Psychological Review
1982, Vol. 89, No. 6, 627-661

Copyright 1982 by the American Psychological Association, Inc.
0033-295X/82/8906-0627$00.75

# A Composite Holographic Associative Recall Model

Janet Metcalfe Eich
University of Toronto, Toronto, Ontario, Canada

In this article, a highly interactive model of association formation, storage, and retrieval is described. Items, represented as sets of features, are associated by the operation of convolution. The associations are stored by being superimposed in a composite memory trace. Retrieval occurs when a cue item is correlated with the composite trace. The retrieved items are intrinsically noisy, may be ambiguous, and may under certain conditions be systematically distorted from their encoded form. A discrete response is selected by matching the retrieved item to all of the items in semantic memory. The model yields several new predictions about errors in single-trial cued recall that depend on similarity relations among the to-be-remembered items, and also about the efficacy of extralist cues. Experiments are presented that test these predictions against human recall. The model is then applied to several well-known results: prototype abstraction, the A–B A–D paradigm—including the independence of the B and D responses—and the Osgood transfer surface.

This article addresses the question of how people associate two ideas and store the association in memory in such a way that at some later time, when only one of the ideas is available, they may use that idea to re-create or retrieve the other. A second question that follows naturally from the first is, How are people able to retrieve such ideas?

The particular answer to these questions that will be investigated and elaborated here has usually been called the holographic hypothesis of associations. It is not the image-like characteristics of holograms that define the holographic hypothesis; rather, it is the formal operations of convolution and correlation that are important. Roughly speaking, association by convolution consists of the interation of all of the parts of one item with all of the parts of another. The result is an interactive new entity. A number of these interactive associations may be stored in a single composite memory trace, somewhat resembling a photograph that has been exposed many times. Retrieval occurs by means of the operation of correlation. The operations of convolution and correlation will be explained in a later section of the article or the reader may refer to Borsellino and Poggio (1973), Murdock (1979), or Stroke (1969). What we get from this scheme—holographic associations stored in a composite trace—is a profoundly interactive memory.

Since its invention by Gabor (1948), the hologram has frequently been cited as a metaphor for human association formation, storage, and retrieval, and the neurological apt-

I am grateful to the many people who contributed to this article. Among his many other contributions throughout the course of this work, Eric Eich provided the name for the model described herein. Peter Liepa first told me about holographic recall, and Bennet Murdock incited and kindled my interest in the hypothesis. I thank Endel Tulving for his support and guidance, Fergus Craik for his many insightul suggestions, and William Batchelder for his numerous contributions to the model. I am grateful to Robert Bjork, Thomas Wickens, Norman Slamecka, Gordon Logan, Paul Kolers, and two anonymous reviewers for thoughtful critiques that helped shape this article. I also appreciate the encouragement of Walter Kintsch. I thank Mary Duncan, Brian McElree, Louis Narens, Henry Roediger, and Eric Eich for facilitating the experiments and the simulations.

This article is based on the author's doctoral dissertation at the Department of Psychology, the University of Toronto. Some of the research was conducted while the author enjoyed the hospitality of the School of Social Sciences at the University of California, Irvine and the Department of Psychology, University of California, Los Angeles. This work was financially supported by Natural Sciences and Engineering Research Council of Canada (NSERC) fellowships to the author, an NSERC grant to Bennet Murdock, and a grant to the author from the Focussed Research Project on Perception and Higher Cognitive Processes, University of California, Irvine, through the auspices of Kenneth Wexler.

Requests for reprints should be sent to Janet M. Eich, who is now at the Department of Psychology, 405 Hilgard Avenue, University of California, Los Angeles, California 90024.

ness of the metaphor has often been noted (Borsellino & Poggio, 1973; Cavanagh, 1976; Heerden, 1963; Julesz & Pennington, 1965; Longuet-Higgins, 1968; Metcalfe & Murdock, 1981; Murdock, 1979, 1982; Pribram, 1971; Pribram, Nuwer, & Baron, 1974; Westlake, 1970; Willshaw, 1981; Liepa, Note 1). In addition, J. A. Anderson (1970) has remarked that the construct of a composite trace, which is an intrinsic part of all holographic models, may be particularly appropriate to human memory because it may give rise to creative behavior. Holographic models, however, have only rarely been directed toward explaining psychological data. One notable exception is Cavanagh's (1976) holographic model of recognition for short lists of items. Also, Metcalfe and Murdock (1981) presented a holographic model of single-trial free recall of unrelated words. Both of these models provide a reasonably good description of the data to which they are applied.

In this article, a traditional psychological variable—similarity—will be represented in a holographic model similar to that proposed by Metcalfe and Murdock (1981). This model will henceforth be referred to as CHARM, which stands for composite holographic associative recall model. The representation of similarity will allow the highly interactive nature of the holographic association itself and of the method of storage to become manifest. It will also allow the model to make some predictions about what people will recall in several experimental situations. Not all of these predictions are intuitively obvious.

CHARM initially will focus on an experimental situation in which people are presented with a list of pairs of words to study and remember. The words may vary in their similarity to each other. The subjects will be instructed to study the pairs as pairs, and it will be assumed that an association is formed between the two words composing each pair. A short time after studying, subjects will be presented with cue words and asked to recall the targets. The reason for focusing on the cued recall of a once-presented list of paired associates is that this is the simplest situation to which CHARM can be applied. It is also the simplest situation in which to address the question of how two ideas may be associated and stored and how retrieval takes place.

Although the initial focus of this article is on implications of the holographic association and the composite trace in single-trial cued recall, several extensions of the model to other paradigms will be presented in a later section. These extensions—to prototype learning, the A–B A–D paradigm, and the Osgood transfer surface—provide some insight into the generality as well as the limitations of CHARM.

## Description of CHARM

CHARM is a fairly simple holographic model of associative recall. A main assumption of the model is that items are represented as patterns of features rather than as discrete indivisible units. These items may vary in their similarity to each other. Two such items are associated interactively by means of the operation of convolution. The result is stored in a composite memory trace that consists of the superimposition of other associations as well. Retrieval occurs by means of the operation of correlation. Finally, the retrieved item is identified as a particular response by being matched to every item in a lexicon representing semantic memory.

CHARM stems from a variety of sources, and each of the main constructs in the model has precedents in the literature. For instance, the representation of items in terms of sets of features is well known (e.g., Bower, 1967; Estes, 1972; Tversky, 1977; Wickens, 1972). The representation of similarity will be in terms of feature overlap. This representation has often been used before (e.g., Reed, 1973; Smith, Shoben, & Rips, 1974; Tversky, 1977), although, to my knowledge, similarity has not previously been studied within a holographic framework.

The holographic association in vector notation, as will be used here, has been specified before (Borsellino & Poggio, 1973). This association appears in a number of other models (Cavanagh, 1976; Heerden, 1963; Julesz & Pennington, 1965; Longuet-Higgins, 1968; Murdock, 1982; Willshaw, 1981; Liepa, Note 1). The model from which CHARM most directly derives was given by Liepa (Note 1), and work on the present model was begun in collaboration with Bennet Murdock (Metcalfe & Murdock, 1981). Murdock (1982) has elaborated a holographic model that stores both item and associative information in a single trace. Murdock's (1982) paper also deals with the issue of noise distributions—an issue that will not be covered in the present article. The main constructs under investigation here—the holographic association and the composite trace—are the same in CHARM as in the above-cited models, and CHARM depends on the theoretical advances made by these investigators.

The construct of a composite trace is an integral part of all holographic memory models but has also been employed in models that do not use the associative operation of convolution and the retrieval operation of correlation (J. A. Anderson, 1970, 1972, 1973, 1977; J. A. Anderson, Silverstein, Ritz, & Jones, 1977; Kohonen, 1977, 1980). Some of the psychological as well as neurological implications of this construct have been delineated by J. A. Anderson (1977), the most important being that it may give rise to creative behavior.

The distinction between semantic and episodic memory is a familiar one (Tulving, 1972). In this article, the composite–holographic trace refers (roughly) to episodic memory. A pattern recognizer corresponding to semantic memory is included as well. These two levels of representation allow CHARM to solve the problem of response ambiguity that occurs in other holographic and distributed models (e.g., Willshaw, 1981; Liepa, Note 1). The construct of a semantic pattern recognizer relates to response availability (e.g., Martin, 1965; Underwood, Runquist, & Schulz, 1959; Underwood & Schulz, 1960). The identification process used in CHARM, by which a particular response is selected, bears a semblance to the resonance process given in Ratcliff's (1978) item-recognition model.

Each part of the model is thus familiar (though some parts may be more familiar than others). What I do in the present article is to mesh these several parts into a model that makes some predictions about recall.

There are, of course, many things one would like such a model to do and explain. The strategy I adopt in the present article is to specify the model sufficiently to allow some of the implications of the interactive nature of the association and the composite nature of the trace to become obvious and testable. Recourse to other constructs (however reasonable) was specifically avoided in an effort to show what sorts of predictions and results could (and could not) be attributed direcly to the composite–holographic trace. In this section of the article, I explain the mechanisms for association formation, storage, and retrieval and describe some results of varying semantic similarity. First, however, the manner in which the individual items are represented needs to be examined.

### *Representation of Individual Items*

In CHARM, items are represented as ordered sets of features. The work of Tversky (1977) indicates that the individual features of items, and not just their point location in multidimensional space (Hutchinson & Lockhead, 1977), may be important for certain tasks, such as judgements of similarities and differences. However, the featural representation, as opposed to a point-location representation, is necessary in the present model because the features themselves participate in both the associative and the retrieval operations.

In contrast to certain other models (see Medin & Schaffer, 1978; Reed, 1973), in which items are represented by a rather small number of features, CHARM assumes that the number of features representing a given item is quite large. Apart from this quantitative difference, the concept of similarity as feature overlap, as it has been developed in earlier models, applies quite readily to CHARM and simplifies the analysis of similarity.

The features in each item are coded as numbers and are assumed to have positive and negative values, with an expected mean of zero. This coding scheme is analogous to the semantic differential. However, Batchelder and Narens (1977) have shown that the identity of particular features, even under highly circumscribed conditions, may not be uniquely determinable. Because this is the case, it seems appropriate to consider the features to be abstract.

The numerical values of the features may vary from item to item, and it is the difference in overall pattern that characterizes each item. The strength of a feature, as given by its absolute value, is an indication of how important that feature is in the representation of the item. In a later section, once some of the implications of similarity have been described, the question, "More important for what?" will be addressed in the context of cued recall. If the average absolute value of all of the features of one item were greater than that of another item, the first item would be said to be stronger than the second.

It is assumed that the numerical value of a particular feature within an item is independent of the values of the other features within the same item. Though possibly unrealistic, this assumption is made for mathematical convenience and presumably could be relaxed.

More critical is the assumption that the features must be ordered. That is, if Feature 15, say, represents a particular attribute in one item, it represents the same attribute in all items. Both the associative formation and retrieval operations depend on the features being ordered.

Unrelated items will be considered to be items in which the feature values are statistically independent. Knowing the value of a particular feature in one item gives no information about the value of that same feature in the rest of the set of unrelated items. We may take the dot product between two items as a measure of the extent to which the two items are similar or unrelated. The dot product is found by multiplying the value of each feature in one item by the value of the corresponding feature in the other item and adding all of the products. The dot product of **F**, coded as $(\cdots, f_{-1}, f_0, f_1, \cdots)$ and **G**, coded as $(\cdots, g_{-1}, g_0, g_1, \cdots)$, is $\mathbf{F}\cdot\mathbf{G}$, given by

$$\mathbf{F}\cdot\mathbf{G} = \sum_{i=-\frac{n-1}{2}}^{\frac{n-1}{2}} f_i g_i ,$$

where $n$ is the number of features in the items. The result is $(\cdots f_{-1}g_{-1} + f_0 g_0 + f_1 g_1 + \cdots)$. Over the set of unrelated items, the dot product between any two items is expected to be zero. In the computer simulations that will follow, unrelated items will be constructed by selecting feature values randomly for each feature of each item from a symmetrical distribution centered on zero. The representational assumptions for unrelated items are the same as those of Metcalfe and Murdock (1981) and differ from those of J. A. Anderson et al. (1977) and Liepa (Note 1) only insofar as the unrelated items are considered to be statistically independent rather than strictly orthogonal.

To go to the other extreme, suppose that two items are identical. Now when the dot product between these two identical items is taken, the result is not zero but rather the sum of squares of the values of the features, which is a positive value. Let us, for convenience, set this value at 1 so that $\mathbf{F}\cdot\mathbf{F} = 1$ for all items. Notice that this value is a measure of the strength of the item and that setting the value to be the same for all items amounts to saying that items begin with equal strength. This is not a necessary assumption, and there are no doubt experimental manipulations that would affect this value. Identity is taken as the most extreme case of similarity. As has often been pointed out, even when the same item is presented twice, it may not be encoded identically on both occasions. The construct of encoding variability will not be used in the present article as an explanation, however, although the model is not incompatible with this idea.

One may, of course, have intermediate degrees of similarity. Suppose, for instance, that half of the features in two items are identical and the other half independent. The independent features, when multiplied and added to form the dot product, will be expected to sum to zero as before, and the identical features will result in a positive value, as before. Because only one half of the total features are identical, the expected value of the dot product will be .5 rather than 1. The dot product thus indicates the extent of feature overlap (beyond independence) of two items.

In this article, similarity will be represented as the proportion of features that two items have in common (i.e., that have identical numerical values). It will be assumed that the number of features in each item is some large constant number and that the dot product of an item with itself gives a value of one. These representational assumptions are very simplified and can probably be relaxed considerably to be more psychologically plausible. For the present purpose, the simplicity of the assumptions makes it easier to see what CHARM does when items are associated, stored, and retrieved.

### Association Formation

It is postulated that two items in consciousness, each represented as an ordered set of features, may be associated by means of the operation of convolution. This method of association formation is fundamental to the holographic hypothesis and is shared by a variety of other models. The operation has other applications as well (see Murdock, 1979, for a review). Suppose that two items **F** and **G** are coded in terms of $n$ features (where $n$ is odd), yielding $\mathbf{F} = (f_{-(n-1)/2}, \ldots, f_{-1}, f_0, f_1, \ldots, f_{(n-1)/2})$ and $\mathbf{G} = (g_{-(n-1)/2}, \ldots, g_{-1}, g_0, g_1, \ldots, g_{(n-1)/2})$. The convolution of F and G, written as F*G, is a $2n - 1$ row vector whose $m$th term is given by

$$\mathbf{T}_m = \sum_{(i,j)\in S(m)} f_i g_j ,$$

where $S(m) = \{(i, j)| -(n-1)/2 \leq i, j \leq (n-1)/2, \text{ and } i + j = m\}$. If $n$ is equal to 3, it is straightforward to compute that $\mathbf{F*G} = (f_{-1}g_{-1}, f_0g_{-1} + f_{-1}g_0, f_1g_{-1} + f_0g_0 + f_{-1}f_1, f_1g_0 + f_0g_1, f_1g_1)$. Figure 1 presents an illustration of convolution for two items each consisting of three features, and a number of other illustrations can be found in Lathi (1968). The associative trace, which is the sum of the products along the central line of intersection, is represented in the box in the center of Figure 1. It is simple to compute the convolution of two items by constructing a matrix, such as the one shown in Figure A1 of the Appendix.

The operation of convolution is itself commutative or symmetric, that is, **F*G = G*F**. This can easily be verified by exchanging the items in Figure 1; the associative trace remains the same. This implies that the forward and backward associative strength between **F** and **G** should be the same if everything else is equal (cf. Asch, 1968; Ekstrand, 1966).

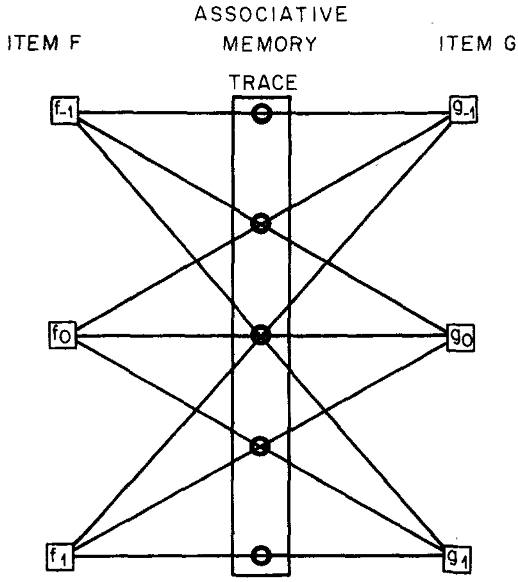


*Figure 1.* An illustration of the convolution of two items, each consisting of three features.

In CHARM, both of the associated items are considered to converge on a higher order associative trace. This trace is a complex interaction between the two items and does not itself resemble, in any obvious way, either of the contributing items. The meaning of the features in the association is not the same as the meaning of the features in the items. For instance, if one were to assign a name to one of the features in the items, that name would not apply (uniquely) to any of the features in the associative trace. Conceptually, this association is nothing like a pathway between the two items (as in, for instance, J. R. Anderson & Bower's, 1973, or J. R. Anderson's, 1976, models). As we shall soon see, the retrieval operation is nothing like a search.

