circumstances to which they refer. This involves the construction of situation models (Johnson-Laird, 1983; Kintsch, 1998; van Dijk & Kintsch, 1983; Zwaan & Radvansky, 1998), and it is these situation models that are accessed during later retrieval (Radvansky & Zacks, 1997). Furthermore, when a set of facts can be easily interpreted as referring to the same situation, people will integrate them into a common situation model. This integration then has an influence on memory retrieval.

According to the situation-model view, when a set of sentences shares a concept but refers to different situations, each sentence is represented by a separate situation model. For example, if the sentences "The welcome mat is in the movie theater," "The welcome mat is in the cocktail lounge," and "The welcome mat is in the hotel" are memorized, a person is more likely to represent these in separate models. Although the sentences share a concept (i.e., welcome mat), it is unlikely that they refer to the same situation. Because information is stored in this way, when a person needs to retrieve any one of these models, the related but irrelevant models can cause interference. When responding to the probe "The welcome mat is in the movie theater," not only is the movie theater model activated but so are the cocktail lounge and hotel models. The greater the number of related but irrelevant models, the greater the interference; the retrieval process slows down accordingly, and a fan effect is observed.

In contrast, when a set of sentences shares a concept and can plausibly be interpreted as referring to the same situation, then the information can be integrated into a single situation model. For example, for the sentences "The oak counter is in the hotel," "The plotted palm is in the hotel," and "The welcome mat is in the hotel," a person is more likely to integrate this information into a single hotel model because these objects can plausibly be in the same location at the same time. When a person is then presented with any of these facts as a memory probe, there is only one model that is accessed. There are no irrelevant representations to produce interference, retrieval time is unaffected by irrelevant models, and so, little or no fan effect is observed.

This pattern of differential fan effects has been observed under a variety of circumstances, including those involving spatial (Radvansky, 1998; Radvansky, Spieler, & Zacks, 1993; Radvansky & Zacks, 1991; Radvansky et al., 1996), temporal (Radvansky, Zwaan, Federico, & Franklin, 1998), and ownership relations (Radvansky, Wyer, Curiel, & Lutz, 1997). The pattern is not affected by the order of the concepts in the study sentences (Radvansky & Zacks, 1991; Radvansky et al., 1996) nor by whether definite or indefinite articles are used in the study sentences (Radvansky et al., 1993).

In sum, it is more accurate to assume that the representations used in this paradigm are situation models (Radvansky, 1999). Thus, the experimental design and the predictions based on it are best understood using the idea that people are forming situation models of the information described by the study sentences, that these situation models are stored in memory, and that there are no strong associations among the different models. Accordingly, the activation and inhibition views of memory retrieval were tested from a situation-model perspective.

Experimental Design

The design of the present experiments is based on other studies reporting a differential fan effect (see, e.g., Radvansky et al., 1993). The design is presented here to clarify how the test of the activation and inhibition views was made. Students memorized sentences about objects in locations of the form "The *object* is in the *location*," such as "The broken window is in the airport." These sentences were generated through a combination of objects and locations using a study-list design that produced one to three associations for each of the concepts. An abstract form of the study-list design is presented in Figure 1A. Figure 1B provides actual object and location concepts that a hypothetical person might receive. An illustration of the associative structure among concepts in a study list based on Figure 1B is provided in Figure 2.

The test of the differential fan effect uses data from cells in which there is one association with one concept and one to three associations with the other. These conditions run along the top row and left column in Figures 1A and 1B. Cases in which a single location is associated with several objects are termed the *single location* condition. Cases in which a single object is associated with several locations are termed the *multiple location* condition. These are the two conditions that are compared in the differential fan-effect analysis. In previous studies (e.g., Radvansky & Zacks, 1991), a fan effect was observed for the multiple location condition but not for the single location condition. There are twice as many items in the 1–1 cell to allow different items to be assigned to the Fan 1 baseline for the single location and multiple location conditions. In Figure 1A, using the notation that a lowercase letter refers to an object concept and an uppercase letter refers to a location concept, the single location condition is composed of sentences aA, bB, eE, fF, gG, and hH. The multiple location condition is composed of sentences cC, dD, iI, jJ, kK, and lL. For the differential fan-effect analysis, those items with multiple associations for both concepts—namely, iG, jH, kE, lF, kG, and lH—are fillers that serve only to provide the single and multiple location items with the appropriate number of associations.

This study-list design also allows for a test of the activation or inhibition views. On the basis of earlier research, one knows that there is a clear fan effect in the multiple location condition but not in the single location condition. The notion of the facts being stored in situation models and that related but irrelevant models can interfere with retrieval is illustrated in Figure 3.

In Figure 3, each box corresponds to a situation model. Because the models here are defined by locations, a location

(A)

Object Fan	Location Fan: 1	2	3
1	aA bB cC dD	eE fF	gG hH
2	II JJ		iG jH
3	kK lL	kE lF	kG lH

(B)

Object Fan	Location Fan: 1	2	3
1	broken window - airport revolving door - city hall ceiling fan - barber shop waste basket - car dealership	wall clock - cocktail lounge fire extinguisher - office building	oak counter - hotel pay phone - high school
2	potted palm - ice cream parlor bulletin board - laundromat		potted palm - hotel bulletin board - high school
3	welcome mat - movie theater cola machine - public library	welcome mat - cocktail lounge coal machine - office building	welcome mat - hotel cola machine - high school

Figure 1. The study-list design used in Experiments 1–3. For Part (A), the lowercase letters refer to object concepts, and the uppercase letters refer to location concepts. Part (B) presents the study-list design using the actual concepts presented in the experiments as a hypothetical participant might receive them.

name is used as a title for each model. Object names are placed in the boxes to convey the idea that each object is in that location. When an object is in different locations, it is repeated in the different models. Relative to a Fan 1 baseline (e.g., "The broken window is in the airport"), retrieval in the single location condition is just as easy because there is also only one model that needs to be accessed. So, when a person is presented with the probe "The oak counter is in the hotel," although the hotel model contains other objects, the oak counter is unique to that location, and the hotel model is the only one activated.

In contrast, retrieval in the multiple location condition is more difficult because several situation models are involved. For example, for the probe "The welcome mat is in the movie theater," although the movie theater has only one object in it, the welcome mat is in two other situation models, namely, the cocktail lounge and hotel models. These related and irrelevant models interfere with the ability to retrieve the desired model. As such, evidence for either activation or inhibition is sought in cases where retrieval requires selection from a set of competitors. This assessment is made by looking at the availability of those competitors on the following trial compared with a more neutral condition. Thus, the trial on which a situation model is selected from a set of competitors serves as the *prime trial* (trial t), and the retrieval of one of those competitor models is the *target trial* (trial $t + 1$). Thus, the multiple location condition, where interference is observed, is the focus of the current study. As a reminder, in the multiple location condition, each sentence is represented by a separate situation model, and the retrieval

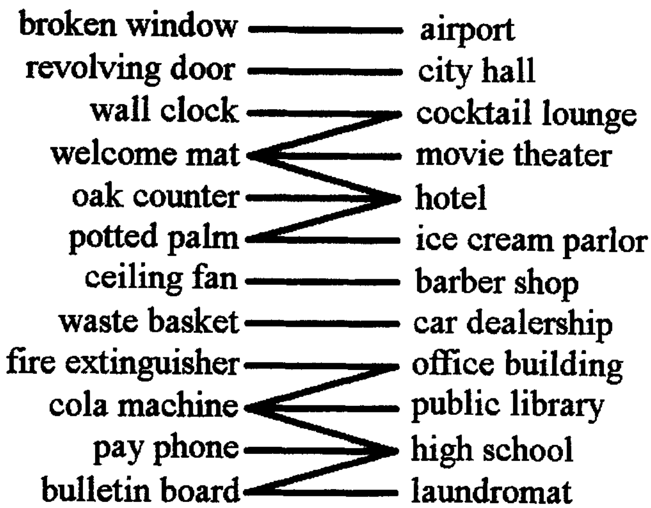

Figure 2. Associative structure among concepts used to form the study sentences.

of one model is impaired by the presence of the other models containing the object mentioned in the memory probe.

Assume that on the prime trial, the sentence is "The welcome mat is in the movie theater." This trial involves the movie theater model as well as the cocktail lounge and hotel models. The question is, how available are the cocktail lounge and hotel models on the target trial? To assess this, on *experimental* trials, people are presented with either the sentence "The wall clock is in the cocktail lounge," "The oak counter is in the hotel," or "The potted palm is in the hotel"—that is, any sentence involving the two related models but without repeating any concept names. Thus, for the experimental trials, this corresponds to the following 16 prime–target pairs: iI–gG, iI–kG, jJ–hH, jJ–lH, kK–eE, kK–gG, kK–iG, lL–fF, lL–hH, lL–jH, kE–gG, kE–iG, lF–hH, lF–jH, kG–eE, and lH–fF.

