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

Janet Metcalfe
Department of Psychology
University of California, San Diego.

Abstract

The idea that compositing or blending may occur in human episodic memory stems from two sources: (1) distributed models of human memory, and (2) studies that have focussed on the distortions and mistakes that occur in eyewitness testimony. In this paper, data that have been uncovered within the eyewitness testimony paradigm are simulated by a distributed memory model--CHARM (composite holographic associative recall memory). Studies done by Loftus have been interpreted as indicating that blending does occur; modification of these experiments conducted by McCloskey and Zaragosa have been claimed to refute Loftus' interpretation. It is shown that both of these results are predicted by the composite-trace model.

Introduction

There has been considerable debate about the nature of human memory storage: whether memories are stored discretely or may interact or even blend with one another. Loftus has argued that the fact that subjects who are given misleading information in a realistic, eyewitness testimony situation may be more inaccurate than are subjects not given the misleading information indicates that subsequent information may distort, erase, or combine with earlier information about the target event. McCloskey and Zaragosa suggest that under the appropriate testing conditions, no evidence for distortion, erasure or blending in memory is found.

The situation of primary interest in this debate is exemplified by a number of experiments by McCloskey and Zaragosa (1985). Subjects saw a series of color slides depicting an incident in which a maintenance man enters an office, repairs a chair, finds and steals $20, and then leaves. Embedded in the sequence was a critical slide in which the man picked up a hammer from a tool kit. After viewing the slide sequence, subjects read a narrative in which the misleading information was embedded, in the experimental condition, and in which neutral information was given in the control condition. In the experimental condition, it was suggested to the subjects that the tool the man had picked up was a screwdriver. In the control condition, a generic term-- tool--was used to refer to the detail in question. At time of test, subjects were asked the following question: "The man slid the calculator beneath a ___________ in his tool box".

The test consisted of a two-alternative forced choice procedure. In what will here be designated the "standard" conditions the label of the actually viewed object (in this case the term "hammer") was contrasted with the suggested objects ("screwdriver"). The pervasive finding in this testing procedure was that correct selection of hammer was impaired in the experimental but not in the control condition. McCloskey and Zaragosa modified this testing procedure such that the correct alternative was pitted against another category member ("wrench") but not against the misleading information itself. We will here designate this testing procedure the "Modified" condition. It was found that under these testing conditions, there was no decrement in performance for the term "hammer" in the experimental conditions.

Table 1. The experimental paradigm.

Standard "Loftus" Conditions			
Presentation	Questionaire	Test	Correct
Control			
Man-Hammer	------	Hammer-Screwdriver	72%
Misled			
Man-Hammer	Man-Screwdriver	Hammer-Screwdriver	37%
Modified "McCloskey" Conditions			
Presentation	Questionaire	Test	Correct
Control			
Man-Hammer	------	Hammer-Wrench	75%
Misled			
Man-Hammer	Man-Screwdriver	Hammer-Wrench	72%

On the basis of these findings, summarized in Table 1, McCloskey and Zaragosa argue that there is no loss or distortion of the initially encoded events. Loftus has usually argued for a blending view of memory (rather than for simple loss or erasure, attributable to the misleading information). About this blending or integration view, McCloskey and Zaragosa state:

> "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 the 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?" (p.15).

In the model, that is outlined below, there is but a single memory trace which consists of the sum of the associations that are entered into it. This composite or superimposed trace is an example of a memory system that produces blended memories. If a cue has been associated with more than one item, that cue will serve to retrieve all of the items with which it has been associated, and they will all be produced together, or in a blend. More explicit computer simulations provide predictions and postdictions about exactly what this composite, or blending model does in the situations outlined above.