To facilitate the exposition of how convolution works in combination with correlation, consider the special case in which an item is convolved with a delta vector, denoted $\delta$. A delta vector is an item that has values of zero on all features except the central feature, which has a value of 1. When an item is convolved with a delta vector, the result is the item itself. This may easily be verified by substituting the values of (0, 1, 0) for either F or G in Figure 1 and multiplying through each connection. The result is the other item. Convolving an item with a zero vector, that is, an item with values of zero on all features, produces a vacuous item—another zero vector. Convolving an item with what will be called an attenuated delta vector (i.e., an item with values of zero on all features except the central one, which has a value greater than zero but less than one) produces the other item with a strength that is equal to the central feature of the attenuated delta vector.

I do not propose that memory items are themselves delta vectors. However, if there were an operation that functionally converted one of the items into a delta vector, the other item would be reconstructed. This is just what the retrieval operation of correlation does.

### *Item Retrieval*

When a retrieval cue is presented and recall of the item with which that cue was associated is required, it is proposed that the cue is correlated with the associative trace. (Correlation, as used here, is not identical to correlation in statistics, which is the dot product, although the two are related.) Correlation may be defined in a manner similar to convolution. When **F** and **G** are coded as $\mathbf{F} = (f_{-(n-1)/2}, \ldots, f_{-1}, f_0, f_1, \ldots, f_{(n-1)/2})$ and $\mathbf{G} = (g_{-(n-1)/2}, \ldots, g_{-1}, g_0, g_1, \ldots, g_{(n-1)/2})$, the correlation of **F** and **G**, written as **F#G**, is a $2n - 1$ row vector whose $m$th term is given by

$$\mathbf{R}_m = \sum_{(i,j)\in S(m)} f_i g_j ,$$

where $S(m) = \{(i, j) | -(n-1)/2 \leq i, j \leq (n-1)/2, \text{ and } i - j = m\}$. It follows that when $n$ is equal to 3, for instance, the result of correlation $\mathbf{F\#G} = (f_{-1}g_1,\ f_0g_1 + f_{-1}g_0,\ f_1g_1 + f_0g_0 + f_{-1}g_{-1},\ f_1g_0 + f_0g_{-1},\ f_1g_{-1})$. As can be seen, the central feature of correlation is $\mathbf{F \cdot G}$.

If an item is correlated with itself, the value of the central feature will be one, because $\mathbf{F \cdot F}$ was set to one. The noncentral features will have an expected value of zero (but will not be exactly zero) because of the independence of the feature values and the overall expected value of zero. Thus, autocorrelation (**F#F**) results in an approximation to a delta vector. If the two items being correlated are similar, the central feature will be equal to the dot product between the two similar items, that is, equal to their similarity. Thus, the correlation of two similar items results in an attenuated delta vector. If the two items being correlated are unrelated, the expected value of the central feature is zero, as it is for all of the other features. Thus, the correlation of two unrelated items approximates a zero vector.

### *Relation Between Operations for Associating and Retrieving*

As was given in the Association Formation section, when an item is convolved with a delta vector, the result is the item itself:

$$\delta * \mathbf{G} = \mathbf{G}. \tag{1}$$

When an item is convolved with a zero vector, the result is a zero vector:

$$0 * \mathbf{G} = 0. \tag{2}$$

The result of convolution with an attenuated delta vector falls between these two extremes.

As was given in the Item Retrieval section, when an item is correlated with itself, an approximation to a delta vector results:

$$\mathbf{F\#F} \simeq \delta. \tag{3}$$

When an item is correlated with an unrelated item, an approximation to a zero vector results:

$$\mathbf{F\#G} \simeq 0, \quad \text{where} \quad \mathbf{F \cdot G} = 0. \tag{4}$$

The result of correlating an item with another similar item falls in between these two extremes, resulting in an approximation to an attenuated delta vector.

The relation between convolution and correlation for unrelated items has been given before (Borsellino & Poggio, 1973; Murdock, 1979; Liepa, Note 1). It will be simpler to give a more general equation here and then derive the result for unrelated items. The equation relating convolution and correlation is given by

$$\mathbf{F\#(F*G)} = \mathbf{(F\#F)*G} + \mathbf{(F\#G)*F} + \text{noise}. \tag{5}$$

The Appendix illustrates the convolution and correlation matrices that produce this summary equation and shows the locus of the two potentially nonnoise components, $(\mathbf{F}\#\mathbf{F})*\mathbf{G}$ and $(\mathbf{F}\#\mathbf{G})*\mathbf{F}$.

If **F** and **G** are unrelated, we may substitute the first four equations into Equation 5 as follows:

$$\begin{aligned}\mathbf{F}\#(\mathbf{F}*\mathbf{G}) &= (\mathbf{F}\#\mathbf{F})*\mathbf{G} + (\mathbf{F}\#\mathbf{G})*\mathbf{F} + \text{noise}\\ &\simeq \delta*\mathbf{G} + 0*\mathbf{F} + \text{noise}\\ &\simeq \mathbf{G} + 0 + \text{noise}\\ &\simeq \mathbf{G}.\end{aligned}$$

If **F** and **G** are unrelated, and **G** is used as a retrieval cue, Equation 5 becomes

$$\begin{aligned}\mathbf{G}\#(\mathbf{F}*\mathbf{G}) &= (\mathbf{G}\#\mathbf{F})*\mathbf{G} + (\mathbf{G}\#\mathbf{G})*\mathbf{F} + \text{noise}\\ &\simeq 0*\mathbf{G} + \delta*\mathbf{F} + \text{noise}\\ &\simeq 0 + \mathbf{F} + \text{noise}\\ &\simeq \mathbf{F}.\end{aligned}$$

Thus, when two unrelated items are associated by convolution, and either of the items is correlated with the resulting associative trace, an approximation to the other item is reconstructed or retrieved.

*Effects of Similarity*

Stimulus and response generalization fall out immediately from the above associative and retrieval scheme. When two unrelated items are associated by convolution, and then an item that is similar to but not identical with one of the associated items is correlated with the result, an approximation to the other item is produced. In this case the two similar items do not form a delta vector with a central feature of precisely one; instead, an attenuated delta vector is formed, with a central feature equal to the similarity (dot product) between the two items. The retrieval operation produces the other item with a strength that is proportional to the dot product. Thus, if **F** and **G** are convolved and then $\mathbf{F}'$ (an item similar to **F**) is correlated with the trace, **F** and $\mathbf{F}'$ will form an attenuated delta vector that will produce or reconstruct **G** with a strength proportional to $\mathbf{F}\cdot\mathbf{F}'$. In a like manner, if $\mathbf{G}'$ was used as a retrieval cue, **F** would be produced with a strength proportional to $\mathbf{G}\cdot\mathbf{G}'$. The fact that stimulus generalization results directly from the method of association formation and retrieval suggests that the model might be directly applicable to a wide variety of experimental paradigms (see Gibson, 1941). The fact that response generalization occurs for the same reason gives rise to the counterintuitive prediction that an item similar to a response will evoke recall of the stimulus—a prediction that will be tested later in this article.

If the two associated items are similar to each other (i.e., $\mathbf{F}\cdot\mathbf{G} > 0$), both of the nonnoise components given in Equation 5 are important. For example, consider the limiting case in which **F** and **G** are identical. Equation 5 becomes

$$\begin{aligned}\mathbf{F}\#(\mathbf{F}*\mathbf{G}) &= (\mathbf{F}\#\mathbf{F})*\mathbf{G} + (\mathbf{F}\#\mathbf{G})*\mathbf{F} + \text{noise}\\ &\simeq \delta*\mathbf{G} + \delta*\mathbf{F} + \text{noise}\\ &\simeq 2\mathbf{G} + \text{noise}\\ &\simeq 2\mathbf{F} + \text{noise}.\end{aligned}$$

The item produced will have twice the strength as would have been the case had it been associated with an unrelated item. (Some complications arise from the internal noise that make this only an approximation, but it is a fairly good one, so long as other associations enter into the composite trace as well as the target association.)

Suppose now that instead of the associated items being identical, they are similar to each other. For the sake of illustration, suppose that $\mathbf{F}\cdot\mathbf{G} = .5$ because about 50% of the features are identical (and the remainder are independent). If these two similar items were correlated, the result would be an attenuated delta vector with a central feature equal to $\mathbf{F}\cdot\mathbf{G} = .5$. Equation 5 then becomes

$$\begin{aligned}\mathbf{F}\#(\mathbf{F}*\mathbf{G}) &= (\mathbf{F}\#\mathbf{F})*\mathbf{G} + (\mathbf{F}\#\mathbf{G})*\mathbf{F} + \text{noise}\\ &\simeq \delta*\mathbf{G} + .5\delta*\mathbf{F} + \text{noise}\\ &\simeq \mathbf{G} + .5\mathbf{F} + \text{noise}.\end{aligned}$$

The retrieved item is a linear combination of the original items **F** and **G**. Each feature in **G** is weighted or multiplied by 1, each feature in **F** is multiplied by .5, and the results are added, feature by feature. This tends to accentuate the features that are common to **F** and **G**, making them stronger than they would otherwise be, and to have an interfering effect on the features that are unique to the target item **G**. The item that is re-created is a systematic distortion of the encoded (target) item **G** such that it is more like the associated item **F** than would be the case had the target been associated with an item that was not similar. The fact that the model systematically distorts the retrieved item from its encoded form when the two associated items are similar, but not when they are unrelated, has experimental implications that will be tested later in the article.

In summary, when two items that are unrelated are associated by convolution and then one of the two is correlated with the trace, an approximation to the other item is retrieved. When an item similar to one of the convolved items is correlated with the trace, the other item is still produced, but with a decreased strength. When two similar items are convolved and then one of the two is correlated with the trace, the item that is retrieved is systematically distorted from its encoded form such that it is more like both of the associated items and less like the target item itself, taken in isolation. The interactive association itself introduces this distortion.

*The Composite Associative Memory Trace*

In CHARM, associations are stored in a composite memory trace—termed *composite* because the composite portraits of Galton (1879), in which photographs of faces were superimposed to yield a prototypical face, provide a good visual analogy to such a memory trace. J. A. Anderson and his colleagues (J. A. Anderson, 1970, 1972, 1973, 1977; J. A. Anderson et al., 1977; J. A. Anderson & Hinton, 1981) have argued for the neurological and psychological plausibility of this same construct.

Figure 2 gives a numerical example of how the composite trace in CHARM is formed. The figure shows how two pairs of items are first convolved and then added into the composite trace. The trace starts out with some numerical values that are the result of previous pairs that have been added into it. When a pair of items (**F** and **G**) is presented, convolution occurs as described earlier. The matrix in the figure is a convenient way of computing the result of convolution. The numerical values in the figure were chosen for their computational simplicity rather than to conform to the independence and strength assumptions in the model. The re-

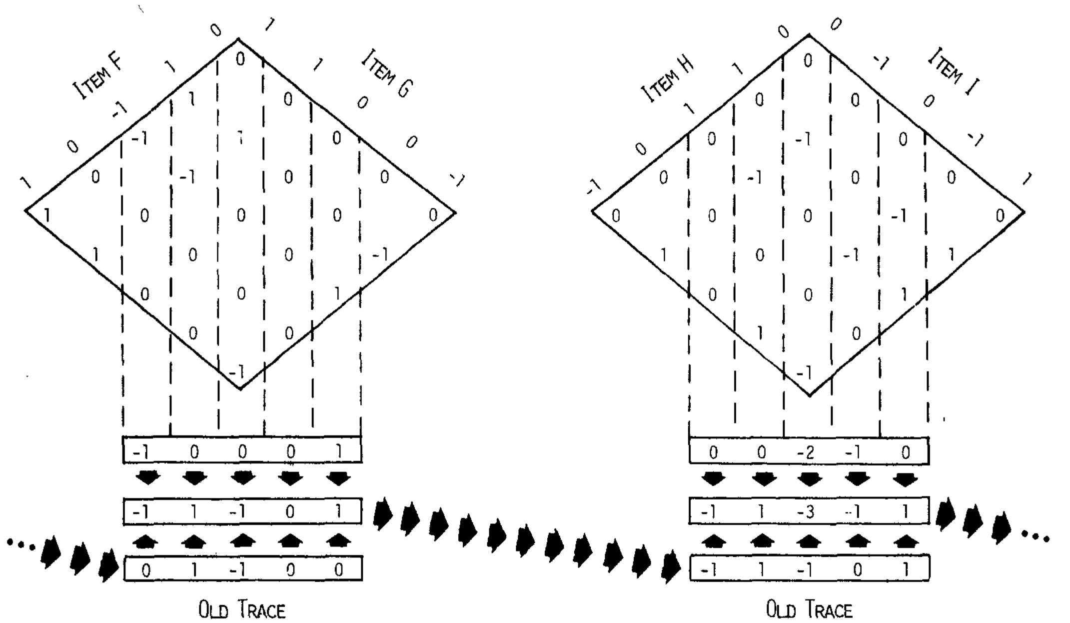


*Figure 2.* A schematic illustration of the formation of the composite memory trace. (Note that the associations are truncated as in the simulations).

sult of convolution is added, feature by feature, into the trace. When the next pair (**H** and **I**) is associated, the result is also added into the composite trace.

How does the composite trace influence what is recalled? When the items of a list are all unrelated, the fact that the trace includes pairs other than the target pair adds noise to the recalled item. Suppose a list consists of two pairs of four unrelated items: **F–G** and **H–I**. When **F** is given as a retrieval cue, the reconstruction can be characterized as follows:

$$\begin{aligned}
&\mathbf{F}\#[(\mathbf{F}*\mathbf{G}) + (\mathbf{H}*\mathbf{I})] \\
&\quad = (\mathbf{F}\#\mathbf{F})*\mathbf{G} + (\mathbf{F}\#\mathbf{G})*\mathbf{F} + \text{noise}_{\mathbf{F}*\mathbf{G}} \\
&\qquad + (\mathbf{F}\#\mathbf{H})*\mathbf{I} + (\mathbf{F}\#\mathbf{I})*\mathbf{H} + \text{noise}_{\mathbf{H}*\mathbf{I}} \\
&\quad = \delta*\mathbf{G} + 0*\mathbf{F} + \text{noise}_{\mathbf{F}*\mathbf{G}} + 0*\mathbf{I} \\
&\qquad + 0*\mathbf{H} + \text{noise}_{\mathbf{H}*\mathbf{I}} \\
&\quad = \mathbf{G} + \text{noise}_{\mathbf{F}*\mathbf{G}} + \text{noise}_{\mathbf{H}*\mathbf{I}}\,.
\end{aligned}$$

In this situation, the effect of the nontarget **H–I** pair is to add noise to the reconstructed item **G**. Quite literally, the composite trace produces interference. Metcalfe and Murdock (1981) present a number of computer simulations that illustrate this buildup of noise due to an increasing number of pairs in the memory trace, when those pairs are unrelated.

Now consider a situation in which all four items are similar to one another. For the sake of illustration, suppose that the dot product of any item with any other item is .5. Going through the reconstruction computations gives

$$\begin{aligned}
&\mathbf{F}\#[(\mathbf{F}*\mathbf{G}) + (\mathbf{H}*\mathbf{I})] \\
&\quad = (\mathbf{F}\#\mathbf{F})*\mathbf{G} + (\mathbf{F}\#\mathbf{G})*\mathbf{F} + \text{noise}_{\mathbf{F}*\mathbf{G}} \\
&\qquad + (\mathbf{F}\#\mathbf{H})*\mathbf{I} + (\mathbf{F}\#\mathbf{I})*\mathbf{H} + \text{noise}_{\mathbf{H}*\mathbf{I}} \\
&\quad = \delta*\mathbf{G} + .5\delta*\mathbf{F} + \text{noise}_{\mathbf{F}*\mathbf{G}} + .5\delta*\mathbf{I} \\
&\qquad + .5\delta*\mathbf{H} + \text{noise}_{\mathbf{H}*\mathbf{I}} \\
&\quad = \mathbf{G} + .5\mathbf{F} + .5\mathbf{I} + .5\mathbf{H} \\
&\qquad + \text{noise}_{\mathbf{F}*\mathbf{G}} + \text{noise}_{\mathbf{H}*\mathbf{I}}\,.
\end{aligned}$$

The main point to note is that in this situation, the composite trace not only produces noise but also causes all of the items in the list to contribute to the pattern that is retrieved from memory.

