Psychological Review

VOLUME 92 NUMBER 1 JANUARY 1985

Levels of Processing, Encoding Specificity, Elaboration, and CHARM

Janet Metcalfe Eich
University of British Columbia, Vancouver, British Columbia, Canada

A model of cued recall called CHARM (composite holographic associative recall model) is applied to several findings that have been investigated within the depth-of-processing framework. It is shown that, given some straightforward, empirically testable assumptions about the representations of the to-be-remembered items themselves, CHARM can account for the main effect of depth of processing, the problem of the negatives, encoding-specificity interactions, and effects—both facilitative and inhibitory—of elaboration. The CHARM model is extended to encompass some depth-of-processing effects found in recognition memory.

One of the most interesting characteristics of human cognition is that when people process information, that processing is such that the information may be changed or transformed. For instance, the verbal context in which a particular item is given is highly influential in determining the interpretation of that item. If a word *cat* is presented in the context of *lions,* the feline qualities of the cat are emphasized; whereas, if the same word *cat* is presented in the context of *dogs,* the domesticated animal characteristics of the cat are liable to be more salient. This result occurs even though the stimulus may be the same in both cases, and so seems to be attributable to the way a person processes and combines information. Although a large number of experiments have shown changes due to context and it is generally agreed that such biasing effects are fundamental to human cognition, our understanding of the processes underlying such interactive effects is far from complete.

The most interesting property of the memory model called CHARM (composite holographic associative recall model; Eich, 1982; Metcalfe & Murdock, 1981; and see also Borsellino & Poggio, 1973; Cavanagh, 1976; Gabor, 1968; Julesz & Pennington, 1965; Liepa, 1977; Longuet-Higgins, 1968; Murdock, 1982, 1983; van Heerden, 1963; Willshaw, 1981, for related models) is that the operations for association formation, storage, and retrieval in this model cause changes in the representation of items from input to memory to output from memory. These changes are sometimes just degradations of the initial input. Under some more interesting conditions, though, the model causes systematic biasing. This biasing may affect the interpretation of items as well as the level of recall.

In the CHARM model, items are characterized as patterns of features. These patterns

This research was facilitated by a Natural Sciences and Engineering Research Council (NSERC) postdoctoral fellowship, NSERC Grant A0505, and a University of California, Los Angeles, intramural computing grant to the author.

I thank James Anderson, Robert Bjork, Patrick Cavanagh, Fergus Craik, Eric Eich, Ronald Fisher, Thomas Nelson, Anne Treisman, Endel Tulving, Thomas Wickens and four anonymous reviewers for their critiques and assistance. I thank Morris Moscovitch for providing data and materials relevant to Simulation 1, and Ronald Fisher for providing the materials as well as the design for Experiment 1.

Requests for reprints should be sent to Janet Metcalfe Eich, Department of Psychology, University of British Columbia, Vancouver, British Columbia, Canada V6T 1Y7.

may vary in their similarity to one another. The representations of items are associated by means of the operation of convolution and stored in a composite memory trace in which other associations are superimposed as well. In order to retrieve, the representation of a cue item is correlated with the composite trace, resulting in a noisy and sometimes distorted item. This retrieved item may be identified as a word—if the task is verbal—by being matched to every item in a semantic lexicon. Figure 1 gives a schematic illustration of the CHARM model. The operations in the boxes in the figure (i.e., convolution, addition, and correlation) are responsible for changes in the nature of the items that depend on context. These operations are discussed in more detail shortly.

In this article I explore some of the memorial consequences of the changes in representation—depending on context—that are predicted by the CHARM model. A number of experiments that show the memorial consequences of varying the context in which items are presented have been conducted within the levels-of-processing framework (Cermak & Craik, 1979; Craik & Lockhart, 1972). Within this framework, the context in which an item is presented is often manipulated by giving subjects different orienting tasks. For instance, a subject might be asked whether a target word rhymed with another word, or alternatively, whether a word was in a particular category. Such differences in context result in pronounced effects on the level of recall of targets. Another program of research (Thomson & Tulving, 1970; Tulving & Osler, 1968; Tulving & Thomson, 1973) showed that it is not only the context at time of study that is important, but also the context at time of test. For instance, if a target item CAT was encoded in the context of *lions,* but is later probed with the cue *dog,* recall is much poorer than if the initial context item

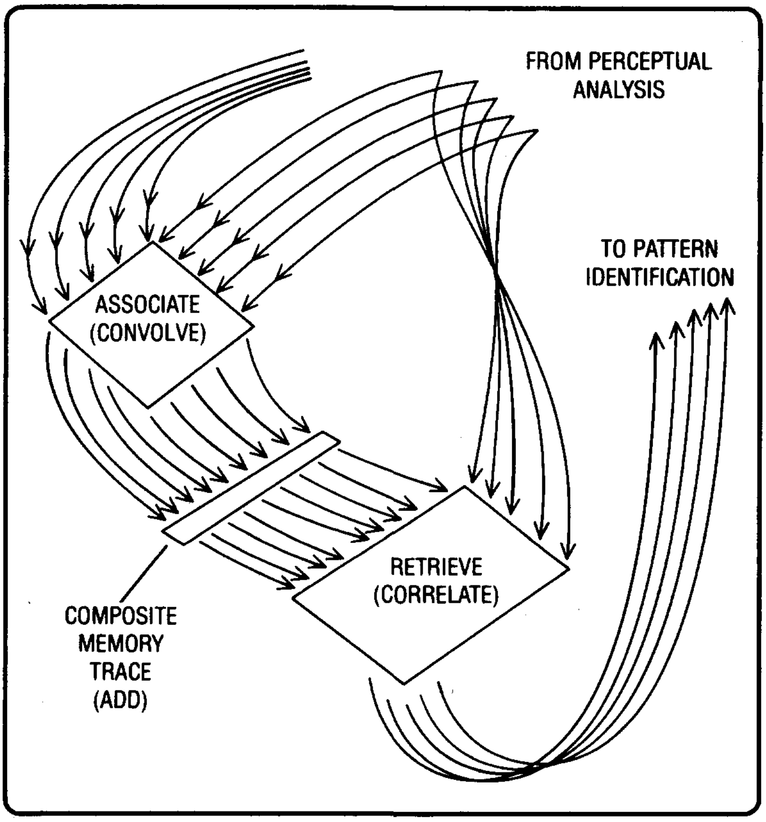

Figure 1. A diagrammatic overview of CHARM.

is used. The nature of the elaborating context in which an item is presented also influences the recall of the item. Sometimes elaboration results in enhanced recall, and sometimes elaboration results in impaired recall. I propose that the memory processes in the CHARM model and the resultant changes in representation are sufficient to account for many of the experimental results of levels of processing, encoding specificity, and elaboration.

Although I propose that many of the levels-of-processing, test context, and elaboration effects are produced by the associative, storage, and retrieval mechanisms in CHARM, the model was not originally designed with these effects in mind. Rather the model was intended to address the general question of how it is that people associate ideas, store those associations in memory, and then later, when given one idea, are able to retrieve another from memory. I have applied the model to a wide range of human learning, memory, and classification phenomena in the cued recall, or paired-associate paradigm (see Eich, 1982), as well as to rehearsal and free recall of unrelated items (Metcalfe & Murdock, 1981). In free recall, many effects appear to depend on factors other than the mechanisms of association formation, storage, and retrieval. For example, the availability of retrieval cues due to the subject's state, and to primary (short-term) memory were important in the holographic free-recall model, in a manner similar to other free-recall models (J. R. Anderson, 1972; Atkinson & Shiffrin, 1968; Gillund & Shiffrin, 1984; Raiijmakers & Shiffrin, 1981). Although of considerable interest and importance in their own right, these factors are nevertheless extraneous to the transformational properties of composite holographic models that are of main interest here, so I mainly consider cued recall. Later in the article, I also sketch out how the main effect of depth of processing could be obtained in recognition, within the CHARM model.

The original levels-of-processing hypothesis stated that the memorability of a particular target item depends on the depth of perceptual analysis performed on that item (Craik & Lockhart, 1972). Items that were analyzed to a semantic level (deeply) were assumed to be more memorable than were items that were analyzed phonemically or orthographically (shallowly). Many experiments have confirmed that the manipulation of an orienting task produces a main effect on the goodness of recall or recognition. Of particular interest in this article are experiments that have produced not only a main effect of orienting task (i.e., level of processing), but also qualifying interactions. The qualifying interactions have given rise to a number of constructs such as elaboration (J. R. Anderson & Reder, 1979; Battig, 1979; Craik & Tulving, 1975; Moscovitch & Craik, 1976; Ross, 1981), cue-trace compatibility (Tulving, 1979; and see also Bransford, Franks, Morris, & Stein, 1979), and distinctiveness (Craik, 1979; Craik & Jacoby, 1979; Eysenck, 1979; Jacoby & Craik, 1979; Klein & Saltz, 1976; Lockhart, Craik, & Jacoby, 1976; Nelson, 1979) that have been evoked instead of or in addition to the original depth hypothesis. However, as Jacoby and Dallas (1981) note, "Although there are currently a large number of experiments showing effects of manipulating an orienting task, there is no generally accepted framework that incorporates the results of those experiments" (p. 309).

In none of the situations that are modeled using computer simulations of CHARM, do I assume a priori that the initial encoding of an item differs from one condition to another. To be more specific, I do not assume that perceptual analysis produces different patterns of features to represent the same word, CAT, for instance, before association formation and retrieval, depending on the experimental condition. Treisman (1979) reviewed the perceptual literature that indicates that the relation between perceptual operations and memory performance may not be as straightforward as was originally supposed. Jacoby and Dallas (1981) argued that if the levels manipulation has its effect by causing different perceptual analyses—for instance, that a shallow task results in the encoding of only a few letters or phonemes, whereas a deep task results in the encoding of the entire word—an effect of level of processing should be found in a perceptual identification task in which the subject must name a tachistoscopically presented word. In contrast to this prediction, they showed that perceptual identification of words that had been presented in a levels task was not differentially affected by the levels manipulation. A sizable levels effect on memory performance, however, was found in the same experiment. Presumably, all of the items, regardless of the levels manipulation, are perceptually analyzed as meaningful, and the memory effects have some other cause.

Because of the interactive properties of the CHARM model, differences in the retrieved patterns of features, that is, in the characteristics of the output from episodic memory, are produced depending on the characteristics of the items that are associated and used as retrieval cues. These differences are produced by means of the operations that are proposed to be the operations that people use to associate ideas, store those associations in memory, and then later to retrieve them.

Description of CHARM

In this section, a brief encapsulation and some illustrative examples of CHARM are given. This model has been described previously (Eich, 1982). Items in the model are represented as ordered sets of features (Tversky, 1977; D. Wickens, 1972). The features are coded as numerical values that may be either positive or negative, but within a particular item are assumed to have values that are independent of one another and to have an expected (in the statistical sense) value of zero. It is assumed that items consist of many such features, and if one wished to, one could consider them as neural units (Estes, 1979). J. A. Anderson and Hinton (1981) suggested that the numerical values in vector representations such as these may correspond to neural firing frequency, with zero being the background frequency. We can also think about the features in more cognitive terms, with a particular feature (or dimension or subset of dimensions) representing a quality. A zero value would then indicate the modal value of that quality. The representation of an item consists of not just one feature, but of a configuration, pattern, or vector of features. Different items have different patterns of numerical values over the set of possible features or dimensions.

Unrelated items are represented such that the feature values in one item are statistically independent of the feature values in other unrelated items. Items that are similar to one another are represented such that the feature values are not independent; that is, there is some feature overlap, beyond independence, among similar items. Figure 2 gives an illustration of feature profiles of four items, **A**, **B**, **C**, and **D**. In the figure, Items **A** and **B** are unrelated; the values on each feature were randomly drawn from a normal distribution centered on zero. I have introduced some similarity between Items **B** and **C**. As can be seen from the figure, Features -10 to -1 have identical values in Items **B** and **C**. Items **C** and **D** are also similar to one another because Features 0 to 10 have identical values. Similarity need not be coded as identical values but, because it will simplify the exposition, I use this coding in the matrices that follow.

The dot product, which is a measure of the similarity between two items, **A** and **B**, is given by

$$\mathbf{A} \cdot \mathbf{B} = \sum_{i=-(n-1)/2}^{(n-1)/2} a_i b_i, \tag{1}$$

where Item **A** is coded as

$$(a_{-(n-1)/2}, \ldots, a_{-2}, a_{-1}, a_0, a_1, a_2, \ldots, a_{(n-1)/2}),$$

B is coded as

$$(b_{-(n-1)/2}, \ldots, b_{-2}, b_{-1}, b_0, b_1, b_2, \ldots, b_{(n-1)/2}),$$

and n is the number of features in the items. If the two items are unrelated, that is, the feature values are independent, the expected value of the product of any two features is zero, and the expected dot product between the unrelated items is also zero. However, the dot product between an item and itself is $(\cdots + a_{-2}^2 + a_{-1}^2 + a_0^2 + a_1^2 + a_2^2 + \cdots)$, which is positive and is set to be one. Consider now, the case where two items are similar to one another. For instance, suppose that $\mathbf{C} = (\cdots, c_{-2}, c_{-1}, c_0, c_1, c_2, \cdots)$, and $\mathbf{D} = (\cdots, d_{-2}, d_{-1}, c_0, c_1, c_2, \cdots)$, the dot product is $(\cdots + c_{-2}d_{-2} + c_{-1}d_{-1} + c_0^2 + c_1^2 + c_2^2 + \cdots)$ with expected values of $(\cdots + 0 + 0 + 1/n + 1/n + 1/n + \cdots)$ or (the number of common features/the total number of features). We could also denote the same two items, **C** and **D**, as $C = (\cdots, c_{-2}, c_{-1},$

$d_0, d_1, d_2, \cdots)$ and $\mathbf{D} = (\cdots, d_{-2}, d_{-1}, d_0, d_1, d_2, \cdots)$.

Convolution

When people associate two items (e.g., the words *ape* and *book*) it is assumed that the ordered sets of features representing those items are combined by the operation of convolution. Convolution, denoted $*$, is defined as

$$(\mathbf{A} * \mathbf{B})_m = \sum_{(i,j)\in S(m)} a_i b_j, \qquad (2)$$

where $S(m) = \left\{(i, j) \middle| -\frac{n-1}{2} \leqslant i, j \leqslant \frac{n-1}{2}, \text{ and } i + j = m\right\}$. The association of two items by means of the operation of convolution is conceptually distinct from the similarity of the items. As an example, two objects (e.g., telephones or computers) may be quite similar to one another without interacting or being connected or associated in any way. The result of association by convolution is an interactive new entity (vector) that does not, in any obvious way, correspond to the items that were so combined.

Figure 3 gives the vector that results from the convolution of Items **A** and **B** in Figure 2. The dimensions in the vector resulting from convolution do not correspond to the features in the item vectors. For instance, suppose that the central (a_0) feature in the items represents some quality such as size

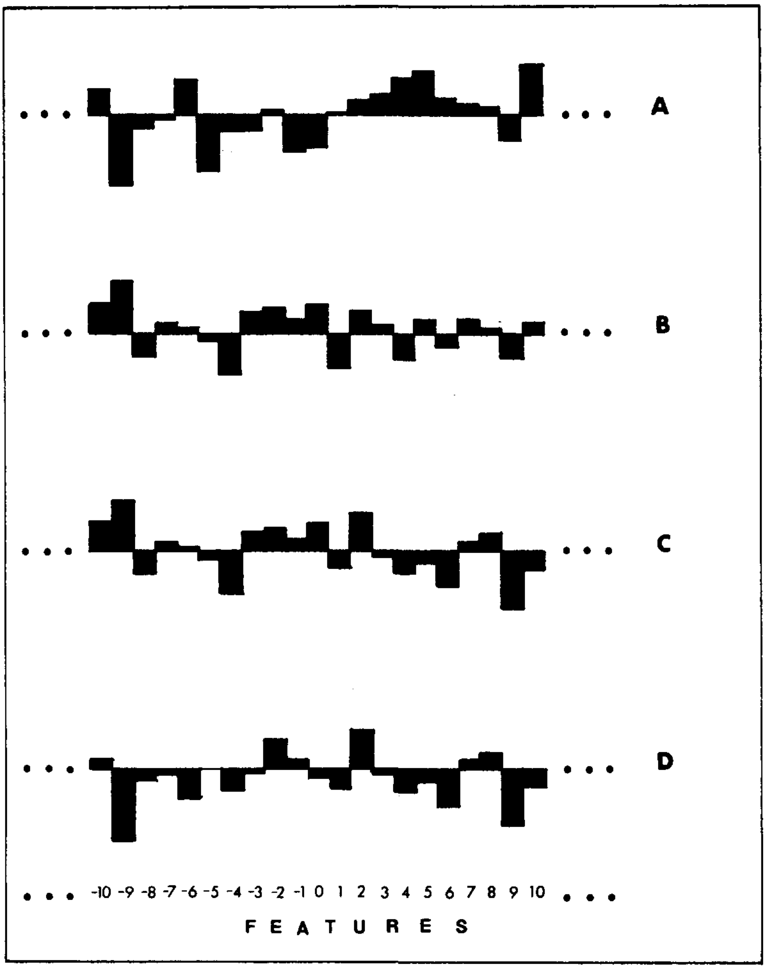

Figure 2. Feature profiles of four items that vary in their similarity to one another.

