

Composite Holographic Associative Recall Model (CHARM) and Blended Memories in Eyewitness Testimony

Janet Metcalfe
University of California, San Diego

The idea that compositing or blending occurs in human episodic memory stems from two sources: (a) distributed memory models and (b) studies on the errors that occur in eyewitness testimony. These two traditions of research—*theoretical and empirical*—have been independent and distinct. Here, data from the eyewitness testimony paradigm are simulated by the distributed model CHARM (Composite Holographic Associative Recall Model). Of focal concern are Loftus's studies, which have been interpreted in favor of the blending hypothesis, and McCloskey and Zaragoza's studies, which have been interpreted as refuting Loftus's position. Both of these seemingly contradictory results, as well as recent findings with yes/no recognition, fall out of the model. Finally, the model predicts empirically found color shifts and provides specifications for when blends and memory impairments will and will not be expected.

Central to our understanding of human memory is the issue of how memories are stored. Traditionally, it has been thought that traces must be stored discretely and separately. However, recently a new conception of memory storage—that the traces may actually be superimposed, composited, or blended—has gained ascendancy. This new idea stems from two sources. First, a number of distributed models that are neural in inspiration have had as a central tenet the construct of superimposed storage. It has been argued (Anderson & Hinton, 1981) that the representations from successive events must be superimposed on one another to make use of the same memory "surfaces." Or, as Kohonen, Lehtiö, and Oja (1981) stated,

Because a neural system is an ensemble of a great number of *collectively* interacting elements, it seems more natural to abandon altogether those physical models of memory in which particular concepts correspond to particular spatial locations (nodes) in the hardware. Instead a physically more plausible approach can be based on the assumption that representations of concepts and other pieces of information are stored as *collective states* of a neural network. (p. 107)

A key theoretical advance made by these models is a variety of learning mechanisms or associative operations that illustrate how retrieval is possible, even though many "traces" are stored on the same layer. Individual patterns that represent the discrete events that were initially encoded are retrievable, despite composite storage (Anderson, 1977; Cavanagh, 1976; Hinton & Anderson, 1981; Longuet-Higgins, 1968; McClelland & Rumelhart, 1986; Metcalfe & Murdock, 1981; Metcalfe & Eich, 1982, 1985; Murdock, 1982, 1983; Rumelhart

& McClelland, 1986). The psychological characteristics of the particular solutions to this retrieval problem vary. Composite trace models make a number of new predictions about the characteristics of the events that are retrieved from memory as a result of blended storage. Some of the predictions from one of these models—the Composite Holographic Associative Recall Model (CHARM)—are explored in this article.

The second source of the idea of blended memories developed quite independently of the recent advances in distributed modeling. Careful study of the systematic distortions and errors that people make when they are asked to be eyewitnesses to important life events, such as when they witness a crime, has also suggested that human memory may be subject to blending of single events into composites (Loftus, 1975, 1977, 1979a, 1979b; Loftus & Loftus, 1980; Loftus, Miller, & Burns, 1978). In a series of studies, Loftus and her colleagues have shown that the remembrances of people who have initially witnessed an important event such as an auto accident or a purse snatching can be altered by the introduction of misleading information that occurs after the event in question. For instance, subjects were shown a series of slides depicting an auto accident. One of the slides showed a car going through a stop sign. Later, some of the subjects were provided with misleading information that indicated that the sign was actually a yield sign, rather than a stop sign. When queried later, the misled subjects were less likely than control subjects to give the correct response concerning the original event. In a different example, subjects viewed a slide sequence in which they saw a green car, in one frame only. Later, it was intimated that the car was blue. When asked to make a judgment about the color of the car, the misled subjects' responses were systematically distorted toward the blue end of the color spectrum. These results, and many others of a similar ilk, have been interpreted by Loftus and her colleagues as indicating either that the initial event is interfered with by the misleading information or that the separate representations of both events are integrated or blended in memory.

Loftus's interpretation of the eyewitness testimony data is not uncontroversial. Indeed, there has been a recent upsurge of debate and experimentation on this eyewitness testimony

This research was facilitated by National Institute of Mental Health Grant 7 R03MH43542-02.

I am grateful for the help and thoughtful suggestions of Gary Cottrell, Earl Hunt, Elizabeth Loftus, Felice Orlich, Jonathon Schooler, and the members of the Parallel Distributed Processing Natural Language Processing Group at the University of California, San Diego.

Correspondence concerning this article should be addressed to Janet Metcalfe, Department of Psychology C-009, University of California, San Diego, La Jolla, California 92093.

paradigm (Bekerian & Bowers, 1983; Belli, 1989; Ceci, Ross, & Toglia, 1987a, 1987b; Ceci, Toglia, & Ross, 1988; Chandler, 1989; Christiaansen & Ochalek, 1983; Geiselman, 1988; Gibling & Davies, 1988; Gudjonsson, 1986; Hammersley & Read, 1986; Kohnken & Brockmann, 1987; Kroll & Ogawa, 1988; Kroll & Timourian, 1986; Lehnert, Robertson, & Black, 1984; Lindsay & Johnson, 1987, 1989; Loftus & Hoffman, 1989; Loftus, Schooler, & Wagenaar, 1985; McCloskey & Zaragoza, 1985a, 1985b; Morton, Hammersley, & Bekerian, 1985; Pirolli & Mitterer, 1984; Register & Kihlstrom, 1988; Schooler, Gerhard, & Loftus, 1986; Sheehan, 1988; Smith & Ellsworth, 1987; Tousignant, Hall, & Loftus, 1986; B. Tversky & Tuchin, 1989; Wagenaar & Boer, 1987; Zaragoza & Koshmider, 1989; Zaragoza & McCloskey, 1989; Zaragoza, McCloskey, & Jamis, 1987). The controversy stems largely from a series of experiments conducted by McCloskey and Zaragoza (1985a) in which (a) they replicated Loftus's basic findings with the misleading-information paradigm and (b) they changed the distractor items at the time of testing. The finding in this second case was that there was no decrement in performance with the target items so long as the distractor items themselves were not presented as lures. On the basis of these findings, they argued that there was no loss or distortion of the initially encoded events. They concluded, correctly, that one could not infer from Loftus's data that the initial events had been lost irretrievably: Under particular conditions (in which, it may be argued, blending would be impossible to detect), there appeared to be no interference with the initial event. This conclusion, though correct, is also just a restatement of the data.

With regard to the blending or integration hypothesis, McCloskey and Zaragoza (1985a) stated,

What sorts of data would, then, support or disconfirm the integration claim? Consideration of this question leads quickly to the realization that what is meant by integration is not at all clear. One might suggest that the claim simply asserts that information from various sources is stored together in memory. Although this answer may be satisfying at an intuitive level, it loses much of its appeal when we ask, What does "stored together in memory" mean? . . . Our point is simply that the integration claim, as it typically appears in the memory literature, is so vague and ambiguous as to be virtually meaningless. Unless the claim is made more specific, we cannot determine whether it is reasonable, what its implications are, or what sorts of data would serve to support or disconfirm it. If we are to progress in our understanding of human memory, we must relinquish vague claims of this sort in favor of specific proposals. (pp. 15-16)

Until now, no distributed model of human memory has been applied to the eyewitness testimony paradigm. Thus if what McCloskey and Zaragoza (1985a) meant by the "memory literature" was only the eyewitness testimony literature, then they were right in asserting that the blending claim is vague. But they were wrong to insist that the idea of integration or blending is, in principle, vague. That claim is specifically instantiated by distributed models, and in fact those models are very precise in delineating exactly what is meant by "stored together in memory." The assumptions of those models have been very carefully laid out, and the operations and mechanisms by which blending occurs have been specified in detail.

The objective of this article is to alleviate the situation of specific formulations existing in one literature—the realms of parallel distributed processing, connectionist networks, and theoretical studies of human memory—but not making contact with a different literature—the area of the eyewitness testimony involving the actual behaviors of people in ecologically valid, and important, situations. In this article, a specific model of human episodic memory in which the construct of composite storage is an inherent, indeed a central, construct is applied to the eyewitness testimony paradigm in which the issue is of both theoretical and practical concern. As will be shown by computer simulations, the data presented by Loftus and those of McCloskey and Zaragoza fall out automatically from this "integration, or composite" model. More recent data (Belli, 1989; B. Tversky & Tuchin, 1989) for which a yes/no paradigm was used, and which have led to even more complicated conjectures (but not to specific mechanistic models) about the nature of human memory, also fall out, with no additions or changes to the model. Finally, the critical finding that people do choose intermediate blended responses (Loftus et al., 1978), when given the opportunity to display these blends, is predicted by the model, as the simulations that follow show.

