Global matching models of recognition memory

Adam F. Osth & Simon Dennis

University of Melbourne

Address correspondence to:

Adam Osth (E-mail: adamosth@gmail.com)

Author Note

Abstract

How does retrieval take place in recognition memory? A number of computational models have been developed that posit that recognition operates by a process of global matching, wherein the cue is compared to each stored memory. These cue-to-memory similarities are then aggregated to produce an index of the global similarity between the cue and the contents of memory which can then be subjected to a decision process. In this chapter, we describe a.) the theoretical rationale and successes of such models, such as their accounts of similarity and list-length effects and generalization to multiple memory tasks, b.) challenges and hurdles they have experienced, including the null list-strength effect, the mirror effect, and the extralist feature effect, and c.) recent developments, such as extensions to predictions about response latency, sources of interference outside of the study list, and integration of more realistic representations.

Keywords: recognition memory; episodic memory; global matching models; memory models

Global matching models of recognition memory

If an informal poll were to be conducted among most researchers in the field as to which are the most commonly used computational models of recognition memory, the most popular responses would likely be signal detection theory (SDT) or the dual process signal detection (DPSD) model. Each of these models has been successful in explaining the form of the receiver operating characteristic (ROC: Wixted, 2007; Yonelinas, 2002) and relations between first and second choice responses in multialternative choice (Kellen & Klauer, 2014; Parks & Yonelinas, 2009). However, the successes of these models are likely due to the specification of the relevant decision variables. SDT, for instance, postulates a set of latent strength distributions, often referred to as "familiarity", for targets and lures that are typically Gaussian in shape. A given trial in a recognition memory experiment is a sample from one of the memory strength distributions – this is compared to a response criterion and an "old" decision is made if this criterion is exceeded, otherwise a "new" decision is made.

A major limitation of the SDT and DPSD models is that they do not specify the representations and retrieval processes that give rise to their decision variables. Consequently, the models are mute with respect to the effects that episodic memory researchers are typically interested in, such as memorial advantages for recent items, spaced repetitions, recognition advantages for low frequency words, and impairments with increasing list-length (Murdock, 2006). We do not mean to imply that such models cannot fit data that emerge from such manipulations. Rather, the models cannot *explain* the data, as they do not specify the psychological processes that give rise to such phenomena.

To give an example, consider data from an experiment which contains a manipulation of list-length, where one study list contains ten items whereas another contains 40 items and it was found that discrimination was poorer in the longer list. A model such as SDT can account for the data by allocating a.) lower familiarity of target items in the longer list, b.) higher familiarity of lures in the longer list, c.) a response criterion that is affected

by the list-length manipulation, or some combination of each of these variables. However, it is not immediately clear as to how or why familiarity would be affected by a list-length manipulation. A number of theoretical explanations could be offered for the result, such as greater interference in the longer list, lateral inhibition among the study list representations, or capacity limits of encoding new information. However, none of these theoretical explanations are central to SDT itself.

A class of models that can explain the results from list-length manipulations in addition to other recognition memory phenomena is the set of global matching models. Global matching has its origins in the theories of Richard Semon (Schacter, Eich, & Tulving, 1978), but became formalized in a number of computational models starting in the 1970's. In these models, retrieval operates via a global matching mechanism where the similarity between the probe item to representations of each of the study list items is computed. These similarity values are then aggregated to produce an index of *global similarity* that is the basis of the recognition decision that reflects how similar the probe item is to the contents of memory. This aggregation often involves summing each of the similarities (e.g., Gillund & Shiffrin, 1984; Nosofsky, 1988; Hintzman, 1988), but sometimes involves the average (Shiffrin & Steyvers, 1997) or the maximum (McClelland & Chappell, 1998). The fact that global similarity is computed from a large number of instances implies that such models can be considered instance theories of episodic memory (e.g., Logan, 2002).

Table 1 provides a brief summary of several of such computational models. Inspection of the table reveals that global matching models differ considerably in their assumptions about learning, representation, association, and the metric for similarity computation. However, as we will discuss, a large number of models make very similar predictions, and in some cases are mathematically identical to each other despite very different psychological interpretations (Humphreys, Pike, Bain, & Tehan, 1989; Kelly, Mewhort, & West, 2017).

The mathematics of global matching are illustrated in a generalized form in Figure 1,

which illustrates a case where four items I have been entered into the contents of memory, where we use M to illustrate a representation in memory. When a target probe is tested (top left), it is compared to each item in memory, with the arrows representing a similarity operation. Similarity operations sometimes involve the calculation of the dot product between the two vectors, but sometimes involve the calculation of distance (e.g., Nosofsky, 1986) or a likelihood ratio representing the odds that M is a target or a lure (e.g., Shiffrin & Steyvers, 1997). We denote these similarities as a draw from a distribution D for matches (D_{match}) and mismatches ($D_{mismatch}$).

It may seem somewhat counterintuitive that similarity can be variable, as similarity is often described as a fixed quantity. However, it is important to keep in mind that this represents variance *across trials*. That is, on one trial, there might be a probe such as "dog" which is compared to a representation of the word "cat" in memory, which might have relatively high similarity. On the next trial, however, the probe word might be "truck", which would have a lower similarity to the representation of "cat." In addition, there are other factors that can produce greater variability in probe-target similarity across trials. Several models postulate that encoding is not all-or-none, but is instead variable. For instance, a variable number of features from each item I may be encoded into its memory representation M , such that some representations are relatively complete and some representations exhibit a small proportion of their features (e.g., Hintzman, 1988; Shiffrin & Steyvers, 1997). This assumption of *encoding variability* is the reason why self-similarity, as represented by the distribution D_{match} in Figure 1, has variability associated with it rather than being a fixed value, as memories encoded with less features will exhibit lower probe-target similarity at retrieval than traces with a greater proportion of features.

Global similarity varies considerably across trials for both targets and lures. Collecting global similarity values across trials in an experiment results in a set of distributions for targets and lures – these are illustrated in the bottom of Figure 1. These distributions are typically Gaussian in shape because even if the distributions of cross-trial

similarities are not Gaussian, the global similarity values should be approximately Gaussian by virtue of the central limit theorem if aggregation across memories is accomplished via the sum or average.

The key factor that results in the difference between targets and lures is the presence of the self-match, or the match between I_1 and M_1 in the figure. The extent to which the mean of D_{match} is greater than $D_{mismatch}$ shifts the target distribution to the right of the lure distribution. In the figure, the variability of the target distribution is larger than that of the lure distribution. This can occur in some models where encoding variability is sufficiently large such that the variability of D_{match} exceeds $D_{mismatch}$ (e.g., Ratcliff, Sheu, & Gronlund, 1992). A number of experiments and modeling analyses in long-term recognition memory have generally found evidence consistent with the proposal that the target distribution contains larger variability than the lure distribution (e.g., Glanzer, Kim, Hilford, & Adams, 1999; Starns & Ratcliff, 2014; Wixted, 2007).

In several models, the global similarity that results from the efforts of retrieval is compared to a criterion to query whether the probe should be accepted as an old or new item, as in SDT. In this respect, global matching models bear a strong resemblance to SDT models, but they differ because the means and variances of the target and lure distributions are not estimated as free parameters but instead are produced by the global similarity between the probe and the study list representations. This produces a more constrained set of predictions for experimental manipulations than in SDT. For instance, global matching models predict that increases in list-length and similarity among the list items should impair performance (see the section "List-Length and Similarity Effects in Recognition Memory"), but this is not predicted *a priori* by SDT. Not all models employ an SDT-like decision mechanism, however; later in the chapter in the section "Predictions about Response Latency" we will describe how some models have leveraged the successes of the state-of-the-art models of decision-making by having global similarity drive an evidence accumulation process, which has the advantage of making predictions about both choice

Figure 1. General illustration of the mathematics behind global matching models. The top left illustrates a target trial, whereas the top right illustrates a lure trial. When global similarities are collected across trials, they result in distributions of global similarity for targets and lures that resemble those from signal detection theory (SDT, bottom). See the text for more details.

and response latency.

We should also note that Figure 1 illustrates a case where each representation in memory is stored separately, as in *multiple trace* models (Hintzman, 1988). However, alternative models exist in which only a single representation is stored that contains a superposition of all learned memories acquired during the study list, which are often referred to as *composite storage* models, where probe items are directly compared to the composite to produce a decision. We further elaborate on how this distinction in the "Composite Storage or Separate Memory Traces?" sub-section, and discuss how several composite storage global matching models can be considered linear versions of neural network or connectionist models.

In addition, in the same sub-section we also describe various computations for how items can be associated to each other, as such computations differ considerably across multiple trace and composite storage models. While item recognition memory does not generally require associations between items, this is not true for paired-associate tasks such as associative recognition, where pairs of items such as $A - B$, $C - D$, and $E - F$ are studied. During the test phase of an associative recognition task, participants are presented with pairs composed of studied words and are asked to discriminate between pairs of words that were both studied together (intact pairs, such as $A - B$) and pairs where both words were studied in different pairs (rearranged pairs, such as $C - F$). A noteworthy feature of the associative recognition task is that knowledge of the items alone is insufficient for the task. Global matching models are capable of associative recognition using the same principles as in item recognition – a representation of the association can be matched against each of the associations present in memory to produce a measure of global similarity for the probe pair.

The chapter will proceed as follows. In the first section, "Rationale and Success of Global Matching Models", we describe why global matching was preferred as a mechanism for retrieval in recognition memory over other mechanisms, such as a serial search through the contents of memory. In addition, while the chapter largely focuses on the case of recognition memory, we also discuss how such models are viable models of other memory tasks such as recall, and have even been extended to several other cognitive tasks, such as categorization and the generation of semantic representations (see the sub-section "Generalization to Other Cognitive Tasks"). In the second section, "Challenges to Global Matching Models" we describe some of the existing challenges to the framework that arose that motivated the extension of existing models or in some cases, the development of newer models.

The final section of the chapter, "Current and Future Directions", summarizes much of the recent progress in these models and the directions such models are likely to be headed

in, including extension to predictions about response times, interference from episodes outside the study list, and incorporation of more realistic representations. This section may be particularly relevant to readers already familiar with such models, as the very thorough existing review of global matching models by Clark and Gronlund (1996) deals with some of the same topics in the first two sections, while the third section summarizes many of the recent developments that have occurred since the publication of their review. The sheer number of global matching models precludes complete mathematical description of each model and thus interested readers should consult the original publications.

