The Effect of List Separation in the Process Dissociation Procedure: The Bind Cue Decide Model of Episodic Memory

Simon Dennis (s.dennis@psy.uq.edu.au)
School of Psychology; University of Queensland
Brisbane, 4072, Australia

Michael Humphreys (mh@psy.uq.edu.au)
School of Psychology; University of Queensland
Brisbane, 4072, Australia

Abstract

The Process Dissociation Procedure as applied to episodic recognition requires subjects to study two lists and then determine which of the words in a test list appeared in the second list (Exclusion condition) or on either list (Inclusion condition). We demonstrate that the dual processing account of episodic recognition (Jacoby 1991) does not account for the effects of manipulating the amount of time between the study lists. In contrast, the Bind Cue Decide Model of Episodic Memory (BCDMEM) is fit to the list separation data.

Introduction

Episodic recognition refers to the task of identifying a stimulus as having occurred within a particular episode or context. In a typical recognition experiment, subjects process a study list of items and are then presented with a test list containing some old items from the study list and some new items which did not appear in the study list. The subject's task is to determine which of the test items appeared in the study list. This basic design can be elaborated in a number of ways by adding additional study lists and requiring subjects to recognize items from individual lists or from all of the lists.

In this paper, we contrast two models of episodic memory: the dual processing model (Jacoby 1991, Yonelinas 1994) and the Bind Cue Decide Model of Episodic Memory (BCDMEM, Dennis & Humphreys in press, Dennis & Humphreys submitted). Firstly, we outline the dual-processing model and describe the process dissociation procedure. Then we outline BCDMEM. Finally, we demonstrate that data on the effect of manipulating study list separation in the PDP are problematic for the dual-processing model, but are well captured by BCDMEM.

The Dual Processing Model and the Process Dissociation Procedure

The Dual Processing model of episodic recognition hinges on the distinction between automatic and controlled processes (sometimes denoted conscious and unconscious processes, intentional and unintentional processes or aware and unaware processes). Automatic processes produce a feeling of **familiarity** which, in the context of episodic recognition, tends to evoke an old response independent of subject control. This familiarity is not specific to an individual list context and is thought to be the same information on which subjects base decisions in implicit tasks such as lexical decision and perceptual identification. In contrast, controlled processes are based on **recollection** retrieving some aspect of the study opportunity that can be used to infer that the item appeared in the list (i.e. "I remembering solving an anagram for this item so this item must have been in the study list"). This source of information is subject to control by the participant and is thought to be vulnerable to disruption at both storage and retrieval.

To separate the influences of familiarity and recollection in episodic recognition, Jacoby (1991) has used a number of procedures. In this paper, we focus on the Process Dissociation Procedure (PDP), which was introduced in the third experiment. In a typical application of the PDP, a subject studies two lists and is asked to make one of two recognition decisions at test. In the *inclusion* condition, they are required to respond "old" if the item was in either of the two lists. In the *exclusion* condition, they are required to respond "old" only if the item appear in the target list (either list 1 or list 2). In Jacoby's third experiment, list one contained read and anagram words and list two contained heard words. Subjects were asked either to recognize words from both lists (inclusion) or from list two only (exclusion).

To estimate the independent contributions of familiarity and recollection, Jacoby assumes that in the inclusion condition the subject will respond old if the item is familiar or if the item is unfamiliar but is recollected:

$$P\,(\text{Inclusion}) = F + (1\text{-}F)\,R$$

In the exclusion condition, when the subject should be responding new, it is assumed that they will mistakenly respond old if the item is familiar but they fail to recollect which list it occurred in:

$$P\,(\text{Exclusion}) = F\,(1\text{-}R)$$

These equations can be solved for R and F giving:

R = P (Inclusion) - P (Exclusion)

F = P (Exclusion) / (1 P (Inclusion) + P (Exclusion))

The estimates of familiarity generated using the equations outlined above have been shown not to vary in a number of manipulations that have an effect on recollection estimates (Jacoby 1991, Yonelinas 1994) suggesting that the procedure was providing estimates of independent familiarity and recollection processes.

