A theory of memory for items and associations

Beige J. Jin, David J. Halpern, and Michael J. Kahana

Department of Psychology, University of Pennsylvania

The authors dedicate this paper to the late Professor Bennet Murdock

Draft, October 6, 2024

Author Note

The authors dedicate this paper to the late Professor Bennet Bronson Murdock. May his memory be a blessing for all who knew him. We also gratfully acknowledge support from National Institutes of Health Grant MH55687. Thanks are due to Masaki Hori for programming Experiment 3 and for analyzing data from that experiment. Jeremy Caplan, Marc Howard, and Dan Rizzuto provided many valuable insights into the analysis of associative memory and successive memory tests. Correspondence concerning this article should be addressed to Michael Jacob Kahana, kahana@psych.upenn.edu.

Abstract

We present a retrieved-context theory of memory for items, associations, and their interaction (CMR-IA). Our theory assumes an evolving representation of temporal context that binds to items and associations, allowing the rememberer to make judgments based on the occurrence of a mnemonic target within a particular context. In addition to the assumptions inherited from prior retrieved-context theories, CMR-IA assumes a conjunctive (Gestalt) representation for paired associates, increased attention to rare items, and variable thresholds for recognition decisions. We apply CMR-IA to key findings concerning recognition of items and associations, including effects of recency, similarity, receiver-operating characteristic curves, word frequency, differential forgettings of items and associations, and contiguity effects for successive probes. We also apply CMR-IA to cued recall phenomena, including serial position effects, distribution of correct responses and errors, contiguity effects, associative symmetry, and similarity effects. Finally, we ask whether CMR-IA can account for the dependencies between successive tests of item and associative memory. We show that combining a Gestalt associative mechanism with retrieved-context theory provides a good account for many empirical phenomena concerning item and associative memory. The analysis of successive memory tests highlights the important role of output encoding in our model.

Keywords: recognition, recall, association, context

A theory of memory for items and associations

Studies of episodic memory have distinguished item-specific information and associative (or relational) information (Humphreys, 1978; Murdock, 1974). Item-specific information represents the encoding in memory of the attributes of an item occurring in a particular context. In contrast, associative information represents the encoding in memory of the relation between two co-occurring items. Associative information underlies the ability to learn that particular words appeared together in a study list, and item-specific information underlies the ability to recognize an item as having appeared in a study list. Although recognition is most often used to probe memory for item information, one can test recognition memory for associations by asking subjects to discriminate between correctly and incorrectly paired items (e.g., Hockley, 1982). The present paper aims to provide a unified theoretical account of memory for items and associations within the framework of retrieved-context theory. We apply our model to classic data regarding memory for items and associations. Further, we report and simulate new data on the interaction between item and associative memory, evaluated using the method of successive tests (Kahana, 2000; Kahana, Rizzuto, & Schneider, 2005).

Items and Associations

A number of experimental variables differentially influence item and associative memory. For example, people forget relational information more slowly than item information (Hockley, 1992). Whereas people exhibit superior recognition performance for rare as compared with common words, this effect does not appear for pairs of rare words as compared with pairs of common words (Clark, 1992), and in some cases, this effect even reverses (Lohnas & Kahana, 2013).

Whereas the distinction between item and associative information pertains to the representation being stored and retrieved from memory, the distinction between recognition and recall pertains to how the information is retrieved. In a recognition task, subjects typically judge whether a target item or association occurred in the context of a just-presented list. In a recall task, subjects attempt to recall a target item based on a specific associated cue item. Both cued recall and associative recognition measure memory for associative information, but cued recall involves a process that leads to the production of an item, sometimes referred to as deblurring. Associative recognition, in contrast, involves a decision process that leads subjects to convert an index of retrieved information to a “yes” vs. “no” response to the question of whether the association occurred in the target list.

Context-Maintenance and Retrieval Model for Items and Associations (CMR-IA)

We develop a model of memory for item and associative information that builds on a class of models that conceive of episodic memory as the interaction between content and context information. Content information is simply the information experienced at a given time and context is a slowly drifting representation of background features that evolve due to our experiences and the associations they evoke (Manning, in press). Experiencing or recalling a memory reinstates its earlier context(s), which in turn updates the present state of context and associates with subsequent experiences. The recursive nature of context allows associations from the distant past to inform new learning and binds items never experienced together (Cohen & Kahana, 2022; Wachter & Kahana, 2024).