Rationale and Success of Global Matching Models

List-Length and Similarity Effects in Recognition Memory

Possibly the most thorough rationale for the global matching mechanism came from Gillund and Shiffrin (1984), who were concerned with extending the search of associative memory (SAM) model, a highly successful model of free recall (Raaijmakers & Shiffrin, 1981), to recognition memory. In SAM, free recall operates by sampling through traces in memory where the sampling probability of a trace is proportional to the similarity between the cues and the memory trace – items are output if this similarity is sufficiently high. Gillund and Shiffrin initially considered this mechanism for recognition memory, where traces in memory would be repeatedly sampled and an "old" decision would only be made if a sampled trace was sufficiently similar to the probe. Such a mechanism resembles serial search models that had been previously popular for recognition, such as Sternberg (1966)'s serial exhaustive model and Murdock (1974)'s conveyor belt model. In these models, an "old" decision is made if a match between the probe and a sampled memory is found and a "new" decision is made if no such match is found after a certain number of retrieval attempts. This decision mechanism makes the prediction that rejections of lures should be slower than rejection of targets. Unfortunately, correct rejections can in some cases be made quickly and with high confidence (e.g., Murdock & Dufty, 1972).

An alternative to serial search models is direct access models, in which a probe item is able to directly access its own representation in memory without searching through the other contents (e.g., Anderson & Bower, 1972; Kintsch, 1967). However, Gillund and Shiffrin noted that such models have difficulty explaining how recognition memory often exhibits *search-like* characteristics. One of the more prominent examples is the list-length effect, wherein both recognition and recall become progressively more impaired as the length of the study list is increased (e.g., Brandt, Zaiser, & Schnuerch, 2019; Cary & Reder, 2003; Strong, 1912). Sequential searches through the contents of memory are naturally impaired by longer lists because the greater number of to-be-searched memory traces makes it more difficult for the representation of a target item to be found.

While hybrid models of search and direct access existed at the time (e.g., Atkinson & Juola, 1974), Gillund and Shiffrin conducted a series of clever experiments to investigate whether there are qualitative differences between early and late retrievals that necessitate a slower search process. In these experiments, they manipulated a number of variables that affect performance, such as list-length and similarity of the lures to the studied items. Critically, each of these manipulations were crossed with an instructional manipulation that forced participants to respond quickly or late in the trial. The logic of the instructional manipulation is that if a search process is operating, it should exert a greater effect later in the decision due to the larger number of search attempts having taken place. Each of the manipulations produced their predicted marginal effects – poorer performance was found for longer study lists, similar lures, and for responses made early in retrieval, the latter of which indicated a speed-accuracy tradeoff. However, there were no interactions between the retrieval time manipulation and the other manipulations that affected performance. If list-length effects were due to the influence of search processes, Gillund and Shiffrin reasoned that list-length impairments should be larger later in the decision than earlier, but the data did not confirm this prediction of hybrid search models.

A global matching mechanism, in contrast, elegantly captures each of these patterns

of data. Rejections can be made with high confidence because lures exhibit considerably lower global similarity to the study list items than target items by virtue of the lure not being within the contents of memory. Rejection becomes poorer for lures that are similar to studied list items because the higher similarity between the probe and such studied items increases the lure’s global similarity and thus increases the false alarm rate. In addition, list-length effects are predicted because each representation in memory contributes additional variance to the global similarity calculation. If summed similarity is the basis of the global similarity calculation, the variance of the resulting distribution of global similarity values is merely the sum of the contributions of each of the list memories. For this reason, variance increases linearly for both targets and lures with increasing list-length. The increase in variance reduces the signal-to-noise ratio (as measured by d'), decreasing performance.

An additional prediction that emerges from the global matching framework is that if the proportion of study list items that are similar to lure probes is increased, false alarm rates should increase. This prediction was tested in category length designs where the number of studied items from a semantic category is manipulated. Experiments consistently confirmed the prediction of higher false alarm rates for categories with a greater number of studied exemplars (Hintzman, 1988; Robinson & Roediger, 1997; Shiffrin, Huber, & Marinelli, 1995), an effect dubbed the *category length effect*. A related and popular finding is that participants often exhibit very high false alarm rates to an item that is highly associated to a set of studied list items in the Deese-Roediger-McDermott (DRM: Deese, 1959; Roediger & McDermott, 1995) paradigm, a finding which is also qualitatively consistent with the predictions of global matching models. The fact that global matching models were simultaneously able to capture the deleterious effects of increasing list-length and category length was considered to be a strong advantage for the framework (Clark & Gronlund, 1996).

Composite Storage or Separate Memory Traces?

Several of the initial global matching models, such as the theory of distributed associative memory (TODAM: Murdock, 1982), SAM, Minerva 2 (Hintzman, 1988), and matrix (Humphreys, Bain, & Pike, 1989; Pike, 1984) models made similar predictions about the manipulations of similarity, list-length, and list-strength, and in some cases the models are even mathematically identical (Humphreys, Pike, et al., 1989). However, that does not imply that the psychological theory behind each model was the same. To the contrary, each model made very different assumptions about how memories are stored and retrieved.

Some models assumed composite storage, in which all memories are stored in a single structure. For instance, in the TODAM model, study list items are represented as vectors denoted by f , and the process of learning involves the summation of each of item vector into a single memory vector m that represents the contents of memory¹:

$$m = \sum_i^L f_i \quad (1)$$

where the subscript i denotes the i th item on the study list of L items. When a probe vector is presented, global similarity s is assessed by calculating the similarity between the probe vector p and the memory vector:

$$s = p \cdot m \quad (2)$$

Another way that TODAM differs from other global matching models is in terms of how it represents associations between items. While associations between items are not commonly required for tasks such as item recognition memory, they are when paired

¹We have depicted a relatively simplified equation for TODAM that omits a forgetting process that occurs as each item is entered into memory. Interested readers are encouraged to consult the original publication (Murdock, 1982).

associates are studied for a task such as associative recognition or cued recall. In TODAM, associations between pairs of items can be learned by calculating the convolution of the constituent item vectors. Formally, the x th element in a convolution between vectors f and g can be expressed as:

$$(f * g)(x) = \sum_i^N f(i)g(x - i) \quad (3)$$

where i indexes the i th element in the item vectors, which contain a total of N elements, and $*$ is the convolution operator. Equation 3 is equivalent to calculating the outer product between two vectors and summing over the diagonals of the resulting matrix in a single direction. The $f * g$ convolution represents a conjunctive binding between the two constituent vectors f and g , and the convolution bears little resemblance to either of its constituent vectors f and g .

An alternative to TODAM that also employs composite storage is the matrix model (Humphreys, Bain, & Pike, 1989; Pike, 1984). Instead of representing associations as convolutions, they are represented as outer product matrices that sum into a single matrix that contains the contents of memory. Formally, a memory matrix M can be defined as the sum of L pairs of constituent items f and g :

$$M = \sum_i^L f_i \otimes g_i \quad (4)$$

Composite storage models can be construed as relatively simple, linear versions of neural network or connectionist models, as the outer product operation is the mathematical formalism of Hebbian learning. Relatively sophisticated neural networks share the assumption of storage of memories in associative matrices which represent the connection weights between layers of neurons, but they deviate through the usage of more

sophisticated learning rules (e.g., Chappell & Humphreys, 1994; Monaco, Abbott, & Kahana, 2007; Norman & O'Reilly, 2003). Composite storage models share much of the same theoretical rationale as connectionist models, such as a commitment to the idea that the brain does not learn by "creating" new units, but instead by sharing capacity between a fixed set of neurons.

An alternative to composite storage models are multiple trace models, in which each experienced event results in the creation of a new memory trace. The Minerva 2 model of Hintzman (1988) is an example, which shares the assumption of TODAM that items are represented as vectors of features, but instead learning results in the acquisition of separate vectors. Specifically, if each vector has N features and a list of L items is learned, in Minerva 2 learning of the study list involves the acquisition of a matrix that is shaped $L \times N$ elements, with each row corresponding to a different memory trace, rather than a single vector of N elements as in the TODAM model.

Hintzman (1986, 1988) made a number of arguments in favor of the multiple trace model over assumptions of composite storage. In particular, participants appear to be sensitive to unique aspects of item presentations and to the spacing between repetitions, suggesting that each presentation leads to unique storage. The principle of unique storage for each item presentation has intuitive appeal and one might be inclined to believe it leads to very different model predictions. However, it should be noted that composite storage and multiple trace models can make mathematically identical predictions under some circumstances. Consider a composite model such as TODAM where a vector M is the sum of L item vectors with N features, versus a multiple trace model that results in a matrix of $L \times N$ elements where each row corresponds to a different memory trace denoted by f . TODAM's global similarity computation (Equation 2) is mathematically identical to the summed similarity between the probe vector p and each of the individual vectors stored in the multiple trace model ($\sum_i^L p \cdot f_i$).

However, the Minerva 2 model is not merely a multiple trace variant of the TODAM

model. Minerva 2 differentiates itself by raising the dot product between the probe and each memory trace to the third power, introducing a non-linearity into the global similarity computation. This similarity function is graphically depicted in Figure 2. This non-linear relationship between the dot product and the resulting similarity makes it such that traces that are highly similar to the probe carry the largest weight in the global similarity calculation. To understand how, consider two vectors that each of have a dot product with the probe vector of .1 and .2. After the cubing process, the psychological similarities are .001 and .008, making them extremely dissimilar to the probe. However, consider a different pair of vectors with dot products of .8 and .9 with the probe vector. After cubing, their similarities are .51 and .73. Not only do these exhibit considerably higher similarity values, but the increase in the dot product from .1 to .2 produced a very small change in similarity, whereas the increase from .8 to .9 produced a very large change.

Figure 2. Graphical illustration of Minerva 2's similarity function, where the dot product between two vectors is raised to the third power to produce the similarity values.

Non-linear transformations of objective similarity into psychological similarity are extremely common in multiple trace models. One of the most noteworthy examples is the generalized context model (GCM: Nosofsky, 1986, 1988). GCM is similar to Minerva 2 in that each study list item is represented separately in memory. However, instead of the dot product, the matching operation in the GCM is a calculation of the distance d between the two representations i and j in a multidimensional space. This distance is subsequently

converted to similarity s through an exponential transformation (e.g., Shepard, 1987):

$$s_{ij} = \exp(-cd_{ij}) \quad (5)$$

where c governs the "steepness" of the exponential transformation. With higher values of c , items in memory that are further from the probe become progressively less similar as c increases such that more similar items will dominate the global similarity computation. Raising the value of c can be considered analogous to raising the power of the dot product between memory vectors in a model such as Minerva 2.

In addition to the aforementioned mechanisms, models such as SAM (Gillund & Shiffrin, 1984) and the retrieving effectively from memory (REM: Shiffrin & Steyvers, 1997) also employ such non-linearity, but by virtue of multiplicative combination across features and/or cues (e.g., Medin & Schaffer, 1978).