The Eyewitness Paradigm

The situation that was modeled in the first series of simulations was given by McCloskey and Zaragoza (1985a). Half of their experiment was based on Loftus et al.'s (1978) design, but they also included a different control condition that they contended undermines Loftus's interpretation of her data. Thus it seems appropriate to consider both halves of the design. McCloskey and Zaragoza (1985a) presented their subjects with a series of 79 color slides depicting an incident in which a maintenance man entered an office, repaired a chair, found and stole \$20, and then left. Embedded in the sequence were a number of critical slides, only one of which is of concern in this discussion, inasmuch as the others were included for purposes of counterbalancing across conditions or for other technical reasons. The critical slide was one in which the subjects saw the man pick up a hammer from the tool kit. After viewing the slide sequence, subjects read a narrative in which the misleading information (in the experimental conditions) was embedded. In the case of interest, it was suggested to subjects that the tool that the man had picked up was a screwdriver. In the control condition, a generic term—*tool*—was used to indicate the event in question.

At the time of testing, subjects were asked to fill in the blank of the following statement: "The man slid the calculator beneath a _____ in his tool box." The alternatives given were, in what is here designated the standard test, "hammer" and "screwdriver." In the modified test, the alternatives were "hammer" and "wrench." A summary of the whole design is shown in Table 1, along with the average percentage correct over six replications. As the table shows, the misled group scored considerably more poorly than did the control (unmisled) group in the standard test (i.e., when the misleading information was presented at the time of testing), qualitatively

Table 1
 Summary of the Experimental Paradigm Modeled in Simulations 1-3

Slides seen	Misleading information	Test alternatives	Result
Standard conditions			
Control: man-hammer	—	hammer vs. screwdriver	72%
Experimental: man-hammer	man-screwdriver	hammer vs. screwdriver	37%
Modified conditions			
Control: man-hammer	—	hammer vs. wrench	75%
Experimental: man-hammer	man-screwdriver	hammer vs. wrench	72%

replicating Loftus's results. However, when the misleading information was not presented at the time of testing, as in the modified conditions, there was no decrement in performance on the target information.

After describing the model and how that model is applied to this experimental situation, I discuss some other findings within the eyewitness testimony paradigm and the predictions that the CHARM model makes with respect to them. However, I now turn to a description of a specific distributed model that incorporates the idea of a composite memory trace (and hence memorial blends) as a fundamental construct.

Description of the CHARM Model

The model used to investigate the eyewitness testimony paradigm is called the CHARM model. This model was not initially formulated to deal with this particular paradigm, and it has, in fact, been applied with some success to a variety of other classic memory situations. The model is associative in nature, based on the idea that items, represented as distributed patterns of features or as vectors, are associated by the operation of convolution. The results are stored by being added into a composite trace and hence "blended" or superimposed. Retrieval occurs by the operation of correlation, which results in a noisy and sometimes systematically distorted output. That output, in recall, is then matched to all of the items in a lexicon, and the best match wins and is given as the recalled item. In the present situation, because only two alternatives were given in the experimental situation, the retrieved item is compared with only those two items, and the best match is said to win or to be the response of the simulated subject.

Although I have described the assumptions and mechanisms in this model in more detail elsewhere (Metcalfe Eich, 1982, 1985), I review them here. An overview of the model is given in Figure 1.

Representation

Items, in the model, are represented as patterns of elements or features. The idea that people's memorial representations consist of features is an old one. It allows models of memory to capture a number of psychologically relevant phenomena,

such as the fact that the similarity of items to one another may vary, that certain aspects may be present in some items and not in others, that items may be decomposed and analyzed into smaller parts, and that the relations among items may be quite complex (see, for instance, A. Tversky, 1977). The fact that, in the model, the entire *pattern* of features is what is important allows for encoding variability over different instances of the same item and for a certain noise resistance, which is also characteristic of other aspects of the model.

It is not necessary, though, for any of the work on the model, to assume that individual features are interpretable as entities such as bars, edges, curves, shape of the eyebrow, length of the nose, or other entities that have been traditionally labeled as features. Rather, without any loss of explanatory power of the model, the elements of the items could represent more holistic properties, such as perhaps principal components or even less nameable components. Because of the possibility that these "features" are themselves distributed, they should perhaps be assigned the more neutral term of *holons*. The term *features* is used in the discussions that follow, but the term *holon* could be used interchangeably with *feature*.

Items are formally represented as random vectors that have an expected value of 0 for all elements and that have actual values that vary randomly around this value. The value of any particular element in an item is statistically independent of the values of other elements in the same item. The dot product between two items provides a measure of their similarity. The dot product of any item with itself is, for mathematical convenience, assumed to be 1. The dot product between unrelated items is 0. The representation is assumed to code phonemic, visual, other sensory, and semantic aspects of the items. In other work, I am investigating the implications of not coding some of the features, for specific items and of implementing manipulations (such as encoding variability) and other representational variables. In this article, however, these manipulations are not considered, and I have no reason to think that this should alter the main results discussed here. I did, in some of the simulations to be described, allow that there is some systematic similarity among items. This may be accomplished in the model by feature overlap between items: The more features two items have in common, the more similar they are. By varying the similarity, one can investigate the implications of the factor of category structure or cohesiveness for the blending results that are found empirically.




Figure 1. A schematic overview of the CHARM model.

Association Formation

When two items occur together or in close proximity, they may be associated in the model. The operation of association formation—convolution—is conceptually distinct from the inherent similarity of the items themselves. Both may have an effect on memory, of course, but they are conceptually different. The operation used for association formation is convolution (denoted $*$) and is given by the following equation, for the m th element of the resultant vector:

$$(F * G)_m = \sum_{(i,j) \in S(m)} f_i g_j, \quad (1)$$

where F and G are the item vectors ($f_{-(n-1)/2}, \dots, f_{-1}, f_0, f_1, \dots, f_{(n-1)/2}$) and ($g_{-(n-1)/2}, \dots, g_{-1}, g_0, g_1, \dots, g_{(n-1)/2}$), respectively, that are being associated. $S(m) = \{(i, j) | -(n-1)/2 \leq i, j \leq (n-1)/2, \text{ and } i + j = m\}$. A numerical

example of convolution is given in Figure 2. As the figure shows, the value of every element of one of the initial items is multiplied by the value of every element in the to-be-associated item; this forms a complete matrix. This matrix is then compressed diagonally, by addition, to form a single vector. Even if the elements at the level of the items were interpretable features, the elements that result from convolution are not. Rather, they consist of completely distributed combinations of elements.

This operation of convolution is symmetric; that is, $A * B = B * A$. It potentially gives rise to *two* signal terms: that of the A item and that of the B item. Under certain interesting conditions, both of these signal terms may be simultaneously produced. When the two items are unrelated to one another, one will produce the other (at time of retrieval under the operation of correlation). When, however, the two items are similar to one another, *both* will be produced at time of

retrieval. The goodness of retrieval and, indeed, what is retrieved will thus depend on what it was that was initially associated. Because of this characteristic of the convolution association, it will produce compatibility and congruence effects between the two items associated: a nearly ubiquitous psychological phenomenon not produced by other associative schemes. In short, then, this form of distributed association does not reduce to a behaviorist-style stimulus-response bond. Instead, the association may be more accurately thought of as an interactive gestaltlike entity in which the compatibilities and similarities between the original items figure large.

Storage

As noted earlier, the result of a single convolution is a single vector. The memory trace, in the model, is also a single vector. Successive convolutions resulting from the various associations that are formed during processing are added into the memory trace vector. (Again, in other work, parameters dynamically specifying the weighting of each new input into the trace, and also the decay of the trace, are being investigated, but these studies are not focal to the work under investigation in this article.) Figure 3 depicts the addition of the results of several associations into the composite memory trace. At any given time, there is only *one* memory trace that combines into a single vector the information from all of the convolved pairs, each of which is itself a vector. This is what is meant by "blending" or "stored together in memory" in the present model.

