7
Differentiation in Episodic Memory

Amy H. Criss and Gregory J. Koop

Global Matching Models and the List Strength Effect

Rich Shiffrin's illustrious career advancing the rigorous study of human memory began with the development of The Modal Model (Atkinson & Shiffrin, 1968). Nearly all subsequent models of memory have their origins in The Modal Model—either developed from elements of the model or developed to contrast with the model. Shiffrin and colleagues continued to develop models of memory within the general framework of The Modal Model, which resulted in the Search of Associative Memory model (SAM; Raaijmakers & Shiffrin, 1981) and the Retrieving Effectively from Memory model (REM; Shiffrin & Steyvers, 1997). In this chapter, we review empirical data demonstrating that one assumption introduced in these later models, termed differentiation, is in fact a fundamental mechanism of memory (Shiffrin, Clark, & Ratcliff, 1990).[1] Before describing differentiation, we review the circumstances that led to its inclusion in memory models.

SAM is a member of the class of models known as global matching models (GMMs). For decades GMMs were highly successful, owing to their ability to account for a wide range of data spanning multiple memory tasks, including recognition, free recall, cued recall, and others (e.g., Humphreys, Pike, & Bain, 1989; Murdock, 1982). The core assumption shared by GMMs is that a memory cue consisting of context and item information is matched to the full contents of memory and a global signal indicating the quality of this match (often called familiarity, strength, or evidence) is used in the decision process. After decades of success, however, the GMMs were categorically disconfirmed with the discovery of the list strength effect (LSE).

The LSE paradigm addresses a fundamental question about memory: how is memory affected by the quality of encoding? A substantial body of empirical evidence supports the informal observation that better encoding through repetition, additional study time, or enhanced processing improves memory (e.g., Craik & Lockhart, 1972; Neisser, 1967). The more interesting theoretical question in the LSE paradigm concerns the mechanism behind this improvement and the consequences of enhanced encoding of an item (or items) on memory for *other* items. The typical LSE experiment (Figure 7.1) includes three conditions: a pure weak list where all items receive some minimal level of study, a pure strong list where all items receive additional encoding (via repetition, duration, or depth of processing), and a mixed list where half of the items are weakly encoded and half are strongly encoded. The key comparison is memory accuracy on items with identical encoding strength but studied in a pure versus a mixed list.

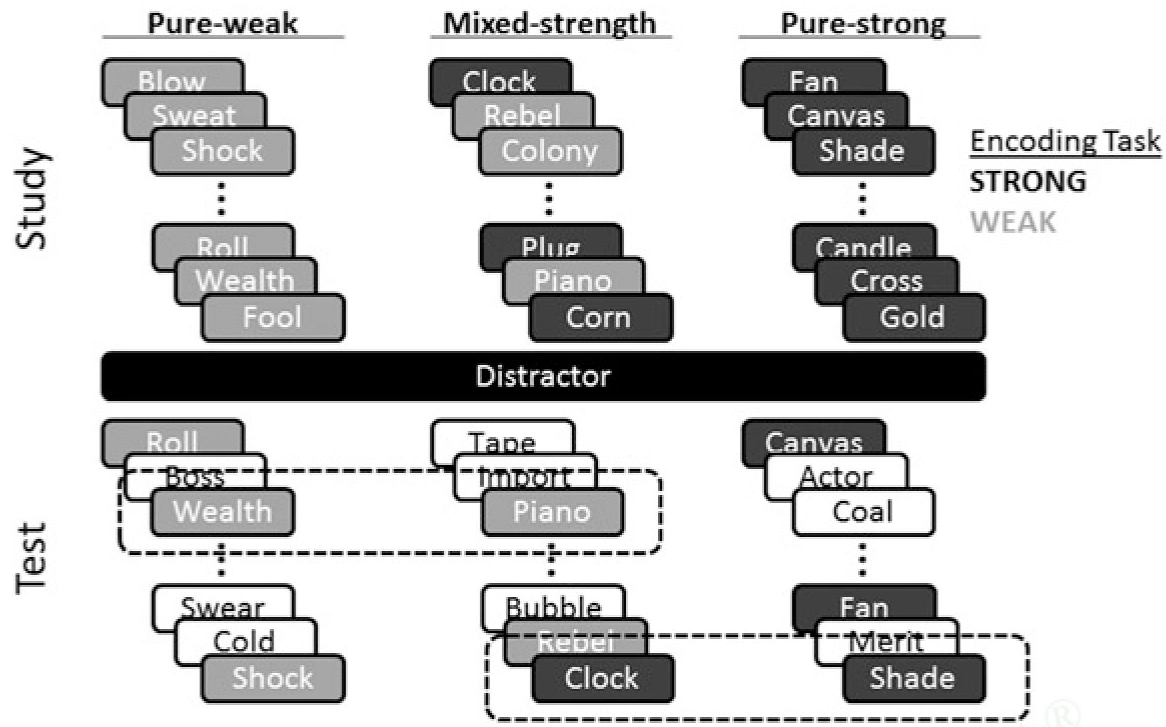

Figure 7.1 Study and test conditions in the typical list strength effect paradigm. Weak encoding manipulations are indicated by light colored boxes, strong encoding manipulations are indicated by dark colored boxes. At test, unstudied foils are indicated by white boxes. Critical comparisons for the list strength effect are between items of identical encoding strength in pure and mixed conditions—here indicated by the dashed rounded rectangles.