To test what is going on in the experimental trials, a *control* condition is needed. For the control trials, the target items are the same as those used on the experimental trials. Furthermore, the primes have the same associative complexity as in the experimental trials. The only difference between the prime and target items on the control trials is that they are unrelated. For example, if the experimental prime–target pair were "The welcome mat is in the movie theater" and "The oak counter is in the hotel," the control prime–target pair would be "The cola machine is in the library" and "The oak counter is in the hotel." Using these constraints for the control trials, this corresponds to the following 16 prime–target pairs: jJ–gG, jJ–kG, iI–hH, iI–lH, lL–eE, lL–gG, lL–iG, kK–fF, kK–hH, kK–jH, lF–gG, lF–iG, kE–hH, kE–jH, lH–eE, and kG–fF.

There are two primary advantages to this design. First, there is no repetition of concept names. This avoids any problems of repetition priming effects. Second, it should be noted that the primes in both the experimental and control trials are the same set of items and have the same fan in both cases. Any difference between these conditions cannot be attributed to the different associative complexity, and thus the retrieval difficulty, of the prime sentences.

The activation and inhibition views predict different outcomes for the experimental trials. For the activation view, because only the competition among activated memory traces causes the fan effect, responses should be faster in the experimental condition than in the control condition. That is, a facilitation or positive priming effect should be observed.

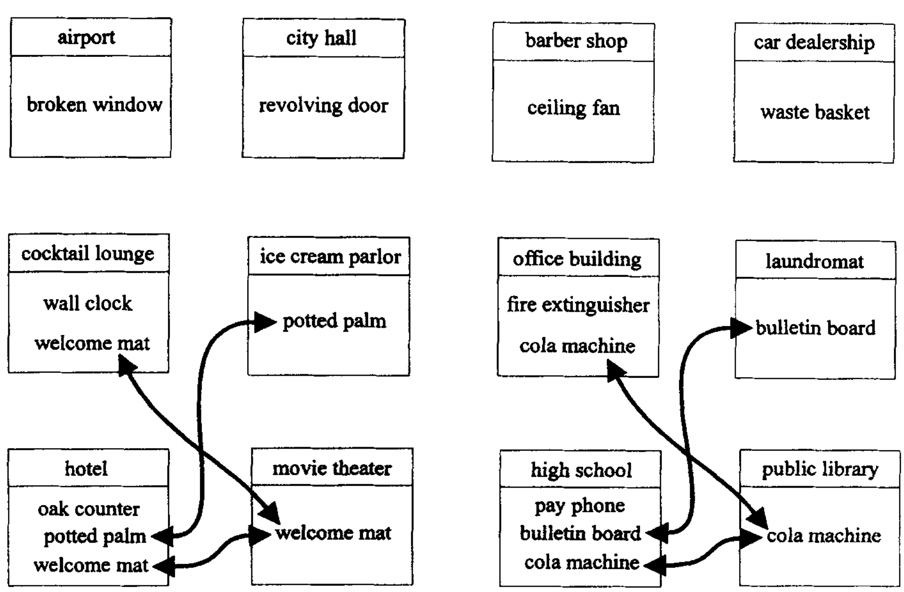

Figure 3. Conceptualization of the organization of information into situation models giving rise to the fan effect. Lines with arrows indicate sources of interference during retrieval for the multiple location condition.

This is generally consistent with a spreading activation view of memory and, in the case of the current study-list design, may be similar to accounts of mediated priming (see, e.g., McNamara, 1992). Assuming the use of situation models in long-term memory retrieval, the explanation would be that all of the related models are being activated and that the residual activation that a model had received from the previous trial allows it to be retrieved faster.

For the inhibition view, when people respond to a memory probe, such as "The welcome mat is in the movie theater," they need to inhibit the cocktail lounge and hotel models to aid the selection of the movie theater model. If a memory probe that requires the activation of either of these irrelevant models is then presented, it should take longer to be retrieved. As such, people should respond more slowly to experimental trials than to control trials. This would essentially be a type of negative priming effect (see, e.g., Tipper, 1985; see Fox, 1995; May, Kane, & Hasher, 1995; and Neill, Valdes, & Terry, 1995, for reviews). The main difference between this study and other negative priming studies is that the to-be-ignored distractors are not physically present and perceptually available in the environment (see also Carlson-Radvansky & Jiang, 1998). Instead, they are conceptually present, existing only in memory.

Experiment 1

Method

Participants. Seventy-one native English speakers were tested. They were recruited from the pool at the University of Notre Dame and given partial class credit for their participation.

Materials. As discussed above, the study facts were a set of 18 sentences combined to create sentences of the form "The *object* is in the *location*" in a fashion similar to earlier studies (e.g., Radvansky et al., 1993). A different random assignment of objects and locations to each condition was used for each participant.

For the recognition test, studied probes were those sentences that were originally memorized. Nonstudied probes were generated from re-pairings of the object and location concepts from within the same cell of the design. For example, if the studied sentences from the same cell were Sentences 1 and 2, the nonstudied sentences would be Sentences 3 and 4.

1. The oak counter is in the hotel.
2. The pay phone is in the high school.
3. The oak counter is in the high school.
4. The pay phone is in the hotel.

Thus, using the notation in Figure 1A, the nonstudied sentences would be dA, aB, bC, cD, fE, eF, hG, hH, jI, iJ, iG, jH, lK, kL, lE, kF, lG, and kH.

This method of generating nonstudied probes avoids the possibility that people may engage in plausibility judgments rather than recognition decisions (see, e.g., Reder & Anderson, 1980). Because the same number of associations was involved for the object and location concepts for the nonstudied sentences, they were assigned to single location and multiple location conditions and were analyzed as such.

Procedure. People memorized a list of 18 sentences by means of a study–test procedure. During memorization, each person was first presented with the study list and instructed to memorize the sentences as efficiently as possible. The sentences were displayed one at a time in white on a black background for 7 s each on a 486sx IBM-compatible computer running in 40-column presentation mode. The sentences appeared halfway down the screen beginning on the left-hand edge. A different random presentation order was used on each cycle. After the list had been presented, a set of test questions was given. The test questions were of the form "Where is the *object*?" and "What is in the *location*?" for each object and location, respectively. The test questions were randomly ordered on each cycle. Accompanying each test question was a number indicating the number of answers to each question (i.e., 1, 2, or 3). People responded by typing their answers into the computer. The computer provided feedback concerning the correctness of each answer. After the appropriate number of answers for a question was given, if there were any incorrect answers, the computer displayed all of the correct answers together for 3 s per answer. After answering all of the questions, the participants returned to the study portion. This study–test procedure continued until a person was able to correctly answer all of the test questions twice in a row. An average of 4.7 (SD = 1.2) study–test cycles was required to memorize the information.

The recognition test was timed and administered on the computer. Each probe sentence was presented 12 times, yielding a total of 432 recognition-test trials. There were 56 prime–target pairs in the recognition test, half experimental and half control (see Experimental Design above for a complete listing of the prime–target pairs). Each prime–target pair was presented twice across two blocks. The reason that there were 56 prime–target trials in Experiment 1 rather than the 64 possible is that some pairs were omitted due to an error in the program that generated the recognition-test files. The omitted experimental prime–target pairs were kG–eE and lH–fF, and the omitted control pairs were lH–eE and kG–fF.

The left button on a computer mouse was pressed to indicate a studied sentence and the right button to indicate a nonstudied sentence. People were encouraged to respond as quickly and as accurately as possible. There was no delay between a response and the presentation of the next memory probe. The order of probe presentation in the recognition test was randomized within each of two blocks with the constraint that prime–target trials follow one another. If an incorrect response was made, immediate feedback was given in the form of a line that read either "*ERROR* SENTENCE STUDIED" or "*ERROR* SENTENCE NOT STUDIED," whichever was appropriate. This feedback was presented for 1 s. A set of 18 practice trials was given to familiarize people with using the mouse buttons. On the practice trials, the computer displayed either "SENTENCE STUDIED" or "SENTENCE NOT STUDIED," and the person pressed the appropriate button. A self-timed break was given in the middle of the recognition test.