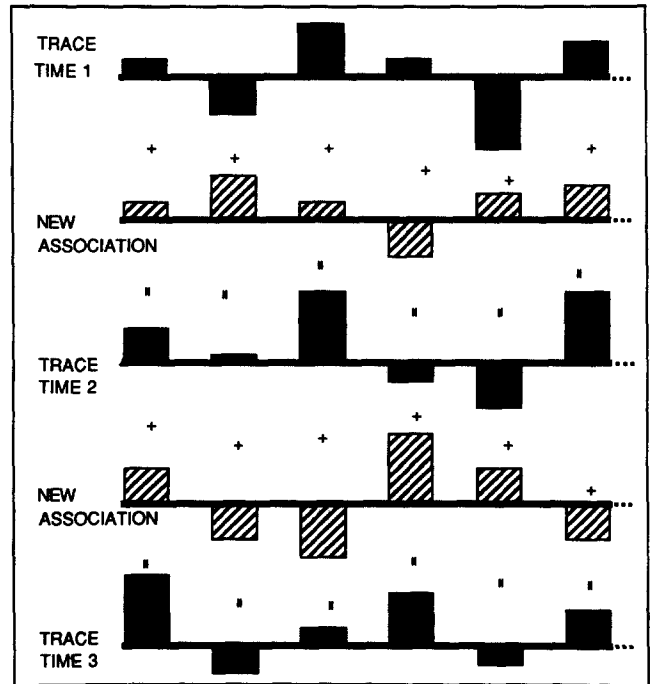



Figure 3. Addition of multiple associations into the composite memory trace. (Bar graphs are of associations and the resultant traces formed by addition.)

$$A = \{ .387, -.632, .317, .447, -.387 \}$$

$$B = \{ -.270, .270, -.50, -.632, .450 \}$$



THE TRACE IS: {-.104, .275, -.450, .037, .641, -.811, .054, .445, -.174}

Figure 2. A numerical example of convolution.

Retrieval

The retrieval operation is correlation, which is defined as

$$R_m = \sum_{(i,j) \in S(m)} f_i g_j \tag{2}$$

where $S(m) = \{(i, j) | -(n - 1)/2 \leq i, j \leq (n - 1)/2, \text{ and } i - j = m\}$. In terms of the model, the cue item F is correlated with the trace, which may contain associations involving F or items related to F . In this case, the items that were associated with the cue F are signal components in the item vector that is the result of retrieval.

The Relation Between Convolution and Correlation

An intuitive summary of the relation between convolution and correlation starts with the observation that the convolution of any vector (say A) with a delta vector results in the vector A itself. A delta vector is a vector composed of zeroes for all elements except for the central element, which has a value of 1. The delta vector is the identity function for convolution:

$$A * \Delta = A.$$

The result of correlating an item with itself, given the constraints outlined earlier in the section on representation, produces an approximation to a delta vector. The central feature of the autocorrelation of, say, $B \# B$ is the dot product between B and B , which is 1. The other features of the autocorrelation function are approximately, but not exactly,

zero. Thus:

$$B\#B = \Delta + \text{error.}$$

If these two relations are combined, one can see that when two unrelated items, A and B , are associated by convolution— $(A*B)$ —and then later one of the items, say A , is used as a retrieval cue and correlated with the trace, the result is the other item, B :

$$\begin{aligned} A\#(A*B) &\approx (A\#A)*B \\ &= \Delta*B + \text{error} \\ &= B + \text{error.} \end{aligned}$$

Similarly, if B is used as the cue, then

$$\begin{aligned} B\#(A*B) &\approx (B\#B)*A \\ &= \Delta*A + \text{error} \\ &= A + \text{error.} \end{aligned}$$

The association thus contains signal components for the two items A and B . Under conditions in which the two items are similar rather than being unrelated, both of these terms may be produced. So, in general, any cue (Q), when correlated with the association, will produce signal and noise terms as follows:

$$Q\#(A*B) = S_{QA}B + S_{QB}A + \text{error}, \quad (3)$$

where Q is the retrieval cue, A and B are items that were associated and entered into the trace, and S_{QA} gives the similarity between the cue and A . Equation 3 can be extended to more complex situations in which multiple associations are added into the composite memory trace. Suppose the trace is $(A*B) + (C*D) + \dots$. The single vector that is the result of retrieval may be broken down into its components, reflecting the original trace and the retrieval cue, as follows:

$$\begin{aligned} Q\#(\text{Trace}) &= Q\#[(A*B) + (C*D) + \dots] \\ &= Q\#(A*B) + Q\#(C*D) + \dots \\ &= S_{QA}B + S_{QB}A + \text{error}_{A*B} \\ &\quad + S_{QC}C + S_{QD}D + \text{error}_{C*D} + \dots \end{aligned}$$

The error terms here have an expected value of zero, and they will not be considered further in this context. Of more interest, for the notion of blends, are the signal terms. The retrieved item vector actually or potentially consists of a variety of signal terms overlaid, superimposed, or added to one another. For instance, suppose that the item that was used as a cue at time of retrieval was associated with two items, B and C , at different points in time. Equation 3 would then become

$$\begin{aligned} A\#(\text{Trace}) &= A\#(A*B) + A\#(A*C) + \dots \\ &= S_{AA}B + S_{AB}A + S_{AA}C + S_{AC}A + \dots \\ &= B + 0 + C + 0 + \text{error} \\ &= B + C. \end{aligned} \quad (4)$$

Most of the potential signal component terms drop out, with expected values of zero, but because the similarity of the cue to itself is 1, both the B and C terms will be retrieved simultaneously and superimposed. This single retrieved vector will match both the original items B and C . However, there

exists the possibility, especially if these two items are similar to one another to begin with or on some kind of continuous scale, that this retrieved blended output from memory will match some *other* item (that may exist) better than it matches either of the two separate signal components. (I suggest later in the article that Loftus's 1977 color-shifting experiment, with the green car and the suggestion that it was blue, may represent a case in point.)

A pictorial analogy to the retrieved item, from a situation in which a cue is associated with two different items that are similar to one another, is given in Figure 4. A slightly different analogy, in which the items are virtually unrelated to one another, is given in Figure 5. These figures show that there may be cases in which the blending inherent in the model will (sometimes, but not all of the time) result in the production of a new entity that may be recognized as a single new and coherent item, combining elements of both of the original items. In other cases, the superposition allows the detection of both items.

Choice of a Response

The response choice stage of the model will depend on the constraints of the experiment. In an unconstrained cued-recall situation, the pattern that is retrieved may be matched to every item in the lexicon. Normally the subject knows that the cue item itself is not a candidate for the target item, and so this possibility is eliminated (see Metcalfe Eich, 1985, for further discussion, and Metcalfe Eich, 1982, for an experiment in which this aspect of the model was tested). In some situations, more than one response may be allowed for a single retrieved item, as in the $A-B/A-C$ paradigm under modified, modified free-recall conditions—that is, when both responses are required. In such a case, the retrieved item is matched to every item in the lexicon, and the *two* items that produce the best match may be given as responses. In a probe-recognition task, in which the subject is asked whether he or she has seen a particular item in a particular context, the output from memory may be matched against the probe.

In the eyewitness testimony test situation given earlier, the output from memory is matched against the representations of only the target and the misleading information in the standard conditions, and the best match is chosen. In the modified conditions, it is matched against the target and a different, nonpresented category member. Other possibilities and constraints exist. For instance, in a later simulation the item retrieved is matched to the probe for a yes/no decision. In Loftus's (1977) color-shifting experiment, the alternatives were an array of 15 color slides. The model chooses from among a similar array. The point here is just that the model uses whatever is in the experiment as the response possibilities.

Alternative Choice Paradigm: Simulation 1a

The first simulation consisted of a computer-modeled experiment of the CHARM model in which an attempt was made to represent the situation portrayed in the misleading-information paradigm. Subsequent simulations provided refinements, extensions, and variants of the first situation.




Figure 4. An analogy to the items that are retrieved from memory. (The two left panels depict single images retrieved under conditions of no specific interference [i.e., in the control conditions]. The right panel depicts a composite, as it is retrieved from memory, in the misled conditions. Here, the blend is coherent. If only the left image were presented for comparison, one would say that the retrieved blend matches it quite well. But in a forced-choice situation in which one had to decide whether the right image [the blend retrieved from memory] matched the left or the center image better, the task would be difficult. There also exists the possibility, with a coherent blend, that a new image might match the retrieved item better than either of the constituent images.)