The non-linearity in such models offers several advantages. A major advantage is that it offers sensitivity to conjunctions of items or features. Consider if a pair of items A and B are learned and are represented by a concatenation of their item vectors ($A \oplus B$), meaning that if each vector has N features, the pair is represented by a vector of $2N$ elements with the B vector to the right of the A vector. In addition, another pair $C - D$ is also entered into memory, represented by the $C \oplus D$ operation. Now, consider if we were to test a participant's memory by presenting the items A and D , but instead of asking if these items were studied, we ask whether or not the two items were studied in the same pair *together* (as in the associative recognition task). Linear matching operations would predict that false alarm rates to $A - D$ pairs should be identical to hit rates to studied pairs such as $A - B$ or $C - D$, producing zero discrimination in the task.

To understand how this is the case, consider if each self-similarity (such as the dot product between A and itself) takes a value of 1 and each inter-item similarity (such as the

dot product between A and B) takes a value of zero. If we adopt the cubing process in Minerva 2, an intact pair such as $A \oplus B$ matches its own representation in memory $(A - B)$ as $(A \cdot A) + (B \cdot B) = 2^3 = 8$ and matches the $C - D$ pair as $(A \cdot C) + (B \cdot D) = 0^3 = 0$, which sum together to produce a global similarity of 8. However, for the rearranged pair $A - D$, this matches the pair $A - B$ as $(A \cdot A) + (D \cdot B) = 1^3 = 1$ and matches the $C - D$ pair as $(A \cdot C) + (D \cdot D) = 1^3 = 1$, resulting a global similarity of 2. Thus, the global similarity for the rearranged pair is considerably lower than for the intact pair. If the cubing process was not used, however, both $A - B$ and $A - D$ would have the same global similarity (2). The cubing process makes it such that the strong match between the intact pair and itself has considerably higher weight than the two relatively weak partial matches to the rearranged pair cue.

The non-linearity in models such as Minerva 2, GCM, SAM, and REM has repeatedly been shown to be advantageous in capturing differential rates of forgetting for exemplars vs. prototypes (e.g., Hintzman, 1986; Nosofsky, Little, Donkin, & Fific, 2011), as well as non-linear effects of study duration on hit rates and false alarm rates to similar lures (Arndt & Hirshman, 1998). While these results may come across as obvious support for separate storage models, a similar non-linearity can be achieved in composite models through the usage of *auto-association*. The term auto-association refers to associations between an item and itself. This may seem somewhat counter-intuitive, but can be understood as the association between an item's features, such as a binding between an item's color and shape. Autoassociation is quite common in neural network models of memory (e.g., Chappell & Humphreys, 1994; Norman & O'Reilly, 2003).

Kelly et al. (2017) demonstrated that if a four-way auto-association is learned, such as a fourth-order tensor (A learned as $A \otimes A \otimes A \otimes A$) or a four-way auto-convolution (A learned as $A \circledast A \circledast A \circledast A$). Under these circumstances, composite models are mathematically equivalent to Minerva 2 in the case of tensor products and approximately equivalent in the case of the convolution operation. The equivalence between the two

frameworks arises because a four-way auto-association achieves the same non-linear transformation of similarity as the cubing of dot products in the Minerva 2 model. These results provide further support for the notion that the different types of global matching models are more similar than they are different. However, there is one further advantage of multiple trace models, which is the ability to modify memory traces after they have entered the contents of memory. In particular, in differentiation models, repetition results in the enhancement of an existing memory trace by updating it with additional features. We will elaborate more on the predictions of differentiation models in the section "The Null List-Strength Effect."

Generalization to Other Cognitive Tasks

Up to this point, this chapter has focused heavily on the case of long-term recognition memory. However, global matching models have also been applied to the Sternberg (1966, 1975) short-term recognition memory task, where a small number of items (usually between 1 and 5 items) are studied and a recognition test on a single probe follows shortly thereafter. Several variants of global matching model have been developed to explain results from the task, including the Noisy Exemplar Model (NEMO: Kahana & Sekuler, 2002) and the Iterative Resonance Model (IRM: Mewhort & Johns, 2005), but the most thorough example was possibly the Exemplar Based Random Walk (EBRW: Nosofsky et al., 2011) model. The EBRW is an extension of the GCM which uses a random walk model to make decisions, enabling the model to make joint predictions about choice and response time (RT). The EBRW has been successfully in capturing both error rates and RTs from several benchmarks in the Sternberg memory task, including manipulations of list-length, serial position effects, and differential rates of forgetting for exemplars and prototypes. The success of the EBRW suggests that the architecture of global matching may be able to extend to all forms of episodic recognition and not just to long-term recognition memory.

The successes of the global matching models are not limited to recognition, however.

Several of the global matching models were able to address recall tasks. A simple illustration of the cued recall task can be seen in the matrix model (Pike, 1984), where pairs such as $A - B$ and $C - D$ would be encoded as outer products and summed into a memory matrix M according to Equation 4. If a pair such as $A-B$ were presented for a recognition trial, the global similarity of the pair s could be derived by matching it against the memory matrix using the dot product operation:

$$s = (A \otimes B) \cdot M \quad (6)$$

However, if only the word A was presented and a participant was tested for cued recall ("What was studied with A ?"), recall can be formalized as:

$$s = A \cdot M \quad (7)$$

where s is now a vector instead of a scalar. The similarity between the s vector and each of the item representations (A , B , C , and D) can be calculated using the dot product, and whichever representation is most similar to the retrieved vector could be chosen for recall. Recall progresses in TODAM in a similar fashion, with the exception that a recall cue is matched against a memory vector using the correlation operation, which results in a vector that can be compared to candidate responses for recall.

The fact that global matching models were capable of addressing both recognition and recall on a common set of representations was a major strength of the framework and differentiated them from earlier models that were task-specific. Collectively, the models enjoyed a great deal of success in their extension to a wide range of memory tasks. In particular, the SAM model served as the dominant model of free recall for quite some time, and its successor, the REM model, has also succeeded in passing several free recall

benchmarks (Lehman & Malmberg, 2013). The TODAM model was extended to serial recall (Lewandowsky & Murdock, 1989) and recency judgments (Murdock, Smith, & Bai, 2001), and the Minerva 2 model was capable of addressing judgments of frequency (Hintzman, 1988) and prototype generation (Hintzman, 1986). In addition, many of the models were capable of addressing paired-associate tasks such as associative recognition and cued recall. The broad generalization of such models is likely due to the fact that the global matching framework is as powerful as it is simple. The models specify the representations and retrieval processes; an extension to a given memory task merely involves a specification of the cues and retrieval operations that are used. It is for this reason that such models can be considered models of episodic memory more generally and are not restricted to recognition².

Such models have enjoyed success in their extension to semantic memory. The architecture of the retrieving effectively from memory (REM) model (Shiffrin & Steyvers, 1997), for instance, was also extended to the lexical decision task, in which participants are presented with letter strings and asked whether each string is a word or a non-word (Wagenmakers et al., 2004). In this model, the letter string probe is globally matched against the contents of word representations in semantic memory, and if there is sufficient global similarity the letter string is considered a word. The REM model was successfully able to account for benchmark phenomena in the task such as repetition and word frequency effects.

Global matching models were also sufficient to generate semantic representations of words from a corpus of natural text, similar to techniques such as latent semantic analysis (LSA: Landauer & Dumais, 1997). In particular, Minerva 2 was extended by Kwantes (2005) and Jamieson, Avery, Johns, and Jones (2018) by adding a trace vector corresponding to each experienced word into the contents of memory along with a representation of the context in which the word was experienced, with the context being

²A lengthier overview of the benchmarks imposed by recall tasks and how global matching models could be amended to achieve them is provided by Osth and Dennis (2015a).

defined as the document that the word was contained in. This was done for each document in the text corpus, with a corpus size typically being tens of thousands of documents to bear some resemblance to the experiences of language learners. Semantic representations are produced "on the fly" using Minerva 2's *echo content* generation process. First, the probe vector is globally matched against the contents of memory, and subsequently a weighted sum of all of the vectors in memory is constructed where the weight of each vector is determined by the similarity to the probe word vector. In Kwantes's formulation, first the probe word is used to construct a set of the retrieved contexts in which the word was associated. In the second stage, the combined word and context cue is matched against the contents of memory, which produces a weighted sum of words that co-occurred in the same contexts as the probe word. These resulting semantic vectors exhibit similarity relationships to other words that resemble human judgments and responses derived from free association norms, similar to other successful models of semantic memory such as LSA. This is an impressive success, as it demonstrates that the storage and retrieval mechanisms present in global matching models might be sufficient to generate representations that could be used in natural language.

Another noteworthy success is the extent to which the global matching mechanism is successful in the domain of categorization. The aforementioned generalized context model (GCM) of Nosofsky (1986) stands as a noteworthy example. In this case, categorization decisions are made on the basis of global similarity constructed from an item cue and a category cue. Consider if a participant studies items a , b , and c as members of category A and d , e , and f as members of category B. When presented with a novel exemplar g and asked which category it belongs to, the global similarity of category A is assessed by calculating the summed similarity of g to members of category A (a , b , and c), and likewise the global similarity of category B is assessed by calculating the summed similarity of g to members of category B (d , e , and f). The strengths of category A and B are subjected to a Luce choice rule and the category with the highest similarity tends to win the competition.

Many subsequent extensions of the GCM to the recognition memory task have proven that the model is capable of addressing both recognition memory and categorization judgments with the same representations and model structure (e.g., Nosofsky, 1988, 1991; Nosofsky et al., 2011).

This review is but a small subset of the successes of these models. Global matching models have also been extended to spoken word recognition (Goldinger, 1998), priming effects in lexical decision tasks (Ratcliff & McKoon, 1988), heuristics and probability estimation in judgment and decision making (Dougherty, Gettys, & Ogden, 1999), hypothesis generation (Thomas, Dougherty, Sprenger, & Harbison, 2008), and artificial grammar learning (Jamieson & Mewhort, 2009). The wide generality of the global matching mechanism in cognitive models is likely due to the fact that it provides a simple and effective means for leveraging structure across a set of memories, and retrieval from memory is often a fundamental component of cognition.

Challenges to Global Matching Models

The Null List-Strength Effect

The zeitgeist of the global matching models in the 1980's in episodic memory was nearly put to a grinding halt by Ratcliff, Clark, and Shiffrin (1990)'s discovery of the null list -strength effect in recognition memory. Most memory models predict that hit rates should improve for items that are strengthened by extra presentations or study time³. But what effect does strengthening an item have on *other* items in memory? That is, if a participant studies items A and B, and we were to strengthen A by extra repetitions or study time, what effect does this have on item B?