Design and analysis. The response-time data were analyzed in two ways. First was in terms of the prime–target pairs. The data from the target trials were of primary consideration. Trials for which response times were shorter than 200 ms or longer than 10 s were eliminated as anticipations and lapses of attention, respectively. In addition, responses that were more than 2.5 standard deviations from a participant's mean in a given cell were eliminated as outliers. These procedures removed 3.4% of the data. In addition, those trials on which a person made an incorrect response on the prime trial were excluded from the analysis. This added constraint removed an additional 2.4% of the trials. The effect of prime (control vs. experimental) was assessed using a repeated-measures analysis of variance (ANOVA).

The other way the data were analyzed was in terms of the differential fan effects of the single location and multiple location conditions. Only the data from probe sentences where there was one association with one concept and one to three associations with the other were analyzed. These data were submitted to a 2 (studied/nonstudied) × 2 (condition: single location vs. multiple location) × 3 (fan) repeated-measures ANOVA. Those sentences from cells in which several objects were associated with several locations were considered as fillers for this analysis. They were not used because they could not be simply and directly assigned to the single location and multiple location conditions. The response-time data were trimmed in the same manner as in the priming analysis. This accounted for 3.5% of the data. Trials on which an error was made were excluded from the response-time analysis. For interested readers, response-time and error-rate data for filler trials for all three experiments are presented in Appendix A. Note that there is some deviation between the means presented for the priming and interference analyses. This is because the priming data, but not the interference data, reflect the consequences of the removal of trials on which an error was made on the prime trial. Unless noted otherwise, $p < .05$ is assumed for all statistical tests.

Results

The results of Experiment 1 support the inhibition view. People responded more slowly to target trials when they were preceded by experimental primes than by control primes. In addition, a differential fan effect was observed, with a fan effect occurring for the multiple location probes but not for the single location probes. This is consistent with the idea that people were using situation models.

Priming analysis. The response-time data for the experimental and control target trials are summarized in Figure 4.[1] Consistent with the inhibition view, people responded more slowly to experimental targets (1,736 ms, $SE = 41$) than to control targets (1,676 ms, $SE = 46$), $F(1, 70) = 4.65$, $MSE = 27{,}617$. Thus, a negative priming effect was observed.[2] The effect on the target trials could not be attributed to a difference in the processing of the prime trials as response times for control primes (1,827 ms, $SE = 46$) were similar to those for the experimental primes (1,813 ms, $SE = 44$), $F < 1$.

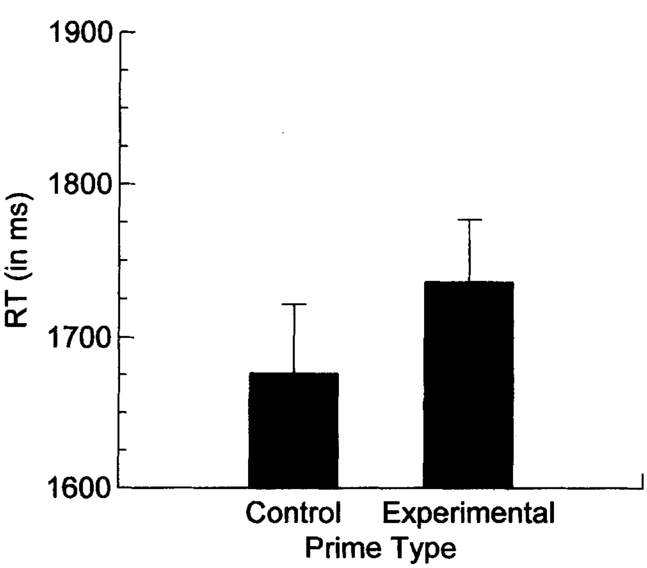

Figure 4. The response-time (RT) data for the target trials in Experiment 1. Error bars reflect standard errors.

The overall error rate for the control and experimental trials was 3.1%. People made similar numbers of errors to control targets (3.4%, $SE = .39$) as to experimental targets (2.6%, $SE = .47$), $F(1, 70) = 2.24$, $MSE = 10$, $p = .14$. Furthermore, people made similar numbers of errors to control primes (2.7%, $SE = .40$) as to experimental primes (2.5%, $SE = .34$), $F < 1$.

Differential fan effects. The response-time and error-rate fan-effect data are summarized in Table 1. As in previous experiments, there were significant main effects of condition, $F(1, 70) = 45.43$, $MSE = 74{,}779$, and fan, $F(2, 140) = 28.63$, $MSE = 92{,}967$, as well as a significant Condition × Fan interaction, $F(2, 140) = 7.23$, $MSE = 105{,}365$. Simple effects tests showed that the fan effect was significant for the multiple location condition, $F(2, 140) = 31.33$, $MSE = 96{,}396$, and the single location, $F(2, 140) = 3.97$, $MSE = 101{,}936$. In addition to the effects of primary interest, there was a significant main effect of studied–nonstudied probes, $F(1, 70) = 120.99$, $MSE = 58{,}493$, with people responding faster to studied probes (1,617 ms) than to nonstudied probes (1,800 ms).

A significant fan effect for the single location condition is unusual in studies such as these. One contributor to this is the fact that some of the trials at Fan Levels 2 and 3 in the studied single location condition were experimental targets for the priming analysis. As such, the increased response times here could be due to inhibition from the previous trial and not to interference that was generated from irrelevant models on the current trial. Because of this, the fan-effect analysis was redone with the experimental target trials removed. This resulted in 17% of the data being removed from the Fan 2 condition (33% in Experiments 2 and 3) and

[1] The data for the experimental and control targets can, in principle, be broken down into the various fan levels. It should be noted, however, that there are often very few possible observations per cell (as few as two in Experiments 2 and 3). Furthermore, there is no theoretical interest in the fan size of the target. Still, for those readers interested, the experimental and control target data divided up by size of fan are presented in Appendix B.

[2] The data from an additional person were dropped from the data presented here. This was done because an examination of the absolute difference between the control and experimental trials revealed an excessively large difference of 1,249 ms (a facilitatory effect) that was more than twice as large as the next largest difference (578 ms, also a facilitatory effect) and was more than 5 *SD* above the mean absolute difference (207 ms, $SD = 191$). With this person included, the experimental trials were slower (1,731 ms, $SE = 41$) than control trials (1,689 ms, $SE = 47$), but this difference did not reach statistical significance, $F(1, 71) = 1.63$, $MSE = 39{,}130$, $p = .21$, presumably due to the excessive amount of variability contributed by this person's data. Overall, 48 participants responded more slowly to the experimental trials than to the control trials, 23 showed the reverse, and 1 showed no difference. This inhibition effect is significant by a sign test, $p = .004$. As such, whatever this one person was doing, it was quite different from the other 71 people, adding an excessive amount of variance to the analysis. The exclusion of this data resulted in no meaningful effect on the differential fan-effect analyses.

Table 1
Differential Fan-Effect Analysis for Experiment 1

	Studied						Nonstudied					
	Fan level 1		Fan level 2		Fan level 3		Fan level 1		Fan level 2		Fan level 3	
Condition	Response time (ms)	Error rate (%)	Response time (ms)	Error rate (%)	Response time (ms)	Error rate (%)	Response time (ms)	Error rate (%)	Response time (ms)	Error rate (%)	Response time (ms)	Error rate (%)
Single location	1,527	2.1	1,582	2.2	1,566	2.2	1,652	2.1	1,808	1.6	1,737	2.2
			1,554	2.8	1,533	2.7						
Multiple location	1,517	2.1	1,747	2.0	1,765	2.5	1,691	2.3	1,932	2.9	1,978	4.0

Note. Response times in milliseconds and error rates in percentages for the differential fan-effect analysis in Experiment 1. Corrected response times for the single location condition are presented beneath the uncorrected times.

50% of the data being removed from the Fan 3 condition. These data are also presented in Table 1. In this analysis, there were main effects of condition, $F(1, 70) = 49.84$, $MSE = 79{,}643$, and fan, $F(2, 140) = 25.33$, $MSE = 95{,}722$, as well as a significant Condition × Fan interaction, $F(2, 140) = 8.64$, $MSE = 103{,}592$. Simple effects tests showed that the fan effect was significant for the multiple location condition, $F(2, 140) = 31.33$, $MSE = 96{,}396$, but was only marginally significant for the single location condition, $F(2, 140) = 3.09$, $MSE = 103{,}381$, $p = .06$. Thus, these data are consistent with previous situation-model analyses of the fan effect. In the corrected analysis, there was also a significant main effect of studied–nonstudied, $F(1, 70) = 131.15$, $MSE = 60{,}190$, with people responding faster to studied probes (1,607 ms) than to nonstudied probes (1,800 ms).