Methods

A lexicon of 90 items was constructed; each item consisted of 63 features, and each feature consisted of a value randomly selected from a truncated Gaussian distribution with a mean value of 0 and a range of -2 to 2 . The items were then normalized so the dot product of each item with itself was 1. The first item in the lexicon is henceforth called “man”; the second item, “hammer”; the 22nd item, “screwdriver”; the 32nd item “tool”; and the 42nd item, “wrench.” In this first simulation, these items were all statistically independent

of one another, but in later simulations I made them related to one another. In addition, a second set of items with the same structure (and that one could think of as the Coke, 7-Up, and Sunkist orange cans, which were alternate critical items in the experiment) replicated this design in the simulation but with different lexical entries, of course. Two different traces were formed (in independent simulations); the control condition trace and the misled condition trace. The logical structure of the control condition trace was

$$T_c = (\text{man} * \text{hammer}) + 2(\text{man} * \text{tool}) \\ + (\text{irrelevant convolutions}).$$




Figure 5. Like Figure 4 except that the left and center images, when superimposed to yield a blend, do not coalesce. (As before, however, if both the left and center images were presented as alternatives pitted against one another, the decision about which one better matched the blend [shown in the right panel] that is retrieved from memory would be difficult. In this case there exists no real-world object that could match the blend, and so it is called a *noncoherent* blend.)

The structure of the misled condition trace was

$$T_M = (\text{man}\cdot\text{hammer}) + 2(\text{man}\cdot\text{screwdriver}) \\ + (\text{irrelevant convolutions}).$$

As before, the symbol \cdot refers to convolution, as defined in Equation 1. The irrelevant convolutions were included to indicate that there were other events stored in the trace, and the number of them—five—is not too important except that the more of these there are, the worse recall will be because they increase the variance of the trace and the output item (see Metcalfe & Murdock, 1981). One of the five associations in each trace was a different pair (e.g., the Coke on the desk) that was set up with the same logic as the critical pair and allows one to double the number of data points on which each observation is based. The data resulting from this second critical pair varied only randomly from those for the first.

Retrieval was simulated by correlating the vector for "man" with the composite trace. The retrieved vector that resulted was compared to "hammer" and "screwdriver" in the standard test conditions or to "hammer" and "wrench" in the modified test conditions. The comparison was simply the value of the dot product of the output vector and the lexical item in question. This value is called the *resonance* between the retrieved item and the response alternative. The match that gave the highest value was the winner and was said to be the choice made on that particular trial. The entire simulation was run twice, first through 100 independent trials (and so each data point is based on 200 observations) and then again through 500 trials to produce points based on 1,000 observations.

Results

The results of this simulation corresponded to the overall pattern of McCloskey and Zaragoza's (1985a) data (see Table 2). In particular, there was no difference in performance in the two modified conditions, whereas in the standard conditions, the misled group showed poorer performance than did the control group.

Table 2
Simulation Series 1

Condition	Percentage correct	
	200 observations/ point	1,000 observations/ point
Unrelated		
Standard control condition	68.0	64.2
Standard misled condition	32.5	34.1
Modified control condition	70.5	64.8
Modified misled condition	62.0	66.3
Moderately similar		
Standard control condition	58.5	63.1
Standard misled condition	38.0	37.9
Modified control condition	67.0	63.1
Modified misled condition	65.0	64.7
High similarity		
Standard control condition	52.5	55.8
Standard misled condition	37.5	42.0
Modified control condition	54.5	56.7
Modified misled condition	58.0	57.0

Note. This simulation included irrelevant noise in the memory vector.

Categorical Structure: Simulations 1b and 1c

Simulation 1a did not mirror the fact that the critical items in the experiment reflected some categorical structure, or within-category similarity, and so although Item 32 was nominally the category prototype, it and everything else were in fact unrelated. Simulation 1a represents the maximum possible category disintegration. In order to make the simulations more realistic to the actual situation, within-category structure was manipulated in Simulations 1b and 1c. Item 32 was the prototype "tool," and Items 2 (hammer), 22 (screwdriver), and 42 (wrench) were rerepresented to be similar to it. In Simulation 1b, the similarity was moderate, and so 40% of the features were selected by a random draw with replacement to have values on Item 2, Item 22, and/or Item 42 that were the same as those of Item 32. Item 32 was hence more similar to Items 2, 22, and 42 than Items 2, 22, and 42 were to one another, though they were no longer unrelated. In Simulation 1c, 80% of the features were randomly chosen in Items 2, 22, and 42 to overlap with Item 32. The results of Simulations 1b and c are shown in the bottom panels of Table 2. As the numbers presented there show, the basic pattern found in the actual data held up very well, and the numbers from the moderate similarity condition closely reflect those produced in McCloskey and Zaragoza's (1985a) experiments.

Effect of Noise: Simulation 2

In the second simulation series, I investigated the effect of noise in the composite trace on the choice outcomes. Here, as in the first simulation series, the misleading information was more heavily weighted than was the original information. The man·hammer convolution was weighted by a factor of 1; the man·screwdriver (misled condition) or the man·tool convolution was weighted by a factor of 2. Aside from different random seeds, everything else in these simulations was the same as in the previous series except that no additional noise was added into the trace. The two series may be directly compared in order to investigate the effects of noise (which represents interference or a temporal delay in which other events would be expected to be added into the trace).

The target in the standard-test misled conditions was very poorly recalled (in comparison with the heavily weighted and encoded lure; see Table 3). The addition of noise actually makes target recall better in this situation (see Table 4) by making the memory of the presented lure less dominant. In all other cases, however, the deletion of the noise term improved recall of the target item.

Yes/No Paradigm: Simulation 3

Recently, additional data about the misleading-information paradigm have come to light in a sequence of articles and commentaries (Belli, 1989; Loftus & Hoffman, 1989; B. Tversky & Tuchin, 1989; Zaragoza & McCloskey, 1989) in which the researchers investigated the basic misleading and control conditions described earlier, but in which the test involved a yes/no decision rather than an alternative-choice situation.

Table 3
Simulation Series 2

Condition	Percentage correct
Unrelated	
Standard control condition	81.3
Standard misled condition	16.9
Modified control condition	82.5
Modified misled condition	83.7
Moderately similar	
Standard control condition	77.8
Standard misled condition	18.5
Modified control condition	76.4
Modified misled condition	78.9
Highly similar	
Standard control condition	69.1
Standard misled condition	29.8
Modified control condition	65.0
Modified misled condition	67.7

Note. There were 1,000 observations per point, as in Simulation 1, but additional noise was eliminated.

Two series of experiments by independent investigators—Belli (1989) and B. Tversky and Tuchin (1989)—have produced the data that were modeled in Simulation Series 3.

These experiments differed in detail, but basically, in both laboratories, subjects were presented with McCloskey and Zaragoza's (1985a) slide sequence. The experimenters then either suggested that an alternative misleading exemplar had been present in the misled conditions or asked a question about the generic term in the control conditions. At the time of testing, subjects were asked to make yes/no recognition decisions. They were given critical sentences embedded in a sequence of noncritical sentences, some of which were true and some false. In the critical sentence, the item in question was underlined. For example, such a sentence could have been "The man picked up a screwdriver from the tool kit?" Subjects had to say yes or no.

The results of the two experiments are presented in Table 4. The original items viewed in the slide sequence (for which the correct answer is "yes") are here called the *original events*; the misleading information (for which the correct answer is "no") are called *misleading alternatives*; and the category exemplars that were not presented in either the slide or the misleading narrative (for which the correct answer is "no") are here called the *novel lures*. Subjects performed better with

Table 4
Results of Yes/No Experiments in Percentage Correct

Condition	Original event	Misleading alternative	Novel lure
Belli (1989), Experiment 1			
Control	46.5	—	56.3
Misled	28.5	—	84.0
Belli (1989), Experiment 2			
Control	62.5	—	70.1
Misled	36.1	—	83.3
B. Tversky and Tuchin (1989)			
Control	65	—	76
Misled	43	38	76

the original events in the control condition than in the misled condition (see Table 4). Both Belli (1989) and B. Tversky and Tuchin (1989) found that subjects were more likely to say "yes" (wrongly) to the misleading alternatives than to the correct original events. Interestingly, in both of Belli's experiments, subjects were more likely to say "yes" (and thus make an error) to the novel lures in the control condition than in the misled condition. B. Tversky and Tuchin, however, failed to find a difference between the two conditions on responses to the novel lures.

B. Tversky and Tuchin (1989) reported the combined probabilities of saying "yes" to more than one critical item (i.e., for saying "yes" to both the slide and the narrative). The joint probabilities were rather high, as might be expected by a blending model. Nearly half of their misled subjects who said "yes" to the original event also said "yes" to the misleading alternative. This joint probability was lower in the control condition. The design of Belli's (1989) experiment did not permit this analysis.