As a final illustration, consider the case where two unrelated target items (**B** and **D**) have been associated with the same cue (**A**) and both associations (**A–B** and **A–D**) have been stored in the composite trace. When the cue is correlated with the trace, the single item that is retrieved is a combination of both of the unrelated targets. This situation, which corresponds to Barnes and Underwood's (1959) A–B A–D paradigm, can be schematized as follows:

$$\begin{aligned}
&\mathbf{A}\#[(\mathbf{A}*\mathbf{B}) + (\mathbf{A}*\mathbf{D})] \\
&\quad = (\mathbf{A}\#\mathbf{A})*\mathbf{B} + (\mathbf{A}\#\mathbf{B})*\mathbf{A} + \text{noise}_{\mathbf{A}*\mathbf{B}} \\
&\qquad + (\mathbf{A}\#\mathbf{A})*\mathbf{D} + (\mathbf{A}\#\mathbf{D})*\mathbf{A} + \text{noise}_{\mathbf{A}*\mathbf{D}} \\
&\quad = \delta*\mathbf{B} + 0*\mathbf{A} + \text{noise}_{\mathbf{A}*\mathbf{B}} + \delta*\mathbf{D} \\
&\qquad + 0*\mathbf{A} + \text{noise}_{\mathbf{A}*\mathbf{D}} \\
&\quad = \mathbf{B} + \mathbf{D} + \text{noise}_{\mathbf{A}*\mathbf{B}} + \text{noise}_{\mathbf{A}*\mathbf{D}}\,.
\end{aligned}$$

The fact that the retrieved item is a combination of more than one item in this situation, as well as in situations where to-be-remembered items are similar to each other, reveals a problem of possible response ambiguity. The construct of semantic memory, detailed below, solves this problem.

To summarize, a quick rule of thumb for saying what will be retrieved from the composite–holographic memory trace is as follows: Take the dot product (i.e., the similarity) of the cue item with each of the items in the list. If the dot product between the cue and a particular item is about zero, then only noise will be produced. If the dot product between the cue and Item X is not zero, then the complement item, with which Item X was convolved, will be retrieved with a strength proportional to the dot product. This calculation may be carried out for every item in the list, giving the strength with which the complement items are retrieved. All of the noise and the retrieved items, weighted by their appropriate strengths, are then added to give an approximation to the single item that is recovered from memory.

### *Semantic Memory*

Because the item that is retrieved from the composite trace is intrinsically noisy, and may be systematically distorted or ambiguous, CHARM requires a pattern recognizer if it is to make predictions about what discrete words will be recalled. This pattern recognizer basically identifies the retrieved item as such and such a word and allows for a discrete response.

There are two parts to the pattern recognizer. First, because the retrieved item is to be assigned a particular label, there is a listing of all the possibilities. This listing of responses corresponds in a general way to what is often called semantic memory, where the term is interpreted in the restricted sense of a lexicon. The usage of the term *semantic memory* seems most similar to that given by Kintsch (1974). Thus, when an item is retrieved from the composite trace, it is identified by being matched to every item in the lexicon.

Second, there is a matching process that serves to select one (or more) of the discrete possibilities as the response(s). The matching process is similar to the resonance operation described in Ratcliff's (1978) item-recognition model. Each feature in the retrieved item is multiplied by the corresponding feature in each lexical item, and the sum of the products for all of the features corresponding to a particular lexical item is taken. In this manner, a separate dot product is computed for every lexical item. The better the match between the retrieved item and a particular lexical item, the higher will be the dot product—or "resonance" value (see Ratcliff, 1978)—for that lexical item. In the simulations that follow, the matching process is allowed to be exhaustive; that is, every feature of every lexical item is matched to the retrieved item. (An alternative approach is to halt the matching process once sufficient information has accumulated—via a random walk to a criterion—for a particular response; see Ratcliff, 1978. Whether the matching process stops at a criterion or is exhaustive probably depends on the demands of the specific task at hand.)

Resonance values allow us to determine what the retrieved item will be identified as and to make contact with verbal cued-recall experiments in which, in most instances, responses are discrete. It is at the identification stage that the strength of features becomes especially important. Features that have high absolute resonance values—that is, strong features—will contribute more to the selection of a response than will weak features. The implementation of a lexicon for pattern recognition is the same as in Metcalfe and Murdock (1981) but is not explicit in any other holographic model. It allows CHARM to address situations involving response ambiguity and thereby enables the model to make predictions concerning errors of recall and to deal elegantly with the A–B A–D transfer paradigm. In the third and fourth experiments reported below, evidence for the separate existence of a semantic memory system is provided.

For simplicity, no occurrence information is stored at the lexical level. The matching process is thus properly considered to be an identification process rather than a recognition-memory process. CHARM is a model of recall, not recognition.

## Experimental Tests of CHARM

In the preceding discussion, I described the workings of the model. I said little, however, about explicit experimental predictions. In this section, I will describe several simple single-trial cued-recall experiments and outline the predictions of the model. I will present computer simulations of the model that demonstrate the predictions in a more concrete fashion. These simulations were run predictively (before the experiments were conducted) and thus are attempts not to fit data but rather to make qualitative predictions. I will present several experiments that seek to determine whether people, when in recall situations that should conform fairly closely to those simulated, produce patterns of recall that correspond to the predictions of the model.

### *Intralist Intrusions*

The first simulation is directed at a situation in which people are presented with a list of cue–target word pairs to study and remember. The list may consist of words that are all unrelated to each other or that are all similar to each other. After studying, subjects are given the cue words as retrieval stimuli, and the question asked is, What will be recalled?

### Simulation 1

*Method.* The entire simulation that I will describe here was run 100 times for each of the two conditions: similar list and unrelated list. The only difference between these runs was that each had a different random assignment for the initial representations of the items. One run of this simulation for the unrelated-list condition will first be detailed. Then the alteration in the program that gives rise to the similar list condition will be described.

A small lexicon consisting of 12 items was constructed. Six of these items were designated to be list items and six were unrelated extralist items. Each feature of each lexical item was a number randomly selected from a unit normal distribution with a mean value of zero. The randomization subroutine RANDR was used and care was taken that the computer specifications match those for which this subroutine was originally designed. Each item consisted of 63 features. The 63-tuple composing each item was then normalized so that the dot product of any item with itself equaled one. Because of the random selection of feature values, the dot product of any item with any other item was approximately equal to zero.

The associative memory trace was formed by convolving Items 1 and 2, Items 3 and 4, and Items 5 and 6. The results of the three convolutions were added into a single 63-tuple made up of the central 63 features of the convolutions. Once this composite trace was formed, the program attempted to recall.

In order to recall, each of Items 1, 3, and 5 was correlated with the composite trace, resulting in three retrieved patterns that were again truncated to the central 63 features. Each retrieved pattern was then matched to each of the 12 lexical items by multiplying each feature in the retrieved item by the corresponding feature in the lexical item and summing the products, yielding a resonance score for each lexical item. For purposes of illustration, the program also chose the lexical item that had the highest resonance score (as long as it was greater than zero) to give as a response. This responding criterion is arbitrary but will illustrate the pattern of recall more concretely than do the resonance scores themselves.

In the similar-list condition, the lexicon was initially constructed in exactly the same way as in the unrelated-list condition. Then the values of every other feature in the first six items were altered so that they were all exactly the same as the feature values originally assigned to Item 1. Thus, the items that were to make up the list had identical values on 32 features and independent values on 31 features. This representation of similarity is crude, but a later simulation on protoypes that uses a more refined representation gives about the same results. The convolution, correlation, and identification procedures were the same in the similar-list condition as in the unrelated-list condition.

*Results.* As evidenced by the results presented in Table 1, the model predicts that the discriminability of the target relative to other list items is worse when the list is composed of similar rather than unrelated items. Thus, a distinctive pattern of intrusion errors should be found when the responses to a list of similar items are compared to those to a list of unrelated items. When the list consists of unrelated items, intralist intrusions should be quite infrequent, whereas when the list items are all similar, the frequency of intralist intrusions should be considerably higher. The resonance scores show this pattern and also show that there should be no difference in the rate of intrusions depending on whether

Table 1
*Mean Resonance and Recall for Each Type of Response in Simulation 1*

| | Type of response | | | | |
|---|---|---|---|---|---|
| Condition | Target | Cue | Nontarget response | Nontarget stimulus | Unrelated extralist |
| Unrelated | | | | | |
| Resonance | .80 (.20) | .00 (.30) | .00 (.20) | .00 (.20) | .00 (.20) |
| Recall | 98.6% | .6% | .6% | .0% | .3% |
| Similar | | | | | |
| Resonance | 1.70 (.60) | 1.50 (.60) | 1.50 (.60) | 1.50 (.60) | .00 (.40) |
| Recall | 42.0% | 14.6% | 21.7% | 21.7% | .0% |
| No. (out of 12) of lexical items corresponding to response type | 1 | 1 | 2 | 2 | 6 |

*Note.* Standard deviations are in parentheses.

a given item had been a stimulus or a response item.

The differential pattern in the two conditions stems from the fact that in the model, the pattern that is retrieved from the composite trace is a combination of all of the items that are stored in the composite trace when those items are similar to each other but consists only of the target item when the list items are unrelated. This intrusion pattern is thus expected even though there may be lexical items in the unrelated-list condition that are equally similar to the target item as are the list items to the target in the similar-list condition.

The table shows the levels of recall given the arbitrary decision criterion of the best matched item. These numbers are not quantitative estimates because no attempt was made to quantify the degree of similarity in the simulation as compared to the experiment that will follow, to equate the number of pairs in the trace, or to estimate the number of features. The recall figures do serve to illustrate that even though the resonance of the target item was higher in the similar-list as compared to the unrelated-list condition, correct recall in the similar-list condition is expected to be worse than in the unrelated-list condition. This occurs because the retrieved item provides little of the information on the nonshared features of the similar items that is necessary to discriminate the target item from the other list items.

The simulation also produced the counterintuitive prediction that there should be cue intrusions—that is, recall of a cue word to itself as a cue—in the similar-list but not in the unrelated-list condition.

### *Experiment 1*

*Method.* In this experiment, people's recall to a list of similar pairs was compared to recall to a list of unrelated pairs.

Four sets of materials were constructed such that each set consisted of two lists: one in which the words were all from a single category and one in which each word was from a different category. Each list consisted of 14 pairs of words. In order to construct the similar list in each set, a category was randomly selected from the Toronto categorized word pool (Murdock, 1976). Words were randomly selected from the category until 14 pairs resulted. The uncategorized list in each set was constructed from the same word pool. Twenty-eight categories were randomly selected (with the exclusion of the category that was used in the similar list in the same set). A word was randomly chosen from each of the 28 categories, and the words were then randomly paired.

Each of the 16 subjects in the experiment (volunteers from the social sciences subject pool at the University of California, Irvine) received two lists of pairs (one set) to study. Four subjects were assigned to each of the four sets of materials. Two subjects studied and were tested on the categorized list first, whereas the remaining two subjects studied and were tested on the unrelated list first. The lists were typed in uppercase letters on sheets that were given to subjects to study for 4 minutes. After studying, subjects counted backwards from a large number by threes for 1 minute and wrote down their final count. They were instructed that the right-hand members of the pairs were the targets or responses and that they would be given the left-hand items as cues.

Cued recall was written in booklets that contained one cue per page. The order of presentation of cues was randomized for every subject. Beside each response, subjects were asked to write their degree of confidence about the correctness of the response on a scale of 1 to 6, where 1 was guessing and 6 was highly confident. Recall was subject paced, and subjects were instructed not to look back to previous responses or forward to upcoming cues until all attempts to respond had been made.

*Results.* The mean number of correct responses, intralist intrusions, and extralist intrusions in each of the two conditions are presented in Table 2, along with the confidence ratings. As can be seen from the table, the entire pattern of recall and intrusions corresponded quite well to the predictions of the model. In particular, the nontarget list items were poorly discriminated from the target items in the categorized-list condition but not in the uncategorized-list condition. An analysis of variance was conducted in which the factors were condition (categorized or uncategorized), response type (correct, stimulus intrusion, response intrusion, and extralist intrusion) and set (1 to 4). There was a significant interaction between response type and condition, $F(3, 36) = 34.47$, $p < .05$. There were no other significant effects or interactions. A chi-square test, where the observations were treated as if they were independent, was also performed. There was a significant relation between the frequency of correct and intralist intrusions and the categorized or uncategorized conditions, $\chi^2(1) = 30.67$.

There was no statistical difference between intrusions that were originally stimulus items and response items, $t(15) = .81$, although the slight tendency was in favor of the stimulus items. This result offers some support for the symmetrical nature of the association in the model.

Table 2
*Mean Number of Correct Responses and Confidence Rating for Each Type of Response in Experiment 1*

| | Type of response | | | |
|---|---|---|---|---|
| Condition | Target | Cue | Other list item | Extra-list item |
| Uncategorized | | | | |
| No. recalled[a] | 11.2 | 0 | .8 | .1 |
| Confidence rating[b] | 5.9 | — | 3.8 | 3.5 |
| Categorized | | | | |
| No. recalled[a] | 7.9 | 0 | 3.1 | .5 |
| Confidence rating[b] | 5.7 | — | 2.6 | 1.4 |

[a] Out of 14. [b] Out of 6.

In summary, the predictions of the model given in Simulation 1 were largely supported by the results of the present experiment. The one prediction that absolutely failed to materialize was that of intrusions of the cue item to itself as a cue: No subject ever said that the target item was A when A was given as a cue. It seems likely that because an item was never paired with itself in the experiment, subjects were able to use a rule to inhibit this particular response. Nevertheless, the prediction of cue intrusions is important, not necessarily for its own sake, but because it is a direct result of the interactive nature of the postulated association.

It is of interest to note that the prediction of differential intrusion errors in the similar as compared to the unrelated condition derives directly from the property of the model that, in the former condition, the composite trace causes the retrieved item to be systematically distorted towards a "prototype" of the list. The results of both the simulation and the experiment indicate that CHARM should be extendable to prototype learning experiments. Further, errors of recall, according to the model, are related to people's ability to form concepts. The extension of the model to prototype learning will be made in a later section of the article.

### *Extraassociative-List Context Effects*

For the same reason that intralist intrusions were predicted in the first experiment, the model predicts that the makeup of the whole list of paired associates, not only the to-be-remembered pair, will influence the goodness of recall. Consider a situation in which a given pair, say NAPOLEON–ARISTOTLE, is embedded in a list of pairs of names of famous people (homogeneous list) or is embedded in a list containing a variety of conceptually different pairs such as RED–BLUE, 14–22, and so forth (heterogeneous list). As was pointed out to the author (Tulving, Note 2), CHARM makes precisely the right predictions about the difficulty of recall under these conditions. NAPOLEON–ARISTOTLE should be more difficult in the homogeneous-list condition because, in that condition, the composite trace will cause a prototype of the entire list rather than just the appropriate target response to be produced. In an experiment that included this comparison, Tulving (Note 3) found that the target association was indeed harder to learn in the homogeneous- than in the heterogeneous-list condition.

### *Cue Intrusions*

The predictions of differential cue intrusions between the similar and unrelated conditions was not confirmed in Experiment 1, probably because subjects were able to use a rule to eliminate these responses. This prediction is, nevertheless, an important one because it bears directly on the structure of the interactive association postulated in CHARM and, further, because it is a counterintuitive or high-risk prediction. The question that will be addressed in the present experiment is whether cue intrusions are more likely to occur when a cue and target are similar rather than unrelated. Tulving (1974) has reported the presence of cue intrusions in a study in which some of the pairs had identical cues and targets, so the inclusion of this type of pair in the list appears to override, at least to some extent, a rule excluding these intrusions.

CHARM makes the prediction that cue intrusions will be more frequent when the cue–target items are similar to each other because the item that is retrieved from memory resembles both the cue and the target. It is because the retrieved item in the similar condition is a combination of both A and A′, where the common features are accentuated and the contrasting features suffer interference, that one expects this retrieved item to be sometimes mistakenly (from the experimenter's point of view) identified as A—the cue itself. When the two associated items A and B are unrelated and A is presented as a cue, the cue item is not itself reconstructed, and so cue intrusions are not expected.

It is not clear that other models make this prediction. In models such as FRAN (J. R. Anderson, 1972), in which the probability of forming an association depends on the similarity of the to-be-associated items, it would seem that the probability of recall might be higher in the similar as compared to the unrelated condition (in contrast to the previous experiment). There is no indication, however, of whether there should be more or fewer cue intrusions in the similar as compared to the unrelated condition. There might be more if a selection of nodes (different from the central node for a given item) occurs, as is suggested by J. R. Anderson (1976). On the other hand, there might be fewer cue intrusions in the similar condition if the subject guesses the cue in the absence of a tagged pathway to a different tagged node (as presumably occurs more frequently when the items are not similar). There is little doubt that either result could be accommodated after the fact, but no a priori prediction is clear. Bower's (1972) multicomponent model makes reference only to the properties of the stimuli and so can make no predictions about phenomena attributable to the relation between the two associated items. The same may be said for Medin and Schaffer's (1978) feature model, which represents similarity but does not consider the responses, only the stimuli. This is also the case for the J. A. Anderson et al. (1977) model, which contains a forward-directional association, although this model is in other respects like CHARM. Because Flexser and Tulving's (1978) gestalten model does not include any representation of similarity, it makes no predictions in this situation. To my knowledge, there are no other models that make the prediction of cue intrusions. If subjects were guessing some of the time when they had no information about the correct response, one would expect that the conditional probability of a cue intrusion given an intrusion should be the same in both conditions.