Summary of the CHARM model

The model that will be used to investigate the blending predictions under the conditions outlined in McCloskey and Zaragosa's experiment is called the CHARM model (composite holographic associative recall model). The model was not devised specifically to apply to this situation, and has, in fact, been applied with some success to a variety of other classic memory situations, such as paired-associate learning, interference as a function of similarity, encoding specificity effects, concept formation, elaboration effects, recognition failure effects, and others (Metcalfe, in press; Metcalfe-Eich, 1982, 1985). The model is associative in nature, based on the idea that items, represented as distributed patterns of features, or vectors, are associated by the operation of convolution. This operation (denoted *) is given by the following equation, for the mth term of the resulting vector:

$$(F*G) = \sum_{(i,j)\varepsilon S(m)} f_i g_j, \qquad (1)$$

where, F and G are the item vectors:

$$(f_1, f_2, f_3, \ldots)\text{and}(g_1, g_2, g_3, \ldots)$$

and

$$S(m) = \left\{ (i,j) \,\middle|\, -\frac{n-1}{2} \le i,j \le \frac{n-1}{2}, \text{ and } i+j=m \right\}.$$

The resulting vector is added into a single vector that is the composite memory trace. As each association is added into this vector the values for each element of the vector may change. Thus, the trace is defined as:

$$T=(G^*G)+(H^*I)+(J^*K)+.... \tag{2}$$

The initial item vectors, F,G,H,I,J, and K may bear any similarity relation to one another, and may vary in terms of their initial strength or length. This summation at time of storage is what is meant by blending or storing together in memory.

The operation that allows retrieval from this composite associative trace is called correlation (denoted #) and is defined as:

$$(F\#T)_m = \sum_{(i,j)\varepsilon S(m)} f_i t_j, \tag{3}$$

where

$$S(m)=\left\{ (i,j) \;\middle|\; -\frac{n-1}{2} \le i,j \le \frac{n-1}{2},\ \text{and } i-j=m \right\}.$$

The result of retrieval is a single vector. But this vector may be broken down into the components that contribute to it as follows:

$$R=F\#T \qquad \text{I(4)}$$

$$=F\#(F^*G)+(H^*I)+\cdots$$

$$=F\#(F^*G)+F\#(H^*I)+...$$

$$=S_{FF}G+S_{FG}F+error_{F*G}+S_{FH}I+S_{FI}H+error_{H*I}+\cdots$$

Here, S is a scalar giving the similarity value between F and F, for example, as measured by their dot product. In the case where two items are associated with a single cue, we see that the single vector that is retrieved by this system will contain components of both of the original items. This output from the model can be simulated, and the result can be assessed within the framework of the Loftus-McCloskey forced-choice paradigm, by simply providing the alternative they allowed in the experiment, and letting the model pick the best match to its retrieved output.

Simulations

A number of simulations were conducted on this and related paradigms. Only one series will be reported here.

Method

A lexicon of 90 items was constructed, where each item consisted of 63 features and each feature consisted of a value randomly selected from a truncated Gaussian distribution with an expected value of zero. The items were then normalized so that the self dot products were 1. The first item in the lexicon we will hereinafter assign the name "man"; the second item "hammer"; the 22nd item "screwdriver"; the 32nd item "tool" and the 42nd item "wrench". In the High Similarity conditions, these exemplars were reassigned feature values so that 80% of their features were the same as the prototype item "tool". In the Moderate Similarity conditions, 40% of these features were reassigned values of the prototype. In the unrelated conditions, the items were statistically independent.

Two different traces were formed to depict the experimental and the control conditions of the experiment. The control trace was:

$$T_c=(MAN*HAMMER)+(MAN*TOOL)+5 \textit{ irrelevant convolutions.}$$

The experimental trace was:

$$T_{e=}(MAN*HAMMER)+(MAN*SCREWDRIVER)+5 \textit{ irrelevant convolutions.}$$

The irrelevant convolutions were included here to indicate that there were other events stored in the trace, and the number is not too important in the present context. (See Metcalfe & Murdock, 1981, for further details on this point).