The dual-processing framework has generated a number of important insights. Firstly, it suggests that episodic recognition involves more than one process and that the application of these processes can vary depending on the conditions, in particular, whether attention is divided.

Secondly, the dual-processing framework has focused attention on multi-list paradigms. Because previous approaches to episodic memory were primarily sensitive to manipulations of the other items in the target list, efforts to test these models have focused on single list paradigms. However, it is only possible to be certain you are dealing with the episodic memory system when multi-list paradigms are employed and so the process dissociation procedure can provide useful data for testing how models approach the use of contextual information.

In the next section, we outline the Bind Cue Decide Model of Episodic Memory.

The Bind Cue Decide Model of Episodic Memory

As the name suggests there are three critical components to BCDMEM - the binding mechanism, the cues employed and the decision rule. The binding mechanism specifies how elements of an episode including items, contexts and other information are associated in episodic memory. The cues are the elements that are used to initiate retrieval. Not all of the information available to a subject need be used as a cue, so part of the theory involves specifying which components are used as cues in a given experimental paradigm. Finally, the decision rule takes the results of retrieval and outputs the required information in the form of an item in recall paradigms or a yes/no decision in recognition paradigms. A complete theory of episodic memory must address these components and in the following sections, the each will be considered in turn.

The Binding Mechanism

Figure 3 depicts the BCDMEM architecture, which consists of three layers of units. Items and contexts (e.g., "the list of words I saw today" or "the list in which I was forming anagrams") are represented as sparse binary distributed vectors.

The components present at a given moment are added to form an input vector (a composite encoding). For instance, if a subject were studying the pair "pencil grass" in the list1 context (i.e. the study list) the input vector would be PENCIL+GRASS+LIST1 where PENCIL and GRASS are sparse vectors representing the items and LIST1 is a sparse vector representing the context (see Figure 1).

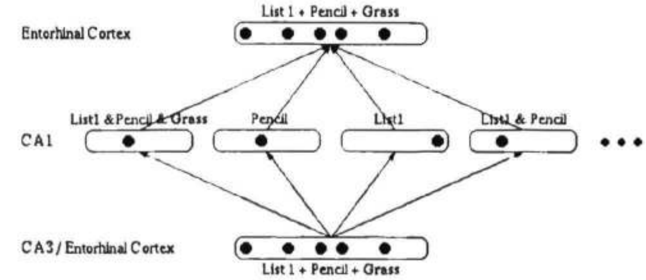

Figure 1: The BCDMEM architecture. The layer names are intended to give an indication of the hippocampal structures that might be involved.

Bindings are formed in the middle layer (the binding or CA1 layer), which is assumed to be a set of pools of competitive units. Each pool will contain a unit representing either a single item or a combination of the input items (a conjunctive encoding - see Figure 1) that will be reactivated if any of those components are presented at a later time. It is assumed that a node will only be reactivated if the current input pattern is very similar to the input pattern to which it responds. So, for instance, the plural form of a word may reactivate the same node as its singular form, but semantically similar items would reactivate different units. In the "pencil grass" example, bindings would be formed representing LIST1 & PENCIL & GRASS, LIST1 & PENCIL, LIST1 & GRASS, PENCIL & GRASS, LIST1, PENCIL and GRASS. The relative proportions of these bindings will depend on factors such as the sparsity of connectivity between CA3/Entorhinal cortex and CA1 as well as the sparsity of the item and context representations.

Finally, the system is assumed auto-associative, so that the input layer is also the output layer. Associative weights connect the binding layer to the output layer forming the memory of the system. For modeling purposes these weights are considered to be binary (either they form or not) and will be present with a probability that depends on the amount of time and attention employed.

In constructing the binding mechanism we have attempted to produce the simplest architecture that captures three key behavioural constraints. The conjunctive pools allow the system to bind three-way information. The bindings are symmetric so that learning the pair AB will allow A to be used to retrieve B and B to be used to retrieve A to approximately the same extent. And, the bindings are formed rapidly rather than as the result of a large number of repetitions. It would be possible to produce a more complicated model that described the sparsity of the connections and provided a more realistic account of learning. Such a model will be necessary to capture the affect on binding formation of training (c.f. Chalmers &

Humphreys in press). However, at this stage, there is insufficient data to constrain the neural mechanism and so we have chosen a simple option that embodies the existing behavioural constraints.