Models based on retrieved-context offer a rigorous account for phenomena observed in studies of free recall where the cue for retrieval is context itself (Lohnas, Polyn, & Kahana, 2015; Healey & Kahana, 2014). Recent work has extended this class of models to account for sequence production and serial recall (Logan, 2021; Lohnas, in press-b). Here, we ask whether a similar context-based theoretical approach can help elucidate recognition and cued recall of item and associative information. As our model extends the Context-Maintenance and Retrieval (CMR) family to item and associative memory we refer to it as CMR-IA.

Our paper is organized as follows. Section 1 presents a selective review of early models of item and associative memory, providing a theoretical backdrop to our work (Osth & Dennis, 2024, provide a comprehsnive review of the recognition modeling literature). Section 2 presents CMR-IA. Section 3 evaluates CMR-IA against key findings concerning item and associative memory. Section 4 presents new data on successive tests of item and associative information. Section 5 evaluates CMR-IA in relation to these new data. Section 6 concludes.

Theoretical Backdrop

Strength theory provided the earliest account of item and associative information and their relations. According to this view, studying items on a list strengthens associations between each list item and some representation of the list itself. Outcomes of item recognition, associative recognition, and cued recall tests depend on the strength of these associations. Consistent with the observation that item recognition is usually easier than cued recall, strength theory held that the essential difference between these tasks is that item recognition could be successfully performed with weaker associations than either cued recall or associative recognition (Kahana, 2012).

Researchers found strength theory's account of the relation between item recognition and cued recall to be deficient, as many experimental variables have opposing effects on these tasks. Subjects easily recognize rare as opposed to common words as having occurred on a recent list, but they find rare words more difficult to recall (Gregg, 1976; Kinsbourne & George, 1974; MacLeod & Kampe, 1996; Lohnas & Kahana, 2013). Other dissociations include the effects of intentionality on encoding (Glenberg & Bradley, 1979; R. M. Schwartz & Humphreys, 1974), the effects of context-change on retrieval, and the effects of medial-temporal lobe damage on recognition of dissimilar items (Holdstock et al., 2002).

Generate-recognize theory quickly emerged as a preferred alternative account (e.g., Bahrick, 1970). This theory posited that recall involves two stages: Subjects first generate possible responses, and then apply a recognition test to decide whether any of the generated responses were on the list. The recognition task lacks a generate stage. A strong version of generate-recognize theory predicts that recallable items will always be recognized, contrary to many studies (e.g. Tulving, 1968; Tulving & Thompson, 1973; Flexser & Tulving, 1978).

A new generation of simulation-based, mathematical models of recognition and recall memory emerged in the 1970s and 1980s. These models often assumed distributed representations of items and associations, and operations that simulated the processes of encoding and retrieval (Kahana, 2020; Cox & Shiffrin, in press). To illustrate how these models simulated data from recognition and recall, let us consider an experiment in which subjects first study a list of randomly paired common words, denoted $\mathbf{f}_i - \mathbf{g}_i$. Later, during a test phase, the experimenter may assess memory in one of several ways: by presenting a single word and asking whether it appeared in any of the studied pairs (item recognition), by presenting a pair of words and asking whether both items in the test pair appeared together as one of the studied pairs (associative recognition), or by asking presenting a single word from one of the studied pairs and asking subjects recall its mate (cued recall).

Defining each item in memory as an N-dimensional vector, the pair of (column) vectors $(\mathbf{f}_i, \mathbf{g}_i)$ would represent the i-th pair of studied items. The following equation represents how a standard matrix model might store these paired vectors in memory:

$$M = \sum_{i=1}^{L} (\mathbf{f}_i + \mathbf{g}_i)(\mathbf{f}_i + \mathbf{g}_i)^\top$$

This model learns each association by first summing the two item vectors in each pair and then multiplying the sum by its transpose to form an outer-product matrix. Rather than storing each matrix separately in memory, the model combines them by simple summation. In this manner, a single memory matrix represents all of the studied pairs in a given list (for the moment, assume that memory begins the experiment as a tabula rasa; following study, it only contains pairs from the encoded list).