A counter-intuitive prediction of the global matching models is that strengthening a

³An exception to this is the TODAM model. Murdock (1982) noted that in TODAM, memory did not improve with repetitions because increasing presentations or study time of a study list proportionally increases the variance of the target and lure distributions, resulting in no net improvement in d' . A number of solutions were proposed for this problem, including variable encoding of item features (Murdock & Lamon, 1988) and prediction-error encoding, where the learning rate for an item is inversely proportional to its similarity to the contents of memory (Lewandowsky & Murdock, 1989; McDowd & Murdock, 1986).

set of items should cause forgetting of the other items in memory. This is because repetitions serve the same functional role as additional item presentations. Recall from earlier that increases in list-length impair performance because each added memory contributes additional variance to the global similarity calculation, increasing the variance of the target and lure distributions and decreasing the signal-to-noise ratio. Repetitions of items harm performance in the same way, as each repeated item that is added to the contents of memory contributes additional variance. In fact, they can cause even *larger* impairments in performance, because the covariance between multiple presentations of the same item increase variance to an even greater degree. In other words, strong memories dominate the global similarity computation at retrieval to the detriment of weak memories.

The list-strength prediction was first tested in the domain of free recall by Tulving and Hastie (1972). In their work, they presented lists of items where all items were presented once and a mixed list of twice presented and once presented items, with the critical comparison being between once presented items in each list condition. Recall of the once presented items was worse for items in the mixed list compared to the condition where all items were presented once. In other words, the twice presented items interfered with the memories of the once presented items for recall.

Ratcliff et al. (1990) conducted similar experiments in the domain of recognition memory, but with three conditions instead of two: a pure weak (PW) condition where all items were relatively weak, a mixed (M) condition where both weak (MW) and strong (MS) items were presented, and a pure strong (PS) condition where all items were strong. Strengthening was conducted by both repetition and additional study time. The global matching models predicted that additional strength should impair performance of other items, so the d' for weak items in the mixed list (MW) should be worse than the PW items ($MW\ d' < PW\ d'$), and the PS list should perform worse than the strong items in the mixed list ($PS\ d' < MS\ d'$). However, a total of seven experiments they conducted did not find any evidence for these predictions. Strength did not appear to produce any

impairment for other items, as evidenced by weak items performing roughly the same across both list types (PW $d' =$ MW d') and likewise for strong items (MS $d' =$ PS d'), despite the fact that strong items performed much better than weak items.

The null list-strength effect has been replicated numerous times across different memory labs (Kahana, Rizzuto, & Schneider, 2005; Murnane & Shiffrin, 1991a, 1991b; Ratcliff, McKoon, & Tindall, 1994; Ratcliff et al., 1992; Yonelinas, Hockley, & Murdock, 1992). It has even been found in other memory tasks, such as associative recognition (Osth & Dennis, 2014), cued recall (Wilson & Criss, 2017), and source memory (Osth, Fox, McKague, Heathcote, & Dennis, 2018). A mathematical analysis by Shiffrin, Ratcliff, and Clark (1990) of the existing memory models found that none of the models at the time yielded the correct predictions about the list-strength paradigm. This failed prediction of the models was extremely unsatisfying – the strongest predictions from global matching models regard the effects of list composition, but how seriously can they be taken if they cannot address something as simple as list items having unequal strengths? If the models could not be revised to address the list-strength paradigm, then there was a risk that they would be abandoned altogether.

As a solution to this problem, Shiffrin et al. (1990) proposed a variant of the SAM model that uses a differentiation mechanism. The SAM model previously assumed that each repetition results in a new memory trace, such that the number of memory traces is increased when items are repeated. In the differentiation variant, in contrast, subsequent repetitions of an item strengthen the trace that was laid down on the first presentation of the item such that strengthening a set of list items does not increase the number of memory traces. In addition, as the item is strengthened, its similarity to other traces decreases, making it such that strengthened memory trace does not respond more strongly to other items. Differentiation models predict null effects of list-strength while still predicting a list-length effect – this is because presentation of novel items does not produce differentiation but instead results in the creation of new memory traces, each of which

contribute interference at retrieval. Differentiation became a major component of the REM model (Shiffrin & Steyvers, 1997), the subjective likelihood in memory (SLiM) model (McClelland & Chappell, 1998), and the model of Cox and Shiffrin (2017), each of which are multiple trace models as they allow for the enhancement of the original traces with additional repetitions.

Differentiation also makes predictions beyond that of the null list-strength effect. As the strength of a set of memories is increased, the memories will be less similar to novel item cues, implying that the false alarm rate should be reduced in conditions of higher list-strength. This prediction has been confirmed across a range of experiments which have all shown reduced false alarm rates for conditions of higher list-strength (e.g., Criss, 2006; Hirshman, 1995; Osth, Dennis, & Kinnell, 2014; Stretch & Wixted, 1998).

Differentiation was not the only solution to the null list-strength effect, however. Other researchers argued that the null effect was indicative of the study list items exerting a relatively small interference contribution relative to the entire contents of memory. We will elaborate more on these models in the "Sources of Interference" section.

The Mirror Effect

The mirror effect refers to any manipulation that produces opposite effects on hit rates and false alarm rates. Many manipulations produce mirror effects, including comparing pictures to words, varying levels of word concreteness, but the most popular of these is likely to be the word frequency mirror effect, wherein low frequency (LF) words exhibit higher hit rates and lower false alarm rates than high frequency (HF) words (Glanzer & Adams, 1985). Some hypothetical data demonstrating the mirror effect are plotted in Figure 3, where across a set of conditions A, B, C, and D, the hit rate (HR) decreases while the false alarm rate (FAR) increases. This figure illustrates how the effect got its name – the false alarm rate appears as a mirror image of the hit rate.

The mirror effect sounds rather intuitive – shouldn't variables that improve

performance increase hit rates and reduce false alarm rates? However, the early global matching models actually predicted that both hit and false alarm rates should move in the same direction. For instance, if we assumed that HF words had more pre-experimental traces than LF words, FAR would be higher for HF words because these pre-experimental traces should increase their global similarity. However, an unfortunate consequence is that such traces should also increase global similarity for targets, increasing the hit rate!

Figure 3. Hypothetical data demonstrating a mirror effect on the hit rates (HR: black) and false alarm rates (FAR: gray).

As a solution to this problem, Gillund and Shiffrin (1984) proposed that HF and LF words additionally elicit different response criteria. That is, HF words may be more familiar overall, which would result in higher hit and false alarm rates for HF words if the same criterion was applied to both classes of words. However, a higher response criterion is applied to HF words, which under the appropriate parameter values can result in a mirror effect. However, a disadvantage of such an approach is that LF word advantages persist even in two-alternative forced choice (2AFC) tests, where a target and a lure is presented and the participant has to choose the studied word (Glanzer & Bowles, 1976). 2AFC testing is conventionally assumed to not involve the usage of a response criterion, as it is a sufficient strategy to simply choose the response option that elicits the highest global similarity. Moreover, in the studies of Glanzer and Bowles, they also included trials where two lures or two targets were presented, where one word is LF and the other is HF. When

two lures are presented, participants choose the HF word, whereas when two targets are presented, the LF word is chosen.

The mirror effect motivated a shift away from a view where decisions are only based purely on direct measures of similarity between the cue and each memory. Rather, it led to a whole new set of models that used a *likelihood ratio* decision mechanism (e.g. Glanzer, Hilford, & Maloney, 2009; Osth, Dennis, & Heathcote, 2017). These models included attention likelihood theory (ALT: Glanzer, Adams, Iverson, & Kim, 1993), the aforementioned REM and SLiM models, and the bind-cue-decide model of episodic memory (BCDMEM: Dennis & Humphreys, 2001). In likelihood ratio models, rather than directly comparing an item’s strength to a decision criterion, the participant constructs an estimate of the strength’s likelihood if the item was a target divided by the strength’s likelihood if the item was a lure.

To understand how this produces the mirror effect, consider if a participant is tested in a condition where weak memory is expected, such as for items that were only briefly presented at study, resulting in a target distribution that is shifted only slightly higher than the lure distribution ($\mu_{lure} = 0$, $\mu_{target} = .5$, both distributions $\sigma = 1$). In this case, if a participant were to sample a strength $x = .25$ and compares it to both the target and lure distribution, the participant would be in an uncertain state because this strength value is equally likely under both distributions. However, consider if the participant enters another condition with relatively strong memory ($\mu_{target} = 3$) and receives another item with the same strength value ($x = .25$). In this case, it is extremely likely that this value is from the lure distribution, because low strength values are unlikely to come from a very strong strength distribution. This results in a lower false alarm rate for items that were studied in a strong condition.

The psychological theory behind likelihood ratio decisions is that decisions are not based directly on an item’s strength, but instead involve a comparison between an item’s strength and its expected strength. Memory models implement the likelihood ratios in

different ways, with several models involving likelihood ratio computations for each feature in a probe stimulus, as in ALT, REM, SLiM, BCDMEM, and the model of Cox and Shiffrin (2017), whereas in the Osth and Dennis (2015b) model, likelihood ratios are calculated on the basis of the item itself. It is interesting to note that each of these models implement wildly different explanations for the word frequency effect, including LF words attracting more attention (ALT), possessing uncommon yet diagnostic features (REM, Cox & Shiffrin), having lower encoding variability (SLiM), or having less interference from pre-experimental representations (BCDMEM, Osth & Dennis). However, each of these models owe their success in capturing the mirror effect to likelihood ratio decisions.

Decision on the basis of likelihood ratios is not without criticism. Hintzman (1994) criticized such theories for not providing a specification for where the expected memory strengths in a given condition come from. That is, how does a participant know to expect poorer performance on high frequency words, and how do they know how much poorer they should actually perform? This becomes especially problematic if one imagines a situation where a different degree of study time is given to each word, with some given very long durations and some very short durations. If a stimulus appears during the test list that was studied for a long duration, employing an expectation of high memory strength already presupposes memory for the word itself.

Several of the models have counter-acted this criticism by only employing *partially-informed* likelihood ratio decisions (Criss & McClelland, 2006). In models such as REM, no knowledge is assumed about whether a memory trace is of high or low frequency or whether the trace is strong or weak. Rather, the memory system only has access to both the number and identity of the features encoded, and this information is sufficient to calculate a likelihood ratio that produces a mirror effect. Likewise, in models such as BCDMEM and the Osth and Dennis model, if a mixed-strength list is studied, at test each of the words is compared to an expected target distribution that reflects the average strength of the words on the study list.

The Extralist Feature Effect

Some of the aforementioned successes of the global matching framework involve their predictions about the impairments of increasing list-length and increasing similarity between the probe and the contents of memory. However, one could argue that these are indirect tests of the models. A more direct test would involve manipulating the feature compositions of the studied list items and lures to evaluate the extent to which global similarity predicts performance for each trial. This was the reasoning behind the work of Mewhort and Johns (2000), who presented short lists of colored shapes for recognition. During the test phase, they constructed lures that exhibited different matches and mismatches to the studied list items on both the color and shape dimensions. For instance, a 1:1 lure would match a single studied item on the color dimension and a single studied item on the shape dimension, whereas a 2:0 lure would match two items on the color dimension and zero items on the shape dimension. Each of the different lure types along with an example study list are depicted in Figure 4.