These simulations were based closely on the first two simulation series. Although other kinds of recognition (based on the explicit retrieval of only the probe term when the probe has been autoassociated and entered into the trace) are currently being formulated in this model, an associative form of recognition is modeled here. In particular, I assumed that subjects would encode the associations as given in the previous simulations. At the time of testing, they would correlate the cue with the composite memory trace to produce a retrieved item as before. They would then check to see whether the item that was retrieved was the underlined word. If the match between the retrieved item and the underlined word was above criterion, then the subjects would answer "yes"; otherwise, the answer would be "no."

Method

Aside from the decision phase, these simulations of the model were the same as previous simulations, except that some irrelevant noise was added. The lexicon, item vectors, similarity relations, association operations, composite trace, and retrieval operations were all unchanged. In the moderate similarity condition, I used 40% feature overlap: in the high similarity condition, 80% overlap. The "noisy" simulation included a vector with a variance of .24, which simply added noise to the trace, whereas the noise-free simulation included only simulated list associations.

The retrieved item was matched to the possible probes through a yes/no decision rule. Thus, for example, if the lexical item for "man" was used as a cue, the output was matched to "hammer," in the original event condition, to "screwdriver" in the misleading alternative test condition, and to "wrench" in the novel lure test condition. The simulation called the item (i.e., "man," "screwdriver," or "wrench") "old" if the magnitude of the dot product between the retrieved vector and the to-be-recognized item in question exceeded .5 in the low-threshold simulations or 1.0 in the high-threshold simulation.

Each simulation was replicated 500 times, and so the simulation data were based on 1,000 observations per point. In addition, the mean resonance scores and their standard deviations for each item were computed.

Results

The percentage-correct results of the simulations are given in Table 5. Although the actual numbers varied (as they did in the experiments), the pattern was clearly the same as that produced in the experiments. In particular, the probability of a correct response was higher in the control condition than in the misled conditions; the tendency to respond "yes" to the misleading alternative was greater than to the original events. The simulation also tended to wrongly respond "old" to novel lures more frequently in the control condition than in the misled condition, as in the data.

Discussion

To make it more clear why the simulation produced the results that people exhibited, the resonance scores between the item retrieved from memory and each of the critical items, as well as for the prototype item, are shown in Table 6. The higher these resonance scores are, the more likely it is that the criterion value of .5 (or, in the high-threshold simulation, of 1.0) was exceeded and the simulation correctly or incorrectly called the item "old." For the sake of illustration, I have also included, in the bottom panel of the table, the resonances that were produced from an analogous simulation in which there was no category structure—that is, in which all of the items were unrelated to one another. This "unrelated" simulation was run for other reasons and will be discussed in the context of a recall experiment shortly. However, the resonances are included here because they help to illustrate how the model functions.

As the bottom panel shows, the items that were explicitly encoded in the trace, shown in boldface, resonated most

Table 5
Simulation Results, in Percentage Correct, Based on 1,000 Observations per Point

Condition	Original event	Misleading alternative	Novel lure
Low threshold: .5			
With noise			
Moderate similarity			
Control	70.7	—	44.6
Misled	60.3	20.8	54.6
High similarity			
Control	84.0	—	19.9
Misled	79.9	15.8	27.8
No noise			
Moderate similarity			
Control	90.6	—	38.2
Misled	80.4	2.8	59.0
High similarity			
Control	99.2	—	3.1
Misled	98.2	1.0	7.5
High threshold: .10			
No noise			
Moderate similarity			
Control	57.9	—	58.3
Misled	48.7	34.0	67.4

Table 6
Resonance Scores From Simulation Series 3

With-noise condition	Original events	Misleading alternative	Novel lure	Prototype
Moderate similarity				
Control	1.44	0.75	0.74	1.83
Misled	0.91	1.62	0.37	1.01
High similarity				
Control	1.95	1.73	1.79	2.18
Misled	1.77	1.99	1.46	1.83
No similarity (comparison condition)				
Control	0.74	-0.03	0.00	1.50
Misled	0.75	1.46	-0.00	0.01

Note. Items in boldface were explicitly encoded into the composite trace.

strongly to the output from retrieval. If everything is unrelated to everything else, that is all that happens. The prototype (in the control condition) and the misleading alternative (in the misled condition) resonated more strongly than did the original events because they were weighted more heavily in the original trace.

When category structure is introduced, however, the pattern becomes more complicated. Recall that the exemplars were constructed from the category prototype, and hence the prototype was more similar to the particular exemplars than the exemplars were to each other, on average. This is only to say that in the simulation, as in natural categories, the prototype was central to the category. It is an uncontroversial claim. (One might be able to choose category exemplars that are more similar to one another than they are to the prototype. For example, Pepsi and Diet Pepsi might be more similar to one another than they are to the prototype soft drink. However, this was not the situation in the experiments under study.) The similarity relations have major implications for the pattern of results produced by the model. Because the model makes use of similarities wherever they occur (see Equation 4), the similarity structure of the categories, the exemplars, and the items that were encoded in the trace are all critical in determining what is retrieved and how strongly. If two things that are similar to one another are both entered into the trace and associated to the same cue, then both will be retrieved. They will bolster the resonances of one another to the extent that they are similar to one another. In some cases, like a forced-choice paradigm in which they are pitted against one another, this bolstering does not necessarily help them. In the yes/no paradigm, however, the increases in resonance due to similarity all work in a way that increases the probability that a so-bolstered item will receive a "yes" response.

In the control conditions of the moderate and high similarity simulations, the fact that a prototype similar to the original event was encoded bolstered the resonance of the original event. This occurred because the retrieval cue produced a blend of the two items that were associated with it: the original event and the prototype, superimposed. Had the original item alone been associated with the cue (in addition to an unrelated

association), its resonance against the retrieved item would have been only 0.74, in comparison with the 1.44 in the moderate similarity condition or the 1.95 in the high similarity condition (see Table 6).

The fact that the misleading information was presented in the misled condition increased the resonance of the retrieved (blended) item with the lexical entry of the original event as well. The resonance value of this item (0.91 in the moderate similarity condition or 1.77 in the high similarity condition) is also above the "unrelated" value of 0.75. The difference between the amount of augmentation shown for the control condition and that shown for the misled condition resulted because the original event and the prototype were more similar to one another than were the original event and the other (misleading) category member. This translates, of course, into a higher probability of responding correctly to the original event in the control condition than in the misled condition.

The fact that the prototype has higher similarity to the category members than do the category members to each other also produces the effect, as shown in Belli's (1989) two experiments, on the novel lure: Subjects were more likely to correctly call the novel lure "new" in the misled condition than in the control condition. Because the prototype was more similar to this novel lure than was the misleading alternative, the fact that the prototype was encoded in the trace in the control condition and the misleading alternative was encoded in the misled condition accounts for this difference. Both gave rise to some resonance on the novel lure, but the prototype gave rise to more. This translates into a lower probability of correct responses (or a higher false alarm rate) for the novel lure in the control condition than in the misled condition. Belli's results, then, were just what is expected by the model.

Belli (1989) suggested that overall control condition performance might be better than overall experimental condition performance. In one of his experiments, the means favored the control condition by 6%. However, in the other experiment, the means favored the misled condition by 5% (not significant). The simulation results vacillated: Sometimes they favored the control condition and sometimes the misled condition.

B. Tversky and Tuchin (1989) found that subjects frequently said "yes" both to the original slide and to the misleading alternative. This was also true in the model. For example, in Simulation 3—moderate similarity with noise—the simulation said "yes" to the slide event 61% of the time, to the misleading alternative 81% of the time, and to both of these 40% of the time.

Recall With No Category Structure: Simulation 4

Zaragoza et al. (1987) conducted an experiment in which recall rather than recognition was the memory test. Their experiment was unusual in several ways. First, although they used the same slide sequence as in the previous experiments modeled here, they did not use the same misleading alternatives. These alternatives, instead, were close to being categor-

ically unrelated. For example, the misleading suggestion was that the man lifted a "sandwich" from the tool box. Second, at the time of recall testing, the misleading alternative was eliminated from consideration by the nature of the question: for example "What *tool* did the man lift out of the box?" In the soft-drink case, it was suggested that the can in question was one of Planter's peanuts, but in the later recall probe, subjects were asked for the brand of soft drink. The results of this experiment showed no difference in recall between the misled and the control conditions.

The setup of the trace in the control and misled conditions was like that of the no-noise conditions in the preceding series, except that no category structure was constructed. So instead of having 40% or 80% overlap, the items were not changed from their original unrelated form. As in previously cited investigations of recall with the model, the item that was retrieved from the composite memory trace was identified by being matched to all items in the lexicon except for the cue itself. Because of the special nature of the constraints in Zaragoza et al.'s (1987) experiment, the misleading alternative was also excluded from consideration as a recall alternative.