The Cues

The cues presented to the memory system at test play an important role in determining the sources from which interference will arise. In BCDMEM, it is assumed that recognition paradigms primarily involve cueing with an item or items (Anderson & Bower 1972, Dennis & Humphreys, in press), while recall paradigms primarily involve cueing with a context. In recognition, interference is generated mainly from the other contexts in which the item has been seen. In recall, interference is generated mainly from the other items from the same context (i.e., in the same list). Of course, circumstances may predispose a subject to use a different cueing strategy (for instance, in some associative recognition experiments subjects may employ recall, Clarke & Gronlund 1996). Making this general distinction, however, provides insight into a number of dissociations between recall and recognition including word frequency effects and the null list-strength effect (Dennis & Humphreys in press, Dennis & Humphreys submitted).

Cueing with the item for contexts in recognition is a major departure from previous approaches, where it is uniformly assumed that the primary source of interference is the other list items. However, we believe that cueing with the item is more consistent with the intuitions behind the dual processing approach. Recollection involves the retrieval of contextual elements associated with the item, as if the item were being used as a cue. Furthermore, familiarity is a property of an item, as if retrieval is starting with the item. So for both familiarity and recollection, the assumption is that subjects focus first on the item, as is the case in BCDMEM.

The Decision Rule

The final essential component of the memory mechanism is the decision rule, which selects the response of the model (Humphreys, Wiles & Dennis 1994). In recall paradigms, the decision rule selects an item for output. In yes/no recognition paradigms, the decision rule selects an old or new response. In BCDMEM, we assume that the neural mechanisms underpinning these decisions will approximate the optimal decision rule in a Bayesian sense (c.f. Anderson & Milson 1989, Glanzer & Adams 1990, Shiffrin & Steyvers 1997, McClelland & Chappell in press). In recognition, the decision can be characterized as the odds ratio:

$$\frac{p(old \mid data)}{p(new \mid data)}$$

where the data is the evidence retrieved from the memory system.

Figure 2 outlines the components of the architecture of BCDMEM relevant to recognition. At study, an item is represented as a single node (i.e. a local code) in the binding layer. Each node at the binding layer is connected via associative weights to the memory layer which contains a distributed representation of the context (with sparsity s and length l) in which the item is being studied. These weights are originally zero, but are set to one with probability r (amount of learning) if the corresponding component of the context vector is one. The item may also have appeared in non-study contexts so that there will also be weights that have been learned during these episodes. The probability that a weight is one as a result of previous learning is called memory noise (p).

In a single item recognition test, the item node is reactivated which in turn reinstates at the memory layer a binary addition (the bit-wise OR operation) of the contexts in which this item has been seen weighted by the amount of learning in those contexts (see Figure 2). In addition, the context vector that was present at study is reconstructed. For a target item, the binary addition will contain the study context whereas for a distracter it will not. There may, however, be components of the study context that are activated even when a distracter is presented as a consequence of the overlap between the study context and the contexts in which the distracter has been seen. Likewise, the target pattern may be missing components of the study context because they were not learned.

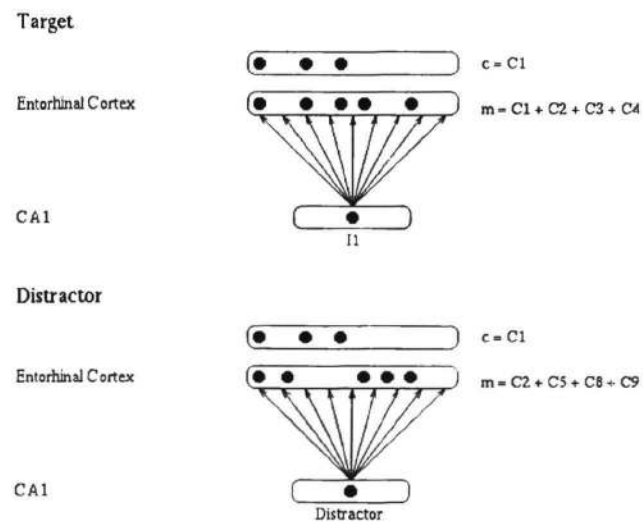

Figure 2: The BCDMEM architecture for recognition.