Figure 4. An example study list of four colored shapes along with the different lure probe types used by Mewhort and Johns (2000).

The fundamental prediction of the global matching models is that rejection of lures

should be based on global similarity, meaning that a strong number of matches should impair performance regardless of the number of mismatches. In other words, a 2:0 lure and a 1:1 lure should be equally difficult to reject, because the match to two studied stimuli on the color dimension in the 2:0 lure is compensated by the mismatch on the shape dimension. However, contrary to these predictions, 2:0 lures were considerably easier to reject than 1:1 lures. Another problematic aspect for global matching models is that 1:0 lures did not show rejection advantages over 2:0 lures. The advantage for non-presented features was dubbed the *extralist feature effect* by Mewhort and Johns (2000) and has been replicated several times since (E. E. Johns & Mewhort, 2002, 2003; Osth, Zhou, Lilburn, & Little, 2019). Mewhort and Johns (2005) later demonstrated that popular models like REM and Minerva 2 were not able to produce the appropriate predictions for the paradigm. In the twenty years since the publication of the extralist feature effect by Mewhort and Johns, newer global matching models have been published without any obvious means of capturing the effect, including the exemplar-based random walk model (EBRW: Nosofsky et al., 2011), the tensor model of Osth and Dennis (2015b), and the dynamic recognition model of Cox and Shiffrin (2017).

Can global matching models be amended to capture the effect? Mewhort and Johns (2005) presented a model called the iterative resonance model (IRM) that qualitatively predicted the extralist feature effect. In the IRM, evidence for "old" decisions is based on global similarity, whereas evidence for "new" decisions is based on mismatches to the probe. Such a model may be able to capture the effect because the mismatch in 1:0 probes can trigger a "no" decision because the mismatch is not combined with the matching information, and therefore the global similarity does not cancel out the negative information present in the mismatch. However, Osth et al. (2019) pursued variants of the EBRW that implemented similar assumptions, and while they were capable of producing the extralist feature effect, they tended to perform poorly on other aspects of the data, particularly distributions of response latency.

Instead, the winning model in the Osth et al. model comparison was a model proposed by Nosofsky et al. (2011) where extra attention is devoted to the extralist dimensions when 1:0 and 2:0 probes are presented. This has the effect of giving that feature extra weight in the global similarity computation, such that 2:0 probes are less similar to the contents of memory than 1:1 probes. However, a disadvantage of such a model that they noted was that it is not obvious how participants can shift attention to and from an extralist feature dimension on a trial-by-trial basis.

The fact that a clear solution to the extralist feature effect has not yet emerged is an obvious remaining challenge for the global matching models. Mewhort and Johns (2000) demonstrated that even words demonstrated extralist feature advantages – words that contained letters not present on the list showed rejection advantages even if other letters were strongly represented. The fact that this challenge has not yet been addressed implies that there may be hidden sources of variability in the ability to reject lures that is not captured by current models.

Current and Future Directions for Global Matching Models

Predictions about Response Latency

Much of this chapter has focused on cases where the decision stage operates in a manner analogous to SDT – the global similarity of a probe is compared to a decision criterion and an "old" response is produced if the criterion is exceeded and a "new" response is made otherwise. But we also mentioned that part of the motivation for global matching came from the fact that participants can at times produce very fast rejection of lures – how can response latency be predicted from this framework?

Initial frameworks assumed that response latency could be predicted from the distance from a response criterion (e.g., Wickelgren & Norman, 1966). However, a difficulty with such models is that they often have difficulty predicting the shapes of response time distributions for correct and error responses. Specifically, when RTs are determined from a

distance-from-criterion model, errors occupy considerably less area in the target and lure distributions than correct responses, predicting that errors should have low RT variability relative to correct responses. However, in recognition memory paradigms error RT variability often equals or exceeds correct RT variability, in stark contrast to predictions from distance-from-criterion models (Murdock & Dufty, 1972; Ratcliff & McKoon, 2008).

An additional SDT framework derived RT predictions through the usage of two criteria instead of one. In the Hockley and Murdock (1987) decision model, a probe is recognized if it exceeds an upper criterion and is rejected if it is below a lower criterion. If a probe's strength is between the two criteria, however, the two criteria converge on the next iteration and the probe is compared to them again. This process continues until the probe is eventually accepted or rejected, and the RT is a quadratic function of the number of iterations. Unfortunately, an analysis of the model by Gronlund and Ratcliff (1991) found the model somewhat lacking in that it predicted multimodal RT distributions and encountered difficulties addressing speed-accuracy tradeoffs. These failures indicate that SDT frameworks may not be well suited for dealing with distributions of response latency.

Gillund and Shiffrin (1984) proposed abandoning the SDT decision rule and instead have global similarity drive a diffusion process (e.g., Ratcliff, 1978). In a diffusion process, evidence oscillates between two response boundaries corresponding to "old" and "new" decisions in a recognition memory experiment until it terminates at one of the boundaries. The boundary that is reached determines the choice and the time taken to reach the boundary is proportional to the RT. Evidence accumulation is noisy, such that even if evidence is generally tending toward the correct boundary, noise can drive it toward the incorrect boundary and produce an error response. A distinct advantage of diffusion models is that they naturally produce right-skewed and unimodal RT distributions. In addition, the parameters of the diffusion model are interpretable. The rate at which evidence accumulates is the drift rate, and increases in this parameter result in more frequent responses and faster latencies. Increases in the response boundary produce slower

decisions because the evidence has more distance to travel, but produces more accurate decisions because it is less likely that noise in the diffusion process will drive the accumulation toward the incorrect boundary.

Gillund and Shiffrin's suggestion was first undertaken in the exemplar-based random walk (EBRW) model (Nosofsky et al., 2011). In the EBRW, the GCM model of Nosofsky (1986, 1988) is used to determine the global similarity between the probe and the contents of memory. This global similarity subsequently drives a random walk process between "old" and "new" decision boundaries. A random walk is similar to a diffusion process, except that on each trial the drift rate, which is a normalization of the global similarity, is a probability of taking a discrete step toward the "old" decision boundary. Such a decision rule was essential, as the EBRW was applied to data from short-term recognition memory tasks (e.g., Sternberg, 1966, 1975) where accuracy can sometimes reach ceiling levels. The EBRW was successful in capturing a wide range of phenomena in such tasks, including effects of recency and list-length, and was able to capture both accuracy and RTs for such manipulations.

In the domain of long-term memory, a similar approach was undertaken by Osth, Jansson, Dennis, and Heathcote (2018), who extended the matrix model of Osth and Dennis (2015b) with a back-end diffusion process. In this model, the likelihood ratios of targets and lures derived from global similarity became the drift rates for the diffusion model. This model was able to capture a wide range of long-term recognition memory phenomena, including list-length, word frequency, study-test delay, list-strength, and decreases in performance with increasing test position, with the model capturing hit and false alarm rates along with the shapes of correct and error RT distributions in each case. Part of the utility of using a back-end diffusion process in this case is that the goal of the modeling was to quantify the extent to which certain model parameters were responsible for decreases in performance with increasing numbers of test trials. Some accounts stated that this phenomenon was due to interference from the learning of test items (e.g., Criss,

Malmberg, & Shiffrin, 2011; Gillund & Shiffrin, 1984), changes in context during the course of retrieval (e.g., Howard & Kahana, 2002a; Osth & Dennis, 2015b), or changes in speed-accuracy thresholds over trials (e.g., Ratcliff, 1978). If changes in speed-accuracy thresholds are occurring over trials, such that some participants are speeding up through the course of testing, these decreases in accuracy would be attributed to memory processes if RTs were not jointly considered. By using a back-end diffusion process, changes in accuracy can be properly partitioned into accuracy decreases that are due to interference from memory vs. those that are due to changes decision-making strategies.

While the usage of a back-end diffusion process can capture RT distributions, there are other phenomena in recognition memory that suggest changes in the information used to cue memory and generate global similarity through the course of a decision. One of the most noteworthy examples is the plurality discrimination paradigm, where participants study singular or plural words and at test are tested on switched plurality lures, such that if a participant were to study the words "cat" and "tires", "cats" and "tire" would be examples of switched-plurality lures. Hintzman and Curran (1994) investigated the switched-plurality task using the signal-to-respond methodology, where participants are given cues to respond at various signal lags after probe presentation, such as 100, 250, 500, 800, or 1,200 ms. Hintzman and Curran found a non-monotonic relationship between the FAR to switched-plurality lures and processing time – early in retrieval, HR and FAR to switched plurality lures were virtually identical, whereas late in retrieval HR and FAR diverged, with the FAR decreasing substantially as the HR rose.

Hintzman and Curran initially attributed their results to a hybrid search process, in which both targets and their switched plurality lures elicit very similar values of global similarity but a recall-like process that occurs late in the trial can be used to distinguish similar lures from targets. However, Brockdorff and Lamberts (2000) argued that such results can be achieved by assuming a *feature sampling* process through the course of the test trial. Specifically, they proposed the feature-sampling theory of recognition

(FESTHER), which resembles the GCM but instead of having all of the probe dimensions available at the beginning of the trial, dimensions are sampled randomly with replacement as the test trial progresses. To capture Hintzman and Curran's data, they assumed that in the switched-plurality paradigm the features that denote whether a word is singular or plural are not sampled until after the other features are sampled. This means that if a probe word "cats" was presented, the letters "c", "a", and "t" would need to be sampled before the plurality feature "s" enters the cue. The gradual accumulation of features that are congruent with both "cat" and "cats" early in the trial results in a rising FAR for switched-plurality lures early in the trial, whereas the late entrance of the critical letter "s" late in the trial results in the decrease in FAR later in the trial.

Recently, Cox and Shiffrin (2017) presented an extension of the feature sampling approach in a dynamic model of recognition memory that bears a strong resemblance to the REM model. Like FESTHER, the model contains a feature sampling process, with the exception being that global similarity is not the direct basis of the decision. Instead, as features are added to the cue, the changes in global similarity drive a set of accumulators, with positive changes producing evidence for "old" decisions and negative changes producing evidence for "new" decisions. The Cox and Shiffrin model extends the FESTHER approach by also addressing a number of other major phenomena, including the late onset of associative information relative to item information (Gronlund & Ratcliff, 1989) and RT distributions for correct and error responses. Both the FESTHER and Cox and Shiffrin models mark a significant departure from the two-stage models, where retrieval drives an evidence accumulation process, in that retrieval and decision-making are fundamentally intertwined.

Collectively, the results of the current models suggest that the current crop of global matching models with more sophisticated decision stages are successful in addressing the major hurdles from response latency data. In particular, they can address the shapes of correct and error RT distributions for each of the usual benchmarks, such as recency,

list-length, list-strength, and word frequency effects, and are also capable of addressing qualitative changes in performance with increasing processing time.