The first replication of the simulation, based on 500 observations, gave a recall score in the control condition of .15, in comparison with .16 in the misled condition. For the second replication the scores were .16 and .18, respectively. The resonance results are those given in the bottom panel of Table 6 and discussed previously. Under these conditions, there was no difference between recall of the original event in the control and in the misled condition, nor was there a difference in the resonance scores. This simulation represents an example of maximum category disintegration, in which the items are unrelated to one another. Of course, had the alternatives been more related to the original event, as in the yes/no recognition experiments simulated in Series 3, recall too would have shown an effect of the misleading suggestion.

A No-Presentation Control

Chandler (1989) conducted an experiment that differed from those described earlier because in the control condition, no prototype or generic term was presented, though the items were similar to one another. I reran Simulation 3, moderate similarity, with noise but without including either a prototype or a novel lure in the trace. The standardized difference in resonance between the original event and the novel lure, in this new control condition, was .55. For the misled condition, it was .37. This difference corresponds to the results of Chandler's first two experiments. When I did the same thing with the high similarity conditions, which had very high standard deviations, no difference was produced. Chandler found a difference favoring the control condition in two experiments but not in two others. The simulations show the positive trend, but the effect is easy to mask.

Color Shifts: Simulation 5

Perhaps the most compelling experimental evidence for blends comes from experiments in which the remembered

item was not either of the items that were explicitly presented during the experimental sessions but rather consisted of a new item that was a compromise between the presented items. Because contributions from both items are necessary to produce a positive (or actually discretely existing) blend, the argument that subjects are using an all-or-none guessing strategy (such as what McCloskey & Zaragoza, 1985a, suggested in the previous context) becomes implausible. In many cases, of course, a positive blend does not exist in the world. There is no such real object as a superimposition of a hammer and a screwdriver, for example. (In the context of the model, one would say that the item retrieved from the composite memory trace may be a superimposition, but if no such entity exists as a lexical entry, then the model will be forced to an either/or, or both, decision.) However, in the case of color biasing, the theoretically very interesting possibility of a positive blend exists. If an unrepresented compromise item is remembered as well as or better than either of the presented items, then the most straightforward and compelling interpretation of the data is that it was a compromise item that was retrieved from memory.

One of the best investigations of this possibility was in a study done by Loftus (1977) in which subjects, during a slide series, saw a photograph of a green car. Later on they were provided the suggestion that the car had been blue. At the time of testing, rather than making only a two-alternative forced-choice judgment, as in other experiments discussed in this article, subjects were given a series of 15 continuously varying color swatches from which to choose the color of the car. Loftus found a significant blue shift in the color judgments when the misled subjects were compared with the control group. The shift (as shown in Figure 6) was very well

behaved and systematic. Misled subjects now preferred a compromise blue-green color.

Although this experiment may seem to be a rather isolated instance within the literature on eyewitness testimony, the paradigm is very similar to prototype formation paradigms in which a variety of specific instances or exemplars are provided to subjects. At the time of testing, subjects are found to recognize the nonpresented prototype very well and sometimes better than any of the presented instances. Most explanations for this phenomenon of prototype emergence depend on the idea that the representations are added or blended at some point in processing, though there is not yet agreement that this is at storage (e.g., Metcalfe & Eich, 1982, 1985) rather than at retrieval (Hintzman, 1986). The color-shifting phenomenon is predicted, in the CHARM model, for the same reason that prototype emergence is predicted: namely, because in the model the "traces" of the events are added, blended, or superimposed in memory.

These simulations differed from the previous simulations in two ways. First, instead of there being two-alternative forced-choice or yes/no experiments, many alternatives were given. Second, the alternatives were not just randomly different from one another but rather were systematically varied in a continuous manner.

Method

A lexicon of 50 statistically independent items, each composed of 63 features, as in previous simulations, was set up. In order to construct a continuous series that would represent the color swatches given at the time of testing in Loftus's (1977) study, Items 20-30 were altered. First, all of the Items 21-29 were reassigned feature

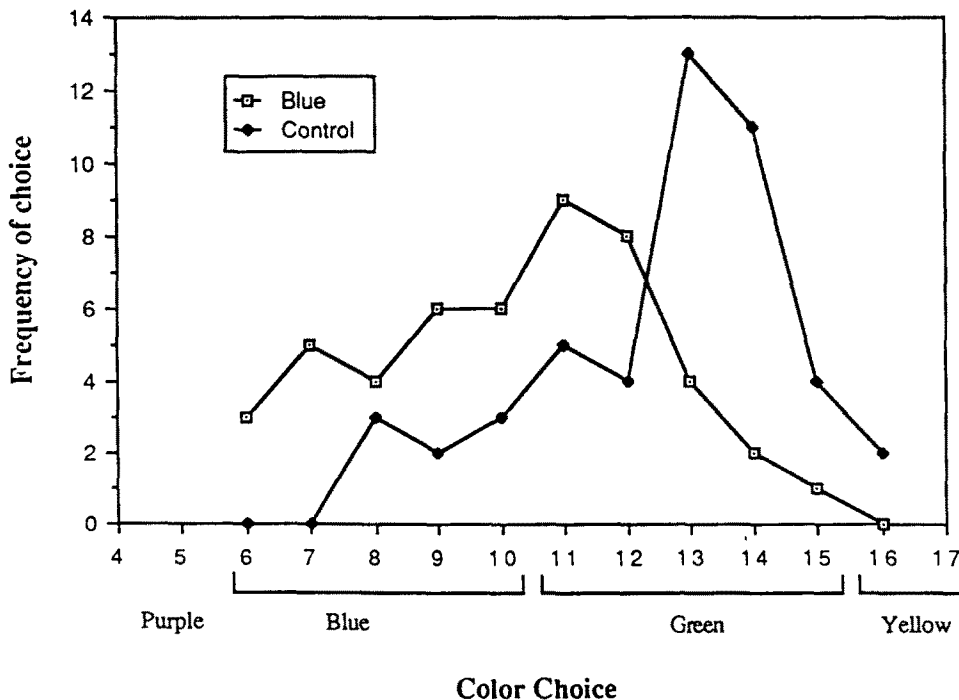


Figure 6. Color choices made by subjects in Loftus's (1977) color-shifting experiment.

values to be identical on all features to Item 20. Then six random features were drawn without replacement. All of Items 21–29 were assigned the values of Item 30 for these six features. A new random selection of six features was drawn, and Items 22–29 were assigned the values of Item 30 for these six. Another six were drawn, and items 23–29 were reassigned the values of Item 30 on these. The drawing of six new features, corresponding to the increment in a lexical item, continued through Item 29. The result was that Items 21–29 varied continuously, and in equal steps, from being highly similar to Item 20 to being highly similar to Item 30. If one wished to make the correspondence to the experimental situation, then Item 20 would be purple and Item 30 would be yellow. Shades of blue, blue-green, and green would fall in between the two statistically independent extremes (Items 20 and 30, or purple and yellow).

I set up the associative memory trace by convolving Item 10 with Item 26 (green, say) and adding in three unrelated convolutions in the control (unbiased) condition. In the experimental condition, in which it was suggested that the car was blue, Item 10 was also convolved with Item 24 (blue, say) and added into the trace. To mimic retrieval, Item 10 was correlated with the trace, and the resonance scores of the output of the correlation against all of the lexical items were computed. The concern here is only with the resonance scores for the 11 related items: lexical Items 20–30.

Results

The resonance scores for both the control and the experimental conditions for the 11 items that were continuously related to the target item, and for one unrelated item are given in Figure 7. By comparing this figure and the previous one, which provides the choice data given by Loftus's subjects, one can see that the resonance scores qualitatively mapped the pattern of the choice data reasonably well. The major differ-

ence between the two figures resided in the overall inflation of the simulation resonance scores in the experimental condition. This occurred because the resonance reflected the similarity (as measured by the dot product) of the retrieved item with the lexical items. When two similar items were entered into the trace, both contributed to this similarity score. Figure 8 shows what happens when one forces a choice in the model by allowing only the item with the highest resonance score to be given as the answer. In this case, the pattern was very close to that shown in the data. In particular, the color between the presented "green" and the suggested "blue" received quite a few votes in the simulation. In addition, the overall distribution in the experimental condition was more even, or squatter, than the distribution in the control condition. The fact that nearby colors in the control condition received some votes was due to pure similarity and the bit of random variability that was inherent in the model. It was not just similarity but rather the composition of the composite memory trace and the nature of the blended item that was retrieved that accounted for the alteration in the distribution found in the experimental condition.