To simulate cued recall, we premultiply the memory matrix, M, with the cue item, $\mathbf{f}_k$. This yields an approximate representation of the studied pair, $\mathbf{f}_k + \mathbf{g}_k$, as given by the following equation:

$$\begin{aligned} M\mathbf{f}_k &= \sum_{i=1}^{L}(\mathbf{f}_i\mathbf{f}_i^\top + \mathbf{f}_i\mathbf{g}_i^\top + \mathbf{g}_i\mathbf{f}_i^\top + \mathbf{g}_i\mathbf{g}_i^\top)\mathbf{f}_k \\ &= (\mathbf{f}_k + \mathbf{g}_k)(\mathbf{f}_k \cdot \mathbf{f}_k + \mathbf{g}_k \cdot \mathbf{f}_k) \\ &\quad + \sum_{l \neq k}(\mathbf{f}_l + \mathbf{g}_l)(\mathbf{f}_l \cdot \mathbf{f}_k + \mathbf{g}_l \cdot \mathbf{f}_k) \\ &= \mathbf{f}_k + \mathbf{g}_k + \text{error} \end{aligned}$$

The match between the retrieved trace and the cue item, given by $M\mathbf{f}_k \cdot \mathbf{f}_k$, then determines the probability of successful cued recall. If the vectors are orthonormal, retrieval is error-free. Because vector addition is commutative, this implementation of cued recall is inherently symmetrical. Even if the strengths of the $\mathbf{f}$ and $\mathbf{g}$ items are differentially weighted, the forward and backward hetero-associative terms would have the same strength. The symmetric nature of these linear models also applies to their nonlinear variants (Hopfield, 1982).

To simulate the recognition process we first cue memory with $\mathbf{g}_i$ and retrieve its associated pattern, as in the cued recall model described above. We then compute the similarity between this retrieved vector and the cue. Using the dot-product as our measure of similarity, as in Murdock's Theory of Distributed Associative Memory (TODAM), we obtain the following expression:

$$M\mathbf{g}_i \cdot \mathbf{g}_i = (\mathbf{f}_i + \mathbf{g}_i) \cdot \mathbf{g}_i + \sum_{k \neq i}(\mathbf{f}_k + \mathbf{g}_k) \cdot \mathbf{g}_i$$

If the resulting similarity, given by the above expression, exceeds a threshold, the model will generate a "yes" response. Or, we can use this expression as the input to a stochastic decision process, which would allow us to obtain predictions regarding response times to various cue types (Hockley & Murdock, 1987).

Summed-similarity and signal-detection models of recognition memory, under certain assumptions, constitute special cases of the matrix model. This model also provides a natural account for interference effects in associative memory, and it predicts greater correlations between cued recall and associative recognition than between cued recall and item recognition. Kahana et al (2005) provide a mathematical analysis of this model and comparisons with several other classes of distributed memory models, including those that use vector convolution as the associative operator (Murdock, 1982; Metcalfe-Eich, 1982; Metcalfe, 1985). By adding some basic auxiliary assumptions, such as probabilistic encoding (Murdock & Lamon, 1988; Murdock & Kahana, 1993a) and output encoding, Kahana et al. (2005) show how each of these classic models can account for a range of recognition and recall phenomena, including the correlations between successive recognition and recall tests as seen in the Tulving-Wiseman relation (Kahana, 2000).

A key limitation, however, of the aforementioned models is the tabula rasa assumption. By resetting memory at the start of each list, these models do not need to grapple with interactions between memories learned in distinct temporal contexts, such as the different lists of a multi-trial experiment. To address this problem, a number of theories introduce the concept of a list context that associates with each presented item and pair (Humphreys, Pike, Bain, & Tehan, 1989; Osth & Dennis, 2015; Osth, Jansson, Dennis, & Heathcote, 2018; Cox, 2024). However, models that include a dynamical representation of spatiotemporal context subsume a broader array of memory phenomena, including the effects of recency and contiguity seen at multiple time scales (see Kahana, 2020, for a review).

2. CMR-IA: a Context Maintenance and Retrieval Model for Item and Associative Memory

Here we develop CMR-IA, a model of memory for item and associative information based upon the ideas of retrieved-context theory embodied in Polyn, Norman, and Kahana (2009); Lohnas et al. (2015); Healey and Kahana (2016); Cohen and Kahana (2022); Wachter and Kahana (2024). In our usage, context refers to a mathematical representation of a psychological construct. At the psychological level, we wish to capture the idea that people form associations between items and their situational, temporal, and/or spatial context (Carr, 1931; McGeoch, 1932; Hollingsworth, 1928; Robinson, 1932; Underwood, 1945). This idea finds its clearest early expression in Bower's (1972) multi-attribute temporal context model. According to Bower's model, contextual representations comprise a multitude of fluctuating features, defining a vector that slowly drifts through a multidimensional context space. These contextual features form part of each memory, combining with other aspects of externally and internally generated experience. Because a unique context vector marks each remembered experience, and because context gradually drifts, the context vector conveys information about the time in which an event was experienced. Bower's model, which drew heavily on the classic stimulus-sampling theory developed by (Estes, 1955), placed the ideas of temporal coding and internally generated context on a sound theoretical footing, and provided the basis for the retrieved context theory developed by Howard and Kahana (1999, 2002).