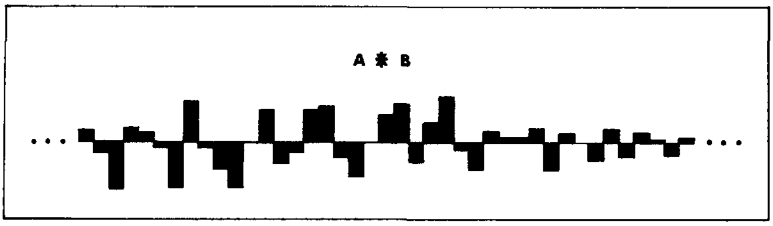

Figure 3. The feature profile that results from the convolution of Items A and B in Figure 2.

(with the numerical values being the sizes of the particular items that are represented). The central feature of the convolution vector is *not* size, but rather is a combination of all of the features in the items that were convolved.

Figure 4 gives an illustration, in matrix form, of the convolution of two items, each coded on five features. The result (t_{-4}, t_{-3}, t_{-2}, t_{-1}, t_0, t_1, t_2, t_3, t_4) is a single vector of dimensionality, $2n - 1$. For the sake of illustration, suppose that the items were two words, Bill and Bob, each coded on features (aggressiveness, first phoneme of name, hair color, age, attractiveness). Both of the items would have a value (not necessarily the same) on each of these features (and on others because it is assumed that there are a large number of features representing each item). The components of the convolution of the two 5-feature items are shown in Figure 4. Substituting into the matrix, it can be seen that the dimensions (t_{-4}, t_{-3}, . . .) in the convolution vector are not the same as are the dimensions in the original items. For instance, the central dimension in the convolution vector is Bill's value on attractiveness × Bob's value on aggressiveness + Bill's value on first phoneme × Bob's value on age + Bill's value on hair color × Bob's value on hair color + Bill's value on age × Bob's value on first phoneme + Bill's value on attractiveness × Bob's value on aggressiveness. The association formed by convolution is complex—an interaction of the original items.

If we consider the patterns of features at the level of the items to be meaningful (and to be consciously interpretable), the pattern

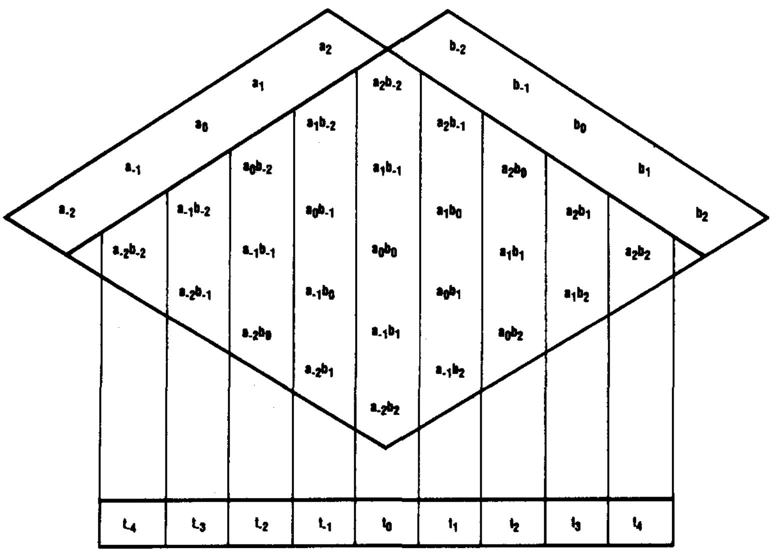

Figure 4. A matrix illustrating the convolution of two items, A = (a_{-2}, a_{-1}, a_0, a_1, a_2), and B = (b_{-2}, b_{-1}, b_0, b_1, b_2).

that results from convolution is not meaningful to the person and is stored in memory. The problem then, of course, becomes one of deconvolving the memory trace so that something meaningful that can be consciously interpreted is recovered. To convert the nonmeaningful and nonconscious memory trace into a form of representation that is meaningful, retrieval occurs. The retrieval operation, that deconvolves the trace, is discussed shortly. The form of the trace, as something not directly available to consciousness, conforms closely to what we generally mean by memory—that which is not conscious. The view of the person that is implied is quite an active one because something (convolution) must be done to associate items, and a different operation (correlation) must be performed to convert the nonmeaningful memory trace into a remembrance that is meaningful and consciously interpretable.

Addition

The results of convolution are added into a composite trace. There is but one associative trace, which is the sum of the various results of association formation. The composite trace is analogous to a photograph that has been exposed many times, except that each "exposure" is itself a complex combination (via convolution) of two items. The trace may be denoted **T**, where

$$\mathbf{T} = \mathbf{A} * \mathbf{B} + \mathbf{C} * \mathbf{D} + \cdots, \tag{3}$$

and **A**, **B**, **C**, and **D** are list items that have been convolved. The idea of a composite trace in which representations are superimposed is used in a number of models (J. A. Anderson, 1973, 1977; J. A. Anderson, Silverstein, Ritz, & Jones, 1977; Cavanagh, 1976; Kohonen, Oja, & Lehtio, 1981; Metcalfe & Murdock, 1981; Murdock, 1982, 1983) and has precidents in the ideas of Galton (1879).

Note that under certain circumstances the results produced by a composite trace are indistinguishable from those of multiple trace models such as those of Hintzman (1983), where the results of multiple retrievals are added. The composite trace conserves storage space and constrains the retrieval mechanism. The search metaphor, or course, cannot apply with a composite trace.

Correlation

When a retrieval cue is given, it is assumed that the representation of the cue item is correlated with the composite trace in order to retrieve an item. The item retrieved, however, is not identical to any item that was initially encoded. The correlation, denoted #, of two vectors is

$$(\mathbf{X}\#\mathbf{Y})_m = \sum_{(i,j)\in S(m)} x_i y_j, \tag{4}$$

where

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

Figure 5 illustrates the correlation of Item **B** with the convolution of two unrelated items, **A** and **B.** I have left the joint convolution–correlation matrix in an expanded form to illustrate the components that reduce to noise (with expected values of zero) and those that produce signal; the signal components have been circled. It can be seen that the central n features of the resulting vector, **R**, produce the Item **A** plus noise. For instance,

$$\begin{aligned} r_{-2} &= a_{-2}(b_{-2}^2 + b_{-1}^2 + b_0^2 + b_1^2 + b_2^2) \\ &\qquad + (\text{noise}) \\ &= a_{-2}(\mathbf{B}\cdot\mathbf{B}) \qquad + \text{noise} \\ &= a_{-2}(1) \qquad + \text{noise}. \end{aligned}$$

In a similar way, r_{-1} is a_{-1} + noise, r_0 is a_0 + noise, and so on for all of the retrieved features. The retrieved item differs from the initial Item **A** insofar as the retrieved item is noisy.

Figure 6 shows the convolution–correlation matrix when **A** is used as a cue. The result is **B** + noise. To simplify the matrix, I have entered 0 for all of the components that, because of the assumption of statistical independence among features, have an *expected* value of zero (even though these values will not be precisely zero, but will be noise). Comparison of Figures 5 and 6 shows that the signal components are different, depending on whether **A** or **B** is correlated with the

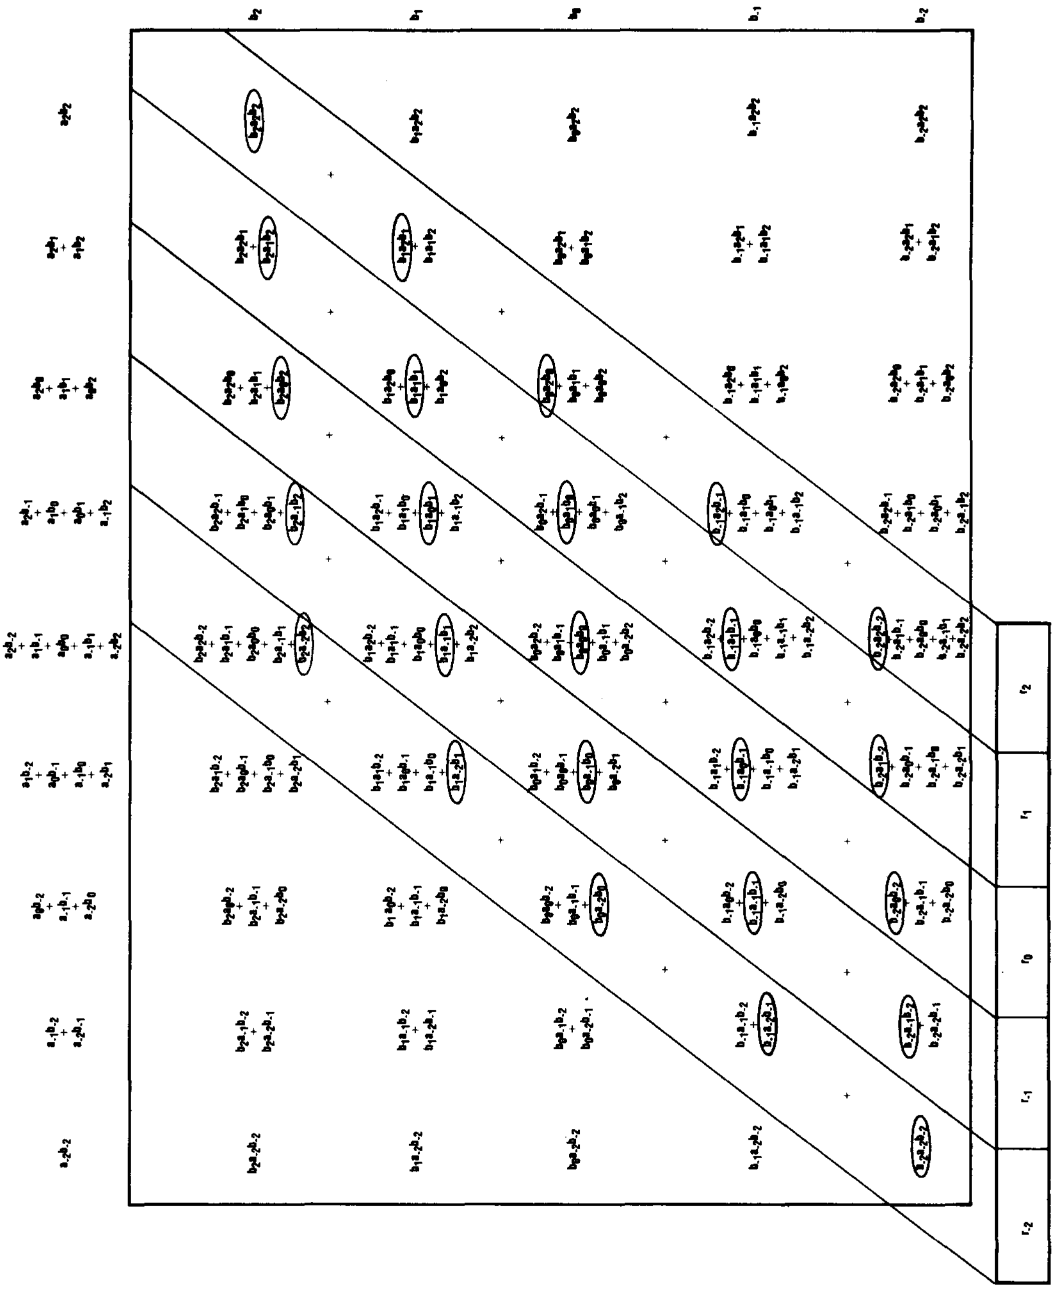

trace. If an item unrelated to either **A** or **B** is correlated with the trace, *all* components in the joint convolution–correlation matrix have expected values of zero and the result **R** is only noise.

Pattern Identification

If the task is verbal, the pattern that is retrieved by correlating a cue with the composite trace must be identified so that a discrete verbal response may occur. In verbal recall, the retrieved item is matched against every item in the lexicon (except the probe or cue itself) and the best match, above a minimum criterion, is the verbal item recalled.

Later in the article, in the section on Recognition and CHARM, I propose that recognition retrieval is similar to recall, except that the item retrieved by the probe is matched only to the probe itself. The matching process involves the computation of the dot product between the retrieved item and the lexical response possibilities (in recall) or the probe (in recognition).

Some Examples

Unrelated associations add noise. Figures 5 and 6 show one association in the composite trace. Suppose, however, that the trace were constructed as

$$\mathbf{T} = \mathbf{A}*\mathbf{B} + \mathbf{C}*\mathbf{D} + \mathbf{E}*\mathbf{F} + \cdots,$$

and all of the items were unrelated. Correlating **A** with **T** gives

$$\begin{aligned}\mathbf{R} &= \mathbf{A}\#\mathbf{T}\\ &= \mathbf{A}\#[(\mathbf{A}*\mathbf{B}) + (\mathbf{C}*\mathbf{D}) + (\mathbf{E}*\mathbf{F}) + \cdots]\\ &= \mathbf{A}\#(\mathbf{A}*\mathbf{B}) + \mathbf{A}\#(\mathbf{C}*\mathbf{D}) + \mathbf{A}\#(\mathbf{E}*\mathbf{F}) + \cdots\\ &= \mathbf{B} + \text{noise}_{\mathbf{A}*\mathbf{B}} + \text{noise}_{\mathbf{C}*\mathbf{D}} + \text{noise}_{\mathbf{E}*\mathbf{F}} + \cdots.\end{aligned}$$

Each unrelated association contributes noise to the item that is retrieved.

Figure 7 gives an analogy to the result produced from the addition of unrelated associations in the composite trace. The analogy is not perfect because in the figure the dots are coded as zeros (white) and ones (black) rather than as having values centered around zero. Thus, in the figure, but not the model, the addition of noise makes the resulting representation darker. Nevertheless, it can be seen from the figure that the addition of noise decreases the discriminability of the signal. Even if only one association were entered in the composite trace, the result of retrieval differs from the initial item insofar as the retrieved item is noisy. The addition of noise is the simplest representational change that results from the holographic association and composite trace in CHARM.

Associative interference—the A–B A–C trade-off. There are no constraints on what items can be associated with one another in the model. The left panel of Figure 8 provides an analogy to the retrieved item when two unrelated items are associated (**A** * **B**) and then **A** is correlated with the trace. The successive panels show the result of retrieval in the following situations:

$$\mathbf{A}\#[(\mathbf{A}*\mathbf{B}) + (\mathbf{A}*\mathbf{C})] = \mathbf{B} + \mathbf{C}\ (+\ \text{noise})$$

$$\mathbf{A}\#[\mathbf{A}*\mathbf{B}) + (\mathbf{A}*\mathbf{C}) + (\mathbf{A}*\mathbf{C})] = \mathbf{B} + \mathbf{C} + \mathbf{C}\ (+\ \text{noise})$$

$$\mathbf{A}\#[\mathbf{A}*\mathbf{B}) + (\mathbf{A}*\mathbf{C}) + (\mathbf{A}*\mathbf{C}) + (\mathbf{A}*\mathbf{C})] = \mathbf{B} + \mathbf{C} + \mathbf{C} + \mathbf{C}\ (+\ \text{noise}),$$

when the items **A**, **B**, and **C** are unrelated. It can be seen from the figure that the increasing number of A–C associations added into the composite trace, causes the Item **C** to become more emphasized in the composite representation that is retrieved by correlating **A** with the trace. This situation, corresponding to the A–B, A–C paradigm (Barnes & Underwood, 1959), has been discussed in more detail in Eich (1982).

Prototype extraction. If items that are similar to one another are each convolved with the same unrelated Item, **X**, and each

Figure 5. The matrix illustrating the correlation of Item $\mathbf{B} = (b_{-2}, b_{-1}, b_0, b_1, b_2)$, with the trace formed by convolution as was shown in Figure 4. (The signal components that are responsible for the retrieval of A have been circled.)