The decision rule must discriminate outputs that contain the target pattern from those that do not. It is in making the discrimination that the majority of the noise is added to the system. Following work by Glanzer & Adams (1990), Shiffrin and Steyvers (1997) and McClelland and Chappell (in press) on the nature of retrieval in recognition, we have specified the discrimination problem as a Bayesian procedure and have derived the expressions for the likelihood of a yes or no response given the history of the item.

The odds ratio, P (old|data)/P (new|data), can be rewritten as follows:

$$\frac{p(old \mid data)}{p(new \mid data)} = \frac{p(old)}{p(new)} \frac{p(data \mid old)}{p(data \mid new)}$$

The data referred to in the likelihood expression are the vector retrieved from memory (m) and the reinstated context vector (c). The probabilities depend on how well these vectors match. The item will be considered old if the probability that it is old given the data is greater than the probability that it is new given the data (i.e. if the odds ratio exceeds one).

In many experiments subjects see equal numbers of targets and distracters, so that it could be assumed that in the absence of specific manipulations of criterion, P (old) = P (new) = 0.5, and the prior probabilities cancel. In this case, the odds ratio is equal to the likelihood ratio (i.e. P (data|old)/P (data|new)).

Since both the context vector and the memory vector are binary there are four types of match (i.e. $c_i=1$ & $m_i=1$, $c_i=1$ & $m_i=0$, $c_i=0$ & $m_i=1$, $c_i=0$ & $m_i=0$). The probability of a given sort of match is independent of the component in which that match occurs, so the data can be summarized by n_{jk}: the number of components in which $c_i = j$ & $m_i = k$.

Now,

$$\begin{aligned} P(data|old) &= \Pi_i P(c_i \,\&\, m_i \mid old) \\ &= P(c_i=1 \,\&\, m_i=1 \mid old)^{n11} P(c_i=0 \,\&\, m_i=0 \mid old)^{n00} \\ &\quad P(c_i=0 \,\&\, m_i=1 \mid old)^{n01} P(c_i=1 \,\&\, m_i=0 \mid old)^{n10} \\ &= [P(c_i=1 \mid old) P(m_i=1 \mid c_i=1 \,\&\, old)]^{n11} \\ &\quad [P(c_i=0 \mid old) P(m_i=0 \mid c_i=0 \,\&\, old)]^{n00} \\ &\quad [P(c_i=0 \mid old) P(m_i=1 \mid c_i=0 \,\&\, old)]^{n01} \\ &\quad [P(c_i=1 \mid old) P(m_i=0 \mid c_i=1 \,\&\, old)]^{n10} \end{aligned}$$

Similarly for P (data | new).

We can now restate the likelihood ratio in terms of the parameters of the model that have been introduced. In summary they are:

Sparsity (s): the probability a component of a context vector is a one.

Learning (r): the probability that the link between an item node and a context component of one is learned during study.

Memory Noise (p): the probability that a component of the memory vector is a one because of other contexts in which this item has been seen. Memory noise incorporates both the number of other contexts in which the item has been seen and the amount of learning in those contexts. However, adding additional contexts is likely to affect the memory noise more than repeating an item within the same context (because the context representations are sparse and chosen independently). Therefore, the number of different contexts in which an item appears should be a sensitive measure of memory noise. However, since the correlation between the number of contexts an item has appeared in and its word frequency is very high (Dennis 1996), we will assume word frequency reflects memory noise under most conditions.

Vector Dimensionality (l): the length of the context and memory vectors. Note $l = n_{00} + n_{01} + n_{10} + n_{11}$.

Substituting into the previous equations, we get:

$$\begin{aligned} P(data|old) = &[s(r+p-rp)]^{n11} [(1-s)(1-p)]^{n00} \\ &[s(1-r)(1-p)]^{n10} [(1-s)p]^{n01} \end{aligned}$$

$$\begin{aligned} P(data|new) = &[s\,p]^{n11} [(1-s)(1-p)]^{n00} \\ &[s(1-p)]^{n10} [(1-s)p]^{n01} \end{aligned}$$

So,

$$P(data|old)/P(data|new) = [(r+p-rp)/p]^{n11} [1-r]^{n10}$$

Note that the 01 and 00 matches have no impact on the likelihood ratio. They affect the P(data|old) in the same way as they affect the P(data|new).