### *Simulation 2*

*Method.* In this simulation, the similarity of the cue-target pairs was varied within a single list. Three levels of similarity were used; identical, similar (50% feature overlap), and unrelated. As in the first simulation, a lexicon of 12 items was constructed such that each item was a random normalized 63-tuple. Items 1 and 2 were reassigned feature values so that they were identical; Items 3 and 4 were reassigned feature values so that 32 of the 63 features were identical and the remaining 31 features were independent; and Items 5 and 6 were left with their original independent feature values. Items 7 through 12 were not included in the study list but were considered to be unrelated extralist possibilities in the lexicon. To make the simulation somewhat more realistic, Item 7 was assigned the same feature values as were Items 1 and 2 on half its features, Item 8 was similarly related to Item 5, and Item 9 was related to item 6.

The associative memory trace was constructed by convolving Items 1 and 2, Items 3 and 4, and Items 5 and 6 and adding the results into a single 63-tuple representing the composite trace. Then Items 1, 3, and 5 were correlated with the trace to produce three retrieved items. These retrieved items were matched to the items in the lexicon by means of the matching process described in Simulation 1. The simulation was run through 100 replications, and the means and standard deviations of the resonances of the lexical items to each of the three retrieved items were computed.

*Results.* The main point to note from the results shown in Table 3 is the difference in the resonances of the cues for the similar and unrelated conditions. When the original two items in a pair were similar, the resonance of the cue was almost as high as was that of the target. In the unrelated condition, the resonance of the cue was essentially zero, whereas the target itself was fairly strongly evoked.

I note parenthetically that if a rule can be used to eliminate the cue as a response, the level of recall is proportional to the resonance scores of the target, in this simulation. Recall will then be a direct function of the similarity between the stimulus and the response. This is the case because, unlike Simulation 1, the only competing response is the cue itself. However, in the experiment that follows, an attempt is made to deliberately inhibit the use of such a rule. Thus, no prediction is made about the level of recall: The only clear prediction is that there should be relatively more cue intrusions in the similar-pair condition than in the unrelated-pair condition.

### *Experiment 2*

*Method.* The lists in the present experiment each consisted of three types of pairs: identical pairs, synonym

Table 3
*Mean Resonance for Each Type of Response in Simulation 2*

| Condition | Type of response: Target | Cue | Other list item | Unrelated Extralist item |
|---|---|---|---|---|
| Unrelated | | | | |
| *M* | .76 | .02 | .00 | .00 |
| *SD* | .20 | .32 | .27 | .22 |
| Similar | | | | |
| *M* | .94 | .85 | .07 | −.02 |
| *SD* | .29 | .41 | .30 | .24 |
| Identical | | | | |
| *M* | 1.42 | 1.42 | −.01 | .02 |
| *SD* | .32 | .32 | .32 | .26 |

*Note.* The unrelated, similar, and identical conditions were all within a single list.

pairs, and unrelated pairs. There were 20 pairs of each type in each list, for a total of 60 pairs. Although the cue–target words varied in their similarity to each other within a pair, there was no obvious relation among the pairs in the list. Four different lists were constructed. The words that constituted the four lists in the experiment were chosen from the 80 highest ranked synonyms of Whitten, Suter, and Frank (1979). The lists were constructed so that the same response term occurred equally frequently as a response in each of the three conditions, over lists, and the mean synonymy was about the same for all lists.

To construct the lists, each of the four pairs that were the highest ranking four synonym pairs given by Whitten et al. (1979) were assigned by means of a Latin square to the four lists in four nominal conditions: identical, synonym, cue of the unrelated pair, target of the unrelated pair. The synonym given first in Whitten et al.'s (1979) listing was the particular member of the synonym pair that was used in these four nominal conditions. In the synonym condition, both words from the synonym listing were used. The assignment of words to conditions and lists was repeated for each successive tetrad of synonym pairs given in the norms until there were 60 pairs in each list. The order of pairs in each list was then randomized. There were two random test orders for each of the four lists.

Thirty-two subjects, who were volunteers from the social sciences subject pool at the University of California, Irvine were tested. Eight subjects were assigned to each list.

The lists were read to subjects at a rate of 10 sec per pair. Following study, subjects wrote the names of the states of the United States in alphabetical order for 6 minutes. They were then given a booklet containing the 60 cue items—always the first item in a pair—and were asked to write their responses along with a rating of the certainty of their responses on a scale of 1 (guessing) to 6 (highly confident). Ten minutes were allowed for cued recall.

*Results.* The prediction that there would be more cue intrusions in the synonym than in the unrelated condition was confirmed. When the probability of a cue intrusion, given an intrusion, was calculated for each subject in both conditions, the resultant mean conditional probability was .60 in the synonym condition and .46 in the unrelated condition, $t(31) = 1.89$, $p < .05$. The same pattern resulted when the data were pooled over subjects. The conditional probability of a cue intrusion was then .50 in the synonym condition and .39 in the unrelated condition.

Table 4 shows the frequency of cue and noncue intrusions for the synonym and unrelated conditions when these were divided into highly confident (with ratings of 4 or higher) and low-confident (with ratings of 3 or lower) intrusions. There was a disproportionate number of highly confident cue intrusions in the synonym conditions, $\chi^2(1) = 5.19$, $p < .025$, but not in the unrelated condition $\chi^2(1) = .00$, *ns*. The confidence-rating results provide converging evidence that these intrusions were not due simply to some complicated guessing strategy.

The level of correct recall was .63 in the synonym condition and .22 in the unrelated condition. When the cue–target pair had originally consisted of identical words, the level of correct responding was .59. I thought that as long as subjects actually associated the identical words with each other, the level of recall would be better in the identical than in the synonym condition rather than show

Table 4
*Comparison of Cue Versus Other Intrusions as a Function of Condition and Confidence in Experiment 2*

| Type of intrusion | Synonym condition: High confidence | Synonym condition: Low confidence | Unrelated condition: High confidence | Unrelated condition: Low confidence |
|---|---|---|---|---|
| Cue | 31 | 28 | 30 | 63 |
| Other | 15 | 34 | 46 | 97 |

*Note.* A total of six intrusions for which no confidence ratings were obtained have been excluded. Of these six intrusions, two occurred in the synonym condition and four in the unrelated condition; within each condition, the excluded intrusions were divided equally according to type (cue versus other).

no difference. One should note that in Tulving's (1974) study, recall of identical pairs was worse than that of similar pairs. In other experiments it has been shown that the level of recall increases with increased similarity between the cue and target (McKoon & Ratcliff, 1979, for instance). It seems that performance in the identical condition, which is not usually included, is most puzzling.

The cue-intrusion prediction was made because the item recreated from the holographic association is, under some conditions, a combination of the cue and target and, hence, is likely to be mistakenly identified as the cue. It would be of some interest to tap experimentally an item that is itself a combination rather than rely on the intrusion probabilities as the only evidence for this combinatorial property. An experiment by Loftus (1979) may be relevant. Loftus showed subjects slides of people going about routine activities. One of the slides showed a person reading a green book. Later, a leading question suggested that the book had actually been blue. When subjects were asked to select the remembered color on a color wheel, their "choices tended to be a compromise between what they actually saw and what they were told on their questionnaire, that is, they selected a bluish-green" (p. 123). One plausible interpretation is that the subjects actually retrieved a superimposed combination of the two events, as both the composite trace and the holographic association in CHARM produce. Gibson (1941) has also noted that in a paired-associate learning task in which subjects are given two different nonsense syllables as responses to the same stimulus, a frequent error was the combination of the two nonsense syllables. Bartlett (1932) noted that in an experiment in which subjects were given name–face pairs and were later asked to describe the face corresponding to a particular name, details from other faces were often incorporated into the descriptions. These results suggest that what is retrieved is not simply a degraded but essentially veridical copy of the target item but may actually be a combination of a number of events.

The finding of differential cue intrusions in the present experiment is an indication that the item that is retrieved from memory may be a combination of the items that were associated when those items are initially similar to each other. This finding was predicted because of the interactive nature of the association in CHARM. It is a counterintuitive finding that is not predicted by other models of recall.

### *Extralist Cuing*

In the preceding experiments, I tested several predictions of the model with respect to errors of recall. In this section, a rather different question will be addressed: the efficacy of extralist cues. These two areas of investigation have in common the fact that the interactive association model yields predictions about both. They are linked not by the experimental paradigm but rather by the concept of the interactive association.

Consider a situation in which a person studies a list of pairs of unrelated words. The subject is told that the right-hand word in each pair is the target word that will later be cued for recall. We now ask the question, What sorts of cues will be effective in eliciting the target items? Three possibilities will be considered, namely, the word that was presented with the target item in the list (i.e., the list cue), a word that is similar to the list cue, and a word that is similar to the target itself. Before presenting the simulation of this situation, let us consider for a moment what might be expected.

If one simply assumes that an association is formed between the list cue and the target item, then it seems likely that the list cue would serve as an effective reminder for the target. Presumably, all associative models, including the present one, are in accord on this point. One might also expect that an extralist cue similar to the list cue should give rise to recall of the target. This is a classic case of stimulus generalization, and although the term *stimulus generalization* does not itself explain the phenomenon, there are a number of feature models (e.g., Bower, 1967; Estes, 1955; Medin & Schaffer, 1978) in addition to the present one that can address it.

Because the phenomenon of stimulus generalization is well known and well substantiated, it is probably intuitive that the item similar to the cue should allow recall of the target. More obscure is the question of what the item that is similar to the target will cause to be recalled. There is a considerable amount of data that show that an item that is preexperimentally related to a list item may facilitate recall of the list item (Bahrick, 1969; Bilodeau, 1967; Bilodeau & Blick, 1965; Light, 1972; Santa & Lamwers, 1974). It is not clear, however, that items that are similar to each other in the sense that they share features are related in this same sense. If they are, then perhaps they should cause the target item to be recalled. However, there are other data (e.g., Santa & Lamwers, 1974; Thomson & Tulving, 1970) that show that even preexperimentally related cues do not necessarily cause the target item to be recalled. Intuition suggests that an extralist word such as BUY might be an effective retrieval cue for a target word such as PURCHASE.

The most interesting prediction of CHARM in this extralist cuing situation is that a synonym of the target word, rather than causing the target word itself to be recalled, will tend to retrieve the unrelated intralist cue. This prediction, as well as the stimulus-generalization prediction, will be demonstrated in the simulation that follows.

### *Simulation 3*

*Method.* In this simulation, the efficacy of extralist cues was examined. A 12-item lexicon was constructed as in the previous simulations. The items consisted of 63-tuples, where each feature value was drawn from a unit normal distribution and then normalized so that the dot product of an item with itself equalled one. The first six items were designated as list items. Item 7 was reassigned feature values so that it showed 84% overlap with Item 1; Item 7 would thus be considered to be like an extralist item similar to the cue. Item 9 was reassigned feature values to show 84% overlap with Item 2 and was thus like an item similar to the target. Item 11 was also used as an extralist cue, but one that was unrelated to either the cue or the target and its feature values were not changed.

Items 1 and 2, 3 and 4, and 5 and 6 were convolved, and the results of the three convolutions were added into a 63-tuple representing the composite trace. Then Items 1 (the list cue), 7 (an extralist item similar to the list cue), 9 (an extralist item similar to the target), and 11 (an unrelated extralist cue) were correlated with the trace. The means and standard deviations of the resonance scores of the retrieved items with each of the lexical items were computed from the 100 runs of the simulation.

*Results.* Table 5 provides a summary of the mean resonances and their standard deviations as obtained from Simulation 3. When the list cue was correlated with the trace, the target was strongly evoked. The extralist cue that was similar to the list cue also retrieved a pattern that had a high resonance with the target, though not as high as that evoked by the list cue. This indicates that stimulus generalization should occur. The extralist cue that was similar to the target produced a pattern that was not at all like the target. Instead of retrieving the target, the extralist cue similar to the target retrieved an approximation to the list cue!

Table 5
*Mean Resonance for Each Type of Response in Simulation 3*

| Cue condition | Response type: List target | List cue | Intralist intrusion | Extralist intrusion |
|---|---|---|---|---|
| List cue | | | | |
| *M* | .75 | .00 | .01 | .00 |
| *SD* | .20 | .30 | .21 | .19 |
| Cue synonym | | | | |
| *M* | .52 | .05 | .00 | .00 |
| *SD* | .19 | .26 | .20 | .20 |
| Target synonym | | | | |
| *M* | .02 | .51 | .01 | .00 |
| *SD* | .23 | .19 | .21 | .19 |
| Unrelated extralist cue | | | | |
| *M* | .02 | .01 | −.02 | .00 |
| *SD* | .20 | .22 | .19 | .18 |

### *Experiment 3*

*Method.* Lists of pairs of unrelated words were presented for study. Subjects were then given as cues for the recall of each right-hand target word either (a) the left-hand cue word that had initially been presented, (b) a synonym of the left-hand cue word, or (c) a synonym of the target.

The design of the experiment included two within-subjects factors: type of cue (list cue, cue synonym, or target synonym) and trials (one to three). There were eight cues of each type in a given list. At recall, eight unrelated extralist cues were also presented as lures.

Three lists were constructed from the 144 highest ranked synonym pairs of Whitten et al. (1979). Of the three highest ranking synonym pairs, one word was randomly selected and then randomly assigned either to List 1, 2, or 3. Thus, at this stage in list construction, each of the three lists contained only one word. Of the next three highest ranking synonym pairs in the norms, one word was again randomly selected and then randomly assigned to either of the three lists. At this stage, each list contained two words, which were then mated to form the initial cue–target pair in each list. These procedures were repeated until a total of 24 pairs had been formed for each list. One half of the pairs in each list were then reversed so that a word that was originally assigned as the cue for a given pair became the target for that pair, and the original target became the cue. This equated the synonymy of the extralist cues to the cue (or stimulus) and target (or response) members of the pairs over the list. The order of presentation of the pairs in each list was then randomized.

For each list, three test sets were constructed. Across the three test sets for a given list, every pair was assigned to every cuing condition. Thus, if in Set A of List 1 the first pair was cued with the list cue, in Set B it might be tested with a synonym of the cue, and in Set C with a synonym of the target. For each test set within each list, there were three random orders.

The entire experiment was counterbalanced so that Lists 1, 2, and 3 occurred equally often on Trials 1, 2, and 3 and the three test sets for each list occurred equally often on every trial.

Subjects were given two set-establishing lists before beginning the experiment. One of these lists consisted of similar pairs and the other of unrelated pairs (see Experiment 1 for more details). Subjects were then given booklets in which three study lists, three blank pages, and three test lists were interleaved. The subjects were instructed to study the first list for 4 minutes, keeping in mind that they would soon be asked to recall the capitalized right-hand words and that the uncapitalized left-hand words should be studied in conjunction with the targets because these might assist later recall. Following study, subjects turned to the next (blank) page of their booklets and wrote the alphabet in reverse order for 1 minute. The subjects were then told that they would be given an equal number of cues that were list cues, synonyms of the list cues, synonyms of the targets, or unrelated extralist cues and that their task was to use these cues to recall the right-hand members of the previously presented pairs (i.e., the target items). On Trial 1, subjects were unaware, while they were studying, of the differences among the various types of cues to which they would soon be exposed during the recall test. On Trials 2 and 3, these differences were known to the subjects at the time of list presentation and study.

Eighteen subjects, who were volunteers from the social sciences subject pool at the University of California, Irvine participated in the experiment. The subjects were tested in small groups.

*Results.* The data resulting from Experiment 3 are presented in Table 6. As can be seen from the table, the level of target recall when a synonym of the list cue was given was the same as when a synonym of the target was given. This finding is in contrast to the prediction of the model, which was that a synonym of the target should produce the original list cue rather than the target itself.

There were a number of other discrepancies between the data and the predictions of CHARM. Cue synonyms and target synonyms produced equal numbers of list-cue responses. According to the model, synonyms of the targets should cause recall of the original list cues, and synonyms of the list cues should not. Again, in contrast to the predictions of the model, a rather large number of responses to the unrelated extralist cues was found.

Table 6
*Percentage of Response Types as a Function of Cue Type in Experiment 3*

| Cue type | Response type: List target | List cue | Intrusion error | Omission error |
|---|---|---|---|---|
| List cue | 34 | 0 | 15 | 51 |
| Cue synonym | 14 | 8 | 9 | 69 |
| Target synonym | 16 | 10 | 9 | 65 |
| Lure | — | — | 13[a] | 87 |

[a] Only 5 of these 58 responses were synonyms of the lures.

In an attempt to rescue CHARM, an analysis of the intercorrelations of each of the three cue-type conditions (list cue, cue synonym, and target synonym) was conducted. Each subject contributed three scores to each correlation, one for each of the three trials. There was no correlation between correct recall in the two synonym conditions, $r(52) = -.01$. There was also no correlation between recall to the target synonyms and recall to the list cues, $r(52) = -.08$. There was, however, a positive correlation between recall to the cue synonyms and recall to the list cues, $r(52) = .44$.

Correlations were also computed between correct (target) recall in the two synonym conditions and responses to the unrelated extralist cues (or lures). Responses to these lures give some indication of guessing. A positive correlation was found between recall to the target synonyms and responses made to the lures, $r(52) = .33$. A negative correlation was found between responses to the lures and correct recall to the cue synonyms, $r(52) = -.22$.