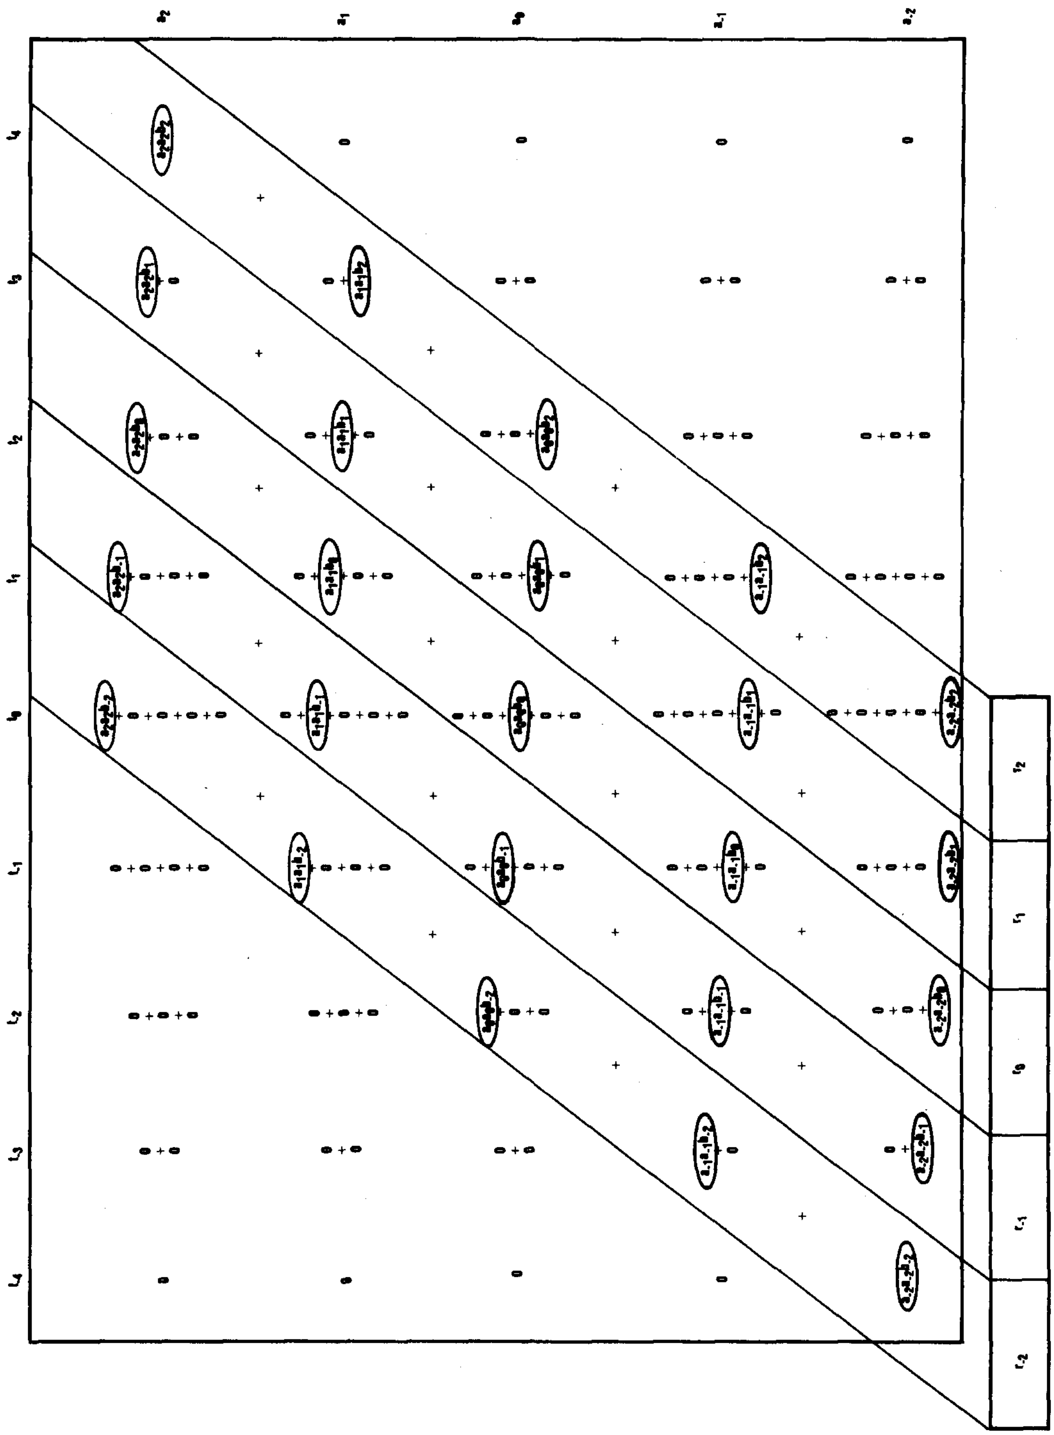

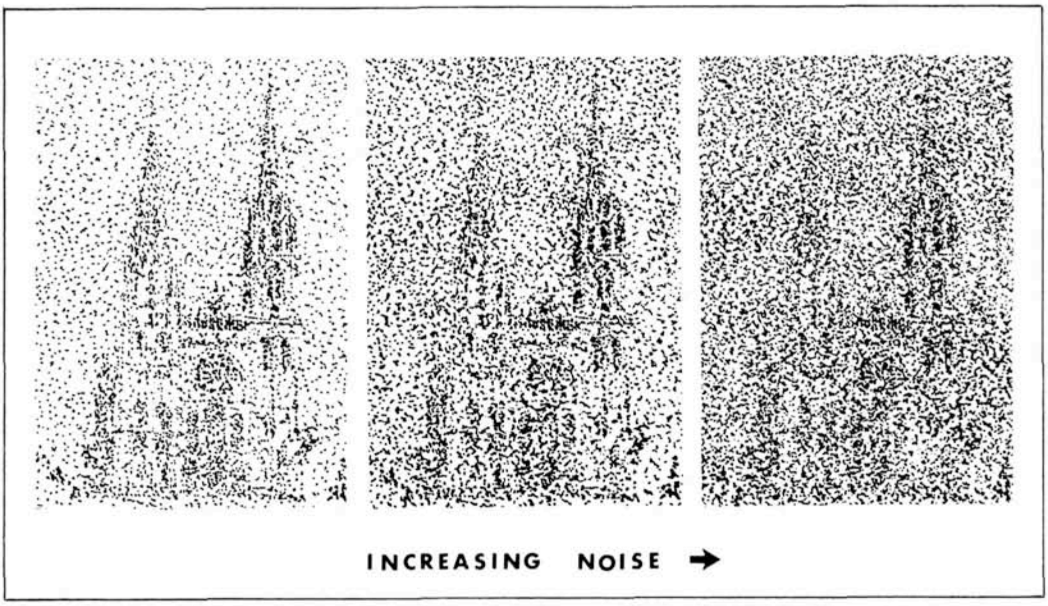

Figure 7. A pictorial analogy to the addition of noise that results when unrelated associations are added into the composite trace.

association is added into the composite trace, the retrieved item will still be composite in form. In this situation, however, the retrieved representation will highlight the features that were common among the items. The top row of Figure 9 gives four items that are similar to one another. The bottom row gives an approximation to the item that is retrieved via correlation as follows from left to right:

$$\mathbf{X}\#(\mathbf{X}*\mathbf{A}) = \mathbf{A},$$

$$\mathbf{X}\#[(\mathbf{X}*\mathbf{A}) + (\mathbf{X}*\mathbf{B})] = \mathbf{A} + \mathbf{B},$$

$$\mathbf{X}\#[(\mathbf{X}*\mathbf{A}) + (\mathbf{X}*\mathbf{B}) + (\mathbf{X}*\mathbf{C})] = \mathbf{A} + \mathbf{B} + \mathbf{C},$$

$$\mathbf{X}\#[(\mathbf{X}*\mathbf{A}) + (\mathbf{X}*\mathbf{B}) + (\mathbf{X}*\mathbf{C}) + (\mathbf{X}*\mathbf{D})] = \mathbf{A} + \mathbf{B} + \mathbf{C} + \mathbf{D}.$$

It can be seen that the idiosyncratic aspects of the separate items are lost and a prototype emerges (see J. A. Anderson, 1977; Eich, 1982; Hintzman, 1983, for further discussion).

We may summarize the result of retrieval **R** as

$$\begin{aligned}\mathbf{R} = \mathbf{X}\#(\mathbf{T}) &= \mathbf{X}\#[(\mathbf{A}*\mathbf{B}) + (\mathbf{C}*\mathbf{D}) + \cdots)] \\ &= S_{\mathbf{XA}}\mathbf{B} + S_{\mathbf{XB}}\mathbf{A} + \text{noise}_{\mathbf{A}*\mathbf{B}} + S_{\mathbf{XD}}\mathbf{C} \\ &\quad + S_{\mathbf{XD}}\mathbf{D} + \text{noise}_{\mathbf{C}*\mathbf{D}} + \cdots, \quad (5)\end{aligned}$$

where **T** is the trace, **A, B, C, D,** and so forth, are list items, $\mathbf{S_{XA}}$ is the similarity between **X** and **A,** and **X** is any cue item. Equation 5 falls directly and immediately out of the method of association formation, storage, and retrieval. It is important for understanding the results that follow.

Intrapair similarity. Suppose, now, that the two items that are associated by convolution are similar to one another. Figure 10 gives an example of the convolution matrix produced by two similar items,

$$\mathbf{A} = (a_{-2}, a_{-1}, a_0, a_1, a_2), \quad \text{and}$$

$$\mathbf{B} = (a_{-2}, a_{-1}, a_0, b_1, b_2).$$

Figure 6. The matrix illustrating the correlation of Item $\mathbf{A} = (a_{-2}, a_{-1}, a_0, a_1, a_2)$, with the trace formed by convolution as was shown in Figure 4. (The signal components that are responsible for the retrieval of **B** have been circled. Components that are only noise with an expected value of zero are indicated as zero in this illustration, in order to simplify the matrix.)

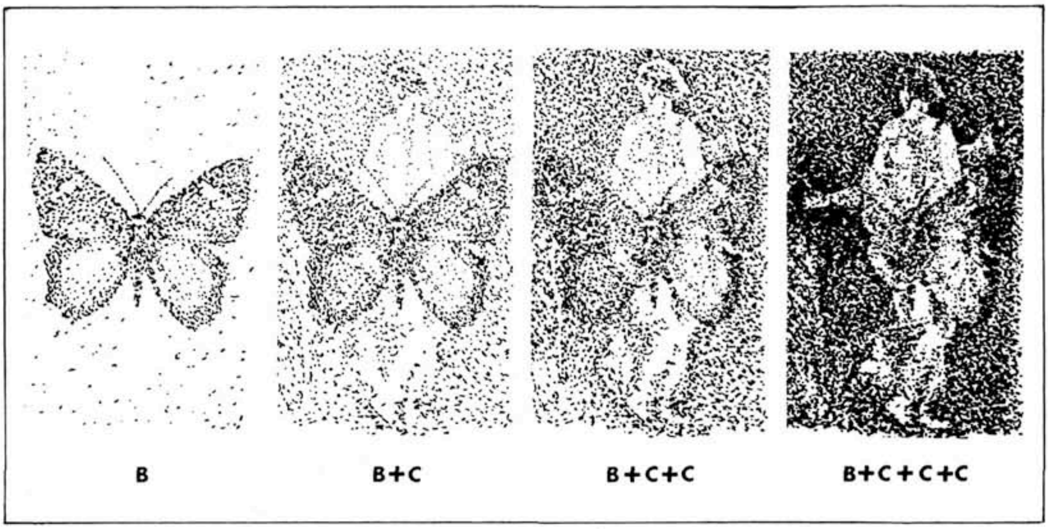

Figure 8. A pictorial analogy to the simultaneous retrieval of more than one unrelated item as in the A–B, A–C paradigm.

Figure 11 shows the joint convolution–correlation matrix that results when **B** is correlated with the association. The components that produce signal are circled, just as in Figures 5 and 6. By comparing Figure 11 to Figure 5, it can be seen that all of the

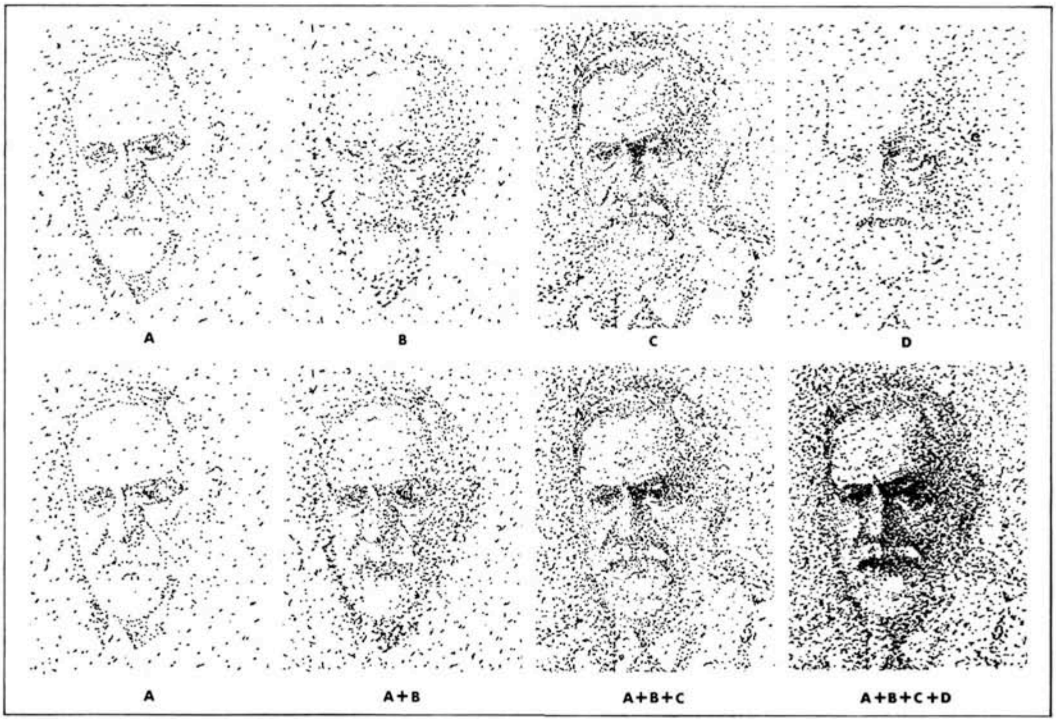

Figure 9. A pictorial analogy to the simultaneous retrieval of highly similar items, which is responsible for prototype abstraction.

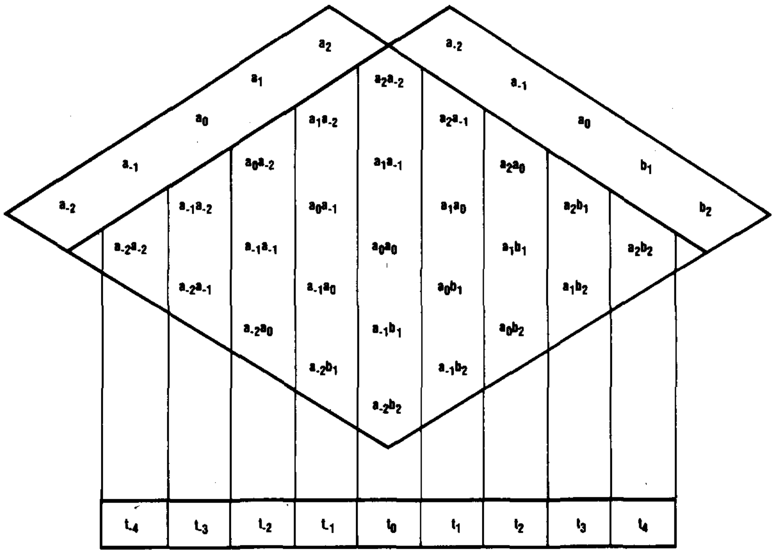

Figure 10. A matrix illustrating the convolution of two similar items, $\mathbf{A} = (a_{-2}, a_{-1}, a_0, a_1, a_2)$, and $\mathbf{B} = (a_{-2}, a_{-1}, a_0, b_1, b_2)$.

components that, in the unrelated case, gave rise to the retrieval of **A** are produced. For the rows of the matrix corresponding to the independent features (i.e., for b_1 and b_2), only the signal components that produce **A** are extracted. For the rows in the correlation matrix corresponding to the common features (a_{-2}, a_{-1}, and a_0), a second set of signal components is extracted. By comparing Figure 11 to Figure 6, it can be seen that these second signal components are the same as are produced when **A** is correlated with $\mathbf{A} * \mathbf{B}$; that is, they are components that produce **B**. A second set of signal components is produced only by the features that are the same in both **A** and **B** (because a_{-2}^2, a_{-1}^2, and a_0^2 have positive values each of expected magnitude $1/n$). The features that are independent in the two items do not extract a second set of signal components (because when multiplied—i.e., $a_1 b_1$ and $a_2 b_2$—they have expected values of zero). Thus, each common feature in the Cue **B**, in addition to producing **A**, also produces **B**. The noncommon features in the Cue **B** only produce **A**.

The fact that **B** is produced by the features that are common to the two initial items (i.e., a_{-2}, a_{-1}, a_0) increases the signal to noise (S/N) ratios. Those common features have been amplified and are less likely to be overwhelmed by noise. The item that is produced in this situation is a compromise between **A** and **B**, such that the features that are common to both have been emphasized. One of the implications of this result is that when the cue and target of a pair are similar to one another, subjects should sometimes give the retrieval cue as a response to itself. This should not happen when the cue–target pair are made up of unrelated items. The experimental results were in accord with this prediction (Eich, 1982, Experiment 2). However, it is only under special experimental conditions that subjects can be induced to make any cue intrusions. Normally, subjects exclude the cue item as a possible response. If, in the model, the cue item is excluded as a possible response, the case in which the cue and target are similar to one another is quite interesting because it results in better recall of the target than would be the case had the cue and target been unrelated. In particular, the S/N ratio on the common features is increased, leading to a stronger item.