A positive LSE describes the finding that strong items are better remembered in a mixed than a pure list whereas weak items are better remembered in a pure than a mixed list (Figure 7.2a). That is, strengthening a subset of items in a list harms memory for the other items in that list: strong items are harmed by the presence of other strong items on a pure strong list and weak items are harmed by the presence of strong items on a mixed list. A positive LSE is found in recall tasks and this finding was predicted by the GMMs (e.g., Ratcliff, Clark, & Shiffrin, 1990).

SAM assumes that free recall is a two-step process of sampling an individual trace for recovery, with sampling being a global process and recovery being an item-specific process (see Huber, Tomlinson, Jang, & Hopper, Chapter 5, this volume). The sampling process is competitive in that a memory trace is selected in proportion to how well it matches the cue compared to all other traces. Repeated items are stored with more information and therefore are more likely to be sampled (and weakly encoded items are relatively less likely to be sampled). Although specifics vary across different models, a general description is that additional encoding increases the variance of the distribution of match values. This mechanism is similar to the GMM account of the list length effect (LLE) which is the finding that accuracy decreases as the number of to-be-remembered items increases. In both cases, adding information to memory (more information about a repeated item in the case of the LSE or additional items in the case of the LLE) does not affect the match between a memory cue and any individual stored trace, instead the global signal decreases due to the noise from that additional information. Essentially the LLE and LSE have the same causal mechanism in the GMMs. Further, because both recall and recognition rely on the global signal, both tasks are predicted to produce an LLE and a positive LSE.

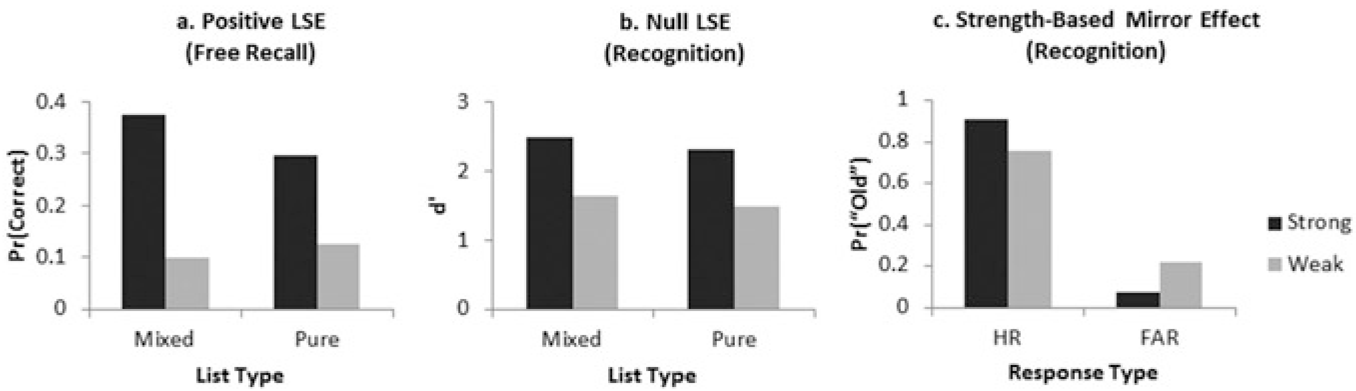

Figure 7.2 A typical positive LSE (a), null LSE (b), and strength-based mirror effect (c) in free recall, recognition, and recognition, respectively. Figures 7.2a and 7.2b are adapted from data in Ratcliff, Clark, and Shiffrin (1990), whereas Figure 7.2c presents data from Koop and Criss (submitted). Note that all three effects are produced by different critical comparisons within the same design (see Figure 7.1).

A series of landmark papers found task-dependent LSEs (Ratcliff, Clark, & Shiffrin, 1990; Shiffrin, Clark, & Ratcliff, 1990): recall tasks resulted in a positive LSE but recognition tasks resulted in a null (or negative) LSE. The null LSE in recognition has been widely replicated, including situations where rehearsal redistribution was unlikely (e.g., Yonelinas, Hockley, & Murdock, 1992). In recognition an LLE occurred simultaneously with a null LSE (e.g., Shiffrin, Huber, & Marinelli, 1995). Taken together, these results indicated that the underlying mechanisms for LLEs and LSEs must differ. In short, these findings were interpreted as a death knell for the GMMs (but see Osth & Dennis, Chapter 8, this volume). The challenge, then, was to provide a theoretical account of an LLE and null LSE in recognition and a positive LSE in recall.