As mentioned previously, when there is no specific manipulation of criterion it is assumed that an item will be called old if the probability that it is old given the data is greater than the probability that it is new given the data, which is true when the likelihood ratio is greater than one. In general, then, as the mean likelihood ratio approaches one, from above in the case of targets and from below in the case of distracters, we expect performance to degrade. We can begin to understand how the above likelihood function simulates performance by looking at how its expected value varies as a function of the parameters (note a full exposition would consider the complete likelihood distribution). Firstly, as memory noise (p), which represents word frequency, approaches one, (r+p-rp)/p approaches one and the expected value of n_{10} approaches zero, so the expected value of P(data|old)/P(data|new) approaches one. In other words performance decreases as word frequency increases. Secondly, as learning (r), approaches zero, (r+p-rp)/p approaches one and 1-r approaches one, so the expected value of P (data|old)/P(data|new) approaches one. So as study time or number of repetitions decreases so does performance.

Adding the Contextual Reinstatement Parameter In the derivations outlined above, it is assumed that the ability to retrieve or otherwise reconstruct the study context at test (contextual reinstatement) is perfect. The context employed at test is identical to that used at study. It seems more likely, however, that features of the original context vectors will be lost. The ability to reinstate context is likely to be compromised by factors such as the length of the list and somewhat by delay.

The contextual reinstatement parameter (d) is the probability that a unit that was a one in the study context is a zero in the reconstructed context. The likelihood ratio can be re-derived taking into account contextual reinstatement (see Dennis & Humphreys in press) to give:

$$P\ (\text{data} \mid \text{old}) / P\ (\text{data} \mid \text{new})$$

$$= [(1\text{-}s\text{+}d\ (1\text{-}r)\ s)/(1\text{-}s\text{+}ds)]^{n00}\ [1\text{-}r]^{n10}\ [(p\ (1\text{-}s)\text{+}d(r\text{+}p\text{-}rp)\ s)/(p\ (1\text{-}s)\text{+}dps)]^{n01}\ [(r\text{+}p\text{-}rp)/p]^{n11}$$

If d is set to zero, indicating that reinstatement is perfect, the likelihood ratio reduces to that previously derived and the 00 and 01 matches have no impact.

Another way in which we might expect the reinstated context to differ from that at study would occur when the test context must encompass more than the study context. For instance, in a dual list design we might expect that the subject would form a reinstated context that incorporates the context vectors from both lists, particularly in the inclusion condition. In BCDMEM, we model the experiment wide context by taking the bitwise OR of the contexts for each study list. The resulting reinstated context will contain more ones (i.e. be less sparse) than either of the study list contexts and the sparsity will decrease if the lists are separated in time (so that the amount of overlap decreases). Likewise, we might expect study list contexts that incorporate different tasks (such as the anagram/read list in the Jacoby, 1991, design) would be less sparse than if only one task was performed.

The Effect of List Separation

One of the important distinctions between BCDMEM and the dual processing approach is the use of a graded temporal context. In BCDMEM, temporal context is represented by a vector and different context vectors can overlap to different degrees. Lists that appear in close temporal proximity will have a greater degree of overlap making it more difficult to determine which list an item appeared in. Under the dual processing assumptions, however, one would not expect the temporal separation of the lists to influence performance. The familiarity of a word is insensitive to whether it appeared in the target or non-target list, so one would not expect the separation of the lists to affect performance when recency is controlled. Recollection is related to the prevailing conditions when a word is being studied. While one might anticipate that decreasing the separation between lists would increase the probability of misidentifying the list in which an item appeared, to allow the solution of the PDP equations it is assumed that recollection is never in error in this way (Dodson & Johnson 1995). So one would not anticipate that the separation of the lists would affect it. Consequently, the temporal separation of the lists in the process dissociation procedure is an important variable to manipulate when distinguishing the two approaches, particularly if this is done in the absence of a study task distinction.