The result that occurs when items within a pair are similar to one another may not be intuitively obvious, so I have prepared some drawings to illustrate this situation. Figure 12 shows dot drawings taken from photo-

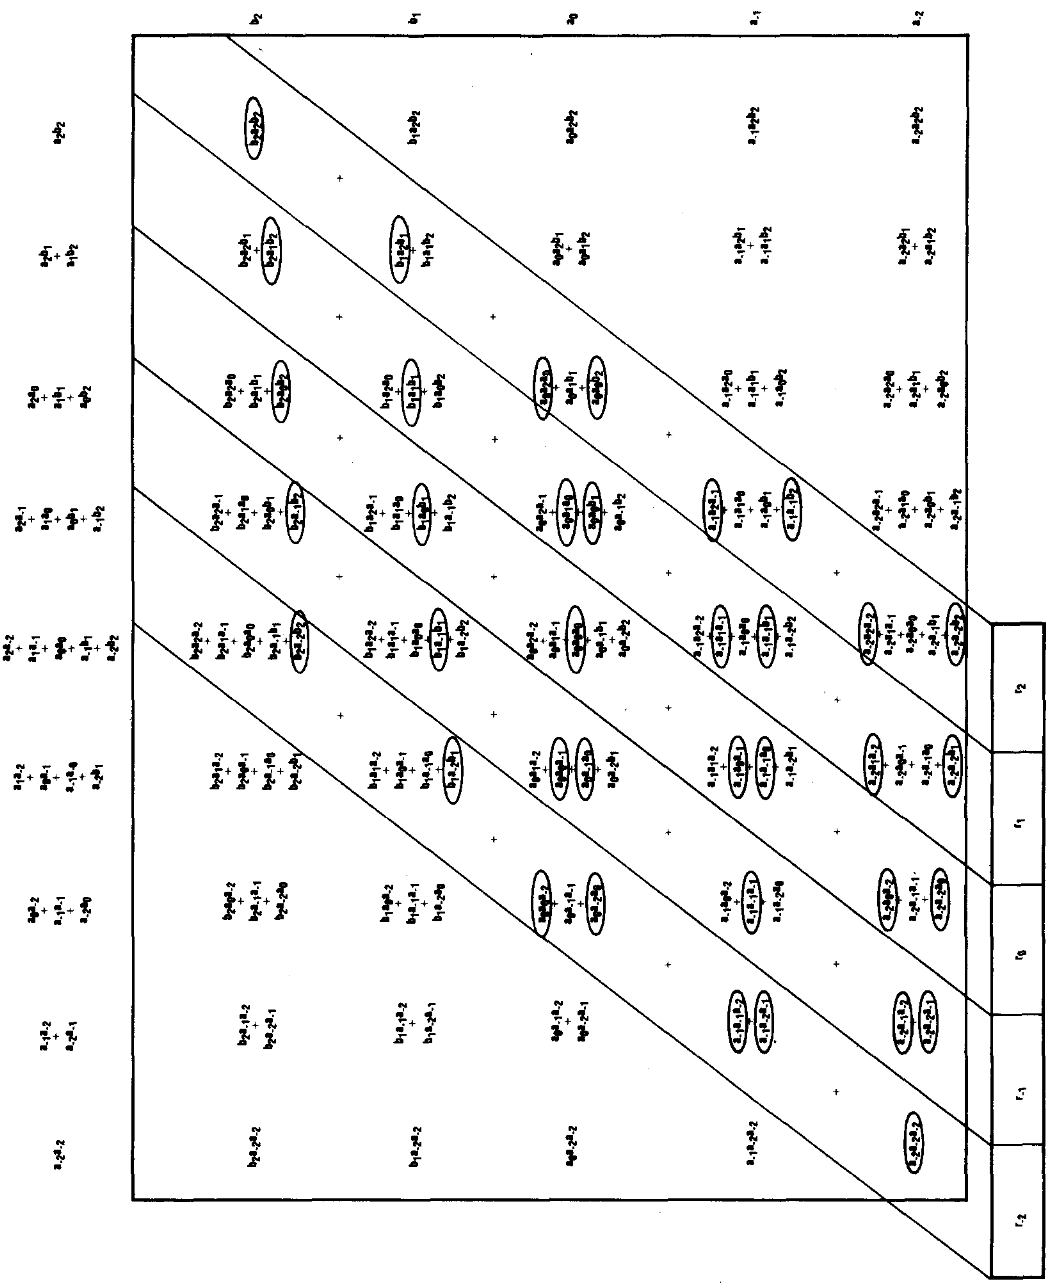

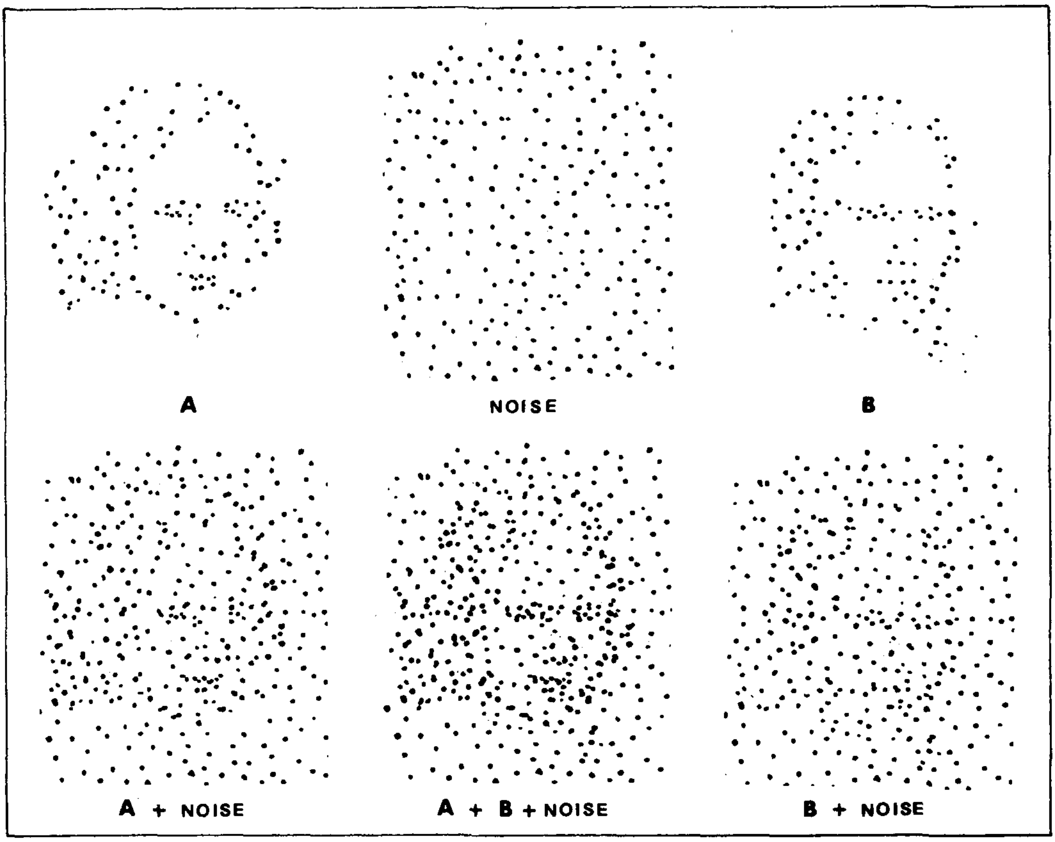

Figure 12. A pictorial analogy to the composite retrieved when two similar items have been associated, is shown in the bottom center panel. (The left and right bottom panels give an analogy to retrieval when the items associated were not similar to one another.)

graphs of two people—Millicent Galton on the top left, and Charles Darwin on the top right. Darwin and Galton were cousins and have a certain family resemblance in face shape, eye spacing, and so on. The top center panel is noise. The bottom left panel shows the result of adding noise to the depiction of Galton alone. Such a result occurs, in the model, under conditions in which the two items convolved are unrelated, and then one item is used to cue the other,

$$\text{X}\#(\text{X} * \text{Galton}) = \text{Galton} + \text{noise}.$$

The bottom center panel is the addition of Darwin, Galton, and noise. This is an approximation to the item retrieved by the model when two similar items are convolved and then one of them is used as a retrieval cue. For instance,

$$\text{Darwin}\#(\text{Darwin} * \text{Galton}) = \text{Galton} + S_{\text{DG}}\,\text{Darwin} + \text{noise}.$$

It can be seen that the figure in the bottom center panel stands out better, against the noise, than does the figure in either of the bottom side panels.

The present explanation of the main effect of levels of processing is based on the testable assumption that often the representations of the deeply encoded context and target items (e.g., *dog*–CAT) are more similar to one

Figure 11. The matrix illustrating the correlation of Item $\mathbf{B} = (a_{-2}, a_{-1}, a_0, b_1, b_2)$ with the trace formed by convolution as in Figure 10. (The signal components have been circled.)

another than are the representations of the shallowly encoded context and target items (e.g., *hat*–CAT). This assumption is discussed in more detail in the next section. Because of the associative and retrieval mechanism in the CHARM model—which produce Equation 5—this difference in similarity, in this situation, results in better recall in the deeply than in the shallowly encoded conditions.

Before proceeding to some evidence about the nature of the interitem similarities in situations like levels experiments, one further comment is necessary. The model does not predict that similarity, under any circumstances whatever, increases recall. To give a few examples, if the similarity of all of the items in one list of paired associates is greater than the similarity of all of the items in a second list, the model predicts recall to be higher in the unrelated than in the highly similar list. This prediction, as well as the corresponding experimental result, was shown in Simulation 1 and Experiment 1 of Eich (1982). If all of the cue items are similar to one another, recall is predicted to be rather bad because all of the items have a tendency to be retrieved, and the response selection is liable to be quite inaccurate. In a like manner, certain combinations in the Osgood (1949) surface with high similarity result in poor recall. The entire Osgood surface is produced by the mechanisms in the CHARM model. It is not possible to reduce the point or hypothesis of the present article to a statement to the effect that similarity improves recall. The more accurate statement is that similarity has an effect on the result of the transformation that occurs by means of storage by convolution and addition, and retrieval by correlation. The particular similarity relation that is usually used in levels tasks corresponds to a situation in the model in which an increase in similarity helps recall.

Similarity, Levels of Processing, and CHARM

A major assumption in applying the CHARM model to levels of processing main effects is that the similarity between the context and the target item tends to be greater in the deep than in the shallow conditions. A task that is frequently given to induce deep processing, is to ask subjects whether a particular item is a member of a certain category. Many researchers investigating the representation of natural and constructed categories have concluded that the similarity among exemplars is of critical importance. For instance, Rosch (1977) discussed categorical structure in terms of the real-world correlational structure of concrete objects. According to Rosch (1977) there is a natural basic level of classification that is the most inclusive level at which lower level objects have (a) a number of attributes in common, (b) motor movements in common, (c) objective similarity in shape, and (d) a class that is identifiable from the averaged shapes of the lower level objects. Note that all four of these aspects reflect what is here called *similarity* among the items. For each of the four aspects of the correlational structure, she has found evidence that category members (subordinates) do indeed overlap with one another and with the more abstract basic level object.

Sokal (1977), in discussing classification from a taxonomic point of view, noted that the process of classification is critically dependent on the perception of similarities. He noted that the distinction may be drawn between monothetic classifications (those with defining features) and polythetic classifications. The latter are those in which "individuals or objects share a large proportion of their properties but do not necessarily agree in any one property . . . A corollary of polythetic classification is the requirement that many properties (characteristics) be used to classify objects" (p. 190). Sokal goes on to argue, after Wittgenstein (1958), that most natural classifications are of the polythetic type. The important point, for the present article, is that classification of an item as a member of a particular category is likely to be based on shared properties (though not necessarily defining shared properties).

The importance of similarity in classification is also evident in data on the acquisition of word meanings among children. The overextensions that children use in applying particular terms strongly suggest that similarity among referents, whether it be perceptual similarity (Clark, 1973), functional similarity (K. Nelson, 1974), or some combination of different sorts of similarities (Bowerman, 1977), is critical in classification judgments even at a young age.

The impact of similarity on classification judgments is not restricted to natural classifications but is also exemplified by studies that have used artificial or constructed stimuli. Posner and Keele (1970) found that items that were large distortions of the learned schema or exemplars were classified less well than items that were small distortions. Items that were highly similar to the constructed category (i.e., small distortions) tended to be classified as category members; items that were less similar (i.e., large distortions) tended to a greater extent to be rejected as category members. It may be inferred that with increasing dissimilarity between the target item and the representation of the category, there would be an increasing tendency to say that an item was not a member of that particular category.

There appears to be rather good agreement that there is high similarity among category members. Most classification models are based on this tenet in one form or another (e.g., Hintzman & Ludlum, 1980; Medin & Schaffer, 1978; Reed, 1973; Rosch & Mervis, 1975; and see Smith & Medin, 1981). There is also high similarity between the category name and its own exemplars (Tversky & Smith, in Smith & Medin, 1981, p. 118). It therefore seems reasonable to suppose that when a subject is given a question in a levels-of-processing task such as "Is it a *fish?*—TROUT," there is considerable feature overlap or similarity between the mental representations of the key ideas, *fish* and TROUT. However, when the person is given a categorical question to which the answer is no, such as "Is it a *musical instrument?*—TROUT," there is rather low similarity between the key ideas. In short, if category inclusion is based on high similarity, category exclusion is based on insufficient similarity.

The similarity between context and target appears to have memorial implications. For instance, Schulman (1974) found that congruous questions (e.g., Is a soprano a singer?) produce better memory performance than do incongruous questions (e.g., Is mustard concave?). Craik and Tulving (1975) tested the idea that it is the congruity that is important and not whether a question receives a positive or negative response. Subjects were asked questions in which the congruity between the central ideas was independent of whether the question should receive a yes or no response. They found, in this case, that there were no memorial differences depending on the type of response.

Semantic properties or features need not be represented differently from sensory properties. For instance, in Rosch's (1977) example, cited above, category members were shown to be similar in shape. Shape, however, would not obviously be considered to be a semantic property, although it may be used to make a semantic judgment. As Lockhart (1979) put it,

> The distinction between sensory and semantic codes has in an important sense been drawn too sharply. There is a tendency to regard the sensory component of a word or picture as a kind of detachable skin within which is a kernel of meaning but that is itself meaningless. Sensory and semantic features do not possess this simple additive relationship. The sensory codes are aspects of analysis of meaning. (p. 82)

Sensory features are here considered to be part of the representation of items, along with what might be called semantic features.

D. Nelson, Fosselman, and Peebles (1971) have shown that memory performance was facilitated when cue–target pairs shared letters in common (e.g., cactus–carrot, instep–influx). This study indicated that sensory, and not just semantic features, have memorial implications. D. Nelson, Wheeler, and Brooks (1976) combined both semantic and sensory similarity in the same experiment and found that both had an effect on performance. It is assumed that even with shallow processing tasks, there is some similarity between the context and target, so long as the answer to the shallow question is yes.

Simulation 1

The first simulation is addressed to the cued recall, rhyme, and category conditions of Moscovitch and Craik's (1976) Experiment 1. Subjects were given questions that required either a positive or a negative response. Examples of the four conditions were (a) category positive (Is it in the category *fruit?* CHERRY), (b) rhyme positive (Does it rhyme with *pond?* WAND), (c) category negative (Is it in the category *office supplies?* LAMP), and (d) rhyme negative (Does it rhyme with *stock?*

CLIPS). At time of test, the subjects were given the encoding questions and were asked for recall of the appropriate (encoded) target words.

The results of this experiment are presented in the top panel of Table 1. As can be seen from the table, there was a main effect of level of processing. Cued recall was better when category questions were asked than when rhyme questions were given. In addition to this main effect, there was a main effect of type of response (yes or no) such that yes targets were recalled better than no targets. Qualifying both of these effects was an interaction between level of processing and type of response. The effect of level of processing was much attenuated with the questions that received a no response. The difference between the rhyme-positive and category-positive conditions was 42%, whereas the difference for the negative conditions was only 7%.

The pattern of recall in the positive conditions, of course, corresponds well to the original depth of processing formulation. However, as Moscovitch and Craik noted, the pattern of results for the negative responses poses some problems. Because it is the case that category questions are assumed to provoke deep processing and rhyme questions shallow processing, regardless of the answers to these questions, there should have been no difference in the recallability of the targets depending on the type of response. To address this problem with the negatives, the ideas of elaboration (Craik & Tulving, 1975), uniqueness or distinctiveness (Moscovitch & Craik, 1976), and congruency (Craik & Tulving, 1975; Moscovitch & Craik, 1976) have been evoked in addition to depth of processing as being determinants of recall. In contrast, in this simulation these results are attributed to the similarity between the context and target items and the mechanisms in CHARM.

It is assumed in this simulation that in the category-positive condition (*fruit*–CHERRY), the context and target are highly similar to one another, and unrelated to the other list items; in the rhyme-positive condition (*pond*–WAND), the context and target are somewhat similar to one another and unrelated to the other list items; in the category-negative condition (*office supplies*–LAMP), the context and target are unrelated to one another and to

Table 1
Percentage Recalled as a Function of Encoding Question and Type of Response

Type of response	Encoding condition: Rhyme (%)	Category (%)
Experiment 1[a]		
Positive	35	78
Negative	11	18
Simulation 1[b]		
Positive	33	80
Negative[c]	13	15

[a] From "Depth of Processing, Retrieval Cues, and Uniqueness of Encoding as Factors in Recall" by M. Moscovitch and F. I. M. Craik, 1976, *Journal of Verbal Learning and Verbal Behavior, 15*, p. 449. Copyright 1976 by Academic Press, Inc. Adapted by permission.
[b] Criterion = 1.0.
[c] The difference in the negatives in the simulation is not systematic.

the other list items; in the rhyme-negative condition (*stock*–CLIPS), the context and target are unrelated to one another and to the other list items. From the evidence outlined, this ordering of similarities seems to be a reasonable guess.[1] The context–target pairs are associated by convolution, stored in a composite memory trace, and retrieved by correlation. The retrieved item is then identified by being matched to every item in the lexicon, except the cue item itself. The lexical item with which the retrieved item most strongly resonates, above a certain minimal criterion, is considered to be the response. In this simulation, the criterion for responding is varied.

Method. A lexicon of 18 items was constructed. All 18 items were assigned numerical values on each feature by randomly selecting each feature value from a unit normal distribution. There were 63 features in each item.