Sources of Interference in Recognition Memory

Many of the discussed models operate under the assumption that forgetting is driven entirely by interference and not by the passage of time (e.g., Waugh & Norman, 1965; Lewandowsky, Oberauer, & Brown, 2008; Oberauer & Lewandowsky, 2008). In global matching models interference operates at retrieval because even if a target cue is presented for recognition, non-target memories will contribute to the global similarity computation, with the degree of interference depending on both the similarity of the memories to the cue in addition to the number of non-target memories.

In recent years however, there has been increasing theoretical interest in which types of memories might be causing the interference that drives forgetting. In their inception, the global matching models assumed that the contents of memory comprised only the items that were learned on the study list. While this assumption was intended to simplify the models, it reflects a theoretical commitment that interference is driven by the studied items in the learning episode and not by memories acquired *prior* to the learning episode. The idea that interference stems from the items in the same learning episode has been retroactively referred to as the *item noise* conception of interference.

An alternative to item noise is the idea that the past contexts in which an item has been experienced is what drives interference, which is referred to as *context noise*. Context noise models, which include the model of Anderson and Bower (1972) and the bind-cue decide model of episodic memory (BCDMEM: Dennis & Humphreys, 2001), posit that at learning items are bound to the context in which they were experienced, where the term "context" is used to refer to a representation of the learning episode. At test, both the context of the study list and the item are jointly used to probe the contents of memory, essentially asking the question "Was this word in the context of interest?"⁴ Furthermore,

⁴It should be mentioned that other models have incorporated context representations, including SAM,

items are assumed to have been associated to past contexts in which the word was experienced to reflect their frequency of experience in episodes prior to the learning episode. If the current context shares elements with previous contexts in which a probe word was experienced, this can lead to confusions as to whether or not the word was experienced in the current context.

To give a concrete example, consider if a word such as "cat" was studied. "Cat" is of relatively high frequency in natural language, and as a consequence it can suffer interference from the many prior contexts in which the word was experienced. Consider if instead a relatively rare word such as "samurai" was studied. Because there are less contexts in which this word has been experienced, it will suffer less interference from prior contexts and performance on this word would be expected to be higher. For this reason, context noise models have a natural explanation for the word frequency in recognition memory, specifically that low frequency words have been associated to fewer contexts and suffer less interference from them. Context noise models posit that the word-frequency effect is essentially a list-length effect that is occurring over the past contexts in which a word has been experienced.

Context noise models also offer a natural explanation of the null list-strength effect. Because only the past contexts associated with an item impair performance and not the other items on the study list, the strength of the other items cannot impair performance. However, an additional consequence is that the number of other list items cannot impair performance either, making them unable to predict impairments with increasing list-length. However, Dennis and Humphreys (2001) countered by arguing that a number of confounds exist in list-length designs. For instance, in some early list-length designs, retention interval was not controlled between the short and long lists, leading to longer retention intervals for earlier list items in the long list condition. In addition, it is possible that attention decreases during the study list, impairing performance for the later items in a

REM, and TODAM (e.g., Murdock, 1997), but the key departure of pure context noise models was the postulation that it is *only* the past contexts in which items were experienced that caused forgetting.

long list (e.g., Underwood, 1978). When each of the proposed confounds are controlled, Dennis and colleagues have routinely observed equal performance for short and long lists (Dennis & Humphreys, 2001; Dennis, Lee, & Kinnell, 2008; Kinnell & Dennis, 2011).

Osth and Dennis (2015b) used an approach where item noise and context noise were estimated within a single model. They also incorporated the assumptions of the Murdock and Kahana (1993) model by incorporating interference from memories beside the cue in other contexts, a type of interference they referred to as *background noise*. Each of these sources of interference can be categorized on the basis of matches or mismatches to the item and context cues employed at retrieval. Specifically, memories that match the context cue (the relevant study list) but mismatch the item cue can be considered item noise, memories that match the item cue but mismatch the context cue (i.e., prior occurrences of the cue at retrieval) can be considered context noise, and memories that mismatch both the item and context cues (other words in other contexts) can be considered background noise. Each of these sources of interference were measured in a matrix model where the similarities between the current and prior contexts in addition to the similarities between the different items were measured. If representations of prior contexts are completely dissimilar to the current context, then memories associated with those prior contexts will not interfere and there will be no context noise or background noise and item noise will comprise the interference at retrieval. If representations of items are highly dissimilar to each other, however, then there will be no item noise and context and background noise will dominate.

Osth and Dennis jointly estimated each of these parameters across ten recognition memory datasets that manipulated list-strength, list-length, study-test delay, word frequency, and stimulus class, and both item and associative recognition were tested with words. The results indicated a dominance of background noise, with almost no context noise for low frequency words and roughly equal contributions of context and background noise for high frequency words. Item noise exerted only a negligible influence on words and natural scenes, while there were larger influences for visually confusable stimuli such as

faces and fractal images. Moreover, in recent years the model was extended to complete RT distributions (Osth, Jansson, et al., 2018) and to source memory (Osth, Fox, et al., 2018) – the introductions of each of these constraints did not change the original conclusions, reinforcing the idea that forgetting in recognition memory is caused by memories acquired prior to the list episode.

These findings are not without controversy. A recent investigation by Brandt et al. (2019) suggested that there were influences of item noise that are often buried by an experimental design. In analyses of manipulations of list-length, Brandt et al. found substantial impairments in performance in long lists when only the first study-test cycle was considered, while finding no differences between conditions in the aggregated data. They argued that when participants are tested on their second study list, interference from the first study list still contributes and this can lead to equal interference in both list-length conditions. To understand why, consider if one participant learns a list of 20 items while the other learns a list of 80 items. However, on the second iteration, the first participant learns a list of 80 items, leading to a total of $20 + 80 = 100$ items in memory. The other participant then learns a list of 20 items, leading to a total of $80 + 20 = 100$ items in memory. Thus, the interference differences between short and long lists are equated, and no performance impairment should be observed. In response, Fox, Dennis, and Osth (2020) modeled the buildup of interference across lists in a list-length paradigm using the Osth, Jansson, et al. (2018) model. The results of the modeling reversed the conclusions of the initial Osth and Dennis work, showing a strong dominance of item noise!

Nonetheless, it is a somewhat undesirable state of affairs that several of these models have measured interference by estimating similarities between items and contexts as free parameters. In the next section, we describe how models could potentially be augmented with more realistic representations of stimuli, such that the similarities between these representations would be naturally specified. Such an approach might offer further constraint in the ongoing debate concerning the sources of interference.

Item Representations

In many of the global matching models, representations of items are often described as "vectors of features." But what are features exactly? Many of the models assume random representations, such as samples from a normal distribution (as in TODAM and the matrix models), or a set of discrete values, such as +1, 0, and -1, as in the Minerva 2 model and samples from a geometric distribution in the REM model. Randomly generated representations such as these are often used for two reasons. The first is that many of these models were applied to data from experiments using words as stimuli, and at the time it wasn't exactly clear what representations of words are. The second is that by making minimal assumptions about representation, it becomes clear that the core assumptions about retrieval are doing most of the work in generating the model predictions rather than representations that might be hand-tailored to the particular problem at hand.

Nonetheless, since the development of the initial global matching models, a number of models have emerged that can learn semantic representations for words in an unsupervised fashion. These include latent semantic analysis (LSA: Landauer & Dumais, 1997), the hyperspace analog to language (HAL: Lund & Burgess, 1996), and the bound encoding of the aggregate language environment (BEAGLE: Jones & Mewhort, 2007), and the semantic models based on Minerva 2 (Kwantes, 2005; Jamieson et al., 2018). Each of these models produces semantic representations for words on the basis of co-occurrence in a large corpus of natural text, producing vector representations for words that are similar to each other to the extent that they co-occur with each other in different linguistic contexts. Specifically, representations of the words "cat" and "dog" may be similar to each other because they co-occur in the same types of documents, such as documents about pet ownership or animals, whereas the words "dog" and "coffee" may be dissimilar because they occur in different types of documents. Each of these models has been shown to demonstrate some rather impressive successes, such as bearing a strong resemblance to human generated similarity ratings and a strong correspondence with free association data,

to name only a couple examples.

An obvious future direction for current models of recognition memory is to substitute the current randomly generated representations with semantic representations from one of the aforementioned models. A disadvantage of random representations is that when such models are being applied to paradigms with similarity effects, typically the models have to invoke additional parameters for either the proportion of shared features between items from the same category (e.g., Criss & Shiffrin, 2004; Shiffrin & Steyvers, 1997) or for the degree of distortion from a prototype (e.g., Arndt & Hirshman, 1998; Hintzman, 1988). If semantic representations for actual words were substituted, the models would be able to make predictions about the effects of similarity without any additional parameters. Such an integration of models seems particularly promising because the semantic representations that emerge from these models are vectors which can be compared using dot products, meaning that the architectures of such models are likely already compatible with such representations.

Such an approach was taken by the recognition-by-semantic-synchronization (RSS) model (B. T. Johns, Jones, & Mewhort, 2012). This model was successful in addressing both false memory effects in the Deese-Roediger-McDermott procedure (DRM: Roediger & McDermott, 1995) and impairments in performance with increasing category length while simultaneously being able to capture a number of benchmarks for recognition models such as the null list-strength effect, the proper shapes of ROCs, and the word frequency effect. The latter demonstration is particularly impressive, as the model was able to capture the advantage for LF words without any additional parameters. Instead, LF words outperform HF words because their semantic representations are less similar to other words than HF words (see also Monaco et al., 2007).

Outside of the RSS model, other current models of long-term recognition memory, such as the models of Osth, Jansson, et al. (2018) and Cox and Shiffrin (2017), have continued to employ random representations of words. Extension of such models with

realistic semantic representations is clearly an important current direction for such models, as it would further constrain accounts of similarity effects with mechanisms that generate predictions about response latency. In the domain of free recall, a number of models have already demonstrated considerable success in utilizing semantic representations (Morton & Polyn, 2016; Polyn, Norman, & Kahana, 2009). This has enabled such models to make transitions between semantically similar words even when presented in lists of unrelated words (e.g., Howard & Kahana, 2002b).

When stimuli other than words are considered, a number of models based on the GCM have been highly successful in capturing similarity effects. In the GCM, representations of simple perceptual stimuli have often been derived from multidimensional scaling (MDS) solutions based on similarity ratings (Nosofsky, 1986). Nosofsky et al. (2011) conducted an experiment manipulating list-length of Munsell colors which can be represented with three dimensions. Their EBRW model was able to account for both choice proportions and mean RT for each individual color tested in their experiment. The success in accounting for variability over items is impressive, as in the EBRW model performance for a color is entirely dependent on its global similarity to other colors on the study list. In other words, an "easy" color is often dissimilar to other colors, whereas a "difficult" color is often highly similar to other colors. These results suggest that the successes of global matching models can be extended beyond the ability to capture qualitative changes in performance with increases in list-length or category length. Instead, models may be able to yield predictions about the memorability of individual items.