In summary, it appeared that correct recall to the cue synonyms was positively correlated with normal cued recall and negatively related to guessing. Correct target recall to the target synonyms was unrelated to normal cued recall and positively correlated with guessing. The regression equations, showing the relative weightings of correct (target) recall in the two synonym conditions in predicting correct recall given the list cues and in predicting guessing to the lures, are given below.

$$z'_{\text{(list cue)}} = +.44z_{\text{(cue synonym)}} - .08z_{\text{(target synonym)}},$$

$$z'_{\text{(responses to lures)}} = -.22z_{\text{(cue synonym)}} + .33z_{\text{(target synonym)}}.$$

These correlations suggest that there might be two different processes or levels of information that combine to give rise to the (unpredicted) equal level of target recall in the cue-synonym and target-synonym conditions as well as the equal level of intrusions of the list cues in these two conditions. To illustrate these ideas, I divided the data into those lists in which subjects made two or more responses to the unrelated extralist cues of lures (high-guessing lists) and those in which no responses were made to the lures (low-guessing lists). The data segregated in this fashion are shown in Figure 3. As can be seen from the figure, the data for the low-guessing lists conformed more closely to the predictions of CHARM than did the data from the experiment as a whole. In particular, in these lists, when a synonym of the target was provided as a retrieval cue, the original list cue tended to be recalled in preference to the target itself. On the high-guessing lists, this trend was dramatically reversed. The list cue was hardly ever recalled to the synonym of the target, whereas the target itself was recalled quite frequently. The holographic association and composite trace in CHARM cannot account for the pattern of recall observed in the high-guessing lists.

The correlations obtained in Experiment 3 suggest that there might be two factors, processes, or levels of information that contribute to the overall pattern of recall found in the experiment. One of these factors seems to produce results consistent with CHARM, whereas the other factor produces results that are at odds with the model. CHARM has nothing to say about this second factor, but the data suggest that it is in some way related to guessing. It seems unlikely that it is just guessing, because the responses that were generated to the lures were rarely synonyms (although the lures also had synonyms). This second factor was nevertheless related to guessing and not to normal recall, where the

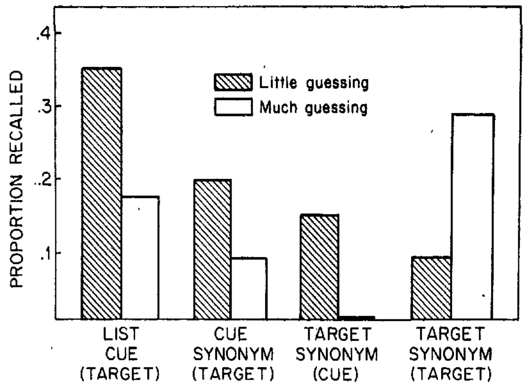


*Figure 3.* Proportion of list cues or target items recalled as a function of cue type and guessing level (Experiment 3). (The abscissa gives the type of cue presented at retrieval—not in parentheses—and the response given—in parentheses.)

stimulus or cue is given and the response or target is recalled. The next experiment was carried out in an attempt to isolate the pattern of recall that is consistent with CHARM.

### *Experiment 4*

*Method.* In the present experiment, as in Experiment 3, subjects studied a list of paired associates and were then given list cues, cue synonyms, and target synonyms as retrieval cues. The question of interest, as before, was what would be recalled to the extralist cues. The most counterintuitive prediction was that extralist target synonyms would result in recall of the list cues rather than the targets.

There were three main differences between the present experiment and the previous one. The first difference was that subjects in the present experiment were not informed that some of the cues would be target synonyms and some would be cue synonyms. (I thought that by not informing subjects, it might be possible to attenuate what was a guessing-related strategy in the previous experiment.) The second difference was the proportion of synonym cues in the test lists. In the present experiment, 16 of the 20 total cues were list cues. Only two cues were critical target synonyms, and only two were cue synonyms. I thought that this change might induce subjects to respond to the extralist synonym cues in the same way that they responded to the list cues. This is, after all, what the simulation did. Finally, no unrelated extralist lures were presented because I felt that their inclusion might disturb the set that the experiment was designed to induce. In summary, Experiment 4 sought to encourage subjects to use only the retrieval operation or strategy that is normally used to recall target items from a list when intralist cues are provided.

The experiment conformed to a 3 × 4 × 2 × 3 mixed design, with cue type (list cue, cue synonym, and target synonym) serving as a within-subjects factor and list, study order, and test order serving as between-subjects variables.

Four lists, each consisting of a unique collection of 20 unrelated cue–target word pairs, were generated via the synonymy norms of Whitten et al. (1979) in a manner similar to that used in Experiment 3. Every subject in the present experiment studied one of these four lists. Two different random presentation or study orders of each list were prepared and assigned randomly to subjects. The cue–target pairs were typed in lowercase and uppercase letters, respectively, on the first page of a three-page booklet. Immediately prior to list presentation, the subjects were read the following instructions:

> In this experiment, you will be given a list of pairs of words. Your task is to remember the right-hand words. However, you should study the right-hand words in conjunction with the left-hand words—trying to form meaningful pairs. You will have three minutes to study the pairs. When the three minutes are up, I will ask you to turn to the last page of your booklet [the second page was blank], where you will find some cues that may help you to recall the right-hand target words. Most of the cues will be the left-hand cue words. You will have four minutes to recall. Do you have any questions?"

A total of 20 cues appeared on the final page of the booklet. The cues were typed in lowercase letters. Of the 20 cues, 14 were original list cues that the subject had seen minutes before. Of the remaining six cues, two were extralist cue synonyms, two were extralist target synonyms, and two were (additional) list cues. These latter six cues provided the experimental data of chief interest and will be referred to as "critical" cues.

On the recall page of the subject's booklet, the top six serial positions were reserved for a corresponding number of "noncritical" list cues, whereas the critical cues were randomly allocated to the bottom 14 positions. On the initial pass through the cues, subjects were asked to read and, if possible, to respond to the cues in their order of presentation from the top of the page to the bottom.

A total of three different test orders were constructed for a given set of six critical cues. These test orders were prepared such that, across all subjects, any given target item that corresponded to a critical cue was probed equally often with (a) a cue synonym, (b) a target synonym, and (c) the original list cue. The critical extralist cue and target synonyms, like all of the words in the lists, were drawn from the synonymy norms of Whitten et al. (1979). Every extralist cue was approximately bidirectionally synonymous with its applicable intralist cue or target item (see Whitten et al., 1979, for details on directional synonymy).

A total of 72 students at the University of Toronto donated a portion of their class time to take part in the experiment. Three subjects were randomly assigned to each of the 24 conditions that were defined by the crossing of the three between-subjects factors (list, study order, and test order). Of the 72 subjects, 69 provided usable data. One subject's data were discarded because of prior knowledge of the hypothesis. The data from two other subjects were also discarded because analysis of their responses to the cue and target synonyms was complicated by the fact that they produced both members of the original pair when probed with the extralist cues.

*Results.* Data obtained via the critical cues are summarized in Table 7 and can be seen to correspond quite closely to the predictions of CHARM. In particular, target synonyms produced recall of the corresponding list cues rather than the targets themselves, whereas the reverse was true of the cue synonyms, $\chi^2(1) = 34.07$, $p < .001$. It has been pointed out to me that if the present experiment was considered in isolation, it is conceivable that subjects used a strategy of not responding with a synonym, even though synonyms were recalled (Slamecka, Note 4). This strategy, though possible, does not seem likely, because it would not account for why the pattern of recall found in the present experiment was also found in the previous experiment, but only when there was evidence that subjects were doing little guessing. When the results of the present experiment are taken in conjunction with those of the previous experiment, it seems more likely that the present experiment largely succeeded in doing what it was designed to do, namely, to attenuate the guessing strategy found in the previous experiment and to induce subjects to use the extralist target synonyms in much the same manner that they use list cues.

The only inconsistency between the data and the predictions of the model was that the extralist synonyms of the target did some-

Table 7
*Proportion of Response Types as a Function of Cue Type in Experiment 4*

| | Response type | | | | |
|---|---|---|---|---|---|
| Cue type | List target | List cue | Intralist intrusion | Extralist intrusion | Omission error |
| List cue | .68 | .00 | .04 | .01 | .27 |
| Cue synonym | .29 | .02 | .03 | .02 | .64 |
| Target synonym | .15 | .27 | .03 | .02 | .53 |

*Note.* Each proportion is based on 138 observations—69 subjects × two presentations of each cue type per subject.

times produce recall of the targets themselves. As Table 7 shows, this happened relatively infrequently. It is possible that despite steps that were taken to dissuade subjects from using a guessing-related strategy like that observed in Experiment 3, such a strategy was not wholly eliminated. On balance, though, the pattern of results found in the present experiment was quite consistent with the predictions of CHARM.

## Extensions of CHARM

In the preceding section, CHARM was used to make several predictions about what would be recalled in a number of simple, single-trial cued-recall situations. These predictions, by and large, were confirmed by the experiments that sought to test them. In the present section, the model will be extended to three more complex areas of investigation—prototype learning, the A–B A–D paradigm, and the Osgood transfer surface—in which the composite–holographic trace might be expected to exert a major influence on performance. These extensions are carried out with some trepidation, because there are, no doubt, factors and processes other than association formation, storage, and retrieval that are important in the data observed in each of the three chosen areas. The aim of the present extensions is exploratory: to determine what effects may be attributable to the holographic association itself and to the composite trace and also what effects do not fall out in a natural way. Therefore, I have specifically avoided embellishing the model by including other constructs such as encoding variability, stimulus or response differentiation, contextual associations, list markers, unlearning, spontaneous recovery, decay, item recognition, response availability, response integration, and the like, even though some of these constructs may be compatible with CHARM and may also be quite appropriate and reasonable. To reiterate, the aim of the present extensions is to illustrate what the current, unembellished version of CHARM does and does not do, without recourse to other constructs.

### *Prototype Learning*

In the discussion of Experiment 1, I noted that the reason the model made the prediction of intrusion errors in the similar-list condition was that the composite–holographic trace produced an item that was a combination of the items in the list. I suggested that in effect, CHARM was producing a list prototype. In this section, an attempt will be made to apply the model in a more formal way to the formation of prototypes.

The experimental paradigm that will be modeled was introduced by Posner and his colleagues (Posner, Goldsmith, & Welton, 1967; Posner & Keele, 1968, 1970) and has been studied by a number of other investigators (Franks & Bransford, 1971; Hintzman & Ludlam, 1980; Homa, Cross, Cornell, Goldman, & Shwartz, 1973; Medin & Schaffer, 1978; Strange, Keeney, Kessel, & Jenkins, 1970). Basically, subjects are given several patterns to study that may be considered to be category exemplars generated from a category prototype. Usually three or four such constructed categories are included in the study list, and subjects are instructed to learn that Exemplar E is an instance of Category C. This paradigm is quite similar to the analysis of natural category learning as outlined by Rosch and Mervis (1975; Mervis & Rosch, 1981) insofar as the exemplars do not specify the category in terms of defining features (i.e., features that must be present) but rather show an overlap or correlation with each other. Subjects, at time of study, are not presented with the prototype itself from which the exemplars were generated. At time of test, subjects are given (a) the old exemplars, (b) the prototype, (c) new exemplars that are formally equivalent to the old exemplars in terms of their similarity to the prototype, and, frequently, (d) new exemplars that are more like the prototype than were the old exemplars. Subjects are asked to sort the materials into different response categories that are provided by the experimenter. In the simulations that follow, the response possibilities will be limited to those given by the experimenter rather than encompassing all of semantic memory.

The simulation will focus on the procedure used by Homa et al. (1973), who varied the number of presented exemplars in each category. They found that when only a few (three) exemplars of a category were given, the old instances were classified best, the classification of the prototypes was considerably worse, and new instances were worse yet. When many category instances were given, however, there was little difference between the correct classification of the old instances and the prototype; the new instances were also classified somewhat better, but correct responding did not improve as much for the new instances as it did for the prototype. It is this approximately Z-shaped interaction (where the prototype gives the crossbar on the Z) that will be simulated below.

### *Simulation 4*

*Method.* I ran the entire simulation three times with three levels of similarity. The high-similarity condition will be described first, and then the changes in the program that produced the medium and low similarity conditions will be detailed.

A lexicon of 12 unrelated items, each consisting of 63 features, was constructed as in the previous simulations. Three of the lexical items (10, 11, and 12) were designated as responses that would be associated with the category exemplars and would be the targets of recall. These three items were left unchanged. Two categories–one with few (one) and one with many (four) exemplars—were constructed as follows: Lexical Item 1 was considered to be a prototype, and Items 2 and 3 were designated to be category exemplars. Items 2 and 3 were reassigned feature values so that 42 of the 63 features were identical to those in Item 1. In order to do this, 42 features were randomly selected and reassigned the same values on Item 2 as those of the corresponding features in Item 1. A different random selection of 42 features were assigned to have identical values in Item 3 as those of the prototype, Item 1.

The large category was constructed in the same way. Forty-two features from each of Items 5 through 9 were randomly selected (with a different random selection for each) to be replaced with the corresponding values from the large-category prototype, that is, Item 4.

In order to construct the memory trace, Item 3 (an exemplar of the small category) was convolved with Response Item 10. Four exemplars of the large category (Items 6, 7, 8, and 9) were convolved with Response Item 11. All five associations (Item 3*Item 10, Item 6*Item 11, Item 7*Item 11, Item 8*item 11, and Item 9*Item 11) were added into a 63-tuple representing the composite trace.

Then Items 1 (the prototype of the small category), 2 (a nonpresented exemplar of the small category), 3 (the presented exemplar of the small category), 4 (the prototype of the large category), 5 (a nonpresented exemplar of the large category), and 6 (a presented exemplar of the large category) were correlated with the composite trace to produce six retrieved patterns. These six patterns were matched to the three response possibilities, namely, Items 10, 11, and 12, by the same procedure used in the previous simulations. The best match of the three was taken to be the response (or classification) that would be given. Means and standard deviations of the resonance scores were calculated for the 100 replications of the simulation.

The entire simulation was run twice more, with only 21 or 11 features being replaced in both categories, to yield the medium- and low-similarity conditions, respectively.

*Results.* The results of the simulation for the three levels of similarity of the exemplars to the prototype are given in Table 8. Results obtained with few presented exemplars show an effect of similarity of the test probe to the encoded item. The presented exemplar was a better cue for the correct response or classification than was the prototype, which in turn was better than the nonpresented exemplar, as Homa et al. (1973) reported. These relations are basically the same as those revealed by Simulation 3, which investigated the efficacy of extralist cues, and may be viewed as manifestations of stimulus generalization. (The prototype was more similar to the presented exemplar than was the nonpresented exemplar by virtue of the way in which the exemplars were generated in the simulation and presumably in the experiment that it was designed to mimic.)

One can also see from the results in Table 8 that when many exemplars were presented, the difference in correct responding between the prototype and the presented exemplar was decreased. In fact, in the high-similarity condition, the prototype was actually more effective at producing the correct response than was a presented exemplar. This is not obvious from the recall data because of a ceiling effect but is readily apparent in the resonance scores. Also, new exemplars were consistently least effective at evoking correct recall. Again, this is not obvious from the percentage-correct scores in the high-similarity condition owing to a ceiling effect but can be deduced from the underlying resonance scores.

On the whole, this simple simulation did a good job of producing the Z-shaped interaction found for correct classification when the type of probe (presented exemplar, prototype, and nonpresented exemplar) was crossed with the number of exemplars presented, as found by Homa et al. (1973).

### *Simulation 5*

*Method.* A second major finding with respect to the prototype learning paradigm is that with a delay in testing, there is a greater loss in the accuracy of classification of the presented exemplars than the prototypes. This

Table 8
*Prototype Abstraction in Simulation 4*

| Probe | Number of exemplars presented: Few | Many |
|---|---|---|
| High-similarity condition | | |
| Presented exemplar | | |
| Mean resonance | .72 (.34) | 1.71 (.41) |
| % recall | 88 | 100 |
| Prototype | | |
| Mean resonance | .40 (.31) | 1.96 (.39) |
| % recall | 62 | 100 |
| Nonpresented exemplar | | |
| Mean resonance | .25 (.34) | 1.29 (.36) |
| % recall | 45 | 100 |
| Medium-similarity condition | | |
| Presented exemplar | | |
| Mean resonance | .76 (.25) | .98 (.34) |
| % recall | 96 | 98 |
| Prototype | | |
| Mean resonance | .24 (.24) | .92 (.33) |
| % recall | 66 | 100 |
| Nonpresented exemplar | | |
| Mean resonance | .05 (.23) | .27 (.30) |
| % recall | 34 | 61 |
| Low-similarity condition | | |
| Presented exemplar | | |
| Mean resonance | .79 (.26) | .84 (.33) |
| % recall | 96 | 100 |
| Prototype | | |
| Mean resonance | .11 (.26) | .49 (.35) |
| % recall | 41 | 79 |
| Nonpresented exemplar | | |
| Mean resonance | .01 (.26) | .03 (.34) |
| % recall | 36 | 36 |

*Note.* The table represents the relation between number of exemplars presented and efficacy of three types of probes for correct classification under three conditions of similarity. Standard deviations are in parentheses.

finding has been reported by Posner and Keele (1970), Strange et al. (1970), and Homa et al. (1973). This difference has usually been interpreted (e.g., Posner & Keele, 1970) as evidence for the differential storage of prototype and exemplar information, although Medin and Schaffer (1978) and Hintzman and Ludlam (1980) have developed models that account for this effect by making use only of exemplar information.