Retrieving Effectively from Memory (REM; Shiffrin & Steyvers, 1997) and the Subjective Likelihood Model (SLiM; McClelland & Chappell, 1998) implemented differentiation (Shiffrin et al., 1990) as a fundamental mechanism of episodic memory and thus became known as the differentiation models. They account for the LLE and null LSE in recognition and other findings as described in the next section. But first, we briefly describe how REM accounts for the positive LSE in free recall. The sampling and recovery processes remain unchanged from the SAM model. The key insight that allows the model to simultaneously predict a null LSE in recognition and a positive LSE in free recall is that context and item features are encoded in different ways and differentially contribute to the two tasks. Free recall is heavily dependent on context as a cue to probe memory because no other information is provided in a typical free recall a cue. In contrast, recognition is more dependent on the item information presented in the memory cue. Malmberg & Shiffrin (2005) found evidence for their one-shot hypothesis: a fixed amount of context information is stored with each presentation whereas the amount of item information that is stored grows with the duration and quality of encoding. They showed that when both additional item and context information is added to a memory trace (e.g., via repetition) a positive LSE occurs. However, when item information but not context information was strengthened (e.g., via depth of processing or study time manipulations), participants showed a null LSE. Additional evidence comes from experiments where participants are asked to indicate whether they "remember" a test item, or whether they simply "know" it was on the study list. Dual process models assume that "remember" responses are generated by the recall of contextual information about the encoded item whereas "know" responses are the result of a global match absent recall. Studies using such a task find a positive LSE for “remember” responses and a null LSE for “know” responses (Diana & Reder, 2005; Norman, 2002), consistent with the idea that the positive LSE is context driven and the null LSE results from differentiation of item information.

Differentiation

Differentiation solved the remaining two challenges, allowing the models to simultaneously predict a null LSE effect and a LLE. Differentiation originated in the perceptual learning, similarity, and categorization literatures and described the decrease in similarity between items as the amount of learning increased (e.g., Gibson & Gibson, 1955; Nosofsky, 1987, 1991; Saltz, 1963). In short, a well-learned item is more distinct from other items. Semantic knowledge develops in a differentiated pattern (Rogers & McClel-land, 2004), providing some rationale that the development of episodic memories should follow a similar trajectory. Specifically, differentiation in episodic memory means that the more that is known about an item in a particular context or episode, the less similar (and thus less confusable) it is with other items (Shiffrin & Steyvers, 1997; McClelland & Chappell, 1998).

REM (Shiffrin & Steyvers, 1997) and SLiM (McClelland & Chappell, 1998) implemented differentiation by assuming that additional encoding results in the updating of a single memory trace. For example, if an item is studied twice on a single study list, the result is a single memory trace that contains a more complete and accurate representation of the item. This assumption accounts for the null LSE because the memory signal increases without a corresponding increase in variance. The assumption that additional encoding results in the updating of a single trace represented a clear departure from the GMMs and separated the net effect of list strength manipulations (updating a trace) from the net effect of list length manipulations (adding a new trace). This mechanism allowed the differentiation models to simultaneously explain the null LSE and LLE in recognition.

Behavioral Evidence for Differentiation

When differentiation was introduced into memory models to account for the null LSE it was not simply the best explanation, it was the *only* explanation (cf. Murdock & Kahana, 1993a, 1993b). In fact, a competitor would not appear in the literature for a half a decade. In the years since, a significant amount of evidence has been collected in support of differentiation.

Murnane and Shiffrin (1991) examined the role context played in initiating trace-updating (rather than creating additional traces) and the nature of subsequent LSEs. Their work helped elucidate the mechanism for updating memory traces—specifically a trace is updated so long as the context has not changed, but a new memory trace is stored when an item is repeated in a new context. One critical experiment contained two conditions of interest where words repeated in the context of sentences. In one condition, full sentences repeated with the hypothesis that such repetition would invoke updating of the traces and a null LSE would be observed. In the second condition, the same words repeated, but in different sentences with the hypothesis that context changed with each sentence. If context changed, then a new trace would be stored and the net result would be a positive LSE. The data confirmed both predictions. Modeling in a slightly different domain clarified that repetition itself is not the causal factor for determining if a trace is updated or a new trace is stored on any given trial. Criss, Malmberg, and Shiffrin (2011) modeled output interference in recognition, the finding that accuracy decreases as a test proceeds. The best-fitting model assumed that when an item is remembered in a specific context, then the best matching episodic trace is updated. If an item fails to be remembered, then a new traces is stored.

More recent work has focused on the strength-based mirror effect (SBME) paradigm as a source for evaluating the role of differentiation in recognition. The first empirical reports focused on the SBME were published contemporaneously with the differentiation models (Stretch & Wixted, 1998) and the finding was attributed to a criterion shift, which, at the time, was the only viable option from a GMM perspective. The deep connection between the list strength findings and the SBME was only later fully explored (e.g., Criss, 2006). The SBME is the finding of a higher hit rate (HR) and lower false alarm rate (FAR) for items tested following a pure strong list compared to items tested following a pure weak list (see panel c in Figure 7.2). Differentiation results in updating of memory traces encoded during the strong list. As memory traces are updated, the match between a foil item and the contents of a strong list become lower than the match between a foil item and the contents of a weak list, which results in a lower FAR for the strong list. This same logic applies to the match between a target item and N-1 other stored traces. However, this is outweighed by evidence generated from the match between that target and its corresponding memory trace, causing a higher HR. Thus the SBME is a direct result of differentiation, at least within the framework of models incorporating this mechanism.

Differentiation models also predict an interaction between list strength and foil similarity. For randomly chosen foils (i.e., dissimilar foils) the match between any given foil and the set of memory traces decreases with increasing list strength for the reasons described above—that is, differentiation occurs. However, moderately and highly similar foils act somewhat like targets (i.e., noisy versions of the target). For mixed lists, differentiation is not sufficient to counteract similarity and the models predict a higher FAR for foils similar to a single strongly encoded target than foils similar to a weakly encoded target. For pure lists, however, differentiation is sufficient to counteract this effect and a lower FAR is predicted for foils that are similar to a strongly encoded target compared to a weakly encoded target. Across multiple studies, Criss (2006) confirmed this predicted interaction between similarity and list strength.