Conclusions

The global matching mechanism provides a simple and effective means of leveraging structure across a set of memories. They not only provide a retrieval mechanism for recognition decisions, they have considerable breadth in their extensions to other episodic memory tasks including recall, semantic memory, and categorization. They have been met

with some considerable challenges over the decades, but the fact that newer models have been developed that provide simple explanations of both the mirror effect and the null list-strength effect along with extensions to response latency suggest that the models are still viable models of episodic memory. While there still exist some challenges such as the extralist feature effect, to date there have been no alternatives to the global matching mechanism that are sufficiently broad as to capture all of the phenomena described in this chapter.

Table 1
List of the global matching models. Columns depict how items are represented, the matching operation between cues and memories, the aggregation method to produce global similarity, and the association mechanism (when applicable).

Model	Representation	Matching operation	Aggregation	Association
Search of Associative Memory (SAM, Gillund & Shiffrin, 1984)	Images	Cue-image Strengths	Sum	Links
Minerva (Hintzman, 1986, 1988)	Vectors	Dot product	Sum	Concatenation
Theory of Distributed Associative Memory (TODAM, Murdock, 1982)	Vectors	Dot product	Sum	Convolution
Matrix (Humphreys, Bain, & Pike, 1989; Pike, 1984)	Vectors	Dot product	Sum	Outer product
Generalized context model (GCM: Nosofsky, 1988)	Vectors	Distance	Sum	n/a
Retrieving effectively from memory (REM: Shiffrin & Steyvers, 1997)	Vectors	Likelihood ratio	Average	Concatenation
Subjective likelihood in memory (SLiM: McClelland & Chappell, 1998)	Vectors	Likelihood ratio	Maximum	Concatenation
Bind-cue-decide model of episodic memory (BCDMEM: Dennis & Humphreys, 2001)	Vectors	Likelihood ratio	n/a	n/a
Recognition by semantic synchronization (RSS: Johns, Jones, & Mewhort, 2014)	Vectors	Cosine	Sum	n/a
Osth & Dennis (2015)	Vectors	Dot product	Sum	Outer product
Cox & Shiffrin (2017)	Vectors	Likelihood ratio	Average	Concatenation

References

- Anderson, J. R., & Bower, G. H. (1972). Recognition and retrieval processes in free recall. *Psychological Review*, 79(2), 97–123.
- Arndt, J., & Hirshman, E. (1998). True and False Recognition in MINERVA2: Explanations from a Global Matching Perspective. *Journal of Memory and Language*, 39, 371–391.
- Atkinson, R. C., & Juola, J. F. (1974). Search and decision processes in recognition memory. In D. H. Krantz, R. C. Atkinson, R. D. Luce, & P. Suppes (Eds.), *Contemporary developments in mathematical psychology: Vol. 1. Learning, memory and thinking*. San Francisco: Freeman.
- Brandt, M., Zaiser, A., & Schnuerch, M. (2019). Homogeneity of item material boosts the list length effect in recognition memory: A global matching perspective. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 45, 834–850.
- Brockdorff, N., & Lamberts, K. (2000). A feature-sampling account of the time course of old-new recognition judgments. *Journal of Experimental Psychology-Learning Memory and Cognition*, 26(1), 77–102.
- Cary, M., & Reder, L. M. (2003). A dual-process account of the list-length and strength-based mirror effects in recognition. *Journal of Memory and Language*, 49(2), 231–248.
- Chappell, M., & Humphreys, M. S. (1994). An auto-associative neural network for sparse representations: Analysis and application to models of recognition and cued recall. *Psychological Review*, 101, 103–128.
- Clark, S. E., & Gronlund, S. D. (1996). Global matching models of recognition memory: How the models match the data. *Psychonomic Bulletin and Review*, 3(1), 37–60.
- Cox, G. E., & Shiffrin, R. M. (2017). A dynamic approach to recognition memory. *Psychological Review*, 124(6), 795–860.
- 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., Malmberg, K. J., & Shiffrin, R. M. (2011). Output interference in recognition memory. *Journal of Memory and Language*, 64, 316–326.
- Criss, A. H., & McClelland, J. L. (2006). Differentiating the differentiation models: A comparison of the retrieving effectively from memory model (REM) and the subjective likelihood model (SLiM). *Journal of Memory and Language*, 55, 447–460.
- Criss, A. H., & Shiffrin, R. M. (2004). Context-noise and item-noise jointly determine recognition memory: A comment on Dennis and Humphreys (2001). *Psychological Review*, 111, 800–807.
- Deese, J. (1959). On the prediction of the occurrence of particular verbal intrusions in immediate recall. *Journal of Experimental Psychology*, 58, 17–22.
- Dennis, S., & Humphreys, M. S. (2001). A context noise model of episodic word recognition. *Psychological Review*, 108(2), 452–478.
- Dennis, S., Lee, M. D., & Kinnell, A. (2008). Bayesian Analysis of Recognition Memory: The Case of the List-Length Effect. *Journal of Memory and Language*, 59, 361–376.
- Dougherty, M. R. P., Gettys, C. F., & Ogden, E. E. (1999). MINERVA-DM: A memory processes model for judgements of likelihood. *Psychological Review*, 106(1), 180–209.
- Fox, J., Dennis, S., & Osth, A. F. (2020). Accounting for the build-up of proactive interference across lists in a list length paradigm reveals a dominance of item-noise in recognition memory. *Journal of Memory and Language*, 110.
- Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review*, 91(1), 1–67.
- Glanzer, M., & Adams, J. K. (1985). The mirror effect in recognition memory. *Memory and Cognition*, 13(1), 8–20.
- Glanzer, M., Adams, J. K., Iverson, G. J., & Kim, K. (1993). The regularities of recognition memory. *Psychological Review*, 100(3), 546–567.
- Glanzer, M., & Bowles, N. (1976). Analysis of the word-frequency effect in recognition

- memory. *Journal of Experimental Psychology: Human learning and memory*, 2(1), 21–31.
- Glanzer, M., Hilford, A., & Maloney, L. T. (2009). Likelihood ratio decisions in memory: Three implied regularities. *Psychonomic Bulletin & Review*, 16(3), 431–455.
- Glanzer, M., Kim, K., Hilford, A., & Adams, J. K. (1999). Slope of the receiver-operating characteristic in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25(2), 500–513.
- Goldinger, S. D. (1998). Echoes of echoes? An episodic theory of lexical access. *Psychological Review*, 105(2), 251–279.
- Gronlund, S. D., & Ratcliff, R. (1989). Time course of item and associative information: Implications for global memory models. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 15(5), 846–858.
- Gronlund, S. D., & Ratcliff, R. (1991). Analysis of the Hockley and Murdock decision model. *Journal of Mathematical Psychology*, 35(3), 319–344.
- Hintzman, D. L. (1986). Schema Abstraction in a Multiple-Trace Memory Model. *Psychological Review*, 93(4), 411–428.
- Hintzman, D. L. (1988). Judgments of Frequency and Recognition Memory in a Multiple-Trace Memory Model. *Psychological Review*, 95(4), 528–551.
- Hintzman, D. L. (1994). On Explaining the Mirror Effect. *Journal of Experimental Psychology-Learning Memory and Cognition*, 20(1), 201–205.
- Hintzman, D. L., & Curran, T. (1994). Retrieval Dynamics of Recognition and Frequency Judgments - Evidence for Separate Processes of Familiarity and Recall. *Journal of Memory and Language*, 33(1), 1–18.
- Hirshman, E. (1995). Decision processes in recognition memory: Criterion shifts and the list-strength paradigm. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21(2), 302–313.
- Hockley, W. E., & Murdock, B. B. (1987). A Decision Model for Accuracy and Response

- Latency in Recognition Memory. *Psychological Review*, 94(3), 341–358.
- Howard, M. W., & Kahana, M. J. (2002a). A distributed representation of temporal context. *Journal of Mathematical Psychology*, 46, 268–299.
- Howard, M. W., & Kahana, M. J. (2002b). When does semantic similarity help episodic retrieval? *Journal of Memory and Language*, 46(1), 85–98.
- 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(2), 208–233.
- Humphreys, M. S., Pike, R., Bain, J. D., & Tehan, G. (1989). Global matching: A comparison of the SAM, Minerva II, Matrix and TODAM models. *Journal of Mathematical Psychology*, 33, 36–67.
- Jamieson, R. K., Avery, J. E., Johns, B. T., & Jones, M. N. (2018). An instance theory of semantic memory. *Computational Brain & Behavior*, 1(2), 119–136.
- Jamieson, R. K., & Mewhort, D. J. K. (2009). Applying an exemplar model to the artificial-grammar learning task: Inferring grammaticality from similarity. *The Quarterly Journal of Experimental Psychology*, 62, 550–575.
- Johns, B. T., Jones, M. N., & Mewhort, D. J. K. (2012). A synchronization account of false recognition. *Cognitive Psychology*, 65(4), 486–518.
- Johns, E. E., & Mewhort, D. J. K. (2002). What information underlies correct rejections in short-term recognition memory? *Memory & Cognition*, 30(1), 46–59.
- Johns, E. E., & Mewhort, D. J. K. (2003). The effect of feature frequency on short-term recognition memory. *Memory & Cognition*, 31(2), 285–296.
- Jones, M. N., & Mewhort, D. J. K. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological Review*, 114(1), 1–37.
- Kahana, M. J., Rizzuto, D. S., & Schneider, A. R. (2005). Theoretical correlations and measured correlations: Relating recognition and recall in four distributed memory models. *Journal of Experimental Psychology: Learning Memory and Cognition*,

31(5), 933–953.

Kahana, M. J., & Sekuler, R. (2002). Recognizing spatial patterns: a noisy exemplar approach. *Vision Research*, 42(18), 2177–2192.

Kellen, D., & Klauer, K. C. (2014). Discrete-state and continuous models of recognition memory: Testing core properties under minimal assumptions. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 40, 1795–1804.

Kelly, M. A., Mewhort, D. J. K., & West, R. L. (2017). The memory tesseract: Mathematical equivalence between composite and separate storage models. *Journal of Mathematical Psychology*, 77, 142–155.

Kinnell, A., & Dennis, S. (2011). The list length effect in recognition memory: An analysis of potential confounds. *Memory & Cognition*, 39, 348–363.

Kintsch, W. (1967). Memory and decision aspects of recognition learning. *Psychological Review*, 74, 496–504.

Kwantes, P. J. (2005). Using context to build semantics. *Psychonomic Bulletin and Review*, 12, 703–710.