In order to mimic a delay in testing, I assumed in this simulation that several unrelated pairs were convolved and entered into the composite memory trace. This amounts to making the assumption that something, though unrelated to the experiment, was stored in the composite trace in the delay interval of about 1 week. In a preliminary simulation, only two random pairs were added, and no forgetting was systematically produced. In the simulation I will describe here, these two intervening pairs were multiplied by five in order to increase the distortion they introduce into the trace.

This simulation was the same as the medium-similarity condition of Simulation 4, except that four additional unrelated lexical items (designated 13–16) were constructed. These four items were unrelated to each other and were also unrelated to any of the other 12 lexical items. The four extra items were simply random normalized vectors, each composed of 63 features.

The composite trace was constructed in the same manner as in Simulation 4. In addition, though, Item 13 was convolved with Item 14 and multiplied by five, Item 15 was convolved with Item 16 and multiplied by five, and the results of these convolutions were added into the trace.

The three probes in each of the large and small categories were correlated with the trace, as before, and the matching and response-selection operations were identical to those used in Simulation 4.

*Results.* The results of this simulation, summarized in Table 9, may be directly compared to those reproduced in the center panel of Table 8, because the two simulations were identical except for the addition of the intervening unrelated associations. As revealed through comparison of the two tables, the presence of the intervening unrelated associations did give rise to forgetting. What did not occur, however, was differential forgetting of the presented exemplars and the prototypes.

In summary, whereas CHARM readily produces a prototype that becomes more dominant as the number of category exemplars increases, it does not readily yield differential forgetting of presented exemplars and prototypes. Such differential forgetting might be obtainable if we assume, in keeping with Posner and Keele (1970) and Strange et al. (1970), that subjects store two types of information—exemplar information and abstracted prototype information—that are lost from memory at somewhat different rates. CHARM abstracts prototypes quite readily (as in Simulation 4). It is possible that there exists another level of information storage at which individual items or exemplars are temporarily stored in a discrete fashion (cf. Hintzman & Ludlam, 1980; Medin & Schaffer, 1978). This possibility is particularly interesting when taken in conjunction with the results of Experiment 3, which provide evidence for a level of memory storage other than that of the composite–holographic trace.

Table 9
*Prototype Abstraction in Simulation 5*

| Probe | No. of exemplars presented: Few | Many |
|---|---|---|
| Presented exemplar | | |
| Mean resonance | .72 (.66) | 1.05 (.71) |
| % recall | 60 | 76 |
| Prototype | | |
| Mean resonance | .25 (.72) | 1.04 (.74) |
| % recall | 44 | 74 |
| Nonpresented exemplar | | |
| Mean resonance | −.02 (.71) | .38 (.77) |
| % recall | 29 | 44 |

*Note.* The table represents the relation between number of exemplars presented and efficacy of three types of probes for correct classification under the delay (medium similarity) condition. Standard deviations are in parentheses.

Elio and Anderson (1981) have provided other evidence in an abstraction paradigm for two types of information.

### *Non-Association-Specific Interference*

In the immediately preceding simulation, the addition of unrelated pairs of items into the composite memory trace produced a decrease in the level of correct recall. This interference was not specific to the associations that had been stored; that is, the interfering associations were in no way related to the stored target associations.

The buildup of interference that results in the model as the number of unrelated pairs added into the composite trace increases has been demonstrated in several computer simulations reported by Metcalfe and Murdock (1981). In essence, these simulations indicate that, everything else being equal, the level of recall is inversely related to the number of unrelated associations that have been stored in the composite memory trace.

The model's property of non-association-specific interference appears to be directly analogous to McGovern's (1964) finding that subjects who learned a C–D list of paired associates following an A–B list showed poorer recall of the B terms than did subjects who studied only the A–B list. The same property might also be related to Underwood's (1957) classic description of an inverse relation between probability of recall of a given serial list and number of lists previously learned.

### *Association-Specific Interference*

Barnes and Underwood (1959) conducted a classic experiment in which a list of A–B pairs was learned to criterion. Following learning, subjects were given a variable number of trials on an A–D list, in which the stimuli were the same as those of the first list and the responses were different. Subjects were then presented with the stimuli and were asked to recall both the B and D terms. The important finding was that when only a few A–D trials were given, recall of B was much higher than recall of D. As more A–D trials were given, recall of D increased and recall of B correspondingly declined. Subsequently, Martin (1971) showed that in the A–B A–D paradigm with modified modified free recall (MMFR) testing (as was used by Barnes & Underwood, 1959), recall of B is independent of recall of D. Over a variety of methods of pooling the data, the relation P(B)P(D) = P(B and D) was found to obtain (Martin, 1971, 1972, 1981).

It is the trade-off between B and D recall, depending on the number of A–D trials, in combination with the independence of B and D responses, that is modeled in the computer simulation described below.

### *Simulation 6*

*Method.* A lexicon of 12 unrelated items was constructed. There were 15 features in each item, and the items were normalized so that the dot product of an item with itself was equal to one. Fifteen features were used, rather than 63 as in the previous simulations, in order to place recall in an appropriate range. (Metcalfe and Murdock, 1981, have shown that with unrelated items, recall is directly proportional to the number of features in the items.)

Item 1 was designated as A, Item 2 as B, and Item 3 as D. Item 1 was convolved with Item 2 and the result was multiplied by four to correspond to four trials on the A–B association. Item 1 was then convolved with Item 3 and multiplied by either 1, 3, 5, or 7 to correspond to an equivalent number of trials on the A–D association. The results of the four repetitions of A∗B, and of the variable number of repetitions of A∗D, were added into a single 15-tuple representing the composite trace.

In order to recall, Item 1 (or A) was correlated with the trace, producing a single retrieved item. This item was matched to each of the 12 lexical items as in prior simulations, producing a resonanace score for each lexical item. Then the program selected as the recalled items the two lexical items with the highest resonance scores, provided that each score exceeded a lower limit of zero. The selection of two items was made in order to mimic the MMFR testing procedure introduced by Barnes and Underwood (1959).

The simulation was replicated 100 times for each of the four A–D repetition conditions. A count was made of the number of trials, in each repetition condition, on which B, D, and both B and D were recalled. Means and standard deviations of the resonance scores were also computed.

*Results.* Results of the present simulation, summarized in Table 10, coincide closely with the actual results reported by Barnes and Underwood (1959): As the number of A–D trials or repetitions increases, recall of D increases whereas the recall of B decreases.

The reason that the model generates Barnes and Underwood's finding becomes apparent when the particular features of B and D are considered. In all cases in the present simulation, Cue Item A produced a combination of B and D. If a feature happened randomly to have the same sign in B and D, then when the item produced by correlation was matched to all of the possibilities in the lexicon, this feature tended to favor the selection of both B and D. On those features where the signs of the values happened randomly to be different between B and D, the number of repetitions became especially important. Suppose, for instance, that the A–B association was repeated more frequently than was the A–D association. In this instance, the sign of the contrasting features tended to be that of the B item. Alternatively, if the A–D association was repeated more frequently, the sign of the contrasting features in the re-created item tended to favor selection of the D item. To give a specific example, suppose that feature $n$ had a value of $-1$ in Item B and $+1$ in Item D. If B had been presented only once whereas D had been repeated four times, the value of this feature in the reconstructed item would have been $+3$—favoring D. Had B been repeated four times and D presented only once, the value would have been $-3$—favoring B.

Because nothing was done to alter the independence of the Lexical Items B and D, independence, as shown by Martin (1971), emerged as a natural consequence. The contingency tables produced by the simulation for the four conditions of A–D repetition are given in Table 11. Inspection of these tables reveals that in all four repetition conditions, the quantity P(A)P(B) was very close to P(A and B), indicating independence.

### *Osgood Transfer Surface*

In the final simulation that will be presented here, the Osgood (1949) transfer surface, as revised by Martin (1965) in his F surface, will be modeled. This surface provides a summary of the transfer relations that are expected, depending on the similarity of the stimuli of two lists and the responses of two lists. The four edges of the surface may be described as follows: (a) When the stimulus terms in the two lists are unrelated and the response terms of the second list vary from being identical to similar to unrelated to those of the first list, no transfer is expected (i.e., A–B C–B = A–B C–B′ = A–B C–D). (b) When the stimulus terms in the second list vary from being identical to similar to unrelated to those in the first list, and the response terms are identical in the two lists, transfer ranges from maximum positive to

Table 10
*A–B A–D Paradigm With MMFR Testing in Simulation 6*

| Response | Number of repetitions of A–D pair: 1 | 3 | 5 | 7 |
|---|---|---|---|---|
| B | | | | |
| Mean resonance | 2.92 (1.01) | 2.92 (1.32) | 2.93 (1.76) | 2.94 (2.26) |
| % recall | 100 | 88 | 59 | 46 |
| D | | | | |
| Mean resonance | .76 (1.11) | 2.26 (1.19) | 3.76 (1.40) | 5.26 (1.69) |
| % recall | 21 | 62 | 91 | 99 |

*Note.* Standard deviations are in parentheses. MMFR = modified modified free recall.

Table 11
*Contingency Tables Illustrating Response Independence in the A–B A–D Paradigm in Simulation 6*

| Item | B | $\bar{B}$ | Total |
|---|---|---|---|
| One repetition of A–D | | | |
| D | 21 | 0 | 21 |
| $\bar{D}$ | 79 | 0 | 79 |
| Total | 100 | 0 | 100 |
| Three repetitions of A–D | | | |
| D | 53 | 9 | 62 |
| $\bar{D}$ | 35 | 3 | 38 |
| Total | 88 | 12 | 100 |
| Five repetitions of A–D | | | |
| D | 52 | 39 | 91 |
| $\bar{D}$ | 7 | 2 | 9 |
| Total | 59 | 41 | 100 |
| Seven repetitions of A–D | | | |
| D | 45 | 54 | 99 |
| $\bar{D}$ | 1 | 0 | 1 |
| Total | 46 | 54 | 100 |

neutral (i.e., A–B A–B > A–B A′–B > A–B C–B, where A–B A–B gives maximum positive transfer and A–B C–B gives no transfer). (c) When the stimulus terms in the two lists are identical and the response terms of the second list range from being identical to similar to unrelated to those of the first list, transfer ranges from maximum positive to maximum negative (i.e., A–B A–B > A–B A–B′ > A–B A–D, where A–D A–B gives maximum positive transfer and A–B A–B yields maximum negative transfer). (d) When the responses are unrelated in the two lists and the stimulus terms of the second list range from being unrelated to similar to identical to those of the first list, transfer ranges from maximum negative to neutral (i.e., A–B C–D > A–B A′–D > A–B A–D, where A–B A–D gives maximum negative transfer and A–B C–D is neutral).

Martin (1965) has argued that the Osgood surface is attributable to the nature of the (forward) associations. He has further argued that two other surfaces may contribute to observed transfer relations. One of these is a backward-association surface—the mirror image of the forward-association, or Osgood, surface. If CHARM can produce the forward-association surface when the stimulus terms are presented as cues, it will also generate the backward-association surface when the response terms are given as cues, by virtue of the fact that the association in the model is symmetric.

The third surface delineated by Martin (1965) was the R surface. In contrast to either the forward- or backward-association surface, Martin attributed the R surface not to the associations themselves but rather to response learning or availability. Because this article is concerned with the problem of how items are associated, stored, and retrieved, and not with other processes or stages of learning, such as response learning or availability, no attempt will be made to model this third, R surface, even though it is not questioned that the stage of learning reflected by this surface may influence observed transfer relations in particular experimental situations.

A problem that I immediately encountered in attempting to simulate the Osgood surface was that if a single association is considered in isolation, it makes no difference whether that association is repeated. A–B A–B does not produce a higher level of recall than does A–B alone. The reason for this anomaly in the model is that when a single association is doubled, for instance, both the signal and the internal noise are doubled, and exactly the same selection errors will occur when the retrieved item is compared to the lexical items as would occur had the association not been doubled. This anomaly occurs only when the composite trace is assumed to start with values of zero on all of its features. It is because there was another association in the composite trace in Simulation 6 that the problem was not obvious in that simulation. When the composite trace does not start out as a blank slate, the problem does not arise. Accordingly, in the simulation that follows, two extraneous unrelated associations, whose purpose was to provide "background noise," were stored in the composite trace in each condition, as were the associations of main interest.

### *Simulation 7*

*Method.* A lexicon of 14 unrelated 31-feature items was constructed as in the previous simulations. Item 1 was designated as A, Item 3 as B, Item 5 as C, and Item 7 as D. In order to mimic the similarity relation between

A and A', Item 2 was reassigned feature values to be identical with those of Item 1 on 15 randomly chosen features. In a similar manner, though with different randomizations, Item 4 was reassigned values on 15 features to be identical with those of Item 3 and, likewise, Item 6 with Item 5 and Item 8 with Item 7.

Nine separate traces were constructed to represent the following conditions: (a) A–B A–B, (b) A–B A–B', (c) A–B A–D, (d) A–B A'–B, (e) A–B A'–B', (f) A–B A'–D, (g) A–B C–B, (h) A–B C–B', and (i) A–B C–D. These nine conditions define a 3 × 3 design in which three levels of stimulus similarity (identical, similar, and unrelated) are factorially combined with three levels of response similarity (identical, similar, and unrelated) to produce the main conditions represented in the Osgood surface.

In all nine conditions, two unrelated pairs—Item 9*Item 10 and Item 11*Item 12—were convolved and added to the 31-tuple representing the composite memory trace. In addition to these two unrelated pairs, the nine conditions were constructed, respectively, as follows: (a) Item 1*Item 3 + Item 1*Item 3, (b) Item 1*Item 3 + Item 1*Item 4, (c) Item 1*Item 3 + Item 1*Item 7, (d) Item 1*Item 3 + Item 2*Item 3, (e) Item 1*Item 3 + Item 2*Item 4, (f) Item 1*Item 3 + Item 2*Item 7, (g) Item 1*Item 3 + Item 5*Item 3, (h) Item 1*Item 3 + Item 5*Item 4, and (i) Item 1*Item 3 + Item 5*Item 7.

To simplify the programming, Item 1 was correlated with each of the nine traces, and recall of Item 3 was compared across conditions. The correlation of Item 1 with each of the nine traces produced nine retrieved items, each of which was matched, as in prior simulations, to all (14) lexical items. The best match (i.e., the most resonant lexical item) was chosen as the item that would be recalled. The simulation was replicated 100 times, using different random values of the lexical items. Means and standard deviations of the resonance scores for each lexical item were computed.

*Results.* The pattern of correct (Item B) recall produced by the simulation is shown in Table 12 and can be seen to correspond closely to the transfer relations represented by the Osgood surface. More specifically, the simulation produced all of the ordinal relations that define the four edges of the Osgood surface: (a) A–B C–B = A–B C–B' = A–B C–D, (b) A–B A–B > A–B A'–B > A–B C–B, (c) A–B A–B > A–B A–B' > A–B A–D, and (d) A–B C–D > A–B A'–D > A–B A–D. I have rerun this simulation using different randomizations, different numbers of extraneous unrelated pairs, and different numbers of features in the items. In all cases, the results were similar to those obtained in Simulation 7, as illustrated in Table 12.

## Relation to Other Models

In the preceding sections, I showed that CHARM makes a number of predictions about what would be recalled in several cued-recall situations. By and large, these predictions were confirmed experimentally. I also showed that the model can account for a variety of well-known results that have been obtained in associative-learning paradigms. In this section I will compare CHARM to several other associative models.

### *Liepa's Holographic Model*

CHARM was developed from the holographic model proposed by Liepa (Note 1; also see Murdock, 1979) and so, in many respects, it is similar to that model. In particular, the convolution-correlation algebra of Borsellino and Poggio (1973) and the construct of a composite memory trace are common to both models. Liepa's model is more

Table 12
*The Osgood Transfer-Retroaction Surface in Simulation 7*

| | Response similarity | | |
|---|---|---|---|
| Stimulus similarity | Identical (B) | Similar (B') | Unrelated (D) |
| Identical (A) | | | |
| Mean resonance | 1.54 (.39) | 1.15 (.36) | .76 (.34) |
| % recall | 98 | 54 | 36 |
| Similar (A') | | | |
| Mean resonance | 1.14 (.35) | .94 (.32) | .77 (.34) |
| % recall | 93 | 73 | 66 |
| Unrelated (C) | | | |
| Mean resonance | .77 (.34) | .76 (.31) | .76 (.32) |
| % recall | 72 | 73 | 72 |

*Note.* Standard deviations are in parentheses.

general than CHARM insofar as it applies to item, associative, and serial-order information, which are stored in a single trace. It is less general than CHARM, however, in that it deals only with orthogonal vectors—similarity is not represented. One reason for the orthogonality constraint in Liepa's model (and other holographic models as well; e.g., Willshaw, 1981) is that when this constraint is relaxed, the retrieved item may be ambiguous.