The overall error rate in the fan-effect trials was 2.3%. An analysis of these data revealed a main effect of condition, $F(1, 70) = 8.12$, $MSE = 9$. This was qualified by a significant Studied–Nonstudied × Condition interaction, $F(1, 70) = 5.90$, $MSE = 10$. Separate analysis showed no difference between the single location and multiple location conditions for the studied probes, $F < 1$, but did show a significant difference for the nonstudied probes, $F(2, 140) = 13.85$, $MSE = 10$, with people making more errors to multiple location probes (3.1%) than to single location probes (2.0%). Finally, the main effect of fan just missed significance, $F(2, 140) = 3.07$, $MSE = 11$, $p = .05$, with people showing a moderate fan effect (Fan Level 1 = 2.1%, 2 = 2.2%, 3 = 2.7%).

Relation between negative priming and fan effects. According to the account offered here, the need for the retrieval process to suppress the related but irrelevant situation models is one of the processes operating on trials that produce a fan effect. To further explore this idea about the relationship between negative priming and fan effects, the data from these two conditions were submitted to a correlation analysis. The size of the fan effect was defined here as the difference between Fan Levels 1 and 3 for the studied multiple location probes, and the negative priming effect was defined as the difference between the control and experimental conditions. These two variables were significantly negatively correlated, $r = -.35$, $p = .003$. Thus, the greater a person's fan effect, the more negative the priming.

Another aspect of the data to be addressed was whether the number of models that needed to be suppressed on the prime trial had an influence on the degree of inhibition observed on the target trial. To explore this, the target trials were broken down depending on whether the prime trials involved the suppression of one or two situation models (see Figure 3). For each person, there were 4 prime–target pairs that contributed to the one irrelevant model condition and 10 that contributed to the two irrelevant models condition. These data were then submitted to a 2 (prime: experimental vs. control) × 2 (set size: one or two irrelevant models) repeated-measures ANOVA. Although the interaction did not reach significance, $F(1, 70) = 2.70$, $MSE = 108{,}248$, $p = .11$, the data suggest that the negative priming effect was larger when there were two irrelevant models rather than one. Simple effects tests showed that for the two irrelevant models trials, responses to experimental targets were slower (1,844 ms, $SE = 52$) than to control targets (1,740 ms, $SE = 49$), $F(1, 70) = 8.09$, $MSE = 47{,}501$. However, for the one irrelevant model trials, responses to experimental (1,935 ms, $SE = 67$) and control trials (1,959 ms, $SE = 67$) did not differ, $F < 1$. This pattern of results—the presence of negative priming when there are two distractors but not one—parallels a visual-selection study by Yee (1991). However, the priming effect for the one irrelevant model trials should be regarded with caution as there were only eight possible observations per person for the one irrelevant model trials.

Discussion

The results of Experiment 1 suggest that suppression is involved in long-term memory retrieval. Thus, they are consistent with the inhibition view outlined above, which states that the processes that result in the fan effect involve the inhibition of related and irrelevant memory representations. In addition to their primary interest, these results provide added support for the idea that people are creating situation models of the information in these facts and that these situation model are used during long-term memory retrieval. This was evidenced by the substantial fan effect in the multiple location condition and the attenuated fan effect in the single location condition. Finally, when performances on the fan effect and negative priming effect were compared, there was some evidence to suggest that these two effects are related.

Experiment 2

Because of the novelty and importance of the inhibition result, a replication was desirable. As such, Experiment 2 was done with this goal in mind. However, some changes were made in the method to make it easier to detect the inhibition. The first change was to include the experimental and control pairs that were inadvertently left out in Experiment 1. This provided a wider array of prime–target pairs that could be assessed. The second change was to use shorter concept names. Some of the concepts used in Experiment 1 were quite long. This increased the encoding time for those items during the recognition test, thereby increasing the variability in the recognition-test response times. This increased variability may have made it harder to detect the negative priming effect.

A third change was to reduce the number of observations collected during the recognition test. Other research on suppression and long-term memory retrieval has suggested that inhibitory effects become less prominent over the course of testing (M. C. Anderson et al., 1997). To test whether this idea was plausible, I divided the data from Experiment 1 into two blocks. An analysis of this data revealed that the negative priming effect was larger in the first block (91 ms) than in the second (11 ms). As such, only one block of recognition test trials was presented for Experiment 2.

Method

Participants. Forty-eight native English speakers were tested. They were recruited from the pool at the University of Notre Dame and given partial class credit for their participation.

Materials and procedure. The same design was used in Experiment 2 as in Experiment 1, with three changes. First, the prime–target pairs that were inadvertently left out of Experiment 1 were included in Experiment 2. Second, a new set of object and location concepts was used that was shorter than those used in Experiment 1. The objects used in Experiment 2 were "bench," "clock," "copier," "desk," "fern," "ladder," "phone," "poster," "radio," "rug," "safe," and "shelf." The locations were "airport," "bank," "bar," "diner," "factory," "hospital," "hotel," "laundromat," "library," "museum," "school," and "theater." Finally, only one block of recognition trials was presented for each person. This resulted in each probe sentence being presented six times for a total of 216 individual recognition-test trials. There was no break given during the recognition test. An average of 4.9 (SD = 1.8) study–test cycles was required to memorize the sentences. For the prime analysis, 3.8% of the target trials were identified as outliers. An additional 4.0% of the data was dropped because of errors on the prime trial. For the differential fan effect, 3.5% of the data was discarded in the data trimming.

Results

The results of Experiment 2 replicated Experiment 1 in showing a negative priming effect. Specifically, relative to an unrelated control, responses to target sentences were slower when they were preceded by a prime sentence in which the object was also in the same location mentioned in the target sentence. In addition, the differential fan effect was again observed, with the fan effect being much larger in the multiple location condition than in the single location condition. This is consistent with the idea that people are using situation models.

Priming analysis. The response-time analyses of the control and experimental target trials are summarized in Figure 5. Consistent with the inhibition view and Experiment 1, people responded more slowly to experimental targets (1,595 ms, SE = 77) than to control targets (1,470, SE = 59), $F(1, 47) = 5.85$, MSE = 63,868. Thus, a negative priming effect was again observed. Response times for control primes (1,733 ms, SE = 62) were similar to those for the experimental primes (1,768 ms, SE = 52), $F < 1$.

The overall error rate for the control and experimental targets was 3.9%. People made similar numbers of errors to control targets (4.0%, SE = .80) as to experimental targets (3.8%, SE = .78), $F < 1$. People also made similar numbers of errors to control primes (4.4%, SE = .72) as to experimental primes (4.2%, SE = .80), $F < 1$.

Differential fan effects. The response-time and error-rate fan-effect data are summarized in Table 2. For the response-time data, there were significant main effects of condition, $F(1, 47) = 5.05$, MSE = 118,552, and fan, $F(2, 94) = 18.39$, MSE = 197,651. However, the Condition × Fan interaction was only marginally significant, $F(2, 94) = 2.49$, MSE = 117,727, $p = .09$. Simple effects tests showed that the fan effect was significant for both the multiple location condition, $F(2, 94) = 13.31$, MSE = 213,844, and the single location condition, $F(2, 94) = 10.65$, MSE = 101,535. In addition to the effects of primary interest, there was a significant main effect of studied–nonstudied probes, $F(1, 47) = 55.17$, MSE = 97,306, with people responding faster to studied probes (1,463 ms) than to nonstudied probes (1,656 ms).

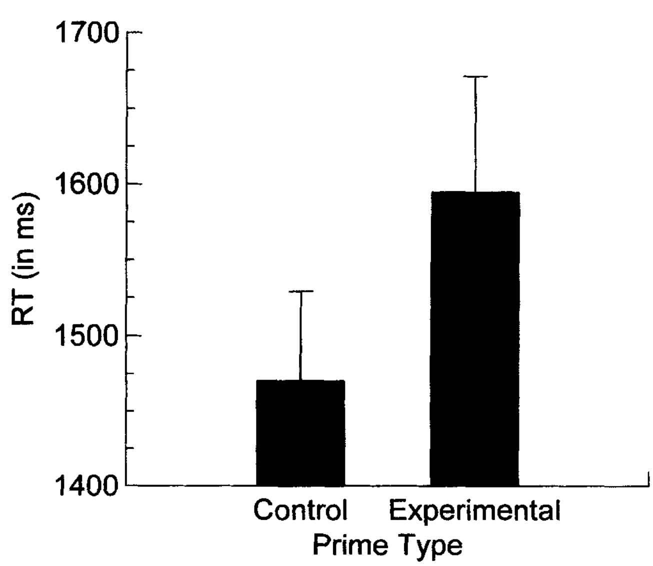

Figure 5. The response-time (RT) data for the target trials in Experiment 2. Error bars reflect standard errors.