[1] Although this guess about similarity seems reasonable it may be incorrect. In particular, the semantic-negative context and target items may be more similar than the rhyme negatives. Among the semantic negatives were questions such as "round object—lamp, cherry, rock, pool, pail, glass," which intuitively seem to exhibit some similarity. The data in several other experiments also suggest a difference in similarity between the two negative conditions. The argument proposed in Simulation 1 is not circular because the similarity values are potentially available empirically, as is exemplified by Experiment 1.

The feature values were normalized for every item (**F**) so that $\mathbf{F} \cdot \mathbf{F} = 1$. Items 1 and 2 were designated as the context and target of the category-positive condition. To represent the similarity between them, 54 features were randomly selected without replacement and were reassigned the same value in Item 2 as those features had in Item 1. Items 3 and 4 were designated as the context and target in the rhyme-positive condition. Twenty-seven features were randomly selected without replacement and reassigned the same values in Item 4 as had the corresponding feature in Item 3. Items 5 and 6 were considered to be the context and target in the category-negative condition; Items 7 and 8 were considered to be context and target in the rhyme-negative condition. These four items were left in their original unrelated form.

The composite trace was constructed as follows:

$$(\text{Item 1} * \text{Item 2}) + (\text{Item 3} * \text{Item 4}) + (\text{Item 5} * \text{Item 6}) + (\text{Item 7} * \text{Item 8}),$$

where $*$ denotes convolution, as before. Only the central n (63) features of the convolutions were used in this simulation and in all simulations in this article.

To retrieve, each of Items 1, 3, 5, and 7 were correlated with the 63-tuple representing the composite trace, to result in four retrieved items. The retrieved items were then matched to every item in the lexicon, except the item that had been used as the retrieval cue. To match, the dot product was calculated between the retrieved items and each lexical item, to result in a resonance value for each lexical item (for each of the four cues). The lexical item that had the highest resonance value above a minimum criterion was considered to be the item that had been recalled. Two criteria were used, 0.5 and 1.0. (The value of 1.0 was chosen because it produces numerical values that are very close to those found in the Moscovitch & Craik, 1976, experiment. The value of 0.5 was chosen to illustrate the effect of lowering the criterion.)

The entire simulation was replicated 100 times with a different randomization for the lexical representations used for each replication. What (if anything) was recalled on each replication for each cue was tabulated, and the means and standard deviations of the resonance values were calculated, for each lexical item.

Results. The frequency of correct recall for the high criterion (1.0) simulation are presented in the bottom panel of Table 1; the relevant data from Moscovitch and Craik's (1976) experiment are shown in the top panel. As can be seen from the table, the simulation produced the main effect of levels—category positive over rhyme positive. It also produced lower recall overall for the negative conditions and an interaction between the levels effect and the type of response. The pattern produced by the model was due to the similarity values among the items, and the fact that given the encoding and retrieval mechanisms in CHARM, the similarities of the items change the results of retrieval.

The resonance values produced by the simulation are shown in Table 2. It can be seen that the target item had the highest resonance value in the high-similarity (category-positive) condition, and the lowest resonance value in the unrelated (category- and rhyme-negative) conditions. The criterion can have a radical effect on responding, however. For instance, the high criterion reduced the rate of recall on the unrelated conditions from about 85% to about 14%. The rate of intrusions was also reduced from 11 total in the low-criterion simulation, to 1 in the high-criterion simulation. A very low (sub-zero) criterion would correspond to a forced recall situation, in which a subject is not allowed the option of not responding. A very high criterion could, of course, eliminate all responding even if the appropriate information were retrieved from memory.

Discussion

The reason the model produced the results given in Table 2 is straightforward. When the context item is used as a retrieval cue, in the CHARM model, recall of the target is an increasing function of the similarity of the context associated with the target, and the target (so long as the rest of the list items are unrelated to the cue). In the case where similarity is high, the retrieval cue recreates not only the target item, but also the context itself. The extent to which the context is

Table 2
Resonance and Recall Scores From Simulation 1

Condition	Resonance: Target	Resonance: Intrusions	Recall criterion (%): 0.5	Recall criterion (%): 1.0
High similarity (category positive)			99	80
M	1.3	0.0		
SD	0.3	0.3		
Medium similarity (rhyme positive)			91	33
M	0.9	0.0		
SD	0.3	0.3		
Unrelated (category and rhyme negative)			85	14
M	0.7	0.0		
SD	0.2	0.2		

recreated depends on the extent of similarity between the context and the target. The reconstruction of the context helps recall of the target by increasing the strength of the features that are common to both the context and the target.

As shown in the next simulation, it is not only the similarity between the encoded context and the target that is important in CHARM, but also, and fundamentally, the nature of the cues that are correlated with the trace at time of retrieval.

Encoding Specificity and CHARM

Tulving has proposed that factors such as depth of encoding, distinctiveness, and elaboration need not be considered separately from or in addition to the compatibility relation between the trace and the cue, that is, that the idea of encoding specificity (Tulving, 1979; Tulving & Osler, 1968; Tulving & Thomson, 1973) applies to levels data. An experiment by Fisher and Craik (1977, Experiment 2) provides an interesting case in point because it showed an encoding specificity interaction and also a main effect of levels of processing.

Fisher and Craik (1977) presented subjects with a list of 62 pairs of words. Four pairs made up a primacy buffer and four were included in a recency buffer that were excluded from further consideration. Of the 54 remaining pairs, half were constructed such that the cue was a high associate of the target, as in *dog*–CAT, and half were pairs in which the cue rhymed with the target, as in *hat*–CAT. A particular target was presented only once in the list, but all of the target items had both a rhyme and a high associate that had previously been equated for the ease of generation of the target from semantic memory. The a priori probability of producing the target items in the absence of any experimental input into episodic memory was .16 for the rhyme cues and .15 for the semantic cues. These probabilities indicate whether the target was one of the words (usually two to four) that were produced within a 10-s interval. Thus, the a priori likelihood of saying *cat* as a response to "rhymes with *hat*" was about the same as the a priori likelihood of saying *cat* to "is associated with *dog*." At time of test, half of the semantically encoded targets were cued by the items with which they had been encoded (compatible cueing conditions), whereas half were cued with the appropriate extralist rhyme cue. For the rhyme encoded items, half were cued with the studied (compatible) rhyme cue and half with the extralist associate cue. The results are shown in the left panel of Table 3. There was a levels-of-processing effect—semantic better than rhyme—but it was obtained only when the list cues were given at time of test.

The finding of a main effect of semantic over rhyme, in the compatible cuing conditions, is consistent with the idea that the semantic encoding led to a deeper (richer, broader, more elaborate, or more distinctive) trace. However, none of the constructs that have been used to account for the better performance with semantic processing are sufficient to explain why this superior semantic trace was not also shown to produce superior recall when cued with the incompatible cues. If the trace was richer in the semantic condition, this richness had no impact when the extralist cues were used.

Tulving has argued that Fisher and Craik's (1977) data are most parsimoniously explained in terms of only the compatibility relation between the cue and the trace. A retrieval cue is proposed to be effective only to the extent that its informational content matches, overlaps, or is compatible with the information content of the trace. In CHARM, the cue and the trace in no sense match or overlap because the meaning of the dimensions in the cue and trace are quite different (see the earlier section of the article, on *Convolution*). The term *compatibility*, however, could be liberalized so as not to implicate a matching mechanism for retrieval. This would be consistent with the mechanisms in CHARM.[2] The model, however, is able to avoid a circularity problem that is implicit in the compatibility explanation. In particular, the compatibility explanation allows no method independent of the to-be-explained recall for assessing what will or will not be a compatible relation. (Note that T. Nelson, 1977, leveled a similar criticism against the depth of processing idea.) CHARM allows that the similarity among items may be ascertained independently of recall.

Table 3
Proportion Recalled as a Function of Encoding Context and Type of Cue

Retrieval cue	Encoding context: Rhyme	Encoding context: Associate
Fisher and Craik (1977) Experiment 2[a]		
Rhyme		
Target recall	.26	.17
A priori probability		.16
Associate		.44
Target recall	.17	
A priori probability	.15	
Simulation 2		
Rhyme		
Target recall	.32	.01
Resonance		
M	.91	.34
SD	.27	.26
Associate		
Target recall	.04	.43
Resonance		
M	.44	.99
SD	.32	.33

[a] From "Interaction Between Encoding and Retrieval Operations in Cued Recall" by R. P. Fisher and F. I. M. Craik, 1977, *Journal of Experimental Psychology: Human Learning and Memory, 3,* p. 707. Copyright 1977 by American Psychological Association. Adapted by permission.

In Simulation 1, it was assumed that the cue and the target were more similar to one another in the semantic (positive) condition than in the phonemic (positive) condition. It was this difference in similarity in combination with the associative, storage, and retrieval mechanisms of CHARM that explained the main effect of depth of processing in Moscovitch and Craik's (1976) experiment. Fisher and Craik's (1977) data are explained in the same way in the simulation that follows. Having proposed that the representation of the items—and their similarity to one another—is an important factor influencing recall, it is possible to obtain an independent measure of the similarity (at least ordinally) without recourse to the recall data. Thus, the model circumvents the circularity implicit in the compatibility explanation. In the experiment that follows, the hypothesis—that the semantic pairs in Fisher and Craik's experiment shared more properties in common than did the phonemic pairs—is tested.

Experiment 1

Subjects were presented with the pairs of items from Fisher and Craik's (1977) experiment and were asked to write down all of the properties that the two items had in common. Note that subjects were not asked to free associate to the items, but were specifically asked for ways in which the items were the same. The hypothesis of the experiment was that subjects would be able to produce more properties in common for the semantic than for the phonemic pairs.

Method. Eighteen subjects were told that

[2] Such liberalization of the compatibility idea may be implicit in Tulving's (1983a) synergistic ecphory proposal. He (Tulving, 1983b) has recently stated that this view is compatible with the correlation retrieval operation.

they would see pairs of items for which they were to write down all the common properties or characteristics shared by the two items in each pair. Four minutes were allowed for each pair. Each subject received 12 pairs—6 of which were phonemic and 6 of which were semantic. The actual pairs were randomly assigned with the constraint that no subject saw the same item in both the phonemic and semantic condition. The 18 subjects provided two complete replications of all of the pair combinations that were used in Fisher and Craik's Experiment 2. An additional 9 subjects provided a third replication. The subjects in the third replication were given 3, rather than 4 min to write down the common properties. Subjects were asked to write each property on a separate line. To score the data, the lines were simply counted; no attempt was made to ferret out redundancies or good or bad properties. Subjects were introductory psychology students at the University of California, Los Angeles. They were run in groups.

Result. The mean number of common properties given for the six semantic pairs was 42.0 and for the six phonemic pairs, 29.4. This difference was significant, $t(26) = 4.83$.

Simulation 2

In this simulation, the representational assumption, consistent with the results of Experiment 1, was made that the semantic pairs in Fisher and Craik's (1977) experiment contained more features in common than did the phonemic pairs. It was further assumed, as in other applications of the model, that when subjects encoded a pair of items, the items were associated by convolution, stored in a composite memory trace, and retrieved by correlation of the cue with the trace. The retrieved item was then designated as a particular response by choosing the lexical item that had the highest resonance score with the retrieved item, above a lower criterion. As in the first simulation, the cue item was not considered as a response possibility.

Method. A lexicon of 18 items was constructed. All 18 items were first assigned numerical values on each feature by randomly selecting each feature value for each lexical item from a unit normal distribution. There were 63 features representing each item. The items were then normalized so that the dot product of an item with itself gave a value of 1.

Four blocks of items were constructed starting from the lexical representations so that the first item in each block was the target item, the second was semantically related to the target item, and the third was phonemically related to the target. Items 1, 4, 7, and 10 were designated as target items, and were left in their original form (i.e., they were unrelated to each other). Items 2, 5, 8, and 11 were designated as semantically related items that were modified so that they were somewhat similar to their respective target items. In order to represent the high similarity, 36 features were randomly selected without replacement and reassigned the numerical values of the corresponding target item on those features. Thus for instance, Item 2 was given the same feature values as Item 1 on 36 features. Items 3, 6, 9, and 12 were designated as phonemically related to the target items, and were assumed to have low similarity to the targets. Twenty-seven features were randomly selected without replacement and reassigned the same value as had the corresponding target items on those features.

The net result of this reassignment of feature values was that the semantic items were more similar to their targets than were the phonemic items. A less obvious result was that there was a small amount of similarity induced between the semantic and phonemic items within a block, by virtue of the fact that they were both similar to the target. By chance, occasionally some of the features selected to be overlapping for the phonemic items were also among the overlapping features in the semantic items.

The composite trace was contructed as

$$(\text{Item } 1 * \text{Item } 2) + (\text{Item } 4 * \text{Item } 5) + (\text{Item } 7 * \text{Item } 9) + (\text{Item } 10 * \text{Item } 12).$$

The items in the first two pairs were semantically related, whereas those in the last two pairs were phonemically related. (Convolution is symmetric and so it does not matter whether Item 1 is convolved with Item 2 or vice versa.) The convolutions were truncated to the central 63 features, and so the composite trace was a 63-tuple.

To mimic retrieval corresponding to Fisher and Craik's (1977) Experiment 2, an intralist semantic cue (Item 2—target is Item 1), an extralist phonemic cue (Item 6—target is Item 4), an intralist phonemic cue (Item 9—target is Item 7), and an extralist semantic cue (Item 11—target is Item 10) were separately correlated with the composite trace to produce four retrieved items. Each of the four retrieved items was then matched to each item in the lexicon, except the cue itself, to produce resonance scores for the lexical items. The lexical item with the highest resonance value (above a lower criterion of 1.0) was selected as the response. The entire simulation was replicated 100 times.

Results. The top panel of Table 3 shows the results of Fisher and Craik's (1977) Experiment 2; the bottom panel gives the results from Simulation 2. It can be seen that, in the simulation as well as in the data, there was a main effect of levels of processing. This main effect, however, was found only in the conditions in which the retrieval cue is the context with which the target item was associated. In those conditions in which an extralist retrieval cue was used, there was no levels effect in either the simulations or the experiment. In the simulation, recall was low in the noncompatible cueing condition. An encoding specificity interaction with levels of processing was shown by both the simulation and the data. The ordering of conditions shown in the data: encode semantic–retrieve semantic > encode phonemic–retrieve phonemic > encode semantic–retrieve phonemic = encode phonemic–retrieve semantic, was reproduced by the simulation.

Discussion

The reason the model produced the levels effect in the compatible cueing conditions is the same as has already been discussed for Simulation 1, and in the description of the model. The lack of difference and low level of recall in the noncompatible conditions may be understood from Equation 5. Although the level of recall is low, qualitative differences in what is predicted to be recalled in the two noncompatible conditions may be interesting. If *dog* and CAT are associated and *hat* is given as a cue, the retrieval of CAT depends mainly on the similarity of *hat* and *dog*. But *dog* is also retrieved to the extent that *hat* and CAT are similar, in this situation. If *hat*–CAT is encoded and *dog* is given as a cue, the retrieval of CAT once again depends mainly on the similarity of *hat* and *dog*. However, *hat* is also retrieved to the extent that *dog* and CAT are similar. There is a trade-off at work with the second (nontarget) item that is retrieved in the two extralist cueing cases. In the first case, the second item, *dog*, is not retrieved as strongly as is the second item, *hat*, in the second case. However, though not retrieved as strongly, *dog* has more features in common with the target CAT than does *hat*, and so its retrieval has a greater tendency to help identification of the retrieved item as being CAT. It is conceivable that these predictions might be testable. The high level of recall with the noncompatible cues that was found in the data (.17) but not the model may be due, to some extent at least, to the a priori (nonepisodic) probability of the responses. This point is discussed in more detail in a section that follows.

Elaboration and CHARM

A number of investigators (J. R. Anderson & Reder, 1979; Battig, 1979; Craik & Tulving, 1975; Moscovitch & Craik, 1976; Ross, 1981) have proposed that levels of processing effects may be explained by means of the construct of elaboration. In the Fisher and Craik (1977) experiment, for instance, it would be argued that when subjects are asked to process *hat*–CAT, they form fewer other associations to CAT than when they are asked to process *dog*–CAT. This interpretation of the term *elaboration* has been most explicitly formalized by J. R. Anderson and Reder (1979). A larger number of potential routes to the target item is assumed to be a concomitant of semantic processing and is assumed to facilitate recall.

An important question to ask, with respect to this explanation of levels effects, is whether having more elaborations necessarily helps recall. The answer to this question is no. There are several studies that show that elaboration may or may not facilitate recall. Fisher and Craik (1980) have shown a facilitative effect, but only under highly circumscribed conditions. Elaboration was defined, in their study, as the complexity of the sentence in which a to-be-remembered word was embedded. When, at test, the target word was presented alone, there was no effect of elaboration. Their second experiment further qualified the elaboration effect, showing that even when highly elaborative sentence frames were presented at time of test, a facilitative effect was obtained only when the elaborations were of a particular kind. The kind of elaboration that did facilitate recognition was called *highly redintegrative,* where the determining characteristic was that the sentences contained a number of items that were high associates of the target items. An example of an elaborative sentence that facilitated recognition was, "He felt *clean* after *washing* in a *hot* BATH." Given the results of Experiment 1, it would appear that the high associates shared many features with the target items. It appears possible that the nature of the items as highly similar to the target (and one another) was critical.