Hall (1996) varied the temporal separation between the lists in a PDP design by including an eight-minute filled interval. In the after lists condition, subjects studied list 1, then list 2, then spent eight minutes solving a puzzle task before being tested. In the between lists condition, the sequence was list 1 – filled interval – list 2 – test. The Jacoby (1991) version of the process dissociation procedure was used with inclusion instructions that covered both lists and exclusion instructions that targeted list two. The critical items were those from list one and the design controls for the recency of these items. In both study lists, subjects were asked to make pleasantness ratings. So, unlike the Jacoby (1991) design, the study task was constant across lists.

Figures 3 and 4 show the results. While the inclusion results for the list one words differed very little as a function of list two placement (0.825 verses 0.865), the exclusion probability of the list one words is much lower when the lists are separated by the filled interval than when they follow each other (0.365 verses 0.590). By varying the placement of list two, Hall (1996) was able to alter both the familiarity and recollection of the list one items (familiarity was 0.68 in the between lists condition and 0.82 in the after lists condition, while recollection was 0.46 in the between lists condition and 0.28 in the after lists condition). This result is inconsistent with the dual processing framework and underlines the importance of temporal context.

To model the manipulation of inter-list interval in BCDMEM we allowed the overlap between the list one and list two context vectors to change. The overlap parameter was defined as the probability that a component is a one in the context vectors of both lists in the PDP paradigm. Placing the filled eight-minute interval between the lists should lead to a lower value of this parameter (i.e. a decrease in the similarity of the list one and list two context vectors).

Bias may also play a role in the Hall (1996) results. All of the exclusion probabilities are below the corresponding inclusion probabilities suggesting the use of a more stringent bias in exclusion. In the exposition of the BCDMEM likelihood ratio we argued that prior odds could be eliminated on the basis that in most experiments the probability of an old word is equal to the probability of a new word. When inclusion instructions include both lists, the probability that a word should be called "old" is higher than in exclusion, which only includes one list. In Hall's (1996) experiment, the actual probabilities of an old word were 0.66 in the inclusion case and 0.33 in the exclusion case. While it is unclear how accurate subjects might be at estimating prior odds, the results suggest they are playing a role. Rather than add two new free parameters to model the prior odds, we chose to set the exclusion probability to 0.33 and allow the inclusion probability to be optimized. A least squares optimization procedure was used to find parameter values and Figures 3 and 4 show the fits to the data for the between lists interval and after lists interval respectively. The parameters of the fit were Inclusion Prior = 0.9093, Exclusion Prior = 0.33, learning rate: r = 0.3540, memory noise: p = 0.1604, sparsity: s = 0.02, vector length: l = 1000, context overlap: o(Between) = 0.0099, o(After) = 0.0152, contextual reinstatement: d = 0.3368. The maximum absolute difference was 0.056 and the correlation between the model and the data was 0.993. So, BCDMEM seems to have captured the effect of temporal separation.

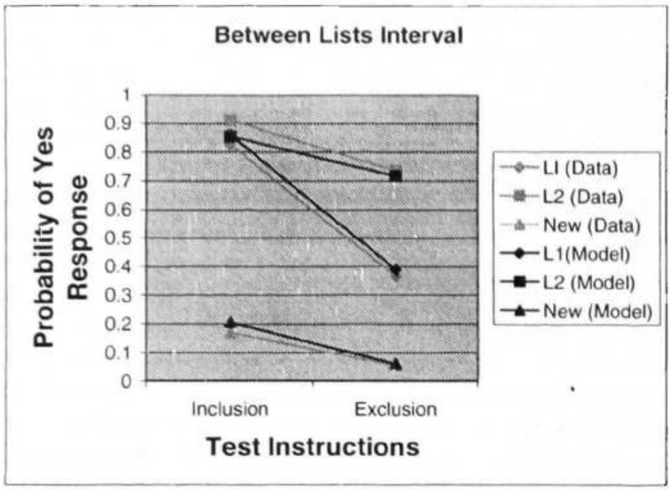

Figure 3: Fit of BCDMEM to Hall (1996) data experiment one for the Between Lists Interval