Table 2
Differential Fan-Effect Analysis for Experiment 2

Condition	Studied						Nonstudied					
	Fan level 1		Fan level 2		Fan level 3		Fan level 1		Fan level 2		Fan level 3	
	Response time (ms)	Error rate (%)	Response time (ms)	Error rate (%)	Response time (ms)	Error rate (%)	Response time (ms)	Error rate (%)	Response time (ms)	Error rate (%)	Response time (ms)	Error rate (%)
Single location	1,387	3.3	1,345	2.8	1,600	2.8	1,513	2.4	1,621	1.4	1,696	3.8
			1,347	2.4	1,502	2.1						
Multiple location	1,316	4.5	1,469	2.3	1,659	3.6	1,534	2.4	1,691	3.1	1,879	5.0

Note. Response times in milliseconds and error rates in percentages for the differential fan-effect analysis in Experiment 2. Corrected response times for the single location condition are presented beneath the uncorrected times.

As in Experiment 1, the fan-effect data were reanalyzed with the inhibited experimental targets removed. In this analysis, there were main effects of condition, $F(1, 47) = 8.41$, $MSE = 111{,}151$, and fan, $F(2, 94) = 16.66$, $MSE = 178{,}166$, and the Condition × Fan interaction was significant, $F(2, 94) = 3.44$, $MSE = 134{,}807$. Simple effects tests showed that the fan effect was significant for both the multiple location condition, $F(2, 94) = 13.31$, $MSE = 213{,}844$, and the single location condition, $F(2, 94) = 5.91$, $MSE = 99{,}129$. However, the fan effect was much smaller in the single location condition. This is interpreted as reflecting the consequences of situation-model integration. In addition to the effects of primary interest, there was also a significant main effect of studied–nonstudied, $F(1, 47) = 60.46$, $MSE = 104{,}216$, with people responding faster to studied probes (1,447 ms) than to nonstudied probes (1,656 ms).

The error rate on the fan-effect trials overall was 3.1%. An analysis of the error-rate data revealed no significant effects.

Relation between negative priming and fan effects. As in Experiment 1, the data for negative priming and fan effect were submitted to correlation analysis. Consistent with Experiment 1, these two variables were significantly negatively correlated, $r = -.43$, $p = .002$. Again, the greater a person's fan effect, the more negative the priming.

In addition, the negative priming data were analyzed in terms of whether they involved the suppression of one or two situation models. With the addition of the prime–target pairs that were omitted in Experiment 1, for each person in Experiment 2, there were 4 prime–target pairs that contributed to the one irrelevant model condition and 12 that contributed to the two irrelevant models condition. These data were then submitted to a 2 (prime: experimental vs. control) × 2 (set size: one or two irrelevant models) repeated-measures ANOVA. The interaction did not reach significance, $F < 1$. To be consistent with Experiment 1, I performed simple effects tests. These showed marginally significant negative priming effects both for the one irrelevant model trials (experimental = 1,830 ms, $SE = 105$; control = 1,681 ms, $SE = 89$; $F(1, 47) = 3.52$, $MSE = 151{,}730$, $p = .07$) and for the two irrelevant models trials (experimental = 1,636 ms, $SE = 72$; control = 1,545 ms, $SE = 61$; $F(1, 47) = 3.06$, $MSE = 64{,}815$, $p = .09$). Note that this pattern is nominally the opposite of that observed in Experiment 1. As such, there does not appear to be any strong relationship between the number of irrelevant models and the amount of observed negative priming. Again, it should be noted that this analysis should be treated with caution as there were only four possible observations per person for the one irrelevant model trials.

Discussion

The results of Experiment 2 were consistent with Experiment 1 in showing that suppression is involved during memory retrieval when there are related and competing memory traces. However, before accepting the inhibition view, I considered an alternative interpretation. The assumption made so far has been that people are creating situation models of the facts they have studied and that these situation models are used during long-term memory retrieval. Although there is a substantial amount of support for this view (Radvansky, 1999), a more traditional analysis deserves some attention. Consider Figure 2, which illustrates the associative relations among the different concept types. Assume that this corresponds to a simplified network representation. Furthermore, assume that retrieval involves the activation of probed-for concepts and the inhibition of concepts that are associated with probed concepts but that are irrelevant.

According to this view, Experiments 1 and 2 have tested only one type of suppression: specifically, the inhibition of unrelated location nodes. For example, based on Figure 2, when a person is presented with the probe "The welcome mat is in the movie theater," the other location nodes associated with the welcome mat node, namely, the cocktail lounge and hotel nodes, are inhibited. So, if a person is then presented with another fact about these locations, such as "The oak counter is in the hotel," they would be slower to respond because all of the related and irrelevant nodes associated with the nodes of the previous probe are inhibited.

Proceeding from this view, the obvious counterpart is to test whether suppression operates in the other direction as well: specifically, the inhibition of irrelevant object nodes. For example, when a person is presented with the probe "The oak counter is in the hotel," the other object nodes associated with the hotel node, namely, the welcome mat and potted palm nodes, are inhibited. So, from this perspective, if a person is then presented with another fact about these objects, such as "The welcome mat is in the movie theater," they should be slower to respond because the object node has been inhibited.

This could be tested easily in the current paradigm by simply reversing the order of the primes and targets in the experimental and control recognition-test pairs. Experiment 3 performed just this test. According to a concept network with inhibition model, one would expect to observe inhibition in this case as well. This sort of analysis assumes that concepts, not situation models, are the primary appropriate unit of analysis.

However, according to a situation-model view, people store facts in separate situation models. When a fact about an object being in a location is retrieved, if the object is in other locations, those locations are inhibited. For example, based on Figure 3, when a person is presented with the probe "The welcome mat is in the movie theater," the other situation models containing the welcome mat, namely, the cocktail lounge and hotel models, are inhibited. However, if there are additional objects in the probed location, these objects are not inhibited because they are part of the model that is being retrieved. It is unlikely that additional locations that contain these unmentioned objects will be inhibited. For example, when a person is presented with the probe "The oak counter is in the hotel," the other objects in that location, such as the welcome mat and the potted palm, are retrieved as being part of that model if they have been integrated. Although these additional objects are also in other situation models, this should have no impact on retrieval because these models are not competing with the retrieval of the target model. This can be seen in the attenuated fan effect in the single location conditions. So, if a person is then presented with another fact about these objects, such as "The welcome mat is in the movie theater," response times should be unaffected relative to an unrelated control condition. Thus, the situation-model view predicts no suppression in this case.

Experiment 3

Method

Participants. Forty-eight native English speakers were tested. They were recruited from the pool at the University of Notre Dame and given partial class credit for their participation. Two additional people replaced 2 of the original participants who had exceptionally slow response times (i.e., more than three standard deviations above the mean).

Materials and procedure. The same materials and design were used in Experiment 3 as in Experiment 2 except that the order of the primes and targets was reversed on the experimental and control pairs of the recognition test. Again, there were no breaks during the recognition test. An average of 4.8 (SD = 1.5) study–test cycles was required to memorize the study sentences. For the prime analysis, 3.7% of the target trials were identified as outliers, along with an additional 2.4% dropped because of errors on the prime trial. For the differential fan effect, 4.0% were discarded in the data trimming.

Results

Consistent with the situation-model view, there was no evidence of negative priming in Experiment 3. People responded to experimental targets just as quickly as they did to control targets. In addition, the differential fan effect was again observed, which is also consistent with the idea that people are using situation models.

Priming analysis. The response-time analyses of the control and experimental target trials are summarized in Figure 6. Consistent with the situation-model view, there was no difference between the control (1,606 ms, SE = 59) and the experimental targets (1,606, SE = 47), $F < 1$. There was also no difference between the response times to the control primes (1,663 ms, SE = 61) and to the experimental primes (1,650 ms, SE = 65), $F < 1$.

The overall error rate for the control and experimental trials was 4.8%. People made a similar number of errors to control targets (4.9%, SE = .98) as to experimental targets (4.7%, SE = .84), $F < 1$. People also made similar numbers of errors to control primes (3.4%, SE = .67) as to experimental primes (2.5%, SE = .52), F = 1.09.

Differential fan effects. The response-time and error-rate fan-effect data are summarized in Table 3. For the response-time data, there was a significant main effect of fan, $F(2, 94)$ = 11.82, MSE = 117,249, and a marginally significant effect of condition, $F(1, 47)$ = 3.67, MSE = 180,819, p = .06, as well as a significant Condition × Fan interaction, $F(2, 94)$ = 4.40, MSE = 88,826. Simple effects tests showed that the fan effect was significant for the multiple location condition, $F(2, 94)$ = 13.78, MSE = 114,075, but not for the single location condition, $F(2, 94)$ = 2.23, MSE = 92,000, p = .11.