Conclusion

There has been a considerable amount of debate about which paradigms and which comparisons are meaningful in determining whether subsequent information may distort the remembrances of previous events. Loftus (1975, 1977, 1979a, 1979b) has argued that the inclusion of the distorting event is appropriate, whereas McCloskey and Zaragoza (1985a, 1985b) have argued against this position. The conclusions

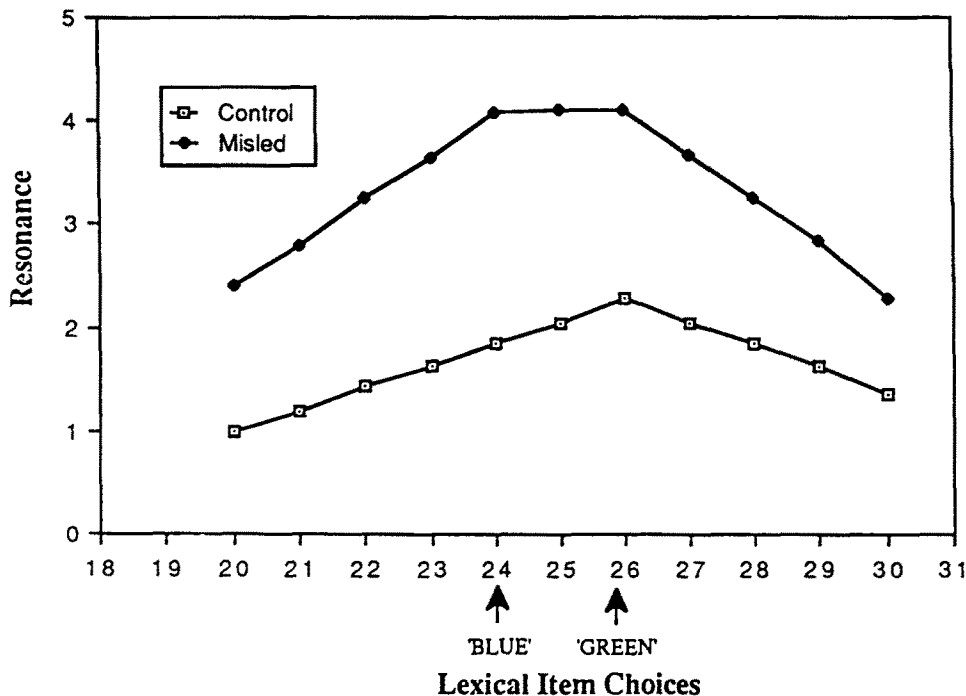


Figure 7. Resonance scores produced by the model in simulations mimicking Loftus's (1977) color-shifting experiment.

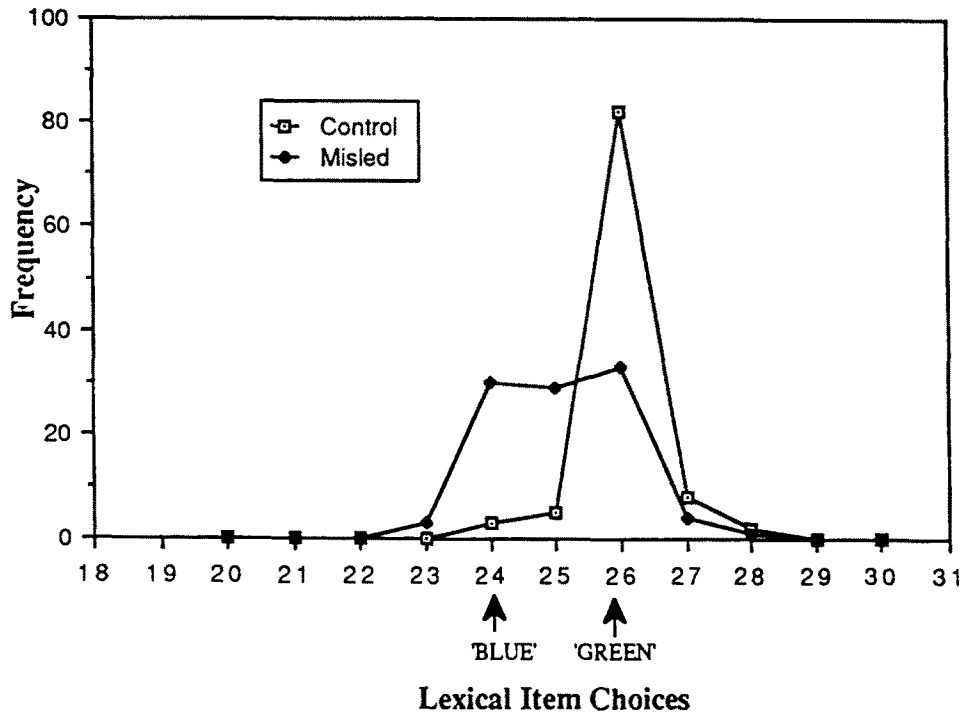


Figure 8. Choice data produced by the model in simulations mimicking Loftus's (1977) color-shifting experiment.

under these two sets of restrictions have been different. As shown by means of the simulations presented in the present article, however, given a well-defined model in which blending unequivocally occurs, both sets of data are predicted. Data in which a yes/no decision rule is used are also predicted. Finally, data collected under different conditions in which a variety of alternatives are possible—the color-shifting experiment—are also predicted. Indeed, the model allows one to make predictions about what will be chosen under any well-defined set of alternatives, and so far those predictions have been borne out. Given the finding of systematic color shifting, those who would argue that memory distortions do not occur are left in the unenviable position of having to argue, in an ad hoc manner, that such shifting did not “really” occur but was due to some kind of complicated deliberation on the part of the subjects. It is hard to imagine that such complicated deliberation could or might be predicted. Certainly the “no distortion” advocates would have been much happier about the experimental results had there been no evidence for blends or color shifts, regardless of the alternatives offered.

What are the implications for real-world memory? There are some situations in which blends might be expected to occur. One prerequisite for such real-world blends is that a real-world object, or the possibility of a real-world object, that closely matches the blend must exist or be capable of existing. In the color-shifting experiment, there are objects that display the graded colors, and so there would be no a priori restriction against the possibility that such a blend occurred. On the other hand, there are no real-world objects that constitute a blend between a stop sign and a yield sign or between a hammer and a screwdriver. Thus a literal blend would be

ruled out immediately, even if such were retrieved from memory (as the CHARM model says it is).

One question that arises concerns the chances of observing blends in face recognition, an area in which the eyewitness testimony paradigm is surely relevant. This is a situation that potentially, at least, is more like the color-shifting experiment than like the stop-sign/yield-sign experiment. Hair color, size and shape of nose, face shape, skin color, and so on are all continuous variables. There thus exists the possibility that there may be compromise alternatives to an actually viewed face and, say, an incidentally viewed face and that these alternatives may provoke a very high resonance. It is possible that a lure that was never viewed before may be “better” recognized than even the actual target face. Furthermore, in a real situation, rather than in a contrived experiment, we are unlikely to have the information to allow us to know that the incorrect but well-recognized compromise face is precisely the misleading lure that should be eliminated from the test. We would choose to provide witnesses with the McCloskey and Zaragoza situation, in which the misleading lure is not present, but, lacking in omniscience, we may be unable to do so.

References

- Anderson, J. A. (1977). Neural models with cognitive implications. In D. LaBerge & S. Samuels (Eds.), *Basic processes in reading* (pp. 27–90). Hillsdale, NJ: Erlbaum.
- Anderson, J. A., & Hinton, G. E. (1981). Models of information processing in the brain. In G. E. Hinton & J. A. Anderson (Eds.), *Parallel models of associative memory* (pp. 9–48). Hillsdale, NJ: Erlbaum.