Criss (2009) capitalized on the assumption that the match between a test probe and the contents of memory corresponds to one's subjective feeling of familiarity for that test item (Shiffrin & Steyvers, 1997; McClelland & Chap-pell, 1998). Criss reasoned that these subjective ratings of familiarity should show differentiation such that ratings of memory evidence should decrease for foils and increase for targets following a strongly encoded study list compared to a weakly encoded study list. Empirical distributions of subjective familiarity were assessed using a direct ratings procedure (Mickes, Wixted, & Wais, 2007) where participants were asked to judge the amount of evidence retrieved from memory on a 1–20 scale (1 indicated strong evidence that the item was not studied and 20 indicated strong evidence that the item was on the study list). As predicted, participants provided more extreme distributions—lower ratings on foil trials and higher ratings on target trials—for test items following strong lists than weak lists. Furthermore, when bias was manipulated via changes to the base rate, participants became more liberal in a yes/no decision, as expected (e.g., Rotello, Macmillan, Hicks, & Hautus, 2006), but the distribution of ratings of memory evidence remained unchanged. Thus, empirical ratings of memory evidence follow the pattern predicted by theoretical predictions of the differentiation models.

An additional piece of evidence that demonstrates the contribution of differentiation comes from the drift diffusion model. The diffusion model (Ratcliff & McKoon, 2008) describes how evidence accumulates towards two decision bounds, "old" or "new" in a recognition task. In this measurement model, memory evidence corresponds to the rate at which evidence accumulates towards either of the two bounds, and is referred to as the *drift rate*. The model has three separate parameters that represent decision bias, only one of which will be described here. The *starting point* indicates how close the evidence accumulation process starts to either decision threshold (Wagenmakers, Ratcliff, Gomez, & McKoon, 2008). As with the direct ratings procedure, Criss (2010) predicted that if differentiation drove the SBME, differences between strong lists and weak lists should be primarily evident in the drift rates. Specifically, the drift rates for targets and foils following a strong list should be more extreme (e.g., have higher absolute values) than those following a weak list. In line with the differentiation account, encoding strength affected drift rates but did not cause a shift in the starting point, whereas base rate manipulations affected the starting point but not the drift rate. Collectively these data present compelling evidence that differentiation is more than a theoretical construct sufficient to account for the null LSE, but instead that differentiation is a fundamental mechanism underlying human memory.

Neural Evidence for Differentiation

Recently, researchers have begun using neuroimaging techniques like fMRI and EEG in hopes of uncovering the unique neural signatures of mnemonic and decisional processes with the ultimate goal of testing process models like REM and SAM. Criss, Wheeler, and McClelland (2013) utilized fMRI to test the differentiation account of the SBME. They first identified brain regions associated with successfully retrieving an item from memory using the traditional method in the fMRI literature of contrasting activity for hits and correct rejections (CRs). Activity in these so-called retrieval success areas (RSA) were then correlated with measures of memory accuracy (*d'*) and response bias (*C*) for individual participants—the idea being that the RSA measurement confounds activity based on memory evidence and bias (see O'Connor, Han, & Dobbins, 2010). These analyses identified significant correlations between retrieval related activity and *d'* in the left angular gyrus, and correlations with *C* in bilateral middle frontal gyrus. After identifying $RSA_{d'}$and RSA_C (respectively), Criss and colleagues tested whether activity patterns in these areas reflected differentiation or criterion shifts. If differentiation drives the SBME, then the evidence for a foil being "new" would be higher following a strong list than a weak list, and activity in $RSA_{d'}$ would accordingly increase more quickly for strong foils than weak foils. Confirming this prediction, the data showed that the rate of accumulation of activation in $RSA_{d'}$ was higher for foil trials after a strong than a weak list. Furthermore, RSA_C did not show any effects of memory strength condition, which suggests that response bias played no role in the SBME, consistent with behavioral evidence documented above. An exploratory voxel-by-voxel analysis (i.e., without identifying RSA) replicated these findings. Testing process models using fMRI is still in its infancy (White & Poldrack, 2013) and no doubt as the field matures the methods and analyses will too (see Turner et al., 2013). The data from Criss and colleagues are a promising indication that differentiation can be measured in the BOLD signal. In fact, reading the literature with the perspective of evaluating differentiation in episodic memory reveals a series of papers with converging evidence (i.e., Daselaar, Fleck, & Cabeza, 2006;Gold and Shadlen 2007; Ploran et al., 2007; Ploran, Tremel, Nelson, & Wheeler, 2011; Yonelinas, Otten, Shaw, & Rugg, 2005).