By allowing for a dynamic representation of temporal context, items within a given list will have more overlap in their contextual attributes than items studied on different lists, or indeed items that were not part of an experiment (Bower, 1972). If the contextual change between lists is sufficiently great, and if the context at time of test is similar to the context encoded at the time of study, then recognition memory judgments should largely reflect the presence or absence of the probe (test) item within the most recent (target) list, rather than the presence or absence of the probe item on earlier lists. This enables a multitrace summed-attribute similarity model to account for many of the major findings concerning recognition memory and other types of memory judgments.

We can implement a simple model of contextual drift by defining a multidimensional context vector, $\mathbf{c} = (\mathbf{c}(1), \mathbf{c}(2), \ldots, \mathbf{c}(N))$, and specifying a process for its temporal evolution. For example, we could specify a unique random set of context features for each list in a memory experiment, or for each experience encountered in a particular situational context. However, contextual attributes fluctuate due to many internal and external variables that vary at many different time scales. An elegant alternative (Murdock, 1997) is to write down an autoregressive model for contextual drift, $\mathbf{c}_i = \rho \mathbf{c}_{i-1} + \sqrt{1-\rho^2}\epsilon$, where ϵ is a random vector whose elements are each drawn from a Gaussian distribution, and where i indexes each item presentation. The variance of the Gaussian is defined such that the inner product $\epsilon_i \cdot \epsilon_j$ equals one for $i = j$ and zero for $i \neq j$. Accordingly, the similarity between the context vector at time steps i and j falls off exponentially with the separation: $\mathbf{c}_i \cdot \mathbf{c}_j = \rho^{|i-j|}$. Thus, the change in context between the study of an item and its later test will increase with the number of items intervening between study and test, producing the classic forgetting curve.

Glenberg and Swanson (1986) suggested that an evolving context representation could explain the phenomena of long-term recency and Mensink and Raaijmakers (1988) demonstrated how a variable context signal in the SAM model could account for the major results obtained in studies of associative interference. Whereas these earlier formulations saw context as a cue for items, they did not explicitly consider items as a cue for context. That is, drifting contextual representations changed the accessibility of items, but item retrieval did not systematically alter the state of context.

Retrieved-context theory (RCT; Howard & Kahana, 1999, 2002) introduced the idea that remembering an item calls back its encoding context which in turn serves as a retrieval cue for subsequent recalls. At its most basic level, this theory proposed that items and context become reciprocally associated during study: Context retrieves items, and items retrieve context. As the learner encodes or retrieves an item, RCT assumes that new associations form both between the item's feature representation and the current state of context (stored in the matrix M^{FC}) and between the current state of context and the item's feature representation (stored in the matrix M^{CF}). These associations follow a Hebbian outer-product learning rule, as in earlier models of associative memory.

Specifically, we can write the evolution of these associative matrices as:

$$\begin{aligned} M_t^{FC} &= M_{t-1}^{FC} + \gamma_{FC}\mathbf{c}_i\mathbf{f}_i^\top \\ M_t^{CF} &= M_{t-1}^{CF} + \gamma_{CF}\mathbf{f}_i\mathbf{c}_i^\top \end{aligned} \tag{1}$$

where γ_{FC} and γ_{CF} parameters determine the learning rates for item-to-context and context-to-item associations, respectively. Through these associative matrices, the presentation or recall of an item evokes its prior associated contexts, and cueing with context retrieves items experienced in overlapping contextual states.