Bransford, McCarrell, Franks, and Nitsch (1977) conducted an experiment in which a large number of elaborations resulted in poorer recall than did a smaller number of elaborations. In the experiment, the minimal elaboration condition consisted of pairs such as "whale–skyscraper (large), mosquito–doctor (draw blood), and lamb–blanket (wool)," and the multiple-elaboration condition consisted of elaborated pairs such as "whale–deer (eat, move, sleep), mosquito–racoon (head, legs, jaws), and lamb–snake (breathe, exist, eventually die)." The similarity relations of the items to each and every item in the list are reasonably complex in this experiment. The main point for the purpose of the present argument—that elaboration does not necessarily help—is that in the more elaborated condition, recall was poorer than in the less elaborated condition.

Reder and Anderson (1980a), in investigating whether elaboration helped recall, reasoned that if elaboration helped, recall for the main points of a text would be superior when the fully elaborated text was presented, as opposed to when an unelaborated summary of the same major points was given. However, it was found that summaries of texts produced performance better than did the original elaborated text (Reder, 1982). As Reder detailed, unelaborated summaries produced better performance regardless of (a) whether direct or indirect questions were asked about the stories; (b) a delay in testing (up to one year); (c) the overall study time being greater for the elaborated conditions; (d) the nature of the test being transfer to the learning of new material, or yes/nor questions, or cued recall (Reder & Anderson, 1982), or free recall (Allwood, Wikstrom, & Reder, 1982); (e) whether subjects studied in the lab or in more naturalistic settings; or even (f) whether subjects were given credit, in the elaborated condition, for recall of noncentral points that were not even presented in the unelaborated condition.

The fundamental assumption that underlies the proposal that levels-of-processing effects are really due to more elaboration in the semantic than in the phonemic conditions, is that elaboration helps performance. Sometimes elaboration does help performance, but sometimes it hurts! It appears that the construct of elaboration is inadequate as an explanation of levels effects. Because the memorial consequences of elaboration vary, it appears rather that some explanation is required for the effects of elaboration.

An experiment by Bradshaw and Anderson (1982) is of particular interest and is modeled shortly with CHARM. Bradshaw and Anderson (1982) found both facilitative and inhibitory effects of elaboration in the same experiment. In the experiment, the extent as well as the kind of elaboration was manipulated. Three conditions are considered. In the unelaborated-control condition, subjects were presented with sentences such as "NEWTON became emotionally unstable and insecure as a child." At time of test, the predicate ("became emotionally unstable and insecure as a child") was presented and recall of the subject (NEWTON) was required. In the elaborated conditions, participants were provided with the same core sentence, and the retrieval cues were the same sentence fragments as in the control condition. In addition, though, they were provided with other statements concerning the target. In what is called the similar-elaborated condition, participants were shown sentences such as "NEWTON became irrationally paranoid when challenged by colleagues," and "NEWTON had a nervous breakdown when his mother died." In the unrelated-elaborated condition, sentences such as "NEWTON was appointed warden of London mint," and "NEWTON went to Trinity College in Cambridge," were given in addition to the core sentence.

Before turning to the data, let us consider some predictions. If it is the case that the number of elaborations determines recall and that the more elaborations there are, the better is recall, then there should be no difference between the similar-elaborated and the unrelated-elaborated conditions. This seems to be the result expected from J. R. Anderson and Reder's (1979) and J. R. Anderson's (1983) model. The pure-elaboration hypothesis indicates that the results should be similar-elaborated = unrelated-elaborated > unelaborated-control.

A propositional network representation of the difference between similar and unrelated elaborations has been given, in a different context, by Reder and Anderson (1980b). It was proposed that when the predicates connected to the same subject are related (i.e., similar), an integrative subnode is formed, intervening between the subject and the predicates that are similar to one another. This particular location of the subnode was required so the model could account for other data that are not of concern here. When the predicates are unrelated, no intervening subnode is formed. According to the propositional network view of retrieval, search begins from the terminal nodes of the fragment presented as a cue, and spreads along the experimentally marked pathways. If an integrative subnode exists between the cue predicate and the subject (target) in the similar-elaborated condition, search should spread from the cue to the subnode and from thence to both the similar elaborations and the target. In short, a fan effect should occur at the locus of the integrative subnode–decreasing the likelihood of recall of the target. Because there is no postulated subnode between the cue predicate and the target in the unrelated-elaborated condition, there should be no intervening fan effect to impair recall, and performance should be the same as would be the case had there been no elaboration. The predictions of a propositional network structure that specifies the difference between the similar- and unrelated-elaborated conditions (Reder & Anderson, 1980b) are unelaborated-control = unrelated-elaborated > similar-elaborated. In contrast to these predictions the results of the experiment were similar-elaborated > unelaborated–control > unrelated–elaborated.

Simulation 3

The essential relationships set up by the experimenters in Bradshaw and Anderson's (1982) experiment may be denoted by considering each sentence to be a subject–predicate pair. Using interference-theory notation, in the similar-elaborated condition, the list was of the form A–B, A–B′; in the unrelated-elaborated condition it was of the form A–B, A–C; in the unelaborated-control condition it was simply A–B. At time of test, B was provided as a retrieval cue in all three conditions, and recall of A was required.

In the simulations that follow, I show the results produced by CHARM when the similarity of the elaborator to the cue varies and also when the number of elaborators varies. Elaborators are either similar or unrelated to the cue. Either one or two elaborators are used. Because Bradshaw and Anderson (1982) included several delay conditions, the simulations give an imitation of delay as well.

Method. A lexicon of 18 items was constructed by randomly selecting feature values for each item from a unit normal distribution. There were 63 features in each item. The vectors were then normalized so that the dot product of an item with itself was 1.0.

The items that would serve as the similar elaborators were constructed. Items 3 and 4 were reassigned feature values on 36 randomly selected features so that they overlapped with the values of Item 2. Items 9 and 10 were reassigned values of 36 features so that they were the same as Item 8 on those features.

The traces for the unelaborated-control, the unrelated-elaborated, and the similar-elaborated conditions were constructed as follows: unelaborated control (Item 1 * Item 2) + (Item 7 * Item 8); unrelated-elaborated (1) (Item 1 * Item 2) + (Item 1 * Item 5) + (Item 7 * Item 8) + (Item 7 * Item 11); similar-elaborated (1) (Item 1 * Item 2) + (Item 1 * Item 3) + (Item 7 * Item 8) + (Item 7 * Item 9).

The extent of elaboration was also varied by using two elaborations for each target rather than just one. The unrelated-control condition was the same as above: (Item 1 * Item 2) + (Item 7 * Item 8). The unrelated-elaborated (2) trace was (Item 1 * Item 2) + (Item 1 * Item 5) + (Item 1 * Item 6) + (Item 7 * Item 8) + (Item 7 * Item 11) + (Item 7 * Item 12). The similar-elaborated (2) trace was (Item 1 * Item 2) + (Item 1 * Item 3) + (Item 1 * Item 4) + (Item 7 * Item 8) + (Item 7 * Item 9) + (Item 7 * Item 10).

Delay was mimicked by adding unrelated noise—(Item 17 * Item 18) × 4—to each of the traces detailed above. Thus, in the unelaborated-control-delay condition, for example, the trace was (Item 1 * Item 2) + (Item 7 * Item 8) + 4(Item 17 * Item 18). The extraneous noise was multiplied by 4 so that the effects would be fairly obvious.

To retrieve, in all conditions, Item 2 was correlated with the trace. The retrieved item was matched, as in the previous simulations, to all items in the lexicon except for the cue Item 2, and the item that showed the highest resonance score above a criterion of 0.5 was designated the recalled item. The simulations were replicated 100 times, and mean resonances, standard deviations, and recall frequences were calculated.

Results. Bradshaw and Anderson's (1982) results are shown in the top panel of Table 4. The center panel shows the results of Simulation 3 when one elaborating pair was used, and the bottom panel shows the results when two elaborators were used. The pattern produced by the simulation mimics the pattern produced by people. In the simulation, as in the data, the unrelated elaborators impaired recall, and the similar elaborators improved recall, relative to the unelaborated control condition. Comparison of the center and bottom panels of the table shows that as the number of elaborators increase, the differences between the elaborated conditions and the control condition also increase, according to the model.

Discussion

The similar-elaborated (2) condition in the simulation produced the best recall because the components extracted by the Cue B as well as those extracted by B'_1 (Item 3) and B'_2 (Item 4) all served to reconstruct the Target A (Item 1). In the unrelated-elaborated condition, the unrelated cues (Item 5 and 6) had no features in common, beyond independence, with the retrieval cue, and so no signal components were systematically extracted from those associations. However, the unrelated associations did produce noise components not present in the unelaborated-control condition. The additional noise in the unrelated-elaborated condition accounts for the poor recall in that condition.

Parameter Values

I discuss here the reasons for the choices of parameter values used in the above simu-

Table 4
Recall Depending on Extent and Kind of Elaboration

Experiment time	Condition: Similar	Unrelated	Control
	Bradshaw and Anderson (1982)[a]		
1 (immediate)	96%	80%	93%
1 (delay)	92%	74%	87%
2 (delay)	75%	45%	61%
3 (delay)	61%	32%	38%
	Simulation 3 (one elaborator)		
Immediate			
Recall	99%	86%	94%
Resonance			
M	1.20	0.75	0.77
SD	0.46	0.22	0.18
Delay			
Recall	69%	48%	53%
Resonance			
M	1.20	0.76	0.77
SD	0.60	0.44	0.43
	Simulation 3 (two elaborators)		
Immediate			
Recall	100%	70%	91%
Resonance			
M	1.60	0.71	0.74
SD	0.59	0.33	0.17
Delay			
Recall	88%	42%	49%
Resonance			
M	1.60	0.70	0.73
SD	0.71	0.53	0.43

[a] From "Elaborative Encoding as an Explanation of Levels of Processing" by G. H. Bradshaw and J. R. Anderson, 1982, *Journal of Verbal Learning and Verbal Behavior, 21,* p. 171. Copyright 1982 by Academic Press, Inc. Adapted by permission.

lations and some of the problems involved in constructing a realistic simulation. I have tried to either keep the values of irrelevant parameters constant, to show that the effects of interest do not interact with the values of certain parameters or to show how effects vary with variations in parameters. The simulations are minidemonstrations of the effects of changing one parameter in particular—the similarity among items. The intent of the article is to demonstrate the changes in this parameter produce interesting levels effects that are often empirically observed. The ordinal differences in the similarity parameter should be open to experimental verification (as in Experiment 1). The magnitude of the differences in similarity among conditions is problematic, though, both because of the measurement problem involved in relating subjects' ratings of similarity to the features in the model and also because the same results can be produced in the model with different numbers or proportions of common features (holding the ordinal relations the same) if other parameters are varied, such as the criterion, or the number of associations in the memory trace, or the number of features in the items. Although a number of parameters have implications for recall in the model, the present article primarily depicts a continuation of an exploration of the repercussions of manipulating one particular parameter—similarity.

There is a real problem in attempting to ascertain appropriate parameter values in the simulation model because the model is sensitive to factors and variables that are impractical to simulate. For example, the model is sensitive to the size of the lexicon. As lexicon size increases, recall falls off. The maximum size lexicon that I have simulated was 350 sixty-three feature items. All of the items were (unrealistically) random vectors. The pattern of results that was obtained with these simulations mirrored the pattern found with a small lexicon of 50 items, except that the level of recall was roughly 20% lower (see Metcalfe & Murdock, 1981, Figure 2). To realistically imitate human recall, one would like to build in a lexicon of several thousand items (words). Of course, one would not want those items to be represented as random vectors. Rather, they should vary in a priori strength (the dot product of an item with itself), in the similarities of items to one another, in the magnitude of particular features, perhaps in intercorrelations among features, and so on.

The model is sensitive to the number of features in the vectors. More features produce better recall. If one keeps the dot product of an item with itself constant, then increasing the number of features results in each feature on average having a smaller value. The overall variance is decreased, and increased accuracy results. The same net result occurs, however, if one lets the dot product increase along with the number of features but keeps the average absolute value of the features constant. In this case, the mean dot product of the retrieved item with its lexical representation increases producing increased accuracy of recall. The number of features per item that I used in the simulations was 63 and was held constant. This number was arbitrary—chosen only because I had arbitrarily chosen this number in some earlier simulations (Eich, 1982). The number may be large enough to suggest that the model is intended to deal with large numbers of features rather than the small numbers used in some feature models. I actually also ran these simulations with 21 features per item as well. Qualitatively the results are the same as those reported here although the variability was greater. A guess at the number of features that a realistic (rather than a demonstration) simulation should use would be a number of the order of magnitude of the number of neurons in, say, the hippocampus. Figure 1 in Metcalfe and Murdock's (1981) article illustrates the linear increase in the level of recall that obtains as the number of features is increased. The work of Tversky (1977) suggests that it might be reasonable, indeed necessary, to represent different items as being coded on different numbers of features (the uncoded features could be assigned a value of zero). The implications for memory of varying the number of features, across items, could be but have not yet been investigated within the CHARM model.

The dot product of an item with itself was always set to 1.0 in the simulations. Changes in this value might reflect differences in effort and arousal. Although it is well known that arousal manipulations may affect memory even within the context of levels experiments

(Krinsky & Nelson, 1981), I have not considered them in the present simulations.

The number of convolutions, or associations, that are added into the composite trace influences the level of recall. If all of the items are unrelated and represented as random vectors with a mean of zero, then recall decreases monotonically as a function of the number of pairs. In particular, if nothing is done to control the variability of the composite trace, then the expected mean dot product of the retrieved item with its correct lexical entry stays the same, but the variance increases linearly with additional associations. If the variance is controlled (by say renormalizing the trace with each entry, or by assigning appropriate weights to each new entry and to the trace, or by allowing decay, attenuation, or inhibition to occur in the model), then the decrease in the level of recall still occurs as more and more entries are added into the trace. However, in this case the poorer performance would be attributed to the mean dot product rather than to the variance.

There should probably be a parameter that gives the probability that a given association was formed and stored at all. There are many precedents for assigning a probability value to the entire process that is modeled by convolution and addition in the present model. If a subject's attention wanders, for example, the to-be-remembered event may not be processed and stored in episodic memory. There may also be neurological insults that affect the entire associative-storage mechanism. I do not rely at all on variations in the probability of association formation and storage or on the probability of enacting the retrieval process in any of the simulations that I have conducted on the CHARM model. This potential parameter is set to 1.0 in the model and so failure to store or to go through the retrieval process is not responsible for any of the effects that are demonstrated by the model in this article.

In the demonstration simulations that have just been presented, I did not set up structured lexicons of 10,000 items or use vectors with 10^6 dimensions or even use the same number of items as were used in the experiments. I used only enough associations to exemplify the designs of the experiments. For these reasons, the numerical values of the parameters that were varied—the similarity of the items to one another and the value of the response criterion—may be different from those that would be used in a more comprehensive simulation.

Simulation 1

Having done a considerable amount of hand waving about why the parameter values should not be taken too seriously and that the pattern of effects shown by the model, but not the absolute levels of recall, are important, I will now procede to admit that I put some effort, in Simulation 1, into ascertaining values of similarity and a value of the response criterion that would produce numbers that correspond to those found in Moscovitch and Craik's (1976) experiment. The major reason for fitting the data from this experiment quite closely is that a casual inspection of Equation 5 might suggest that differences in the level of recall that exceed a 2:1 ratio could have been impossible to produce with the CHARM model. Moscovitch and Craik's results exceed the 2:1 ratio. However, as the first simulation shows, the model can be fit to these data fairly readily. To do so, I used a high-response criterion. This threshold was treated as a free parameter, and the results of lowering it are shown in Table 2. The high threshold has the additional benefit over the low threshold, of almost completely eliminating intrusion errors. People do not typically produce many intrusions, so this is a pleasing spin-off of the high threshold.

Simulation 2

In Simulation 2 I left the similarity parameter of the rhyme condition at the same value it had had in the first simulation. There was about a 5% difference in the level of recall in this same condition from Moscovitch and Craik's (1976) to Fisher and Craik's (1977) experiment. I do not know why this difference occurs in the two experiments, or whether it is a meaningful difference. The materials were no doubt somewhat different, and it may have been the case that Fisher and Craik's rhyming pairs were slightly more similar to one another than were Moskovitch and Craik's rhymes. It would have been very easy to modify the value of the similarity parameter in the second experiment to produce a more precise result in the rhyme–rhyme condition, but I saw no clear justification for doing so. On the other hand, there was no logical reason for making the free-association condition in the Fisher and Craik (1977) experiment equivalent to the category condition in Moscovitch and Craik's (1976) experiment.

What the model does not do is almost as interesting as what it does, in the situation set up by Fisher and Craik (1977). The levels of recall in the noncompatible-cuing conditions were .17 and .17. As can be seen from Table 3, the simulation produces a level between .01 and .04—considerably lower than that shown by subjects. A retrieval cue that is not similar to the cue associated with the target simply does not retrieve the target item in the model. Manipulating parameters does not substantively change this result in the model. Unless the similarity between the encoded cue and the retrieval cue is quite high, poor target recall results. Because the model does not produce a high level of target recall to the extralist cue, it is necessary to claim that the high (.17) level of correct responding in the noncompatible-cuing conditions in Fisher and Craik's (1977) experiment is due to some mechanism other than episodic retrieval from the composite memory trace. Alternatively, the mechanisms in the model might be wrong. Fortunately for the model, it is defensible to claim that the responding in the noncompatible-cueing conditions is not purely episodic, and that the model, which is currently explicit only about episodic retrieval, should not account for the high level of responding. In particular, Fisher and Craik (1977) determined the a priori level of response production under conditions where there was no episodic input at all. This level was .15 in the associative-cue case, and .16 for the rhyme cues. These levels are very close to the .17 levels shown in the experiment. However, there is probably something other than pure guessing occurring to produce the .17 levels found in the experiment because the .15 and .16 levels refer to the probability that the target is one of the two to four words produced as a response to the rhyme or associative question and not just the sole response. It seems plausible to speculate that the .17 levels of responding might have resulted from some combination of guessing and semantic priming. The semantic priming could have occurred because the target word had been recently presented (but in a different context).