Retrieval was simulated by correlating the vector for MAN with the composite trace. The retrieved vector that resulted was then compared to HAMMER and SCREWDRIVER, in the standard conditions, or to HAMMER and WRENCH, in the modified conditions. The comparison consisted of taking the dot product of the retrieved vector with that of the lexical item in question. The match that gave the highest value was the winner and was said to be the choice that was made on that particular trial. The entire sequence of simulations was run twice, the first time producing 200 replications or observations per point, and the second time 1000 observations per point.

Results

The pattern of results produced by the simulations is shown in Table 2. As can be seen, in each of the three manipulations of similarity, the model produced the basic pattern shown in McCloskey and Zaragosa's data. In particular, under Standard testing conditions, the misleading information in the Experimental condition resulted in poorer performance than did the neutral information in the Control conditions, whereas when the Modified test situation was simulated, there was no difference between the Control and the Misled conditions.

Conclusions

It is clear that the simple blending model is able to generate data that have been construed as indicating that there are distortions in memory and also data that have been rallied to reject the blending idea. What are the implications for real-world memory? There are some situations in which one might expect blends to occur. One prerequisite in the model for the appearance of evidence for such blends is that there must exist a lexical representation that depicts or at least is very similar to the composite blended entity that is retrieved from memory. Figure 1 shows an example where two objects that are unlike one another are superimposed. But there is no real world object that could correspond to the blend shown in the far right panel. Figure 2 shows a second example that was created in exactly the same way as the first example. However, in this case the items themselves were highly similar to one another, and the blended entity could plausibly be a real-world entity.

Positive blends-- where a composite model will predict a compromise between the presented and suggested items-- are difficult to find. However, Loftus (1977) has

Table 2. Simulation results.

Category Similarity	Test Condition	Information Condition	Percentage Correct	
Unrelated	Standard *(hammer-screwdriver)*	Control	68.0	64.2
		Misled	32.5	34.1
	Modified *(hammer-wrench)*	Control	70.5	64.8
		Misled	62.0	66.3
Moderate	Standard	Control	58.5	63.1
		Misled	38.0	37.9
	Modified	Control	67.0	63.1
		Misled	65.0	64.7
High	Standard	Control	52.5	55.8
		Misled	37.5	42.0
	Modified	Control	54.5	56.7
		Misled	58.0	57.0

provided one such example in which a car that had in fact been green in a slide sequence was guessed most frequently as having been a blue-green color after misleading blue information about its color was given. In this color-shift experiment, the intermediate colors could and do exist in the real world and so there would be no a priori restriction against the possibility that such a color had occurred. In many other cases, however, there are no real-world objects that comprise a blend. For instance, there is no real world object that consists of a blend between a screwdriver and a hammer. Thus, a literal blend could be ruled out immediately, even if such were retrieved from memory.

Face recognition poses an interesting puzzle, and one that may have practical significance. As Figure 2 illustrates, there may be cases in which the superimposition of two faces could produce a plausible blend. In such a situation even if we were to eliminate from the testing alternative the face that was used as the misleading information the possibility exists that a third face (or actually, in the model, a whole family of intermediate faces) might nevertheless be accepted by the subject as plausible interpretation of the blend that is retrieved from memory. In conclusion, then, the composite model does a good job of predicting the data from both the McCloskey and the Loftus testing conditions. It also makes further predictions that may be of both practical and of theoretical importance.