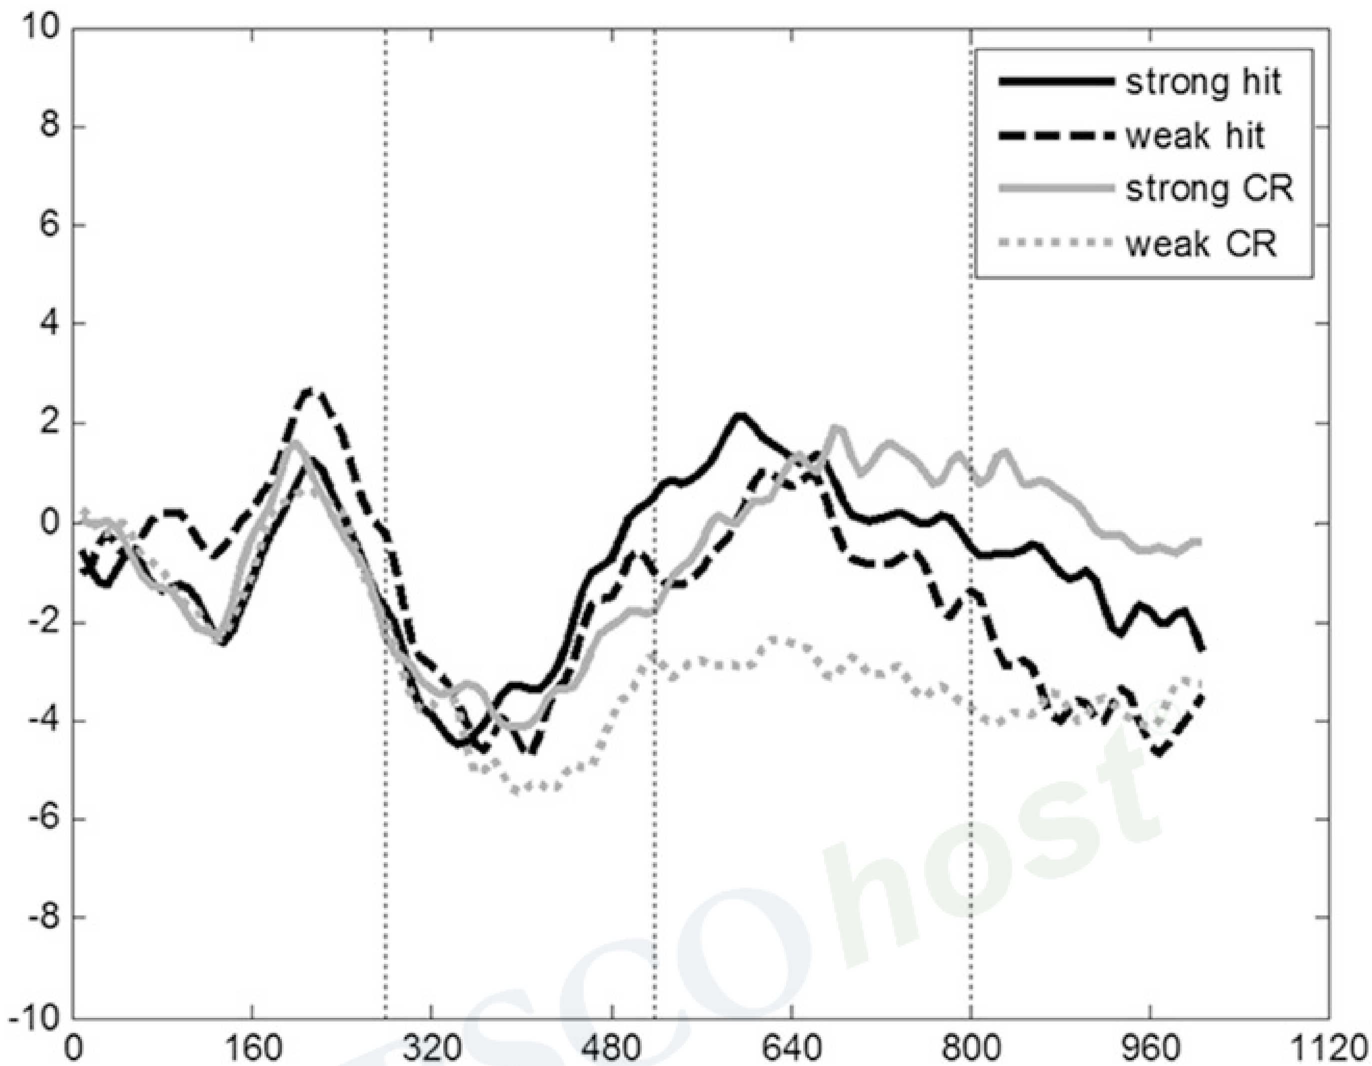

Figure 7.3 Evoked potentials time-locked to the stimulus onset in a strength-based mirror effect paradigm (see Figure 7.1). CR refers to a trial with a correct rejection.

Data from event-related potentials (ERPs) suggest similar conclusions. Analyses of retrieval ERPs typically show an "old/new effect" (e.g., Warren, 1980) in which hits produce more positive deflections in both early (N400) and late (late parietal positivity, or LPP[2]) components of the waveform than do CRs. Hemmer, Criss, and Wyble (2011) examined the effects of encoding strength on the ERP during retrieval. In addition to replicating the typical old/new effect in the FN400, they also demonstrated that the LPP uniquely discriminated between foils following strong lists and those following weak lists (see Figure 7.3). That is, CRs from a strongly encoded list (called strong CRs) elicited a more positive deflection of the waveform than did CRs from a weak list (called weak CRs), in accord with a differentiation mechanism (cf. Finnigan, Humphreys, Dennis, & Geffen, 2002 but see Herron, Quayle, & Rugg, 2003; Windmann, Urbach, & Kutas, 2002). Whereas the above data evaluate ERPs related to retrieval, other recent work has focused on the relationship between retrieval ERPs and encoding ERPs. Chen, Lithgow, Hemmerich, and Caplan (2014) found that late components reflect the distinctiveness of targets and foils and this distinctiveness increases with encoding strength. In other words, late components of the ERP reflect differentiation; a finding that directly fits with the interpretation provided by Hemmer and colleagues (2011). In summary, there exists substantial evidence in the form of both behavioral and neural data consistent with a differentiation mechanism.