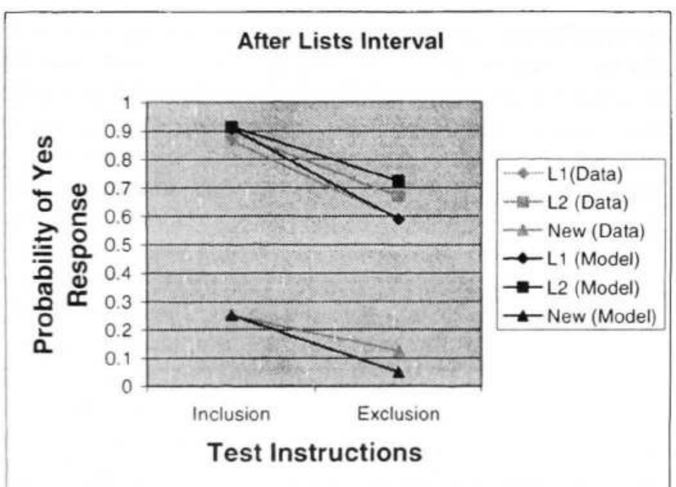

Figure 4: Fit of BCDMEM to Hall (1996) data experiment one for the After Lists Interval

Conclusion

We have demonstrated that the dual processing approach to episodic recognition is unable to account for the effects of manipulating the temporal separation of lists in the process dissociation procedure. BCDMEM is able to model these effects primarily because it proposes that contexts be modeled as vectors and allows these vectors to overlap to varying degrees. Some construct of this nature would seem necessary in order to provide an account of the effects of context in episodic recognition.

References

Anderson, J. A. & Bower, G. H. (1972). Recognition and retrieval processes in free recall. Psychological Review, 79(2), 97-123.

Anderson, J. R., & Milson, R. (1989). Human memory: An adaptive perspective. Psychological Review, 96 (4), 703-719.

Chalmers, K. A. & Humphreys, M. S. (in press). Role of generalized and episode specific memories in the word frequency effect in recognition. Journal of Experimental Psychology: Learning, Memory, and Cognition

Clark, S E. & Gronlund, S D. (1996). Global Matching Models of Recognition Memory: How the Models Match the Data. Psychonomic Bulletin & Review. 3(1). 37-60.

Dennis, S. (1996). The effect of the environment on memory: A connectionist model. Noetica. 2(1), http://psy.uq.edu.au /CogPsych/Noetica/.

Dennis, S., & Humphreys, M. S. (in press). Cueing for Context: An alternative to global matching models of recognition memory. In M. Oaksford & N. Chater (Eds.), Rational Models of Cognition.

Dennis, S., & Humphreys, M. S. (submitted). Integrating global matching and dual processing approaches to episodic recognition: The Bind Cue Decide Model of Episodic Memory.

Dodson, C. S. & Johnson, M. K. (1996) Some problems with the Process Dissociation Approach to Memory. Journal of Experimental Psychology: General. 125(2), 181-194.

Glanzer, M., & Adams, J. K. (1985). The mirror effect in recognition memory. Memory & Cognition, 13 (1), 8-20.

Hall, M. (1996) Process Dissociation Procedure: The effects of manipulating inter-list discriminability on estimates of familiarity. Honors thesis from the School of Psychology, University of Queensland.

Humphreys, M. S., Wiles, J., & Dennis, S. (1994). Towards a theory of human memory: Data structures and access processes. Behavioral and Brain Sciences, 17(4), 655-692.

Jacoby, L.L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. Journal of Memory and Language, 30, 513-541.

McClelland, J. L. & Chappell, M. (in press). Familiarity breeds differentiation: A subjective likelihood approach to the effects of experience in recognition memory. Psychological Review.

Shiffrin, R. M., & Steyvers, M. (1997). A model for recognition memory: REM – retrieving effectively from memory. Psychonomic Bulletin & Review. 4(2), 145-166.

Yonelinas, A. P. (1994). Receiver-Operating Characteristics in Recognition Memory: Evidence for a Dual-Process Model. Journal of Experimental Psychology: Learning, Memory and Cognition, 20 (6), 1341-1354.