Howard and Kahana (2002) further suggested that the mechanism of retrieved context was the driver of context evolution itself. Rather than context evolving randomly, they proposed that the currently active item, $\mathbf{f}_i$, retrieves (a weighted sum of) its previous contextual states, via the M^{FC} matrix, which in turn serves as the input to the context evolution equation: $\mathbf{c}^{IN} = M^{FC}\mathbf{f}_i$. Normalizing this input yields the recursion:

$$\mathbf{c}_i = \rho_i\mathbf{c}_{i-1} + \beta\frac{M^{FC}\mathbf{f}_i}{\|M^{FC}\mathbf{f}_i\|} \tag{2}$$

where ρ_i is a constant ensuring that $\|\mathbf{c}_i\| = 1$, and where model parameter β governs the rate of contextual drift (a large value of β causes context states to decay quickly). During list recall, or when items are repeated, context recursively updates based on the past contextual states associated with the recalled or repeated item[1]. Because context is always of unit length, it can be thought of as a point on a hypersphere, with β determining the distance it travels with the newly presented item and $M^{FC}\mathbf{f}_i$ determining the direction of

[1] Prior to proposing the recursion in Equation 2, Howard and Kahana (2002) considered the simpler idea that items themselves form the input to context, as given by $\mathbf{c}^{IN} = \mathbf{f}_i$. Although this formulation had several desirable properties, it failed to capture the idea that context includes more than the perceptual and conceptual features of an item, but rather embodies its history of prior associated states. For example, when studying the word "chocolate", one would be expected to update context with its past associated mental representations. Similarly, when recalling "chocolate", both the encoding context of the study list and the past associations should come to mind.

travel.

In modeling free recall, we can use the current state of context $\mathbf{c}_t$ as a cue to retrieve items, via associations stored in the M^{CF} matrix:

$$\mathbf{f}_t^{\mathbf{IN}} = M^{CF}\mathbf{c}_t \tag{3}$$

The resulting $\mathbf{f}_t^{\mathbf{IN}}$ gives the degree of support, or activation, for each item in the model's lexicon. Different implementations of retrieved-context theory have used different models to simulate the recall process. Several early papers (e.g., Howard and Kahana (1999)) used a probabilistic choice rule similar to the sampling rule in SAM, but then one must propose a specific rule for recall termination. Sederberg, Howard, and Kahana (2008) modeled the recall process as a competition among racing accumulators (one for each candidate recall), based on the decision model proposed by Usher and McClelland (2001). The first word whose accumulator crosses its threshold, prior to the end of the recall period, wins the competition. The contextual retrieval process in RCT produces the temporal and semantic organization seen in the dynamics of memory retrieval. Much more can be said about these models and we refer the interested reader to Kahana (2020) and Manning (in press) for accessible reviews.

Data on the dynamics of free recall has provided a rich testbed for the development of retrieved-context theory (Kahana, 2020; Lohnas, in press-a). The cue-dependent nature of sequential retrievals in free recall lends itself particularly well to this type of theory. In recent work, Logan (2021; 2023) has applied a very similar modeling framework to explain data on various types of serial recall tasks. Here, we evaluate whether a retrieved-context theory imbued with a Gestalt model of associations can help us understand memory for items, associations, and their interactions.

Using Context-Similarity to Model Item Recognition

A recognition test probe will retrieve its associated context state in exactly the same way that items retrieve their associated contexts during encoding and retrieval. Comparing this retrieved context to the current state of context provides a measure of the match between the current context and the encoding context of the item, which can then drive a recognition decision (Healey & Kahana, 2016). Subjects judge an item $\mathbf{f}_i$ as "old" when the similarity between the context retrieved by $\mathbf{f}_i$ and the current context $\mathbf{c}_t$ passes a threshold c_{recog}:

$$M^{FC}\mathbf{f}_i \cdot \mathbf{c}_t > c_{recog} \tag{4}$$

where c_{recog} may vary across different conditions of an experiment. If endorsed as "old", the item's associated context is integrated into the current state of context (see Equation 2). Healey and Kahana (2016) assumed that context updating did not occur for "new" judgments. Here we generalize their approach by assuming context updating for both "old" and "new" judgments (in the simulations below, we assume the same rate of contextual drift for both types of judgments, but this could be generalized if needed).