- Bekerian, D. A., & Bowers, J. M. (1983). Eyewitness testimony: Were we misled? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 9, 139-145.
- Belli, R. F. (1989). Influences of misleading postevent information: Misinformation interference and acceptance. *Journal of Experimental Psychology: General*, 118, 72-85.
- Cavanagh, P. (1976). Holographic and trace strength models of rehearsal effects in the item recognition test. *Memory & Cognition*, 4, 186-199.
- Ceci, S. J., Ross, D. F., & Toglia, M. P. (1987a). Age differences in suggestibility: Narrowing the uncertainties. In S. J. Ceci, M. P. Toglia, & D. F. Ross (Eds.), *Children's eyewitness testimony* (pp. 79-91). New York: Springer.
- Ceci, S. J., Ross, D. F., & Toglia, M. P. (1987b). Suggestibility of children's memory: Psycholegal implications. *Journal of Experimental Psychology: General*, 116, 38-49.
- Ceci, S. J., Toglia, M. P., & Ross, D. F. (1988). On remembering . . . more or less: A trace strength interpretation of developmental differences in suggestibility. *Journal of Experimental Psychology: General*, 117, 201-203.
- Chandler, C. C. (1989). Specific retroactive interference in modified recognition tests: Evidence for an unknown cause of interference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15, 256-265.
- Christiaansen, R. E., & Ochalek, K. (1983). Editing misleading information from memory: Evidence for the coexistence of original and postevent information. *Memory & Cognition*, 11, 467-475.
- Geiselman, R. E. (1988). Improving eyewitness memory through mental reinstatement of context. In G. M. Davies & D. M. Thomson (Eds.), *Memory in context: Context in memory* (pp. 245-266). New York: Wiley.
- Gibling, F., & Davies, G. (1988). Reinstatement of context following exposure to post-event information. *British Journal of Psychology*, 79, 129-141.
- Gudjonsson, G. H. (1986). The relationship between interrogative suggestibility and acquiescence: Empirical findings and the theoretical implications. *Personality and Individual Differences*, 7, 195-199.
- Hammersley, R., & Reed, J. D. (1986). What is integration? Remembering a story and remembering false implications about the story. *British Journal of Psychology*, 77, 329-341.
- Hinton, G. E., & Anderson, J. A. (1981). *Parallel models of associative memory*. Hillsdale, NJ: Erlbaum.
- Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review*, 93, 411-428.
- Kohnken, G., & Brockmann, C. (1987). Unspecific postevent information, attribution of responsibility, and eyewitness performance. *Applied Cognitive Psychology*, 1, 197-207.
- Kohonen, T., Lehtio, P., & Oja, E. (1981). Storage and processing of information in associative memory systems. In G. E. Hinton & J. A. Anderson (Eds.), *Parallel models of associative memory* (pp. 105-144). Hillsdale, NJ: Erlbaum.
- Kroll, N. E. A., & Ogawa, K. H. (1988). Retrieval of the irretrievable: The effect of sequential information on response bias. In M. M. Gruneberg, P. E. Morris, & R. N. Sykes (Eds.), *Practical aspects of memory: Current research and issues* (Vol. 1, pp. 490-495). Chichester, England: Wiley.
- Kroll, N. E. A., & Timourian, D. A. (1986). Misleading questions and the retrieval of the irretrievable. *Bulletin of the Psychonomic Society*, 24, 165-168.
- Lehnert, W. G., Robertson, S. P., & Black, J. B. (1984). Memory interactions during question answering. In H. Mandl, N. L. Stein, & T. Trabasso (Eds.), *Learning and comprehension of text* (pp. 355-369). Hillsdale, NJ: Erlbaum.
- Lindsay, D. S., & Johnson, M. K. (1987). Reality monitoring and suggestibility: Children's ability to discriminate among memories from different sources. In S. J. Ceci, M. P. Toglia, & D. F. Ross (Eds.), *Children's eyewitness memory* (pp. 92-121). New York: Springer-Verlag.
- Lindsay, D. S., & Johnson, M. K. (1989). The eyewitness suggestibility effect and memory for source. *Memory & Cognition*, 17, 349-358.
- Loftus, E. F. (1975). Leading questions and eyewitness report. *Cognitive Psychology*, 7, 560-572.
- Loftus, E. F. (1977). Shifting human color memory. *Memory & Cognition*, 5, 696-699.
- Loftus, E. F. (1979a). *Eyewitness testimony*. Cambridge, MA: Harvard University Press.
- Loftus, E. F. (1979b). The malleability of memory. *American Scientist*, 67, 312-320.
- Loftus, E. F., & Hoffman, H. G. (1989). Misinformation and memory: The creation of new memories. *Journal of Experimental Psychology: General*, 118, 100-104.
- Loftus, E. F., & Loftus, G. R. (1980). On the permanence of stored information in the human brain. *American Psychologist*, 35, 409-420.
- Loftus, E. F., Miller, D. G., & Burns, H. J. (1978). Semantic integration of verbal information into a visual memory. *Journal of Experimental Psychology: Human Learning and Memory*, 4, 19-31.
- Loftus, E. F., Schooler, J., & Wagenaar, W. A. (1985). The fate of memory: Comment on McCloskey and Zaragoza. *Journal of Experimental Psychology: General*, 114, 375-380.
- Longuet-Higgins, H. C. (1968). Holographic model of associative recall. *Nature*, 217, 104.
- McClelland, J. L., & Rumelhart, D. E. (1986). *Parallel distributed processing* (Vol. 1). Cambridge, MA: MIT Press.
- McCloskey, M., & Zaragoza, M. (1985a). Misleading postevent information and memory for events: Arguments and evidence against memory impairment hypotheses. *Journal of Experimental Psychology: General*, 114, 1-16.
- McCloskey, M., & Zaragoza, M. (1985b). Postevent information and memory: Reply to Loftus, Schooler, and Wagenaar. *Journal of Experimental Psychology: General*, 114, 381-387.
- Metcalfe, J., & Murdock, B. B. (1981). An encoding and retrieval model of single-trial free recall. *Journal of Verbal Learning and Verbal Behavior*, 20, 161-189.
- Metcalfe, J., & Eich, J. (1982). A composite holographic associative recall model. *Psychological Review*, 89, 627-661.
- Metcalfe, J., & Eich, J. (1985). Levels of processing, encoding specificity, elaboration, and CHARM. *Psychological Review*, 92, 1-38.
- Morton, J., Hammersley, R. H., & Bekerian, D. A. (1985). Headed records: A model for memory and its failures. *Cognition*, 20, 1-23.
- Murdock, B. B. (1982). A theory of the storage and retrieval of item and associative information. *Psychological Review*, 89, 609-626.
- Murdock, B. B. (1983). A distributed memory model for serial order information. *Psychological Review*, 90, 316-338.
- Pirolli, P. L., & Mitterer, J. O. (1984). The effect of leading questions on prior memory: Evidence for the coexistence of inconsistent memory traces. *Canadian Journal of Psychology*, 38, 135-141.
- Register, P. A., & Kihlstrom, J. F. (1988). Hypnosis and interrogative suggestibility. *Personality and Individual Differences*, 9, 549-558.
- Rumelhart, D. E., & McClelland, J. L. (1986). *Parallel distributed processing* (Vol. 2). Cambridge, MA: MIT Press.
- Schooler, J. W., Gerhard, D., & Loftus, E. F. (1986). Qualities of the unreal. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13, 501-518.
- Sheehan, P. W. (1988). Confidence, memory, and hypnosis. In H. Pettinati (Ed.), *Hypnosis and memory* (pp. 96-154). New York: Guilford.

- Smith, V. L., & Ellsworth, P. C. (1987). The social psychology of eyewitness accuracy: Misleading questions and communicator expertise. *Journal of Applied Psychology, 72*, 294-300.
- Tousignant, J. P., Hall, D., & Loftus, E. F. (1986). Discrepancy detection and vulnerability to misleading post-event information. *Memory & Cognition, 14*, 329-338.
- Tversky, A. (1977). Features of similarity. *Psychological Review, 84*, 327-352.
- Tversky, B., & Tuchin, M. (1989). A reconciliation of the evidence on eyewitness testimony: Comments on McCloskey and Zaragoza. *Journal of Experimental Psychology: General, 118*, 86-91.
- Wagenaar, W. A., & Boer, H. P. A. (1987). Misleading postevent information: Testing parameterized models of integration in memory. *Acta Psychologica, 66*, 291-306.
- Zaragoza, M. S., & Koshmider, J. W. (1989). Misled subjects may know more than their performance implies. *Journal of Experimental Psychology: Learning, Memory and Cognition, 15*, 246-255.
- Zaragoza, M. S., & McCloskey, M. (1989). Misleading postevent information and the memory impairment hypothesis: Comment on Belli and reply to Tversky and Tuchin. *Journal of Experimental Psychology: General, 118*, 92-99.
- Zaragoza, M. S., McCloskey, M., & Jamis, M. (1987). Misleading postevent information and recall of the original event: Further evidence against the memory impairment hypothesis. *Journal of Experimental Psychology: General, 113*, 36-44.

Received March 13, 1989

Revision received August 4, 1989

Accepted August 7, 1989 ■