The Criterion-Shift Hypothesis

Although the differentiation models continue to be the dominant class of process models of episodic memory, other accounts continue to embrace the claim that the criterion (not the foil distribution) changes with list strength. Hirshman (1995) first proposed this criterion shift account of the null LSE. He considered three models for criterion placement based on the strength of items at study: a means model, range model, and probability-matching model. In the means model, the criterion is placed at a fixed location between the means of the target and foil distributions; when the target distribution increases (e.g., with a strongly encoded list), the criterion also increases (i.e., becomes more conservative). The range model assumes that the criterion is placed some proportion of the way between the estimated high and low end points of familiarity. The probability matching model assumes that the criterion is set such that the probability of saying “old” is constant across conditions. Implementation of these different assumptions within a signal detection model showed that the range model best fit the data. Hirshman presented participants with mixed and pure weak study lists, each followed by a test of foils and weak targets (i.e., strong targets were not tested for the mixed list). The FAR differed for pure and weak lists, which Hirshman interpreted as evidence that the criterion is placed on the basis of the study list. In other words, strong targets affected the criterion placement despite not being present at test. Within a criterion-change account, this finding requires that the criterion be set on the basis of the study list.

Stretch and Wixted (1998) came to a similar conclusion based on a failure to change performance through test manipulations. They presented a mixed study list and a mixed test list but color-coded targets and foils with respect to encoding strength (e.g., red indicated that an item, if studied, was repeated and green indicated no study repetition). If, as expected, participants adopted a criterion reflecting the expected strength of the test item, then the FAR for green and red foils should differ. Despite many attempts, Stretch and Wixted were never able to elicit different FARs for foils expected to be strong than for foils expected to be weak.

In contrast to the assumption that the criterion was set based on the expectations about memory derived from the study list, Verde and Rotello (2007) concluded that expectations were based on the initial test trials. They had participants study mixed lists but the test was blocked by strength of the targets. They found no difference in FARs across test blocks, indicating that participants did not change their criteria in response to changes in the difficulty of the test. However, they found lower a FAR when the first test block contained strong targets and concluded that the criterion was based on the strength of the initial test items.

All of the above data were discussed within a criterion-shift framework with the preferred model being one in which the criterion is set at the beginning of the test (albeit the factors determining this setting remain controversial) and remains fixed over the duration of the test list. However, each of the findings is equally consistent with differentiation. Specifically, if the memory probe is compared to the entire contents of episodic memory then the FAR need not differ as a function of expected target strength (as was found in Verde and Rotello [2007] and Stretch and Wixted [1998]). Further, the FAR should differ for tests following pure weak versus mixed study lists because mixed lists are partially differentiated. Specifically, the FAR should be lowest for pure strong and highest for pure weak conditions with mixed lists falling between the two.

Starns and colleagues took an approach similar to those above but went so far as to suggest that there is no need to assume a differentiation mechanisms exists (e.g., Starns, Ratcliff, & White, 2012;

Starns, White, and Rat-cliff, 2010). The key to their approach has been to give participants mixed study lists and test with either strong or weak targets (à la Hirshman, 1995) among various other conditions. Critically, Starns and colleagues inform participants of the strength of the target items and require participants to repeat this information before continuing the experiment. Under these circumstances, participants show a lower FAR for a test containing strong targets compared to a test containing weak targets. The REM framework (and indeed all memory models that we can think of) has no problem with the supposition that participants may shift their criterion when instructed to do so, as is the case in the Starns and colleagues. However, evidence for a change in criterion is not evidence against differentiation.

In summary, a viable criterion shift account of the effects of list strength in recognition is the following: participants set their criterion in response to some experimental factor and that criterion is adjusted in response to instructions provided by the experimenter. Whereas participants are willing and able to set the criterion prior to beginning a test, they do not alter the criterion placement during test (except when feedback is provided [Verde & Rotello, 2007], but not always [Starns et al., 2010]). Note that a criterion shift accounts provide no explanation at all for the positive LSE in recall.

Importantly, differentiation may provide a more parsimonious explanation for changes in FAR (or lack thereof) than relying on criterion shifts. Criterion shifts in response to experimental manipulations are fickle. Some studies have shown that participants change criteria in response to the base rate of targets at test (Criss, 2006, 2009; Koop, Criss, & Malmberg, in press; Rotello et al., 2006), but not others (Cox & Dobbins, 2011; Koop, Criss, & Malmberg, in press). Participants sometimes adjust their criteria in response to the encoding strength of test items (Verde & Rotello, 2007, Experiment 5), but not always (Starns, White, & Ratcliff, 2010; Stretch & Wixted, 1998; Verde & Rotello, 2007). This is but a small sampling of the literature but it clearly indicates that the hypothesis that participants readily and easily adjust decision criteria in the response to experimental manipulations is dubious at best. A serious attempt to provide a formal account of the mechanisms underlying criterion shifts and the situations in which a criterion shift is induced must be developed for memory in order to make any progress (e.g., Brown, Steyvers, & Hemmer, 2007; Cox & Shiffrin, 2012; Turner, Van Zandt, & Brown, 2011).