Simulation 3

In this simulation I attempted to show that the model would produce the right pattern of results and that increasing the extent of the elaboration manipulation, by increasing the number of pairs, would exaggerate the effects, both positive and negative. Because the conditions in this experiment were quite different from those in the previous two experiments, it did not seem necessary to maintain constant parameter values across simulations.

It may be interesting that the same pattern of facilitation and inhibition would result in the situation set up by Bradshaw and Anderson (1982) even if the events were represented somewhat differently. If each proposition were treated as a vector and the target were treated as part of that vector, then one may perform an autoconvolution ($\mathbf{A} * \mathbf{A}$) at time of storage. At time of retrieval a fragment of **A** (the predicate) suffices to reconstruct all of **A** including the target, by correlation. The goodness of reconstruction of the missing target fragment will be a function of the similarity of the elaborating propositions. The possibility of autoconvolution as an encoding operation is pursued in more detail in the section that follows on recognition memory.

Recognition and CHARM

Many experiments within the levels-of-processing framework have used recognition as the memory test, rather than recall. Recognition experiments, like recall experiments, show sensitivity to levels manipulations (e.g., Craik & Tulving, 1975; Jacoby & Dallas, 1981). Several researchers have shown that the preexperimental associative value of paired items (which, as Experiment 1 shows, is confounded with similarity) influences recognition (Kinsbourne & George, 1974; D. Nelson, Brooks, & Wheeler, 1975; Rosenberg, 1968; Underwood, 1976), although these results are not always straightforward (e.g., Fisher, 1979). In this section, I give a thumbnail sketch of a composite holographic autoassociative recognition model, and show that the recognition model produces changes in recognition depending upon context. My intention is only to show that some standard levels of processing effects can be handled in recognition within the CHARM model. A more thorough and encompassing exposition of recognition may be presented in a future article.

The CHARM recognition model outlined here had its inception in a conversation with R. M. Shiffrin (personal communication, February 22, 1983), and bears a resemblance in some respects to the recognition model of Gillund and Shiffrin (1984), as well as to the holographic recognition model of Cavanagh (1976). Following these theorists, it is assumed that recognition memory is, to some large extent, based on the autoassociation of the list items. The model bears a resemblance to those of Mandler (1980) and Murdock (1982) insofar as two kinds of information are implicated. However, neither Mandler's nor Murdock's recognition model use the convolution–correlation storage and retrieval operations. The model differs radically from Mandler's (1980) recognition model insofar as Mandler proposes a search as the retrieval process. The idea of searching is incompatible with the retrieval mechanisms in CHARM.

Autoassociation is easily written in the CHARM model as autoconvolution (i.e., $\mathbf{A}*\mathbf{A}$). Suppose that a subject is presented an A–B pair. To handle recognition as well as recall, it is assumed that not only is **A** convolved with **B**, but also that **A** is convolved with itself, and that **B** is convolved with itself. The trace, **T**, is $(\mathbf{A}*\mathbf{B}) + (\mathbf{A}*\mathbf{A}) + (\mathbf{B}*\mathbf{B})$. Very likely, there should be parameters attached to the probability of auto and interitem convolution, but they are not discussed here. It may be noted that the addition of autoassociations in the composite trace does not change the applicability of Equation 5, except that the autoconvolutions must be included, in **T**, of course.

It is assumed that when a probe is presented in recognition, the probe is correlated with the trace to produce a retrieved item. If the probe retrieves itself, then positive recognition occurs; if the probe does not retrieve itself, then recognition does not occur. If an interitem association as well as an autoassociation was formed, the trace also supports recall or a recall check in a recognition task. I deal only with the more straightforward recognition process here. The mechanics of retrieval in recognition are identical to those in recall. The decision process, however, is different. In recognition, once an item has been retrieved, the question is whether the retrieved item is the same as or different from the probe. If the probe item has been present in the list, then its autoassociation should have been added into the composite trace. If the autoassociation was added into the trace, then correlating the probe with the trace produces an approximation to the probe item (plus noise, and possibly plus other items that may be more or less similar to the probe). If the autoassociation was not added into the trace, the correlation of the probe with the trace should (if the similarity of the probe to the list items is zero) produce *only noise.* The decision process, therefore, assesses whether the item retrieved from memory is similar to or different from the probe item.

To decide whether the retrieved item is like the probe, the central n features of the retrieved item are compared with the features of the probe. If the product of a Feature f of the retrieved item and Feature f of the probe is a positive value, the item can be said to match on this feature and a positive accumulator is incremented by the value of the product. If the product is negative, the features are contrasting, and a negative evidence accumulator is incremented by the value of the product. The positive and negative incrementers race to positive and negative criteria and the one that reaches its criterion first determines whether the response is yes or no. Decision processes, similar to this one, have been used previously in recognition memory models (J. A. Anderson, 1973; Ratcliff, 1978; and see also Link, 1975; T. Wickens, 1982).

To obtain reasonable results, the criterion for the positive accumulator (for a yes response) must have a greater absolute value than the criterion for the negative accumulator (for a no response). This requirement results from the fact that the expected dot product between the item retrieved by an extralist probe and the probe is zero. If the yes and no criteria were the same, the probability would be 0.5 that a new unrelated item would incorrectly be called old. In a similar manner, if only one accumulator were used, in which both positive and negative values were added, the model would have serious problems in accounting for no, and especially fast no responses, because, even for unrelated items that were not in the list, the expected dot product is zero, rather than a negative value.

Simulation 4a

The point of this simulation was to show that a main effect of level of processing would obtain in recognition memory with the similarity of the representations coded (as indicated earlier in the article), and with the encoding, storage, retrieval, and decision processes outlined above. The simulation described below is somewhat like Simulation 1 insofar as three conditions are used: highly similar (category positive), moderately similar (rhyme positive), and unrelated (category or rhyme negative). In addition to each context–target association, each item is also convolved with itself and entered into the composite trace.

Method. A lexicon of eighteen 63-feature items was constructed by randomly selecting feature values for each item from a unit normal distribution, and then normalizing each item so that $\mathbf{F} \cdot \mathbf{F} = 1$. A highly similar pair was constructed by randomly selecting 54 features and replacing the values in Item 2 with those in Item 1. Two items, moderately similar to one another, were constructed by randomly selecting 27 features and replacing in Item 4, the values in Item 3.

The trace was (Item 1 $*$ Item 2) + (Item 3 $*$ Item 4) + (Item 5 $*$ Item 6) + (Item 1 $*$ Item 1) + (Item 2 $*$ Item 2) + (Item 3 $*$ Item 3) + (Item 4 $*$ Item 4) + (Item 5 $*$ Item 5) + (Item 6 $*$ Item 6). Thus, as well as interitem associations, as in the previous simulations, in this simulation each item was also autoassociated.

To retrieve, Items 1, 3, 5, and 7 were correlated with the trace. Item 1 is the probe for the semantic or category-positive condition; Item 3 for the shallow rhyme-positive condition; Item 5 is the probe for the rhyme- or category-negative condition; and Item 7 is an unrelated extralist lure that should be given a no response. Each of these four probes retrieved a pattern that was then compared to the probe that had been used to produce the retrieved pattern.

Table 5
Results of Simulation 4a: Recognition

Condition	Responses (%) Yes	No
Old item		
High-similarity pair	100	0
Low-similarity pair	88	12
Unrelated pair	71	29
New item (Unrelated)	17	83

There were two accumulators, one for the yes and one for the no responses. The features in the probe and the retrieved item were compared serially (presumably with each feature comparison requiring some unit of time). If the product of Feature f in the probe and Feature f in the retrieved item was positive, the yes accumulator was incremented by the value of the product. If the product was negative, the no accumulator was incremented by the absolute value of the product. The criterion of the yes accumulator was set at 0.5, whereas the no criterion was set at 0.25. The response was determined by the accumulator that reached its criterion first. The simulation was replicated 100 times. The number of yes and no responses for each of the four probes was tallied, as were the number of features compared for each response.

Results. The frequency of yes and no responses for each probe is shown in Table 5. It can be seen that the items that were members of high-similarity pairs (i.e., deeply processed items) were recognized better than were items that were members of the low-similarity pairs (i.e., shallowly processed items). The lures were usually correctly rejected. This pattern of results corresponds to experimental results. It might be possible, in some future development of the recognition model, to relate the number of features compared to reaction time (as in J. A. Anderson's, 1973, model). There were a considerable number of *no* responses with relatively few feature comparisons (i.e., fast noes), so this kind of response does not appear to be a particular problem.

A Comparison to Murdock's Recognition Model

The recognition system that I have proposed here is different from that proposed by Murdock (1982, as well as by J. A. Anderson, 1973). Murdock proposed that storage for later recognition conforms roughly to a matched-filter model. No convolution of the to-be-recognized items is conducted. Furthermore, recognition is not a retrieval process using the retrieval operation of correlation, as I propose here, but rather consists of taking the dot product between the probe and the trace. In the recognition model proposed here a representation that is an entity like the input items is retrieved from memory. Nowhere in Murdock's recognition model is an item produced that is itself a vector like the input vectors. Rather, only a spike or a signal (a scalar) is produced that varies in magnitude depending on the similarity of the probe to the trace.

I chose to model the autoassociative, retrieval-plus-decision scheme for recognition over the matched-filter scheme used by Murdock for the following reasons. First, the convolution operator can be thought of as a transformation that converts the items. When items are associated by convolution, the meaning of the features in the resultant entity is different from the meaning of the features in the initial items. In Murdock's model, features in the transform domain—the associations—are added into the composite memory trace, just as they are in the CHARM model. However, vectors that have not been so transformed and that consequently have a different status and dimensionality are also added into the same trace. These latter vectors support item recognition. The incompatibility of the associative and item information in the same trace makes Murdock's model conceptually complicated. It is not clear what the advantages are, if any, to this scheme. The first reason, then, for preferring the present proposal over Murdock's is that the information that supports recognition is of the same kind and dimensionality as is the information that supports recall. The most important reason for preferring the present scheme for recognition, over Murdock's proposal, however, is that the present model can account for the levels of processing data in recognition, whereas Murdock's recognition model does not produce these results. In particular, Murdock's model, as it currently stands, or even if similarity were represented, could not distinguish between an orienting question that is directed to the correct response and one that is directed to the wrong response. Consider List A, category animal–sheep? category flower–tulip? as compared to List B, category animal–tulip? category flower–sheep? If everything were appropriately counterbalanced, Murdock's recognition model predicts the same level of recognition for both lists whereas the CHARM model predicts better recognition in the first than in the second list. In short, the associations make no difference (except for the addition of noise) to the recognition judgment in Murdock's model. In the experiments where this situation has been examined, a difference is found. Because the CHARM recognition model proposed here has not been extensively tested in other situations, I do not yet know whether this success in accounting for the data will hold up in general or not.

Simulation 4b

This final simulation was conducted to demonstrate that the inclusion of the autoconvolutions in the composite memory trace does not qualitatively change the results obtained in recall.

Method. The lexicon and the episodic trace in this simulation were constructed in exactly the same way as in Simulation 4a. The retrieval and identification mechanisms were the same as in Simulation 1.

Results. The recall, resonance, and standard deviations of the resonance scores are shown in Table 6. As can be seen, the recall still reflects the levels manipulation, just as it did in Simulation 1, which is similar to the present simulation except that the autoassociations were not included. The resonance scores are much inflated (as are the standard deviations). These increases are due to the autoassociations. However, the ordinal pattern

Table 6
Simulated Recall Using the Recognition Memory Trace From Simulation 4

Type of pair	%	*M*	*SD*
High similarity			
Recall	100		
Resonance		3.78	0.93
Low similarity			
Recall	92		
Resonance		2.26	0.89
Unrelated			
Recall	30		
Resonance		0.87	0.75
Extralist lure[a]			
Intrusion	3		
Resonance		0.05	0.54

[a] This condition gives an approximation to the intrusion rate in the simulation for an extralist cue retrieving a particular lexical item.

is still the same. The proportion recall results look rather like the low criterion version of Simulation 1, and as in that simulation, there were a good number of intrusions. The inclusion of the autoconvolutions in the composite trace does not qualitatively alter the biasing effect due to context, that is found in recall. Nor does it qualitatively alter the differences in the goodness of recall that are attributable to this biasing.

Simulations 4a and 4b provide a demonstration that the transformational characteristics of the holographic trace, which give rise to levels-of-processing effects in recall, could also produce those effects in recognition. In the next section I discuss the relation of the CHARM model to a number of other constructs that have been proposed to facilitate our understanding of levels-of-processing effects.

General Discussion

Comparison of CHARM to Other Theoretical Ideas

Depth of processing. The view proposed in the CHARM model is very different from that proposed by Craik and Lockhart (1972). In CHARM there is a level or layer of representation at which events are or may be available to consciousness; there is another layer that is the (composite) episodic memory trace; and, there are defined operations for getting from one layer, state, or stage, to another. CHARM is thus very much a stage theory of memory. It seems to be in the same general class as the stage theories Craik and Lockhart (1972) argued against, and to which they proposed the depth-of-processing theory as a conceptual alternative.

In CHARM it is assumed that a number of sensory and perceptual processes occur before the event in question reaches the form of representation (at the item level) necessary for consciousness and for association formation. The preconscious operations are not assumed to impact on episodic memory. This model fits nicely with the empirical studies of Jacoby and Dallas (1981), and Jacoby and Witherspoon (1982), although not necessarily with their interpretation. Craik and Lockhart (1972) placed considerable stress on perceptual and memorial processes. In CHARM, too, the processes used for associating, storing in memory, and retrieving (i.e., convolution, addition, and correlation) are of prime importance. However, it is not considered to be the case in CHARM, as seems to be implied by the depth of processing theory, that all mental operations are memorial operations. CHARM has nothing to say about many cognitive functions such as mental arithmetic, sentence parsing, syllogism construction, or indeed, even the memorial impact of similarity, category, or rhyme decision processes. CHARM deals only with the operations considered necessary for association formation, storage, and retrieval. My argument has been that for a large number of levels effects, that is all one needs to consider.

The episodic memory system in CHARM is envisaged as a rather specialized faculty. This system may interact with other cognitive systems; the changes and distortions in representation that are introduced by the memory system, and that have been the major concern of this article, probably impact on functions other than episodic remembering. Nevertheless, the episodic memory system can be thought of as functionally, and perhaps anatomically, distinct. Though holographic models are in the class of distributed models, this in no way precludes localization of the association or retrieval functions themselves, as Figure 1 perhaps suggests. The CHARM model is compatible with *function*-al or *operation*-al localization where the operations are convolution and correlation. It is incompatible with the idea that particular discrete episodic memory traces of individual items or events are localized in the sense of being stored separately. Though perceptual analysis is, of course, assumed by the CHARM model, it is not considered to be the case that episodic memory is only an automatic by-product of perceptual processing.

Elaboration. The CHARM model does not implicate elaboration as an explanation of levels effects as do propositional network models. Rather, elaboration, if and when it occurs, is (like rehearsal) considered to be data requiring explanation (Cavanagh, 1976; Metcalfe & Murdock, 1981). The question with elaboration, as with recall and rehearsal, is how people are able to retrieve information that is not in consciousness, regardless of whether that retrieval occurs at time of study or test. A quote from Nisbett and Wilson (1977) helps clarify the point.

> An example of the confusion of intermediate output with process was provided by an acquaintance . . . who was asked to introspect about the process by which he had just retrieved from memory his mother's maiden name. "I know just what the process was," he said. "I first thought of my uncle's last name, and since that happens to be my mother's maiden name, I had the solution." This only pushes the process question back a step further, of course, and our acquaintance's answer would appear to reflect a confusion of intermediate results with the process by which the final result was obtained. (p. 255)

Distinctiveness. A number of researchers have stressed that the distinctiveness of the to-be-remembered events may influence the memorability of those events (Craik & Jacoby, 1979; Eysenck, 1979; Klein & Saltz, 1976; Lockhart, Craik, & Jacoby, 1976; D. Nelson, 1979). Distinctiveness is not, in and of itself, considered to be an explanation in CHARM. Though distinctiveness sometimes helps performance, the model predicts that it sometimes hurts performance. For instance, if the cue and target are distinctively different from one another, the model predicts worse recall than if the cue and target are less distinctive and more similar to one another. Because associations are stored in a composite memory trace, the model provides a mechanism whereby the similarity or lack of similarity of items, not only within a pair but also between pairs, can have memorial repercussions. The representations of the items in combination with the associative, storage, and retrieval mechanisms in CHARM potentially lead to what have been called distinctiveness effects.

Compatibility. The explanation of depth of processing given in the present article bears a relation to the idea of cue-trace compatibility (Tulving, 1979), congruence (Craik & Tulving, 1975; Schulman, 1974), and similarity (D. Nelson, 1979). The model also produces the results subsumed under the construct of encoding specificity (Tulving & Thomson, 1973). In agreement with these ideas, the nature of the materials and the context in which the to-be-remembered event is encoded and retrieved are crucial. The CHARM model is firmly footed in this interactive approach to human memory and contributes a mechanism whereby the similarity intrinsic in the presented events manifests memorial consequences.