In addition to the effects of primary interest, there was a significant main effect of studied–nonstudied, $F(1, 47)$ = 78.12, MSE = 97,609, with people responding faster to

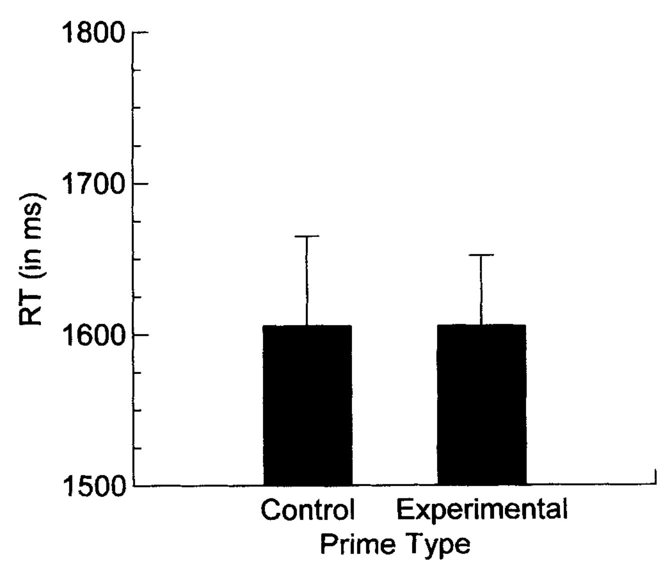

Figure 6. The response-time (RT) data for the target trials in Experiment 3. Error bars reflect standard errors.

Table 3
Differential Fan-Effect Analysis for Experiment 3

	Studied						Nonstudied					
	Fan level 1		Fan level 2		Fan level 3		Fan level 1		Fan level 2		Fan level 3	
Condition	Response time (ms)	Error rate (%)	Response time (ms)	Error rate (%)	Response time (ms)	Error rate (%)	Response time (ms)	Error rate (%)	Response time (ms)	Error rate (%)	Response time (ms)	Error rate (%)
Single location	1,327	3.0	1,357	2.1	1,433	1.0	1,577	1.9	1,583	1.6	1,647	3.1
Multiple location	1,367	3.0	1,464	3.0	1,489	3.7	1,465	1.4	1,699	1.6	1,846	3.5
			1,456	2.1	1,503	3.1						

Note. Response times in milliseconds and error rates in percentages for the differential fan-effect analysis in Experiment 3. Corrected response times for the multiple location condition are presented beneath the uncorrected times.

studied probes (1,406 ms) than to nonstudied probes (1,636 ms). There was also a significant Studied–Nonstudied × Fan interaction, $F(2, 94) = 3.26$, $MSE = 45{,}589$, and a significant Studied–Nonstudied × Condition × Fan interaction, $F(2, 94) = 4.10$, $MSE = 63{,}992$. To break this interaction down, I did separate analyses of single location and multiple location condition data. There was a significant Studied–Nonstudied × Fan interaction for the multiple location condition data, $F(2, 94) = 7.65$, $MSE = 52{,}683$, with a larger fan effect for the nonstudied probes relative to the studied probes, $F(2, 94) = 2.67$, $MSE = 74{,}925$, $p = .07$. However, for the single location condition data, the Studied–Nonstudied × Fan interaction was not significant, $F < 1$.

As in Experiments 1 and 2, the data were submitted to a corrected analysis. Because of the way the priming trials were constructed, now only the multiple location condition was affected. As can be seen in Table 3 and consistent with the absence of a negative priming effect, this correction had little effect on the results. Because the results of the corrected analysis were the same as those of the uncorrected analysis, they are not reported.

The error rate on the fan-effect trials overall was 2.4%. An analysis of the error-rate data revealed no significant effects.

Relation between negative priming and fan effects. As in Experiments 1 and 2, the data for negative priming and fan effect were submitted to correlation analysis. Consistent with the absence of a negative priming effect and very little fan effect in the single location condition, these two variables were not correlated, $r = .02$, $p = .92$.

In addition, the negative priming data were analyzed in terms of whether one or two irrelevant objects were in the situation model. These data were then submitted to a 2 (prime: experimental vs. control) × 2 (set size: one or two irrelevant objects) repeated-measures ANOVA. The interaction did not reach significance, $F(1, 47) = 1.54$, $MSE = 161{,}085$. To be consistent with Experiments 1 and 2, I performed simple effects tests. These showed no significant negative priming effects either for the one irrelevant object trials (experimental = 1,809 ms, $SE = 93$; control = 1,890 ms, $SE = 99$; $F < 1$) or for the two irrelevant models trials (experimental = 1,688 ms, $SE = 60$; control = 1,625 ms, $SE = 54$; $F(1, 47) = 1.23$, $MSE = 76{,}256$, $p = .27$).

Discussion

Consistent with the situation-model view, no clear evidence of inhibition was observed in Experiment 3. Inhibition of this type of information is seen only when an object in the probe sentence is in another, irrelevant location. Suppression is not observed when a location contains additional irrelevant objects. This is consistent with the idea that suppression is confined to separate representations in memory and that, in the present case, these representations are identified by location. Thus, the suppression mechanism that is operating here appears to be focused primarily on situation models rather than on individual concepts.

General Discussion

The analysis of the prime–target pairs suggests that the selection of a situation model in memory involves the inhibition of related but irrelevant models because responses on experimental trials were slower than responses on control trials. As such, this is a negative priming effect (see, e.g., Tipper, 1985).

The current results fall under the rubric of retrieval-based inhibition (see, e.g., M. C. Anderson & Spellman, 1995). Although the paradigm used by M. C. Anderson and colleagues assessed the availability of target items given the repeated practice of related items, the current experiments used negative priming to show that retrieval-based inhibition can be observed as a consequence of having retrieved a related representation on the previous trial. Thus, the current results further support the idea that the retrieval of information can involve suppression. This occurs when competing memory traces are sufficiently related to the target trace that they produce interference. So, the fan effect and the negative priming effect found here appear to be related. Memory retrieval is an attentional task (Carrier & Pashler, 1995). People are not able to access large amounts of information in memory without penalty. As more information becomes heavily involved in the memory-retrieval situation, it is more likely that deficits will be observed. Hence, it should not be too surprising that general attention-based effects are found in memory retrieval.

Consistent with previous research (e.g., Radvansky & Zacks, 1991), a differential fan effect was observed, suggesting the use of situation models. For the single location condition, the fan effect was absent or attenuated; for the multiple location condition, a clear fan effect was observed. It should be noted that the fan effect in the single location condition was more pronounced in Experiments 2 and 3 than in previous similar research. A likely explanation for this has to do with the plausibility of some of the study items. The primary motivation in selecting the object and location concepts in Experiments 2 and 3 was that they be shorter than those used in Experiment 1. Though many of the study sentences that were generated using these concepts were quite plausible, such as "The copier is in the library," others may have been less so, such as "The copier is in the bar." As a result, this likely compromised the students' identification of several sentences as referring to the same situation or compromised their willingness to draw the inference that multiple items were coreferential. Thus, across all students, integration may have been less complete and some information stored in separate models. However, it should also be noted that evidence for a fan effect in the single location condition is largely confined to the difference between the Fan 2 and Fan 3 probes. This is consistent with the simple idea that the integration of two pieces of information is relatively easy, whereas the integration of three is more difficult. The point of the differential fan effect is that there is a substantial difference in the size of the fan effects in the two conditions. The point is not that no fan effect will be observed in one condition (see Radvansky, 1999, for a more extensive consideration of fan effects in conditions such as the single location condition).

More importantly, this does not compromise the main point of this article. Specifically, the focus of this article is how retrieval is affected when information is stored in separate situation models rather than when this information is integrated into a common model. The fact that there was less integration observed in Experiments 2 and 3 than in Experiment 1 is of less concern.

Alternative Account of Negative Priming

The current experiments have been described from an inhibition view, primarily because this is the dominant interpretation of negative priming (see, e.g., Tipper, 1985). However, this is not the only framework that has been used to describe this effect. A prominent alternative in the visual attention literature is the episodic retrieval hypothesis (for comparisons of the inhibition and episodic retrieval views, see Neill & Valdes, 1996; Tipper & Milliken, 1996).

According to the episodic retrieval hypothesis (see, e.g., Neill & Valdes, 1992; Neill, Valdes, Terry, & Gorfein, 1992), the processing of a target item involves not only information presented during that trial but also the implicit and automatic use of the episodic information from the prime trial. The basic idea is that the negative priming effect is produced when information from the prime and target trials is in conflict. Specifically, the representation for the target item is marked as irrelevant in the memory trace for the prime trial but as relevant in the target trial. The added time needed to resolve this conflict gives rise to the negative priming effect.