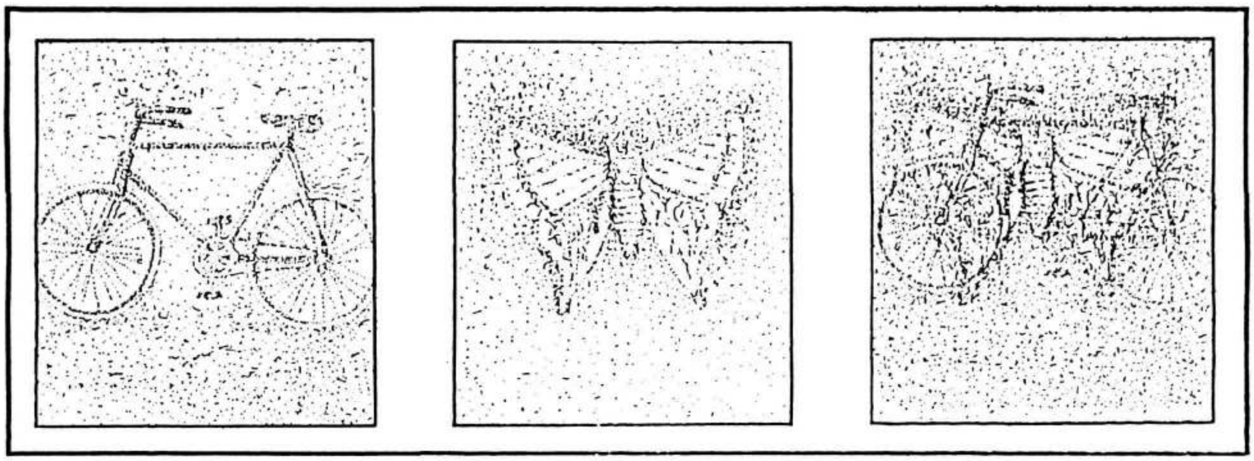

Figure 1. The superposition of two nonintegrable objects that do not yield a positive blend.

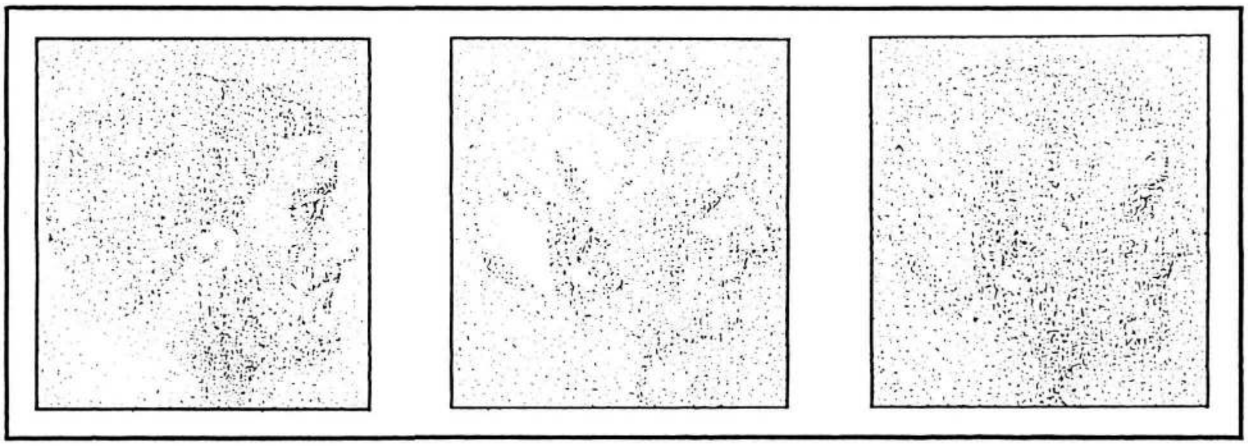

Figure 2. The superpositon of two integrable objects producing a positive blend.

References

Loftus, E. F. (1977). Shifting human color memory. *Memory & Cognition, 5,* 696,699.

McCloskey, M., & Zaragosa, M. (1985). Misleading postevent information and memory for events: Arguments and evidence against the memory impairment hypothesis. *Journal of Experimental Psychology: General, 114,* 3-18.

Metcalfe, J. A composite holographic associative recall model (CHARM) and blended memories in eyewitness testimony. *Journal of Experimental Psychology: General* , in press.

Metcalfe, J. & Murdock, B. B., Jr. (1981) An encoding and retrieval model of free recall. *Journal of Verbal Learning and Verbal Behavior,20,* 161-189.

Metcalfe-Eich, J. (1982). A composite holographic associative recall model. *Psychological Review, 89,* 627-661.

Metcalfe-Eich, J. (1985). Levels of processing, encoding specificity, elaboration, and CHARM. *Psychological Review, 92,* 1-38.