Conclusion

The concept of differentiation is an elegant and parsimonious explanation for a large set of empirical effects including the positive LSE in recall, null LSE in recognition, and the SBME. Direct evidence of differentiation includes both behavioral and neural data. The differentiation approach is not without criticism, and a criterion shift account has been raised as an alternative to differentiation in recognition. Differentiation models include a criterion—thus findings of a criterion shift are not evidence against differentiation. To further advance theoretical development, detailed mechanisms for criterion shifts should be elucidated for testing against the explicitly described differentiation models.

Acknowledgements

This material is partially based upon work supported by the National Science Foundation under Grant

Number 0951612.

Notes

1. Despite the significant impact of his work, Rich often reminded his lab members, myself included (A.H.C.), that we shouldn't take ourselves too seriously because all models are wrong. Rather, modelers are developing simple approximations of the truth. This lesson is so ingrained that it feels somewhat inappropriate to be writing a chapter extolling one of the ways in which Rich's theoretical models were in fact approximately true. But, if the shoe fits …

2. This component is also sometimes known as the "late parietal component" or "late positive complex" (Finnigan et al., 2002).

References

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: Volume 2* (89–195). New York: Academic Press.

Brown, S., Steyvers, M., & Hemmer, P. (2007). Modeling experimentally induced strategy shifts. *Psychological Science*, *18*, 40–45.

Chen, Y. Y., Lithgow, K., Hemmerich, J. A., & Caplan, J. B. (2014). Is what goes in what comes out? Encoding and retrieval event-related potentials together determine memory outcome. *Experimental Brain Research*, *232*(10), 3175–3190.

Cox, G. E., & Shiffrin, R. M. (2012). Criterion setting and the dynamics of recognition memory. *Topics in Cognitive Science*, *4*(1), 135–150.

Cox, J. C., & Dobbins, I. G. (2011). The striking similarities between standard, distractor-free, and target-free recognition. *Memory & Cognition*, *39*, 925–940.

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

Criss, A. H. (2006). The consequences of differentiation in episodic memory: Similarity and the strength based mirror effect. *Journal of Memory and Language*, *55*, 461–478.

Criss, A. H. (2009). The distribution of subjective memory strength: List strength and response bias. *Cognitive Psychology*, *59*, 297–319.

Criss, A. H. (2010). Differentiation and response bias in episodic memory: Evidence from reaction time distributions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *36*, 484–499.

Criss, A. H., Malmberg, K. J., & Shiffrin, R. M. (2011). Output interference in recognition memory. *Journal of Memory and Language*, *64*, 316–326.

Criss, A. H., Wheeler, M. E., & McClelland, J. L. (2013). A differentiation account of recognition memory: Evidence from fMRI. *Journal of Cognitive Neuroscience*, *25*, 421–435.

Daselaar, S. M., Fleck, M. S., & Cabeza, R. (2006). Triple dissociation in the medial temporal lobes: Recollection, familiarity, and novelty. *Journal of Neurophysiology*, *96*, 1902–1911.

Diana, R. A., & Reder, L. M. (2005). The list strength effect: A contextual competition account. *Memory & Cognition, 33*, 1289–1302.
Finnigan, S., Humphreys, M. S., Dennis, S., & Geffen, G. (2002). ERP 'old/new' effects: Memory strength and decisional factor(s). *Neuropsychologia*, *40*, 2288–2304.
Gibson, J. J., & Gibson, E. J. (1955). Perceptual learning: Differentiation or enrichment? *Psychological Review, 62,* 32–41.
Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience, 30*, 535–574.
Hemmer, P., Criss, A. H., & Wyble, B. (2011, November). *Assessing a neural basis for differentiation accounts of recognition memory*. Poster presented at the Psychonomic Society meeting, Seattle, WA.
Herron, J. E., Quayle, A. H., & Rugg, M. D. (2003). Probability effects on event-related potential correlates of recognition memory. *Cognitive Brain Research, 16*, 66–73.
Hirshman, E. (1995). Decision processes in recognition memory: Criterion shifts and the list strength paradigm. Journal of Experimental Psychology: Learning Memory and Cognition, *21*, 302–313.
Humphreys, M. S., Bain, J. D., & Pike, R. (1989). Different ways to cue a coherent memory system: A theory for episodic, semantic, and procedural tasks. *Psychological Review, 96,* 208–233.
Koop, G. J., & Criss, A. H. (submitted). The Response Dynamics of Recognition Memory: Sensitivity and Bias.
Koop, G. J., Criss, A. H., & Malmberg, K. J. (in press). The role of mnemonic processes in pure-target and pure-foil recognition memory. *Psychonomic Bulletin & Review*.
Malmberg, K. J., & Shiffrin, R. M. (2005). The "one-shot" hypothesis for context storage. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*, 322–336.
McClelland, J. L., & Chappell, M. (1998). Familiarity breeds differentiation: A subjective-likelihood approach to the effects of experience in recognition memory. Psychological Review, 105, 724–760.
Mickes, L., Wixted, J. T., & Wais, P. E. (2007). A direct test of the unequal-variance signal detection model of recognition memory. *Psychonomic Bulletin & Review, 14*, 858–865.
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. and Kahana, M. J. (1993a). Analysis of the list strength effect. Journal of Experimental Psychology: Learning, Memory and Cognition, *19*, 689–697.
Murdock, B. B. and Kahana, M. J. (1993b). List-strength and list-length effects: Reply to Shiffrin, Ratcliff, Murnane, and Nobel. Journal of Experimental Psychology: Learning, Memory and Cognition, *19*, 1450–1453.
Murnane, K., & Shiffrin, R. M. (1991). Interference and the representation of events in memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*, 855–874.
Neisser, U. (1967). *Cognitive psychology*. East Norwalk, CT: Appleton-Century-Crofts. Norman, K. A. (2002). Differential effects of list strength on recollections and familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28*, 1083–1094.
Nosofsky, R. M. (1987). Attention and learning processes in the identification and categorization of integral stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*(1), 87–108.
Nosofsky, R. M. (1991). Tests of an exemplar model for relating perceptual classification and recognition memory. *Journal of Experimental Psychology: Human Perception & Performance, 17* (1), 3–27.