In the episodic retrieval hypothesis, one factor that is thought to mediate the occurrence of negative priming is the distinctiveness of the representations used during the prime trial. The more distinctive the irrelevant information is, the better it is processed. The better it is processed, the more it can interfere on the subsequent trial. This idea was supported by findings in visual attention experiments showing that the length of the interval before the prime trial affects the size of the negative priming effect (Neill et al., 1992). The larger time interval makes the memory for the prime trial more distinctive and thus easier to retrieve, thereby producing more interference.

For the current experiments, this idea of the distinctiveness of irrelevant distractors can be translated into the number of irrelevant situation models associated with the prime memory probe (see Neill et al., 1992). In Experiments 1 and 2, in addition to the relevant model, there were either one or two irrelevant situation models that contained the object concept in the memory probe. The prime representations should have been more distinct when there was only one irrelevant model, because there was no additional representation to compete with it. Therefore, the episodic retrieval hypothesis would predict that negative priming is greater when there is only one irrelevant model rather than two.

The data relevant to this point were presented in the *Results* sections of Experiments 1, 2, and 3. Although there was nominally greater negative priming in the one than in the two irrelevant models condition in Experiment 2, this difference was not significant and was not observed in Experiments 1 and 3. Thus, although across Experiments 1 and 2 there was 24 ms more negative priming in the two irrelevant models condition, no clear result emerged. Still, overall, there is not the sort of support for the episodic retrieval hypothesis that one would expect.

J. R. Anderson and Reder (1999)

The current research suggests that, to effectively capture the processes involved in memory retrieval, one needs to take into account the operation of suppression. The current experiments have the most obvious implications for the ACT family of models (J. R. Anderson, 1976, 1983, 1993), which are often used to interpret fan-effect data. However, the ACT models are not alone in this regard. There are many other memory models that lack such a suppression mechanism (see M. C. Anderson & Bjork, 1994).

Recently, J. R. Anderson and Reder (1999) tested for inhibition within the ACT-model framework. This was a fan-effect study that was modeled on the retrieval-practice paradigm of M. C. Anderson and Spellman (1995), although there were some major differences. Specifically, the information was learned experimentally, the items were sentences not concepts, and the repetition occurred during recognition, not study. People memorized 48 sentences, such as "The biker is in the tower," with either two or four associations with each concept. During recognition, items were presented either five times or once per block of trials, of which there were three.

Suppose a person had memorized the following sentences as part of a list: "The biker is in the tower," "The biker is in the factory," and "The writer is in the factory." Furthermore, assume that the item "The biker is in the tower" was repeated. Responses to repeated items, such as "The biker is in the tower," were facilitated. There was also a modified fan effect. For example, the biker–factory association of the once-presented item grew weaker compared with the biker–tower association of the repeated item. Importantly, J. R. Anderson and Reder (1999) suggested that if suppression were operating, then the concept "factory" should be inhibited. Thus, responses to other "factory" facts, such as "The writer is in the factory," should be slowed. However, they found no evidence of such inhibition.

J. R. Anderson and Reder's (1999) results differ from those reported here. They did not find evidence of inhibition, whereas I did. A solution to this contrast may lie in a consideration of several issues. First, their study indirectly tested whether suppression is involved in the fan effect through the overall effects of repetition. A more direct way to test this is to assess whether related information is suppressed as a consequence of retrieving certain items, as was done here.

A second issue is the length of the recognition test. Some studies have found that as memory testing progresses, inhibition effects get weaker. As reported in Experiment 1, negative priming was greater in Block 1 than Block 2. Also, M. C. Anderson et al. (1997), using a repeated practice paradigm, reported greater inhibition effects during earlier test blocks. The cause for this decline is unclear at this time. Nevertheless, because J. R. Anderson and Reder (1999) averaged across three blocks of trials, this may have reduced the ability to observe inhibition.

Third, there were a large number of breaks during testing (i.e., 11, one after every 36 trials). These breaks may have given any developing inhibition an opportunity to dissipate. If so, this constant disruption served to make the detection of inhibition more difficult.

Fourth, there is the issue of how the recognition test was structured. Items were presented in a random order within each block. As such, it was possible for the critical once-presented items to occur before all of the repeated items were presented. If the inhibition is long-lasting, then this is less of a problem. However, as stated above, later blocks of trials are less likely to show evidence of inhibition. If a way of correcting for this is to look at the data from earlier blocks, this sort of randomization compromises such an effort.

Finally, J. R. Anderson and Reder (1999) used an analysis at the concept level, although other research has shown that analyses at the situation-model level can be more informative (Radvansky, 1999; Radvansky & Zacks, 1991). For information that can readily be interpreted as referring to a situation, people, when they have the time and resources to do so, create situation models (see, e.g., Radvansky & Zacks, 1997; Zwaan, 1994; Zwaan & Radvansky, 1998).[3] J. R. Anderson and Reder (1999) used sentences about people in locations that can contain several people at once. Situation models are configurations of different types of information (see Zwaan & Radvansky, 1998). This opens the possibility for the type of information used by J. R. Anderson and Reder to be organized using either the location or the person components (much like a story protagonist). Radvansky et al. (1993) suggested that for sentences like these, both location-based and person-based organizations are plausible (a location can contain several people at once, and people are agents that can go from place to place) and no clear organizational bias is observed (see also Taylor & Tversky, 1997). There is no way to tell which participants were using which organization. Thus, it is not possible to see if related but irrelevant situation models were inhibited. Their "inhibited" items sometimes corresponded to an inhibited situation model and sometimes not. This weakens the ability to detect the operation of suppression.

In summary, given the evidence for inhibition in memory retrieval from several studies (e.g., M. C. Anderson et al., 1994, 1997; M. C. Anderson & McCulloch, 1999; M. C. Anderson & Spellman, 1995; May et al., 1997) including the current one, the null result of J. R. Anderson and Reder's (1999) one experiment seems to be an exception. There appear to be some methodological issues that may have compromised their test for inhibition.

Conclusion

There remain unexplored a number of related issues having to do with the boundary conditions of retrieval-based inhibition, such as the time course of suppression and shifts in suppression over the course of testing. However, the current experiments establish that suppression is involved in memory retrieval and that this inhibition can operate at the level of situation models. These two points need to be taken into account in the further development of models of human memory retrieval.

[3] M. C. Anderson and Spellman (1995) also used the concept level as their unit of analysis. This is appropriate in their case as they used individual concepts during study, not sentences. For example, it would be difficult to form a situation model from "lettuce."

References

Anderson, J. R. (1974). Retrieval of propositional information from long-term memory. *Cognitive Psychology, 6,* 451–474.

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

Anderson, J. R. (1983). *The architecture of cognition.* Cambridge, MA: Harvard University Press.

Anderson, J. R. (1993). *Rules of the mind.* Hillsdale, NJ: Erlbaum.

Anderson, J. R., & Reder, L. (1999). The fan effect: New results and new theories. *Journal of Experimental Psychology: General, 128,* 186–197.

Anderson, M. C., & Bjork, R. A. (1994). Mechanisms of inhibition in long-term memory: A new taxonomy. In D. Dagenbach & T. H. Carr (Eds.), *Inhibitory processes in attention, memory, and language* (pp. 265–325). San Diego, CA: Academic Press.

Anderson, M. C., Bjork, R. A., & Bjork, E. L. (1994). Remembering can cause forgetting: Retrieval dynamics in long-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 1063–1087.

Anderson, M. C., de Kok, D., & Childs, C. (1997, November). *Retrieval-induced forgetting on a test of recognition memory.* Poster presented at the 38th Annual Meeting of the Psychonomic Society, Philadelphia, PA.

Anderson, M. C., & McCulloch, K. C. (1999). Integration as a general boundary condition on retrieval-induced forgetting. *Journal of Experimental Psychology: Learning, Memory and Cognition, 25,* 608–629.

Anderson, M. C., & Neely, J. H. (1996). Interference and inhibition in memory retrieval. In E. L. Bjork & R. A. Bjork (Eds.), *Memory: Handbook of perception and cognition* (2nd ed., pp. 237–313). San Diego, CA: Academic Press.

Anderson, M. C., & Spellman, B. A. (1995). On the status of inhibitory mechanisms in cognition: Memory retrieval as a model case. *Psychological Review, 102,* 68–100.

Bjork, R. A. (1970). Positive forgetting: The noninterference of items intentionally forgotten. *Journal of Verbal Learning and Verbal Behavior, 9,* 255–268.