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato’s problem: The Latent Semantic Analysis theory of the acquisition, induction, and representation of knowledge. *Psychological Review*, 104, 211–240.

Lehman, M., & Malmberg, K. J. (2013). A buffer model of memory encoding and temporal correlations in retrieval. *Psychological Review*, 120(1), 155–189.

Lewandowsky, S., & Murdock, B. B. (1989). Memory for serial order. *Psychological Review*, 96, 25–57.

Lewandowsky, S., Oberauer, K., & Brown, G. D. A. (2008). No temporal decay in verbal short-term memory. *Trends in Cognitive Sciences*, 13, 120–126.

Logan, G. D. (2002). An instance theory of attention and memory. *Psychological Review*, 109, 376–400.

Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical

- co- occurrence. *Behavior Research Methods Instruments & Computers*, 28(2), 203–208.
- McClelland, J. L., & Chappell, M. (1998). Familiarity breeds differentiation: A subjective-likelihood approach to the effects of experience in recognition memory. *Psychological Review*, 105(4), 724–760.
- McDowd, J. M., & Murdock, B. B. (1986). Mathematical models of memory and the problem of stimulus variation: A comparison of Minerva 2 and TODAM. *Acta Psychologica*, 62, 177–188.
- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, 85(3), 207–238.
- Mewhort, D. J. K., & Johns, E. E. (2000). The extralist-feature effect: Evidence against item matching in short-term recognition memory. *Journal of Experimental psychology: General*, 129(2), 262–284.
- Mewhort, D. J. K., & Johns, E. E. (2005). Sharpening the echo: An iterative-resonance model for short-term recognition memory. *Memory*, 13, 300–307.
- Monaco, J. D., Abbott, L. F., & Kahana, M. J. (2007). Lexico-semantic structure and the word-frequency effect in recognition memory. *Learning & Memory*, 14, 204–213.
- Morton, N. W., & Polyn, S. M. (2016). A predictive framework for evaluating models of semantic organization in free recall. *Journal of Memory and Language*, 86, 119–140.
- Murdock, B. B. (1974). *Human memory: Theory and data*. Erlbaum.
- Murdock, B. B. (1982). A theory for the storage and retrieval of item and associative information. *Psychological Review*, 89(6), 609–626.
- Murdock, B. B. (1997). Context and mediators in a theory of distributed associative memory (TODAM2). *Psychological Review*, 104(4), 839–862.
- Murdock, B. B. (2006). Decision-making models of remember-know judgments: Comment on Rotello, Macmillan, and Reeder (2004). *Psychological Review*, 113(3), 648–656.
- Murdock, B. B., & Dufty, P. O. (1972). Strength theory and recognition memory. *Journal*

- of Experimental Psychology*, 94(3), 284–290.
- Murdock, B. B., & Kahana, M. J. (1993). Analysis of the list-strength effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 19(3), 689–697.
- Murdock, B. B., & Lamon, M. (1988). The Replacement Effect - Repeating Some Items While Replacing Others. *Memory & Cognition*, 16(2), 91–101.
- Murdock, B. B., Smith, D., & Bai, J. (2001). Judgments of frequency and recency in a distributed memory model. *Journal of Mathematical Psychology*, 45(4), 564–602.
- Murnane, K., & Shiffrin, R. M. (1991a). Interference and the representation of events in memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17(5), 855–874.
- Murnane, K., & Shiffrin, R. M. (1991b). Word repetitions in sentence recognition. *Memory & Cognition*, 19(2), 119–130.
- Norman, K. A., & O'Reilly, R. C. (2003). Modeling hippocampal and neocortical contributions to recognition memory: A complementary learning systems approach. *Psychological Review*, 110(4), 611–646.
- Nosofsky, R. M. (1986). Attention, similarity and the identification-categorization relationship. *Journal of Experimental psychology: General*, 115, 39–57.
- Nosofsky, R. M. (1988). Exemplar-based accounts of relations between classification, recognition, and typicality. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 700–708.
- Nosofsky, R. M. (1991). Tests of an exemplar model for relating perceptual classification and recognition memory. *Journal of Experimental Psychology: Human Perception and Performance*, 17(1), 3–27.
- Nosofsky, R. M., Little, D. R., Donkin, C., & Fific, M. (2011). Short-term memory scanning viewed as exemplar-based categorization. *Psychological Review*, 118(2), 280–315.
- Oberauer, K., & Lewandowsky, S. (2008). Forgetting in immediate serial recall: Decay,

- temporal distinctiveness, or interference? *Psychological Review*, 115, 544–576.
- Osth, A. F., & Dennis, S. (2014). Associative recognition and the list strength paradigm. *Memory & Cognition*, 42(4), 583–594.
- Osth, A. F., & Dennis, S. (2015a). A prospective for a unified model of episodic memory. In J. G. W. Raaijmakers, A. H. Criss, R. L. Goldstone, R. M. Nosofsky, & M. Steyvers (Eds.), *Cognitive Modeling in Perception and Memory: A Festschrift for Richard M. Shiffrin*. Psychology Press.
- Osth, A. F., & Dennis, S. (2015b). Sources of interference in item and associative recognition memory. *Psychological Review*, 122(2), 260–311.
- Osth, A. F., Dennis, S., & Heathcote, A. (2017). Likelihood ratio sequential sampling models of recognition memory. *Cognitive Psychology*, 92, 101–126.
- Osth, A. F., Dennis, S., & Kinnell, A. (2014). Stimulus type and the list strength paradigm. *Quarterly Journal of Experimental Psychology*, 67(9), 1826–1841.
- Osth, A. F., Fox, J., McKague, M., Heathcote, A., & Dennis, S. (2018). The list strength effect in source memory: Data and a global matching model. *Journal of Memory and Language*, 103, 91–113.
- Osth, A. F., Jansson, A., Dennis, S., & Heathcote, A. (2018). Modeling the dynamics of recognition memory testing with an integrated model of retrieval and decision making. *Cognitive Psychology*, 104, 106–142.
- Osth, A. F., Zhou, A., Lilburn, S., & Little, D. R. (2019). The extralist feature effect revisited: A challenge for global matching models of recognition memory. *PsyArxiv Preprint*.
- Parks, C. M., & Yonelinas, A. P. (2009). Evidence for a memory threshold in second-choice recognition memory responses. *Proceedings of the National Academy of Sciences*, 106, 11515–11519.
- Pike, R. (1984). Comparison of Convolution and matrix distributed memory systems for associative recall and recognition. *Psychological Review*, 91(3), 281–294.

- Polyn, S. M., Norman, K. A., & Kahana, M. J. (2009). A context maintenance and retrieval model of organizational processes in free recall. *Psychological Review*, *116*(1), 129–156.
- Raaijmakers, J. G. W., & Shiffrin, R. M. (1981). Search of Associative Memory. *Psychological Review*, *88*(2), 93–134.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*, 59–108.
- 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*(2), 163–178.
- Ratcliff, R., & McKoon, G. (1988). A retrieval theory of priming in memory. *Psychological Review*, *95*, 385–408.
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, *20*, 873–922.
- Ratcliff, R., McKoon, G., & Tindall, M. (1994). Empirical generality of data from recognition memory: Receiver-operating characteristic functions and implications for the global memory models. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*(4), 763–785.
- Ratcliff, R., Sheu, C. F., & Gronlund, S. D. (1992). Testing global memory models using ROC curves. *Psychological Review*, *99*(3), 518–535.
- Robinson, K. J., & Roediger, H. L. (1997). Associative processes in false recall and false recognition. *Psychological Science*, *8*, 231–237.
- Roediger, H. L., & McDermott, K. B. (1995). Creating false memories: Remembering words not present in lists. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *21*, 803–814.
- Schacter, D. L., Eich, J. E., & Tulving, E. (1978). Richard Semon’s theory of memory. *Journal of Verbal Learning and Verbal Behavior*, *17*, 721–743.
- Shepard, R. N. (1987). Toward a universal law of generalization for psychological science.

- Science*, 237, 1317–1323.
- Shiffrin, R. M., Huber, D. E., & Marinelli, K. (1995). Effects of category length and strength on familiarity in recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21(2), 267–287.
- Shiffrin, R. M., Ratcliff, R., & Clark, S. E. (1990). List-strength effect: II. Theoretical mechanisms. *Journal of Experimental Psychology: Learning Memory and Cognition*, 16(2), 179–195.
- Shiffrin, R. M., & Steyvers, M. (1997). A model for recognition memory: REM - retrieving effectively from memory. *Psychonomic Bulletin & Review*, 4(2), 145–166.
- Starns, J. J., & Ratcliff, R. (2014). Validating the unequal-variance assumption in recognition memory using response time distributions instead of ROC functions: A diffusion model analysis. *Journal of Memory and Language*, 70, 36–52.
- Sternberg, S. (1966). High-speed scanning in human memory. *Science*, 153(3736), 652–654.
- Sternberg, S. (1975). Memory scanning: New findings and current controversies. *Quarterly Journal of Experimental Psychology*, 27, 1–32.
- Stretch, V., & Wixted, J. T. (1998). On the differences between strength-based and frequency-based mirror effects in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24(6), 1379–1396.
- Strong, E. K. J. (1912). The effect of length of series upon recognition memory. *Psychological Review*, 19, 447–462.
- Thomas, R. P., Dougherty, M. R. P., Sprenger, A. M., & Harbison, J. I. (2008). Diagnostic hypothesis generation and human judgment. *Psychological Review*, 115, 155–185.
- Tulving, E., & Hastie, R. (1972). Inhibition Effects of Intralist Repetition in Free Recall. *Journal of Experimental Psychology*, 92(3), 297–304.
- Underwood, B. J. (1978). Recognition memory as a function of the length of study list. *Bulletin of the Psychonomic Society*, 12, 89–91.

- Wagenmakers, E.-J., Steyvers, M., Raaijmakers, J. G. W., Shiffrin, R. M., van Rijn, H., & Zeelenberg, R. (2004). A model for evidence accumulation in the lexical decision task. *Cognitive Psychology*, *48*(3), 332–367.
- Waugh, N. C., & Norman, D. A. (1965). Primary memory. *Psychological Review*, *72*, 89–104.
- Wickelgren, W. A., & Norman, D. A. (1966). Strength models and serial position in short-term recognition memory. *Journal of Mathematical Psychology*, *3*, 316–347.
- Wilson, J. H., & Criss, A. H. (2017). The list strength effect in cued recall. *Journal of Memory and Language*, *95*, 78–88.
- Wixted, J. T. (2007). Dual-process theory and signal-detection theory of recognition memory. *Psychological Review*, *114*(1), 152–176.
- Yonelinas, A. P. (2002). The nature of recollection and familiarity: A review of 30 years of research. *Journal of Memory and Language*, *46*(3), 441–517.
- Yonelinas, A. P., Hockley, W. E., & 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.