O'Connor, A. R., Han, S., & Dobbins, I. G. (2010). The inferior parietal lobule and recognition memory: Expectancy violation or successful retrieval? *Journal of Neuroscience*, 30, 2924–2934.

Ploran, E. J., Nelson, S. M., Velanova, K., Donaldson, D. I., Petersen, S. E., & Wheeler, M. E. (2007). Evidence accumulation and the moment of recognition: Dissociating perceptual recognition processes using fMRI. *Journal of Neuroscience, 27*, 11912–11924.

Ploran, E. J., Tremel, J. J., Nelson, S. M., & Wheeler, M. E. (2011). High quality but limited quantity perceptual evidence produces neural accumulation in frontal and parietal cortex. *Cerebral Cortex, 21*, 2650–2662.

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

Ratcliff, R., Clark, S. E., & Shiffrin, R. M. (1990). List-strength effect: I. Data and discussion. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16*, 163–178.

Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation, 20*, 873–922.

Rogers, T. T., & McClelland, J. L. (2004). *Semantic cognition: A parallel distributed processing approach.* Cambridge, MA: MIT press.

Rotello, C. M., Macmillan, N. A., Hicks, J. L., & Hautus, M. (2006). Interpreting the effects of response bias on remember-know judgments using signal-detection and threshold models. *Memory & Cognition, 35*, 254–262.

Saltz, E. (1963). Compound stimuli in verbal learning: Cognitive and sensory differentiation versus stimulus selection. *Journal of Experimental Psychology*, *66*, 1–5.

Shiffrin, R. M., Clark, S. E., & Ratcliff, R. (1990). List-strength effect: II. Theoretical mechanisms. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16*, 179–195.

Shiffrin, R. M., Huber, D., & Marinelli, K. (1995). Effects of category length and strength on familiarity in recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 267–287.

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

Starns, J. J., Ratcliff, R., & White, C. N. (2012). Diffusion model drift rates can be influenced by decision processes: An analysis of the strength-based mirror effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *38*, 1137.

Starns, J. J., White, C. N., & Ratcliff, R. (2010). A direct test of the differentiation mechanism: REM, BCDMEM, and the strength-based mirror effect in recognition memory. *Journal of Memory and Language*, *63*, 18–34.

Stretch, V., & Wixted, J. T. (1998). On the difference between strength-based and frequency-based mirror effects in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*, 1379–1396.

Turner, B. M., Forstmann, B. U., Wagenmakers, E., Brown, S. D., Sederberg, P. B., & Steyvers, M. (2013). A Bayesian framework for simultaneously modeling neural and behavioral data. *NeuroImage*, *72*, 193–206.

Turner, B. M., Van Zandt, T., and Brown, S. (2011). A dynamic, stimulus-driven model of signal detection. Psychological Review. 118(4), 583–613.

Verde, M. F., & Rotello, C. M. (2007). Memory strength and the decision process in recognition memory. *Memory & Cognition*, *35*, 254–262.

Wagenmakers, E. J., Ratcliff, R., Gomez, P., & McKoon, G. (2008). A diffusion model account of criterion shifts in the lexical decision task. *Journal of Memory and Language, 58*, 140–159.

Warren, L. R. (1980). Evoked potential correlates of recognition memory. *Biological Psychology, 11*, 21–35.

White, C. N., & Poldrack, R. A. (2013). Using functional neuroimaging to constrain theory testing. *Perspectives in Psychological Science, 8*, 79–83.

Windmann, S., Urbach, T. P., & Kutas, M. (2002). Cognitive and neural mechanisms of decision biases in recognition memory. *Cerebral Cortex, 12*, 808–817.

Yonelinas, A. P., Hockley, W., & Murdock, B.B. (1992). Tests of the list-strength effect in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18*(2), 345–355.

Yonelinas, A. P., Otten, L. J., Shaw, K. N., & Rugg, M. D. (2005). Separating the brain regions involved in recollection and familiarity in recognition memory. *Journal of Neuroscience, 25*, 3002–3008.