Bjork, R. A. (1989). Retrieval inhibition as an adaptive mechanism in human memory. In H. L. Roediger & F. I. M. Craik (Eds.), *Varieties of memory and consciousness: Essays in honour of Endel Tulving* (p. 309–330). Hillsdale, NJ: Erlbaum.

Blaxton, T. A., & Neely, J. H. (1983). Inhibition from semantically related primes: Evidence of a category-specific inhibition. *Memory & Cognition, 11,* 500–510.

Cantor, J., & Engle, R. W. (1993). Working-memory capacity as long-term memory activation: An individual differences approach. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19,* 1101–1114.

Carlson-Radvansky, L. A., & Jiang, Y. (1998). Inhibition accompanies reference-frame selection. *Psychological Science, 9,* 386–391.

Carrier, L. M., & Pashler, H. (1995). Attentional limits in memory retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 1339–1348.

Dagenbach, D., & Carr, T. H. (1994). Inhibitory processes in perceptual recognition: Evidence for a center-surround attentional mechanism. In D. Dagenbach & T. H. Carr (Eds.), *Inhibitory processes in attention, memory, and language* (pp. 327–357). San Diego, CA: Academic Press.

Fox, E. (1995). Negative priming from ignored distractors in visual selection: A review. *Psychonomic Bulletin & Review, 2,* 145–173.

Gerard, L. D., Zacks, R. T., Hasher, L., & Radvansky, G. A. (1991). Age deficits in retrieval: The fan effect. *Journal of Gerontology, 46,* 131–136.

Gernsbacher, M. A., & Faust, M. E. (1991). The mechanism of suppression: A component of general comprehension skill. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17,* 245–262.

Hasher, L., & Zacks, R. T. (1988). Working memory, comprehension, and aging: A review and a new view. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 22, p. 193–225). San Diego, CA: Academic Press.

Johnson-Laird, P. N. (1983). *Mental models: Towards a cognitive science of language, inference, and consciousness.* Cambridge, MA: Harvard University Press.

Kintsch, W. (1998). *Comprehension: A paradigm for cognition.* Cambridge, England: Cambridge University Press.

May, C. P., Einstein, G. O., & Knight, S. (1997, November). *Retrieval-induced forgetting and prospective memory.* Poster presented at the 38th Annual Meeting of the Psychonomic Society, Philadelphia, PA.

May, C. P., Kane, M. J., & Hasher, L. (1995). Determinants of negative priming. *Psychological Bulletin, 118,* 35–54.

McNamara, T. P. (1992). Theories of priming: I. Associative distance and lag. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 1173–1190.

Myers, J. L., O'Brien, E. J., Balota, D. A., & Toyofuku, M. L. (1984). Memory search without interference: The role of integration. *Cognitive Psychology, 16,* 217–242.

Neill, W. T., & Valdes, L. A. (1992). Persistence of negative priming: Steady state or decay? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 565–576.

Neill, W. T., & Valdes, L. A. (1996). Facilitatory and inhibitory aspects of attention. In A. F. Kramer, M. G. H. Coles, & G. D. Logan (Eds.), *Converging operations in the study of visual selective attention* (pp. 77–106). Washington, DC: American Psychological Association.

Neill, W. T., Valdes, L. A., & Terry, K. M. (1995). Selective attention and the inhibitory control of cognition. In F. N. Dempster & C. J. Brainerd (Eds.), *Interference and inhibition in cognition* (pp. 207–261). San Diego, CA: Academic Press.

Neill, W. T., Valdes, L. A., Terry, K. M., & Gorfein, D. S. (1992). Persistence of negative priming: II. Evidence for episodic trace retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 993–1000.

Neumann, E., & DeSchepper, B. G. (1992). An inhibition-based fan effect: Evidence for an active suppression mechanism in selective attention. *Canadian Journal of Psychology, 46,* 1–40.

Radvansky, G. A. (1998). The organization of information retrieved from situation models. *Psychonomic Bulletin & Review, 5,* 283–289.

Radvansky, G. A. (1999). The fan effect: A tale of two theories. *Journal of Experimental Psychology: General, 128,* 198–206.

Radvansky, G. A., Spieler, D. H., & Zacks, R. T. (1993). Mental model organization. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19,* 95–114.

Radvansky, G. A., Wyer, R. S., Curiel, J. C., & Lutz, M. F. (1997). Situation models and abstract ownership relations. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23,* 1233–1246.

Radvansky, G. A., & Zacks, R. T. (1991). Mental models and the fan effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17,* 940–953.

Radvansky, G. A., & Zacks, R. T. (1997). The retrieval of situation-specific information. In M. A. Conway (Ed.), *Cognitive models of memory* (pp. 173–213). Cambridge, MA: MIT Press.

Radvansky, G. A., Zacks, R. T., & Hasher, L. (1996). Fact retrieval in younger and older adults: The role of mental models. *Psychology and Aging, 11,* 258–271.

Radvansky, G. A., Zwaan, R. A., Federico, T., & Franklin, N. (1998). Retrieval from temporally organized situation models. *Journal of Experimental Psychology: Learning, Memory and Cognition, 24,* 1224–1237.

Reder, L. M., & Anderson, J. R. (1980). A partial resolution of the paradox of interference: The role of integrating knowledge. *Cognitive Psychology, 12,* 447–472.

Simpson, G. B., & Kang, H. (1994). Inhibitory processes in the recognition of homograph meanings. In D. Dagenbach & T. H. Carr (Eds.), *Inhibitory processes in attention, memory, and language* (pp. 359–381). San Diego, CA: Academic Press.

Taylor, H. A., & Tversky, B. (1997). Indexing events in memory: Evidence for index dominance. *Memory, 5,* 509–542.

Tipper, S. P. (1985). The negative priming effect: Inhibitory priming by ignored objects. *Quarterly Journal of Experimental Psychology, 37A,* 571–590.

Tipper, S. P., & Milliken, B. (1996). Distinguishing between inhibition-based and episodic retrieval-based accounts of negative priming. In A. F. Kramer, M. G. H. Coles, & G. D. Logan (Eds.), *Converging operations in the study of visual selective attention* (pp. 337–363). Washington, DC: American Psychological Association.

van Dijk, T. A., & Kintsch, W. (1983). *Strategies of discourse comprehension.* New York: Academic Press.

Yee, P. L. (1991). Semantic inhibition of ignored words during a figure classification task. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology, 43A,* 127–153.

Zacks, R. T., Radvansky, G. A., & Hasher, L. (1996). Studies of directed forgetting in older adults. *Psychology and Aging, 22,* 143–156.

Zwaan, R. A. (1994). Effect of genre expectations on text comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 920–933.

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

Appendix A

Filler-Trial Data

	Fan level					
	3–2		2–3		3–3	
Experiment & condition	Response time (ms)	Error rate (%)	Response time (ms)	Error rate (%)	Response time (ms)	Error rate (%)
Experiment 1						
Studied	1,893	4.4	1,963	3.9	2,060	4.2
	1,832	4.8			2,023	4.2
Nonstudied	2,011	3.8	2,175	3.6	2,207	3.7
Experiment 2						
Studied	1,729	6.1	1,682	6.1	1,824	5.2
	1,677	5.7			1,762	5.6
Nonstudied	1,879	4.7	1,975	4.2	2,016	4.9
Experiment 3						
Studied	1,713	5.4	1,823	7.8	1,880	6.1
			1,784	8.3	1,844	6.0
Nonstudied	1,859	4.0	2,025	2.8	2,122	5.6

Note. Response times (in milliseconds) and error rates (in percentages) for the filler trials in Experiments 1–3. The notation X–Y is used where the first number corresponds to the fan from the location concept and the second number corresponds to the fan from the object concept. Corrected response times for the single location condition are presented beneath the uncorrected times.

Appendix B

Response Times by Fan Level

Experiment & condition	Fan level 2–1	3–1	3–2	3–3	1–2	1–3	2–3	3–3
Experiment 1								
Experimental	1,586	1,583	1,921	2,053				
Control	1,498	1,551	1,808	2,003				
Experiment 2								
Experimental	1,345	1,526	1,743	2,032				
Control	1,285	1,471	1,610	1,623				
Experiment 3								
Experimental					1,478	1,523	1,748	1,892
Control					1,506	1,456	1,744	2,132

Note. Response times (in milliseconds) for the experimental targets broken down by fan level. The notation X–Y is used where the first number corresponds to the fan from the object location and the second number corresponds to the fan from the object concept. As a reminder, there is some deviation between the means presented in the analyses of target trials and the interference analyses. This is a result of the fact that the target trial data reflect the consequences of the removal of trials on which an error was made on the prime trial.

Received April 9, 1997
Revision received December 23, 1997
Accepted January 19, 1999 ■