Although the effects of similarity demonstrated in this article were facilitative, similarity in the model does not always produce improved performance. I do not use similarity, per se, as an explanation. Rather, the representation of each item is considered to be intrinsic to that item. Without some model saying how representations interact, or fail to interact, it would not be possible to predict what kind of effects similarity should have. (It is easy to construct models that do not produce the effects shown in the data.) CHARM provides an interactive model that produces the results found in the data. Predictions are derivable from the model and so it is open to falsification. By providing a working mechanism, CHARM is able to sidestep the circularity of the compatibility idea, while at the same time maintaining and formalizing the interactiveness that is its cornerstone.

Transformation of Information

I have argued in this article that many of the levels effects that have been reported in the literature can be explained by means of the operations used for association formation, storage and retrieval in the CHARM model, and by the a priori representations of the to-be-remembered and context items. These two threads were tightly woven during the development of the argument. The fact that the CHARM model would produce differences in the level of recall depending on interpair and, especially, intrapair similarity led me to look for situations in which these similarities might be a factor in the observed performance. Many levels of processing experiments, as well as encoding specificity and elaboration experiments, seem to exemplify these differences. Even though the two threads of the argument—the processes postulated by the model and the representational structure of the information—are both necessary in some form, it is possible to tease them apart. In this final section of the article, I try to do so.

Process. It seems reasonable to suppose that there exists a class of models, perhaps involving different mechanisms, that like CHARM may produce the results of interest in this article, given the representational constraints that I have argued should be placed on the to-be-remembered items. Probably the most important condition that must be met by a model in this class is that it produce Equation 5 or something reasonably close to it. The association is bidirectional and symmetrical. It may produce either or both of the associated items. All of the items that are produced are added in a single representation.

At the present time, to my knowledge, there are no other models of human memory (except Murdock's TODAM model, for recall but not for recognition) that automatically produce Equation 5. However, it would be possible to modify some extant memory models to do so.[3] For example, reduplicating the directional association that is used in other models (e.g., J. A. Anderson et al., 1977; Hintzman, 1983; Medin & Schaffer, 1978) by adding a second operation that stores in memory a backward as well as the forward association, might allow those models to approximate Equation 5. I do not advocate this step because I think that it fundamentally undermines a basic concept that is implicit in the directional models, namely, their directionality. It seems most likely that there are some mental processes of types or information processing for which a unidirectional association is appropriate and even required. It would be interesting to use the directional versus symmetrical associative properties and the implications of these properties as a litmus test to assist in taxonomizing different tasks and memory systems. Allowing ad hoc mechanisms to be added to elegant unidirectional associative mechanisms obscures their proper predictions, and should, I think, be avoided. However, I do not deny the possibility that such a modification could produce Equation 5. Nor do I deny the more interesting possibility that it may be possible to develop more elegant or more empirically defensible symmetrical associative schemes than those used in the CHARM model, that might produce Equation 5 and results like those outlined here.

Structure. The nature of the to-be-remembered items is of considerable importance in the CHARM model. If the thesis of this article is correct, it is also crucial for many levels-of-processing, encoding specificity, and elaboration experiments. I have proposed that the content of what is associated, stored, and retrieved (in combination with the processes in the model) may result in different levels of recall.

Just as there may be several different processes that are capable of producing Equation 5, there is no doubt also a class of different types of representations that support this equation. In the CHARM model, feature-set representations are used. These representations are probably the simplest and most easily understood means of representing mental events for which Equation 5 is meaningful. However, it is possible to represent items more complexly—as matrices, for example—and still maintain the properties of interest. It does not seem possible to represent items as undifferentiated points in multidimensional space (as has been done in some node-search models) and still show the facilitative and inhibitory effects of similarity in a manner consistent with the present formulation. The CHARM model, and Equation 5, depends for its predictive ability on adding the values of certain features to result in stronger or weaker feature values that may influence pattern identification and recall. Addition of feature values is fundamental to

[3] One example of this strategy has been advanced and brought to the attention of the author following the writing and acceptance of the present article.

the model. The concept of addition does not seem to be meaningful, or at least not meaningful in a way that makes predictions about recall under the representational assumption of undifferentiated unitary concept nodes that exist as points in multidimensional space. Thus, the present explanation of levels-of-processing, encoding specificity, and elaboration data allows for some flexibility in the way items are represented but is incompatible with node-search conceptualizations of human memory.

Summary

In this article I detailed some implications of a model of human episodic memory that combines and modifies events. The result of retrieval differs from the item that was originally encoded. The transformations that occur depend on the similarities among the to-be-remembered items, the association and storage of items by convolution and addition, and the representations of the cues that are used to retrieve by correlation. One important point to be taken from the model is this: Both the intrinsic similarity among items—the representational structure—and the associative operations that subjects perform to interact items—the cognitive process—are important and neither of these two ideas is reducible to the other.

It was proposed that the representational transformations, which occur because of the episodic memory mechanisms in the holographic model, CHARM, have several consequences. Among the most studied are some recall and recognition effects that have often been attributed to the depth of processing of the items. These effects, as well as encoding-specificity effects and both facilitative and inhibitory effects of elaboration, were explained by the intrinsic similarities among items and the highly interactive associative, storage, and retrieval mechanisms in the CHARM model.

References

Allwood, C. M., Wikstrom, T., & Reder, L. M. (1982). Effects of presentation format on reading retention: Superiority of summaries in recall. *Poetics, 11,* 145–153.

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

Anderson, J. A. (1977). Neural models with cognitive implications. In D. LaBerge & S. J. Samuels (Eds.), *Basic processes in reading* (pp. 27–90). Hillsdale, NJ: Erlbaum.

Anderson, J. A., & Hinton, G. E. (1981). Models of information processing in the brain. In G. E. Hinton & J. A. Anderson (Eds.), *Parallel models of associative memory* (pp. 9–48). Hillsdale, NJ: Erlbaum.

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

Anderson, J. R. (1972). FRAN: A simulation model of free recall. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 5, pp. 315–378). New York: Academic Press.

Anderson, J. R. (1983). A spreading activation theory of memory. *Journal of Verbal Learning and Verbal Behavior, 22,* 261–295.

Anderson, J. R., & Reder, L. M. (1979). An elaborative processing explanation of depth of processing. In L. S. Cermak & F. I. M. Craik (Eds.), *Levels of processing in human memory* (pp. 385–404). Hillsdale, NJ: Erlbaum.

Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. In K. W. Spence & J. T. Spence (Eds.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 2, pp. 89–195). New York: Academic Press.

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

Battig, W. F. (1979). The flexibility of human memory. In L. S. Cermak & F. I. M. Craik (Eds.), *Levels of processing in human memory* (pp. 23–44). Hillsdale, NJ: Erlbaum.

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

Bowerman, M. (1977). The acquisition of word meaning: An investigation of some current concepts. In P. N. Johnson-Laird & P. C. Wason (Eds.), *Thinking* (pp. 239–253). Cambridge: Cambridge University Press.

Bradshaw, G. H., & Anderson, J. R. (1982). Elaborative encoding as an explanation of levels of processing. *Journal of Verbal Learning and Verbal Behavior, 21,* 165–174.

Bransford, J. D., Franks, J. J., Morris, C. D., & Stein, B. S. (1979). Some general constraints on learning and memory research. In L. S. Cermak & F. I. M. Craik (Eds.), *Levels of processing in human memory* (pp. 331–354). Hillsdale, NJ: Erlbaum.

Bransford, J. D., McCarrell, N. S., Franks, J. J., & Nitsch, K. E. (1977). Toward unexplaining memory. In R. E. Shaw & J. D. Bransford (Eds.), *Perceiving, acting, and knowing: Toward an ecological psychology* (pp. 431–466). Hillsdale, NJ: Erlbaum.

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

Cermak, L. S., & Craik, F. I. M. (1979). *Levels of processing in human memory.* Hillsdale, NJ: Erlbaum.

Clark, E. (1973). What's in a word? On the child's acquisition of semantics in his first language. In T. E. Moore (Ed.), *Cognitive development and the acquisition of language* (pp. 65–110). New York: Academic Press.

Craik, F. I. M. (1979). Levels of processing: overview and closing comments. In L. S. Cermak & F. I. M. Craik (Eds.), *Levels of processing and human memory* (pp. 447–462). Hillsdale, NJ: Erlbaum.

Craik, F. I. M., & Jacoby, L. L. (1979). Elaboration and distinctiveness in episodic memory. In L.-G. Nilsson (Ed.), *Perspectives on memory research: Essays in honor of Uppsala University's 500th anniversary* (pp. 145–166). Hillsdale, NJ: Erlbaum.

Craik, F. I. M., & Lockhart, R. S. (1972). Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior, 11,* 671–684.

Craik, F. I. M., & Tulving, E. (1975). Depth of processing and the retention of words in episodic memory. *Journal of Experimental Psychology: General, 104,* 268–294.

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

Estes, W. K. (1979). 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* (pp. 35–60). Hillsdale, NJ: Erlbaum.

Eysenck, M. W. (1979). Depth, elaboration, and distinctiveness. In L. S. Cermak & F. I. M. Craik (Eds.), *Levels of processing and human memory* (pp. 89–118). Hillsdale, NJ: Erlbaum.

Fisher, R. P. (1979). Retrieval operations in cued recall and recognition. *Memory & Cognition, 7,* 224–231.

Fisher, R. P., & Craik, F. I. M. (1977). Interaction between encoding and retrieval operations in cued recall. *Journal of Experimental Psychology: Human Learning and Memory, 3,* 701–711.

Fisher, R. P., & Craik, F. I. M. (1980). The effects of elaboration on recognition memory. *Memory & Cognition, 8,* 400–404.

Gabor, D. (1968). Holographic model of temporal recall. *Nature, 217,* 584.

Galton, F. (1879). Composite portraits, made by combining those of many different persons into a single resultant figure. *Journal of the Anthropological Institute, 8,* 123–144.

Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review, 91,* 1–67.

Hintzman, D. L. (1983, June). *"Schema abstraction" in a multiple-trace memory model.* Paper presented at the Conference on Priority of the Specific, Elora, Ontario, Canada.

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

Jacoby, L. L., & Craik, F. I. M. (1979). Effects of elaboration of processing at encoding and retrieval: Trace distinctiveness and recovery of initial context. In L. S. Cermak & F. I. M. Craik (Eds.), *Levels of processing and human memory* (pp. 1–22). Hillsdale, NJ: Erlbaum.

Jacoby, L. L., & Dallas, M. (1981). On the relationship between autobiographical memory and perceptual learning. *Journal of Experimental Psychology: General, 110,* 306–340.

Jacoby, L. L., & Witherspoon, D. (1982). Remembering without awareness. *Canadian Journal of Psychology, 36,* 300–324.

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

Kinsbourne, M., & George, J. (1974). The mechanism of the word-frequency effect on recognition memory. *Journal of Verbal Learning and Verbal Behavior, 13,* 63–69.

Klein, K., & Saltz, E. (1976). Specifying the mechanisms in a levels-of-processing approach to memory. *Journal of Experimental Psychology: Human Learning and Memory, 2,* 671–679.

Kohonen, T., Oja, E., & Lehtio, P. (1981). Storage and processing of information in distributed associative memory systems. In G. E. Hinton & J. A. Anderson (Eds.), *Parallel models of associative memory* (pp. 105–143). Hillsdale, NJ: Erlbaum.

Krinsky, R., & Nelson, T. O. (1981). Task difficulty and pupillary dilation during incidental learning. *Journal of Experimental Psychology: Human Learning and Memory, 7,* 293–298.

Liepa, P. (1977). *Models of content addressable distributed associative memory (CADAM).* Unpublished manuscript, University of Toronto.

Link, S. W. (1975). The relative judgment theory of choice reaction time. *Journal of Mathematical Psychology, 12,* 114–135.

Lockhart, R. S. (1979). Remembering events: Discussion of papers by Jacoby and Craik, Battig, and Nelson. In L. S. Cermak & F. I. M. Craik (Eds.), *Levels of processing and human memory* (pp. 77–88). Hillsdale, NJ: Erlbaum.

Lockhart, R. S., Craik, F. I. M., & Jacoby, L. L. (1976). Depth of processing, recognition, and recall: Some aspects of a general memory system. In J. Brown (Ed.), *Recall and recognition* (pp. 75–102). London: Wiley.

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

Mandler, G. (1980). Recognizing: The judgment of previous occurrence. *Psychological Review, 87,* 252–271.

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

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

Moscovitch, M., & Craik, F. I. M. (1976). Depth of processing, retrieval cues, and uniqueness of encoding as factors in recall. *Journal of Verbal Learning and Verbal Behavior, 15,* 447–458.

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

Murdock, B. B., Jr. (1983). A distributed memory model for serial-order information. *Psychological Review, 90,* 316–338.

Nelson, D. L. (1979). Remembering pictures and words: Appearance, significance, and name. In L. S. Cermak & F. I. M. Craik (Eds.), *Levels of processing and human memory* (pp. 45–76). Hillsdale, NJ: Erlbaum.

Nelson, D. L., Brooks, D. H., & Wheeler, J. W., Jr.

(1975). Sensory and meaning features in stimulus recognition and associative retrieval. *Journal of Experimental Psychology: Human Learning and Memory, 1,* 711-719.

Nelson, D. L., Fosselman, J., & Peebles, J. (1971). Words as phonological sequences. *Journal of Experimental Psychology, 87,* 361-366.

Nelson, D. L., Wheeler, J. W., Jr., & Brooks, D. H. (1976). Meaning and the elimination of sensory interference. *Journal of Experimental Psychology: Human Learning and Memory, 2,* 95-102.

Nelson, K. (1974). Concept, word, and sentence: Interrelations in acquisition and development. *Psychological Review, 81,* 267-285.

Nelson, T. O. (1977). Repetition and depth of processing. *Journal of Verbal Learning and Verbal Behavior, 16,* 151-172.

Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review, 84,* 231-259.

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

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

Raiijmakers, J. G. W., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review, 88,* 93-134.

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

Reder, L. M. (1982). Elaborations: When do they help and when do they hurt. *Text, 2,* 211-224.

Reder, L. M., & Anderson, J. R. (1980a). A comparison of texts and their summaries: Memorial consequences. *Journal of Verbal Learning and Verbal Behavior, 19,* 121-134.

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

Reder, L. M., & Anderson, J. R. (1982). Effects of spacing and embellishment on memory for the main points of a text. *Memory & Cognition, 10,* 97-102.

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

Rosch, E. (1977). Classification of real-world objects: Origins and representations in cognition. In P. N. Johnson-Laird & P. C. Wason (Eds.), *Thinking* (pp. 212-222). Cambridge, England: Cambridge University Press.

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

Rosenberg, S. (1968). Associative facilitation in the recall and recognition of connected discourse. *Journal of Experimental Psychology, 78,* 254-260.

Ross, B. H. (1981). The more, the better?: Number of decisions as a determinant of memorability. *Memory & Cognition, 9,* 23-33.

Schulman, A. I. (1974). Memory for words recently classified. *Memory & Cognition, 2,* 47-52.

Smith, E. E., & Medin, D. L. (1981). *Categories and concepts.* Cambridge, MA: Harvard University Press.

Sokal, R. R. (1977). Classification: Purposes, principles, progress, prospects. In P. N. Johnson-Laird, & P. C. Wason (Eds.), *Thinking* (pp. 185-198). Cambridge, England: Cambridge University Press.

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

Treisman, A. (1979). The psychological reality of levels of processing. In L. S. Cermak & F. I. M. Craik (Eds.), *Levels of processing in human memory* (pp. 301-330). Hillsdale, NJ: Erlbaum.

Tulving, E. (1979). Relation between encoding specificity and levels of processing. In L. S. Cermak & F. I. M. Craik (Eds.), *Levels of processing in human memory* (pp. 405-428). Hillsdale, NJ: Erlbaum.

Tulving, E. (1983a). *Elements of episodic memory.* Oxford, England: Clarendon Press.

Tulving, E. (1983b). Ecphoric processes in episodic memory. *Philosophical Transactions of the Royal Society of London, 302*(B), 361-371.

Tulving, E., & Osler, S. (1968). Effectiveness of retrieval cues in memory for words. *Journal of Experimental Psychology, 77,* 593-601.

Tulving, E., & Thomson, D. M. (1973). Encoding specificity and retrieval processes in episodic memory. *Psychological Review, 80,* 352-373.

Tversky, A. (1977). Features of similarity. *Psychological Review, 84,* 327-352.

Underwood, B. J. (1976). Recognition memory for pairs of words as a function of associative context. *Journal of Experimental Psychology: Human Learning and Memory, 2,* 404-412.

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

Wickens, D. D. (1972). Characteristics of word encoding. In A. W. Melton & E. Martin (Eds.), *Coding processes in human memory* (pp. 191-215). Washington, DC: Winston.

Wickens, T. D. (1982). *Models for behavior: Stochastic processes in psychology.* San Francisco: W. H. Freeman.

Willshaw, D. (1981). Holography, associative memory, and inductive generalization. In G. E. Hinton & J. A. Anderson (Eds.), *Parallel models of associative memory* (pp. 83-104). Hillsdale, NJ: Erlbaum.

Wittgenstein, L. (1958). *Philosophical investigations* (2nd ed.; *G.E.M.* Anscombe, Trans.). Oxford, England: Blackwell.

Received September 12, 1983
Revision received May 15, 1984 ■