In CHARM, the orthogonality assumption is abandoned and similarity is represented in a traditional psychological way. The resulting ambiguity in retrieval produces some of the most psychologically interesting properties of the model: that certain types of errors of recall will be made, that prototypes will be formed, and that interference will occur. To be sure, CHARM is not a maximally efficient memorizing device in terms of producing a veridical copy of the original items or events. However, it is precisely because CHARM transforms and combines events that the model is psychologically interesting.

A second major difference between CHARM and Liepa's model is that the former includes a semantic-memory-pattern recognizer whereas the latter is a single-layer model. This pattern recognizer is essential for CHARM to be applied to experimental situations in which the items retrieved by correlation are either ambiguous or represent combinations of two or more items, and yet, the responses emitted by subjects are discrete. Further, as Martin (1965) has noted and schematized in his R surface, there appears to be a stage of learning that is separable from association formation itself. The results of Experiments 3 and 4, and also those of Simulation 5 concerning prototype retention, point to the importance of a level of information distinct from that represented by the composite trace. Thus, the construct of a semantic-memory-pattern recognizer is formally necessary if CHARM is to be broadly applicable to psychological data and also appears to be independently supported by the data themselves.

### *J. A. Anderson's Neural Model*

CHARM uses the idea of a composite memory trace, and this idea derives, in part, from the neural model given by J. A. Anderson et al. (1977). Both models account for the formation of prototypes by allowing for the superimposition of items in the composite trace.

However, the associations in the two models are different. The association in the J. A. Anderson et al. (1977) model is directional rather than symmetric, as it is in CHARM. As such, the J. A. Anderson et al. (1977) model cannot account for or predict backward intrusions (Experiment 1), differential cue intrusions (Experiment 2), or recall of the cue when an item similar to the target is given as a retrieval cue (Experiments 3 and 4). J. A. Anderson (1977) has noted that there is a problem in associating two responses to a single stimulus in his model. It is not obvious that the feedback mechanism Anderson proposed in conjunction with the association in his model solves this problem. The semantic-memory-pattern recognizer included in CHARM does allow for the removal of ambiguity of responses in situations where two responses are associated with a single cue.

### *Multicomponent or Feature Models*

A number of feature models of recall have been proposed for purely psychological reasons (e.g., Bower, 1972; Medin & Schaffer, 1978; Underwood, 1969). CHARM, too, uses the idea that items are composed of sets of features, making explicit the link between the vector elements in neural models and the features that have been proposed for psychological reasons, as was suggested by Estes (1979). Similarity is represented in CHARM in much the same way as it is in feature models—that is, in terms of feature overlap—and I showed that this variable produces a number of psychologically interesting results, such as the Osgood transfer surface, prototype formation, and intrusions in recall.

CHARM differs from other feature models in that all of the associations are stored in a single composite trace. However, it may not be possible to distinguish psychologically the construct of a composite trace from the idea that associations are stored separately but are combined at time of retrieval. Mathematically, it makes no difference to the end result whether associations are combined at time of storage or at retrieval. Thus, there is a formal similarity between CHARM and models such as those proposed by Medin and Schaffer (1978) and Hintzman and Ludlam (1980) that add the results of a variety of retrieved items at time of retrieval. Similarly, features have usually been coded as positive values, with the absence of a feature coded as zero. In CHARM, features are coded as positive and negative values around zero, which permits representation of common as well as contrasting features (Tversky, 1977). However, it may be that the apparently different views of feature coding will in some sense prove equivalent. Nevertheless, the combination of the construct of a composite trace and the zero-centered feature coding makes it easy to see certain relations among phenomena that might not otherwise be obvious. For instance, interference and the abstraction of prototypes both depend on the addition of features and differ only in the expected resultant values. The same combination of constructs also allows for the observation of certain conceptual differences. For instance, a loss of discriminability in CHARM reflects an increase in the randomness of retrieved items and is not conceptually equivalent to an increase in the similarity of retrieved items (as it is in Medin and Schaffer's, 1978, model of prototype learning).

CHARM uses a symmetric association, whereas the association in other feature models is directional. The association in CHARM allows it to address phenomena that depend on the similarity relations between the stimulus and response terms of a single pair. It also gives rise to the predictions of equal frequencies of stimulus and response intrusions (Experiment 1), cue intrusions (Experiment 2), and the efficacy of extralist items similar to the target to generate the cues (Experiments 3 and 4). These predictions are not made by other feature models.

### *Martin's Encoding-Variability Model*

Martin (1972) proposed that associative recall depends critically on recognition of the functional stimulus, reviving Hoffding's earlier view of recall (see Rock, 1962). This is in marked contrast to CHARM, which does not consider or include item recognition. When convolution and correlation serve as the associative encoding and retrieval operations, it is not necessary that cue or stimulus recognition precede target or response recall (although a pattern-identification process that might be related to recognition is included in CHARM).

Martin has also argued that the construct of variable encoding of stimulus terms is necessary for a person to recall multiple responses to a single stimulus, as in the A–B A–D situation. As Martin (1972) has stated, "If $R_1$ and $R_2$ are two distinct behaviors and if $S_1$ is the presumed stimulus for $R_1$, then any assertion that $S_1$ can also be the stimulus for $R_2$ is a mistake. . . . A learner cannot retain both A–B and C–D where A and C are identically encoded" (p. 71). This "mistake" is precisely what was committed in Simulation 6 that modeled the A–B A–D paradigm. The result of that simulation was not only a good approximation to the trade-off between B and D recall but also showed independence of recall of the two responses. The implication, then, is that it is not necessary to invoke the assumption of encoding variability in order to account for the capacity of a single nominal stimulus to control two or more distinct behaviors or responses. This does not imply, however, that the representations in CHARM are incompatible with the idea of encoding variability. Indeed, CHARM goes beyond the assumption of encoding variability by providing a memory mechanism that may actually change the representation of an item from encoding to retrieval, depending on the nature of the associated item and also the nature of the list.

### *Node/Network Theories*

CHARM is radically different from node/network theories of associative recall (e.g., J. R. Anderson, 1972, 1976; J. R. Anderson & Bower, 1973). To cite just a few obvious differences, network theories do not consider the question of how new associations are formed (see Postman, 1975), which is a matter of prime importance in CHARM. Also, CHARM places no emphasis whatever on item recognition, whereas this process is quite important for recall in network theories. The representation of items as conceptual nodes in network theories differs from the feature representations in CHARM. Further, the idea that retrieval is a search or that it consists of spreading activation is metaphorically far removed from the filterlike operation of correlation proposed in CHARM. And the single-layer semantic memory given by network theories differs from the dual memory system—composite-holographic trace and semantic-memory-pattern recognizer—employed in CHARM.

The basic assumptions, or the approach to the problem of recall, are also different. It seems that one assumption made about the nature of the association in network theories

Table 13
*Predictions and Applications of* CHARM

| Prediction or phenomenon | Simulation | Favorable: Yes | Favorable: No | Evidence |
|---|---|---|---|---|
| Interaction between correct recall and intralist intrusions depending on whether paired-associate list consists entirely of similar or of unrelated items | 1 | X | | Experiment 1 |
| Equal frequency of stimulus- and response-term intralist intrusions | 1 | X | | Experiment 1 |
| Greater proportion of cue intrusions when cue and target items are similar than when they are unrelated | 2 | X | | Experiment 2 |
| Stimulus generalization | 3 | X | | Experiment 4; Gibson, 1941. |
| Extralist item similar to target evokes recall of cue rather than target | 3 | X | | Experiment 4 |
| Responses of a particular pair learned less well in a homogeneous than in a heterogeneous list | — | X | | Tulving, Note 3 |
| Interaction between number of category exemplars presented and cue efficacy of prototype as compared to exemplars | 4 | X | | Homa, Cross, Cornell, Goldman, and Schwartz, 1973 |
| Prototype more resistant to forgetting than presented exemplars | 5 | | X | Posner and Keele, 1970; Strange, Keeney, Kessel, and Jenkins, 1970 |
| Non-association-specific interference: A–B C–D gives worse recall than A–B alone | Metcalfe and Murdock, 1981 | X | | McGovern, 1964 |
| Trade-off between B and D recall in A–B A–D paradigm with MMFR testing | 6 | X | | Barnes and Underwood, 1959 |
| Independence of B and D responses in the A–B A–D paradigm | 6 | X | | Martin, 1971 |
| All of the transfer or retroaction relations given in the Osgood surface | 7 | X | | Osgood, 1949; Martin, 1965 |

*Note.* MMFR = modified modified free recall.

is that if the association between two items A and B has been tagged and the tag is still present at time of retrieval, A will automatically lead to B. When this assumption is made about the nature of associations, the problem then becomes one of inferring, whether, and with what probability, the tagged association exists. An additional problem is that of delineating what is associated with what; for instance, is Item A associated with context, which meaning node of A is associated with B, and so on.

In CHARM, the problem of recall is not formulated in terms of inferring what is associated to what and with what probability. Rather, it was assumed early on that the experimenter had control over what was associated with what. Thus, if an A–B pair was presented by the experimenter, it was allowed that A and B were associated. The explanatory burden in CHARM was placed on how associations are formed, stored, and retrieved.

The approach taken in CHARM of focusing on the microstructure of association formation, storage, and retrieval seems to have several advantages over the "higher level" approach taken in network theories. For instance, in discussing how human associative memory (HAM) accounts for some of the basic findings of interference theory, Anderson and Bower (1973) appeal to context, recognition, a construct roughly equivalent to unlearning, spontaneous recovery, rehearsal differences in some conditions, mediated semantic associations, and both forward and backward associations between stimuli and responses. In contrast, as has been demonstrated, CHARM can account for a considerable number of paired-associate learning results by detailing only how the pairs presented by an experimenter are associated, stored, and retrieved. In addition, CHARM can account for list-context effects, associative-context effects, and the abstraction of prototypes. Further, CHARM makes predictions about errors of recall (see Experiments 1, 2, 3, and 4) that network theories certainly do not predict and may not be able to accommodate. The simplicity with which both well-known and counterintuitive (but experimentally confirmed) findings are produced by CHARM indicates the value of attempting to specify at a "micro" level how associations are formed, stored, and retrieved.

## Conclusion

In this article, a composite holographic associative recall model was proposed as a solution to the problem of how it is that people associate pairs of items, store the associations in memory, and then later, when given a retrieval cue, use that cue to evoke recall. The model yielded predictions about what people would recall in several cued-recall experiments. It was also applied in an unelaborated form to a variety of well-established phenomena that ostensibly depend on associations. The results of the model are summarized in Table 13. As can be seen from the table, the model seems to be doing something similar to what people do when they associate, store, and retrieve from memory.

CHARM is a highly interactive model of human association formation, storage, and retrieval. Nearly all of the predictions and applications depend on the interactive nature of the holographic association and of the composite trace. The events stored in such a memory combine and interfere with one another so that the output from memory is different from the input to it. Under certain conditions, as when unrelated items are associated, the difference from input to output may be that the retrieved items are noisier than the encoded items. Under other conditions, as when the items are similar to one another, the retrieved items are systematically transformed from their encoded form.

It is an old notion that memory is not just a passive store for holding ideas without changing them but may in fact transform those ideas. The Gestalt psychologists gave a number of demonstrations showing that what is retrieved may differ from what was initially given. It seems intuitive that people do not simply take ideas and passively store them. Rather, people seem all the time to be altering and mentally transforming what was given. The holographic hypothesis of associations may have many implications for the study of human memory. It may turn out that the most important of these is that it may not only increase our understanding of how ideas are associated, stored, and retrieved but also provide some insight into the question of how it is that people are able to take old ideas and transform them into ideas that are new.

## Reference Notes

1. Liepa, P. *Models of content addressable distributed associative memory (CADAM).* Unpublished manuscript, University of Toronto, 1977.
2. Tulving, E. Personal communication, August 1981.
3. Tulving E. *Organization and interference.* Paper presented at the meeting of the American Association for the Advancement of Science, New York, December 1967.
4. Slamecka, N. J. Personal communication, December 2, 1981.

## References

Anderson, J. A. Two models for memory organization using interactive traces. *Mathematical Biosciences,* 1970, *8,* 137–160.

Anderson, J. A. A simple neural network generating an interactive memory. *Mathematical Biosciences,* 1972, *14,* 197–220.

Anderson, J. A. A theory of recognition of items from short memorized lists. *Psychological Review,* 1973, *80,* 417–438.

Anderson, J. A. Neural models with cognitive implications. In D. LaBerge & S. J. Samuels (Eds.), *Basic processes in reading.* Hillsdale, N.J.: Erlbaum, 1977.

Anderson, J. A., & Hinton, G. E. Models of information processing in the brain. In G. E. Hinton & J. A. Anderson (Eds.), *Parallel models of associative memory.* Hillsdale, N.J.: Erlbaum, 1981.

Anderson, J. A., Silverstein, J. W., Ritz, S. A., & Jones, R. S. Distinctive features, categorical perception, and probability learning: Some applications of a neural model. *Psychological Review,* 1977, *84,* 413–451.

Anderson, J. R. FRAN: A simulation model of free recall. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 5). New York: Academic Press, 1972.

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

Anderson, J. R., & Bower, G. H. *Human associative memory.* Washington, D.C.: V. H. Winston, 1973.

Asch, S. E. The doctrinal tyranny of associationism: Or what is wrong with rote learning. In T. R. Dixon & D. L. Horton (Eds.), *Verbal behavior and general behavior theory.* Englewood Cliffs, N.J.: Prentice-Hall, 1968.

Bahrick, H. P. Measurement of memory by prompted recall. *Journal of Experimental Psychology,* 1969, *79,* 213–219.

Barnes, J. M., & Underwood, B. J. "Fate" of first-list associations in transfer theory. *Journal of Experimental Psychology,* 1959, *58,* 97–105.

Bartlett, F. C. *Remembering: A study in experimental and social psychology.* Cambridge, England: Cambridge University Press, 1932.

Batchelder, W. H., & Narens, L. A critical examination of the analysis of dichotomous data. *Philosophy of Science,* 1977, *44,* 113–135.

Bilodeau, E. A. Experimental interference with primary associates and their subsequent recovery with rest. *Journal of Experimental Psychology,* 1967, *73,* 328–332.

Bilodeau, E. A., & Blick, K. A. Courses of misrecall over long-term retention intervals as related to strength of preexperimental habits of word association. *Psychological Reports,* 1965, *16*(3, Pt. 2), 1173–1192.

Borsellino, A., & Poggio, T. Convolution and correlation algebras. *Kybernetik,* 1973, *13,* 113–122.

Bower, G. H. A multicomponent theory of the memory trace. In K. W. Spence & J. T. Spence (Eds.), *The psychology of learning and motivation* (Vol. 1). New York: Academic Press, 1967.

Bower, G. H. Stimulus sampling theory of encoding variability. In A. W. Melton & E. Martin (Eds.), *Coding processes in human memory.* Washington, D.C.: V. H. Winston, 1972.

Cavanagh, P. Holographic and trace strength models of rehearsal effects in the item recognition test. *Memory & Cognition,* 1976, *4,* 186–199.

Ekstrand, B. R. Backward associations. *Psychological Bulletin,* 1966, *65,* 50–64.

Elio, R., & Anderson, J. R. The effects of category generalizations and instance similarity on schema abstraction. *Journal of Experimental Psychology: Human Learning and Memory,* 1981, *7,* 397–417.

Estes, W. K. Statistical theory of distributional phenomena in learning. *Psychological Review,* 1955, *62,* 369–377.

Estes, W. K. An associative basis for coding and organization in memory. In A. W. Melton & E. Martin (Eds.), *Coding processes in human memory.* Washington, D.C.: V. H. Winston, 1972.

Estes, W. K. On the descriptive and explanatory functions of theories of memory. In L.-G. Nilsson (Ed.), *Perspectives on memory research: Essays in honor of Uppsala University's 500th anniversary.* Hillsdale, N.J.: Erlbaum, 1979.

Flexser, A. J., & Tulving, E. Retrieval independence in recognition and recall. *Psychological Review,* 1978, *85,* 153–171.

Franks, J. J., & Bransford, J. D. Abstraction of visual patterns. *Journal of Experimental Psychology,* 1971, *90,* 65–74.

Gabor, D. A new microscopic principle. *Nature,* 1948, *161,* 777–778.

Galton, F. Composite portraits, made by combining those of many different persons into a single resultant figure. *Journal of the Anthropological Institute,* 1879, *8,* 132–144.

Gibson, E. J. Retroactive inhibition as a function of degree of generalization between tasks. *Journal of Experimental Psychology,* 1941, *28,* 93–115.

Heerden, P. J. van. A new optical method for storing and retrieving information. *Applied Optics,* 1963, *2,* 387–392.

Hintzman, D. L., & Ludlam, G. Differential forgetting of prototypes and old instances: Simulation by an exemplar based classification. *Memory & Cognition,* 1980, *8,* 378–382.

Homa, D., Cross, J., Cornell, D., Goldman, D., & Schwartz, S. Prototype abstraction and classification of new instances as a function of number of instances determining the prototype. *Journal of Experimental Psychology,* 1973, *101,* 116–122.

Hutchinson, J. W., & Lockhead, G. R. Similarity as distance: A structural principle for semantic memory. *Journal of Experimental Psychology: Human Learning and Memory*, 1977, *3*, 660–678.

Julesz, B., & Pennington, K. S. Equidistributed information mapping: An analogy to holograms and memory. *Journal of the Optical Society of America*, 1965, *55*, 604.

Kintsch, W. *The representation of meaning in memory.* Hillsdale, N.J.: Erlbaum, 1974.

Kohonen, T. *Associative memory: A system-theoretical approach.* Berlin, Federal Republic of Germany: Springer-Verlag, 1977.

Kohonen, T. *Content-addressable memories.* Berlin, Federal Republic of Germany: Springer-Verlag, 1980.

Lathi, B. P. *Communication systems.* New York: Wiley, 1968.

Light, L. L. Homonyms and synonyms as retrieval cues. *Journal of Experimental Psychology*, 1972, *96*, 255–262.

Loftus, E. F. *Eyewitness testimony.* Cambridge, Mass.: Harvard University Press, 1979.

Longuet-Higgins, H. C. Holographic model of associative recall. *Nature*, 1968, *217*, 104.

Martin, E. Transfer of verbal paired associates. *Psychological Review*, 1965, *72*, 327–343.

Martin, E. Verbal learning theory and independent retrieval phenomena. *Psychological Review*, 1971, *78*, 314–332.

Martin, E. Stimulus encoding in learning and transfer. In A. W. Melton & E. Martin (Eds.), *Coding processes in human memory.* Washington, D.C.: V. H. Winston, 1972.

Martin, E. Simpson's paradox resolved: A reply to Hintzman. *Psychological Review*, 1981, *88*, 372–374.

McGovern, J. B. Extinction of associations in four transfer paradigms. *Psychological Monographs*, 1964, *78*(16, Whole No. 593).

McKoon, G., & Ratcliff, R. Priming in episodic and semantic memory. *Journal of Verbal Learning and Verbal Behavior*, 1979, *18*, 463–480.

Medin, D. L., & Schaffer, M. M. Context theory of classification learning. *Psychological Review*, 1978, *85*, 207–238.

Mervis, C. B., & Rosch, E. Categorization of natural objects. *Annual Review of Psychology*, 1981, *32*, 89–115.

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

Murdock, B. B., Jr. Item and order information in short-term serial memory. *Journal of Experimental Psychology: General*, 1976, *105*, 191–216.

Murdock, B. B., Jr. Convolution and correlation in perception and memory. In L.-G. Nilsson (Ed.), *Perspectives on memory research: Essays in honor of Uppsala University's 500th anniversary.* Hillsdale, N.J.: Erlbaum, 1979.

Murdock, B. B., Jr. A theory for the storage and retrieval of item and associative information. *Psychological Review*, 1982, *89*, 609–626.

Osgood, C. E. The similarity paradox in human learning: A resolution. *Psychological Review*, 1949, *56*, 132–143.

Posner, M. I., Goldsmith, R., & Welton, K. E. Perceived distance and the classification of distorted patterns. *Journal of Experimental Psychology*, 1967, *73*, 28–38.

Posner, M. I., & Keele, S. W. On the genesis of abstract ideas. *Journal of Experimental Psychology*, 1968, *77*, 353–363.

Posner, M. I., & Keele, S. W. Retention of abstract ideas. *Journal of Experimental Psychology*, 1970, *83*, 304–308.

Postman, L. Verbal learning and memory. *Annual Review of Psychology*, 1975, *26*, 291–335.

Pribram, K. H. *Language of the brain: Experimental paradoxes and principles in neuropsychology.* Englewood Cliffs, N.J.: Prentice-Hall, 1971.

Pribram, K. H., Nuwer, M., & Baron, R. J. The holographic hypothesis of memory structure in brain function and perception. In D. H. Krantz, R. C. Atkinson, R. D. Luce, & P. Suppes (Eds.), *Contemporary developments in mathematical psychology* (Vol. 2). San Francisco: Freeman, 1974.

Ratcliff, R. A theory of memory retrieval. *Psychological Review*, 1978, *85*, 59–108.

Reed, S. K. *Psychological processes in pattern recognition.* New York: Academic Press, 1973.

Rock, I. A neglected aspect of the problem of recall: The Hoffding function. In J. Scher (Ed.), *Theories of the mind.* New York: Free Press, 1962.

Rosch, E., & Mervis, C. B. Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 1975, *7*, 573–603.

Santa, J. L., & Lamwers, L. L. Encoding specificity: Fact or artifact. *Journal of Verbal Learning and Verbal Behavior*, 1974, *13*, 412–423.

Smith, E. E., Shoben, E. S., & Rips, L. J. Structure and processes in semantic memory: A featural model for semantic decisions. *Psychological Review*, 1974, *81*, 214–241.

Strange, W., Keeney, T., Kessel, F. S., & Jenkins, J. J. Abstraction over time of prototypes from distortions of random dot patterns. *Journal of Experimental Psychology*, 1970, *83*, 508–510.

Stroke, G. W. *An introduction to coherent optics and holography.* New York: Academic Press, 1969.

Thomson, D. M., & Tulving, E. Associative encoding and retrieval: Weak and strong cues. *Journal of Experimental Psychology*, 1970, *86*, 255–262.

Tulving, E. Episodic and semantic memory. In E. Tulving & W. Donaldson (Eds.), *Organization of memory.* New York: Academic Press, 1972.

Tulving, E. Recall and recognition of semantically encoded words. *Journal of Experimental Psychology*, 1974, *102*, 778–787.

Tversky, A. Features of similarity. *Psychological Review*, 1977, *84*, 327–352.

Underwood, B. J. Interference and forgetting. *Psychological Review*, 1957, *64*, 49–60.

Underwood, B. J. Attributes of memory. *Psychological Review*, 1969, *76*, 559–573,

Underwood, B. J., Runquist, W. N., & Schulz, R. W. Response learning in paired-associate lists as a function of intralist similarity. *Journal of Experimental Psychology*, 1959, *58*, 70–78.

Underwood, B. J., & Schulz, R. W. *Meaningfulness and verbal learning.* New York: Lippincott, 1960.

Westlake, P. R. The possibilities of neural holographic processes within the brain. *Kybernetik*, 1970, *7*, 129–153.

Wickens, D. D. Characteristics of word encoding. In A. W. Melton & E. Martin (Eds.), *Coding processes in human memory.* Washington, D.C.: V. H. Winston, 1972.

Willshaw, D. Holography, associative memory, and inductive generalization. In G. E. Hinton & J. A. Anderson (Eds.), *Parallel models of associative memory.* Hillsdale, N.J.: Erlbaum, 1981.

Whitten, W. B., II, Suter, W. N., & Frank, M. L. Bidirectional synonym ratings of 464 noun pairs. *Journal of Verbal Learning and Verbal Behavior,* 1979, *18,* 109–127.

## Appendix

## Computational Illustrations of Convolution and Correlation

The convolution of two items $\mathbf{A} = (\cdots, a_{-1}, a_0, a_1, \cdots)$ and $\mathbf{B} = (\cdots, b_{-1}, b_0, b_1, \cdots)$ is conveniently computed by constructing a matrix as shown in Figure A1. The trace **T** that results from summing across the positive diagonals is given by $\mathbf{T} = (\cdots, t_{-2}, t_{-1}, t_0, t_1, t_2, \cdots) = (\cdots, a_{-1}, b_{-1}, a_0b_{-1} + a_{-1}b_0, a_1b_{-1} + a_0b_0 + a_{-1}b_1, a_1b_0 + a_0b_1, a_1b_1)$, which conforms to the definition of convolution given in the text in the Association Formation section.

In order to correlate one of the items, say **B**, with the trace **T**, a new matrix is formed as shown in Figure A2. The retrieved item is $\mathbf{R} = (r_{-1}, r_0, r_1)$. The central $n$ features of this retrieved item provide an approximation to **A**, provided that the initially convolved items are unrelated.

To illustrate the nonnoise components in the resultant item **R**, it is necessary to expand the joint convolution–correlation matrix. In the expanded matrix, given in Figure A3, the components that produce the retrieved item **A** have been underlined. Regardless of the dimension of the vectors (i.e., the number of features in the items), if the items are unrelated, there is one signal component in each cell of the joint convolution–correlation matrix for the central $n$ features of **R**. It can be seen that

$$\begin{aligned} r_{-1} &= b_1a_{-1}b_1 + b_0a_{-1}b_0 + b_{-1}a_{-1}b_{-1} + b_1a_1b_{-1} \\ &\quad + b_1a_0b_0 + b_0a_0b_{-1} \\ &= a_{-1}(b_1^2 + b_0^2 + b_{-1}^2) + b_1a_1b_{-1} \\ &\quad + b_1a_0b_0 + b_0a_0b_{-1} \\ &= a_{-1}(\mathbf{B}\cdot\mathbf{B}) + b_1a_1b_{-1} + b_1a_0b_0 + b_0a_0b_{-1} \\ &= a_{-1} + \text{noise}. \end{aligned}$$

Similar computations can be carried out for all of the central features of **R** and one will see that the $r_0 = a_0 +$ noise and $r_1 = a_1 +$ noise. Thus, $\mathbf{B}\#(\mathbf{A}*\mathbf{B}) = \mathbf{A} +$ noise.

Had **A** been correlated with $\mathbf{T} = \mathbf{A}*\mathbf{B}$, the retrieved item would be an approximation to **B**. In that case, $r_{-1}$ is as follows:

$$\begin{aligned} r_{-1} &= a_1a_1b_{-1} + a_1a_0b_0 + a_1a_{-1}b_1 + a_0a_0b_{-1} \\ &\quad + a_0a_{-1}b_0 + a_{-1}a_{-1}b_{-1} \\ &= b_{-1}(a_1^2 + a_0^2 + a_{-1}^2) + a_1a_0b_0 \\ &\quad + a_1a_{-1}b_1 + a_0a_{-1}b_0 \\ &= b_{-1} + \text{noise}. \end{aligned}$$

Similar computations can be carried out for $r_0$ and $r_1$, and they too produce the corresponding fea-

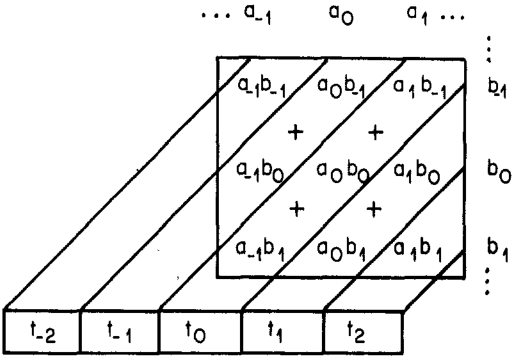


*Figure A1.* A computational illustration of the convolution of two items $\mathbf{A} = (a_{-1}, a_0, a_1)$ and $\mathbf{B} = (b_{-1}, b_0, b_1)$ to produce an association $\mathbf{T} = (t_{-2}, t_{-1}, t_0, t_1, t_2)$.

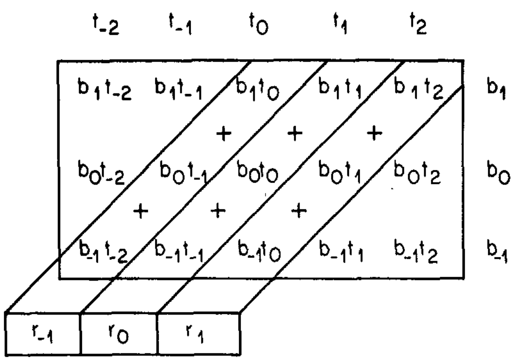


*Figure A2.* A computational illustration of the correlation of an item $\mathbf{B} = (b_1, b_0, b_1)$ with the association given in Figure A1 to produce a retrieved item $\mathbf{R} = (r_{-1}, r_0, r_1)$.

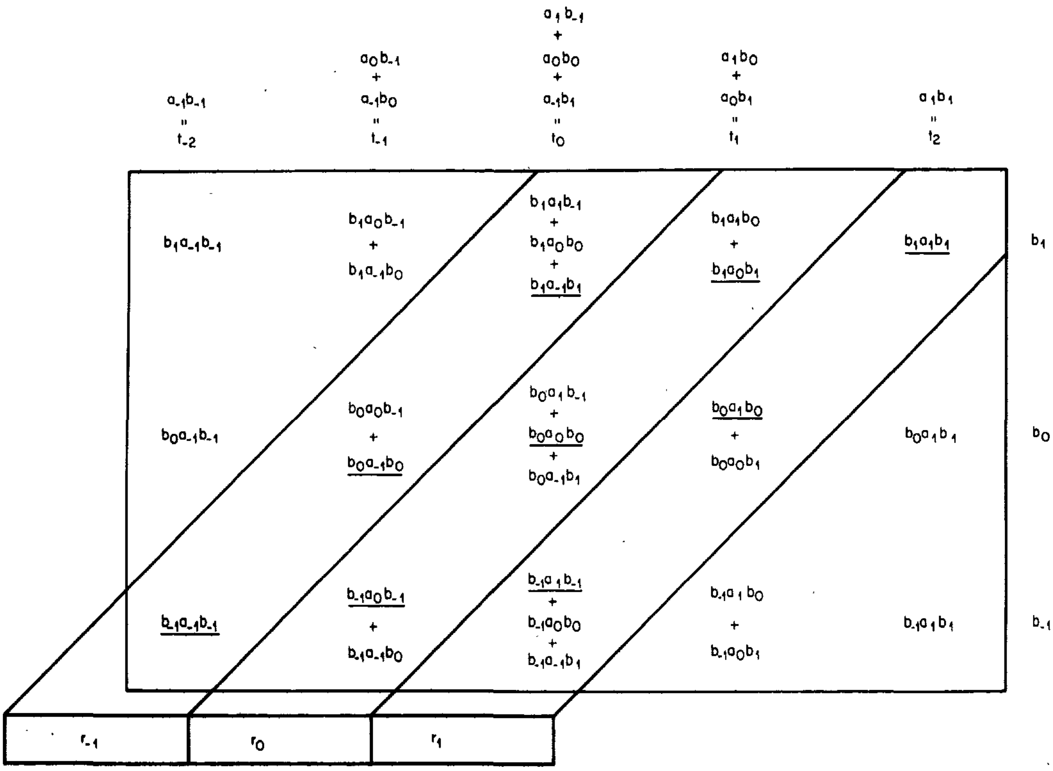


*Figure A3.* An expanded version of the joint convolution–correlation matrix that was given in Figure A2. (The underlined components produce the retrieved item **A** when the initially convolved items **A** and **B** are unrelated.)

tures of **B**, namely, $b_0$ and $b_1$. Thus, **A#(A*B)** = **B** + noise.

Suppose **A** and **B** are convolved, as diagrammed in Figure A1, and then an item similar to **B** is correlated with the result. A "weak" approximation to **A** is retrieved. For instance, suppose that the retrieval cue $\mathbf{B}' = (b_{-1}, b_0, x_1)$ so that two features are identical to the corresponding features in **B** and one feature has a value independent of the value of the corresponding feature in **B**. In this case, the bottom two rows of the joint expanded convolution–correlation matrix will be identical to those given in Figure A3. However, no nonnoise components will be systematically produced in the top row, that is, for the $x_1$ feature. Thus, instead of producing **A** + noise, **B′** will retrieve ($n$ common/$n$)**A** + noise. It may be noted that all of the features of **A** are still retrieved, but with a decreased strength.

If **A** = **B**, and **B** is correlated with the resultant trace,

$$r_{-1} = a_{-1}(\mathbf{B}\cdot\mathbf{B}) + b_1a_1b_{-1} + b_1a_0b_0 + b_0a_0b_{-1}$$

$$= a_{-1} + b_{-1}(a_1^2 + a_0^2) + b_1a_0b_0\,.$$

The term $(a_1^2 + a_0^2)$ results in a positive value but is not quite equal to $\mathbf{A}\cdot\mathbf{A}$. In particular, there is one missing component: in this case $a_{-1}^2$. The missing component stems from the cases, in the original convolution matrix, where the a and b subscripts were the same (i.e., from the negative diagonal of the original convolution matrix). There is thus exactly one missing component from the dot product $\mathbf{A}\cdot\mathbf{A}$ in each reconstructed feature. However, as the number of features in the convolved items becomes large, the difference between $(n - 1/n)$ and 1 becomes vanishingly small, and so

$$r_{-1} = a_{-1} + b_{-1} + \text{noise}$$

$$= 2a_{-1} + \text{noise}.$$

A similar result may be computed in a like manner for all features.

Received October 19, 1981
Revision received April 5, 1982 ■