Consider the case of subjects who studied a list of unrelated items. We would model these as orthogonal vectors (basis functions) on feature space. In this case, the following summation defines the similarity between the current context and the context retrieved by $\mathbf{f}_i$:

$$\sum_j \mathbf{c}_j \mathbf{f}_j^\top \mathbf{f}_i \cdot \mathbf{c}_t = \sum_j (\mathbf{c}_j \cdot \mathbf{c}_t)(\mathbf{f}_j \cdot \mathbf{f}_i) = \mathbf{c}_i \cdot \mathbf{c}_t$$

For the case of a list with a single repeated element occurring in positions i and k, we would have two matching context terms, $\mathbf{c}_i$ and $\mathbf{c}_k$, but the state of $\mathbf{c}_t$ will also be altered due to the context evolution equation. Specifically, $\mathbf{c}_t$, in this case, will be more similar to the earlier of the two retrieved contexts than it would in a list without repetitions, thus producing a greater increase in similarity than would be expected on the basis of the repetitions themselves (see Siegel & Kahana, 2014). The case of similar items should fall between that of repeated and unrelated items.

Modeling Associative Recognition and Cued Recall

Extending CMR to associative recognition and cued recall requires us to posit a mechanism for associative learning. Here, there are two logical possibilities: 1) Treat the association as a list of two items, each associated with its own context, but allow context to drift more between pairs than within pairs. This is basically the associative isolation hypothesis of Caplan (2005). 2) Treat the pair as a single item that combines features of the two constituent elements. Here, the model would essentially create a new features vector for a pair, unitizing the elements into a single memory associated with a single context. This unitization approach accords with the Gestalt conceptions of associative learning and findings of associative symmetry (Kahana, 2002).

Although logically distinguishable, the two approaches outlined above are by no means mutually exclusive. One could imagine that sometimes subjects construct a new item out of two constituents, and at other times they treat the pair as if it were a two-item list. Further, although studies of cued recall generally find symmetric retrieval of forward and backward associations, associative recognition may not exhibit associate symmetry (Yang et al., 2013) (and Experiment 2 of this paper). Finally, subjects can remember the order of a pair of studied items, suggesting that temporal structure does find its way into the memorial representation of a pair. Here we adopt the second hypothesis, mainly for parsimony, but we remain open to both possibilities.

In CMR-IA the two items constituting a given pair share a single context representation. Specifically, CMR-IA treats the two items, $\mathbf{f}_{i_1}$ and $\mathbf{f}_{i_2}$, as a whole by combining them into one vector representation:

$$\mathbf{f}_i = \mathbf{f}_{i_1} + \mathbf{f}_{i_2} \tag{5}$$

The model then updates the context using Equation 2 and updates M^{FC} and M^{CF} using Equation 1. The formal implementation of CMR-IA also incorporates the assumption of a primacy gradient (greater learning of earlier list items) when updating M^{CF} as in prior work (Sederberg et al., 2008; Polyn et al., 2009; Lohnas et al., 2015; Healey & Kahana, 2016) (see Appendix A).

The retrieval process for associative recognition and cued recall also differs from item recognition and free recall. To recognize or recall associations, CMR-IA first updates the context with one of the items in the test pair (in associative recognition) or the cue item (in cued recall), and then uses the updated context to perform recognition or recall. Let's consider associative recognition. Facing a pair of probes, $\mathbf{f}_{i_1}$ and $\mathbf{f}_{i_2}$, our model first updates the current context with $\mathbf{f}_{i_1}$ using Equation 2. We denote the updated context as $\mathbf{c}'_t$. Then, the model compares the context input of $\mathbf{f}_{i_2}$ with the updated context. If $M^{FC}\mathbf{f}_{i_2} \cdot \mathbf{c}'_t > c_{recog}$ (in a similar fashion to Equation 4), it would give a "yes" response. Separate parameters control the threshold for item and associative recognition, denoted as c^{item}_{recog} and c^{assoc}_{recog} respectively. For cued recall, CMR-IA first updates the current context with the cue using Equation 2. Then, the model uses the updated context to retrieve items from memory in the same way as free recall (see Equation 3). In other respects, CMR-IA inherits the assumptions of the CMR2 model developed by Lohnas et al. (2015) and Healey and Kahana (2016).

CMR-IA allows for the possibility that subjects will learn context-to-item associations and item-to-context associations for correctly recognized items and pairs. Similarly, successful cued recall of an item can lead to the storage of item-to-context associations. Although output encoding has also been envisioned in prior CMR implementations, it has only rarely been introduced into simulations. Here we use output encoding in all simulations involving repeated tests, as described below. This mechanism plays an essential role in modeling the correlations between successive memory tasks.

3. Evaluating CMR-IA's Account of Benchmark Phenomena

Here, we evaluate CMR-IA's ability to account for a set of key benchmark phenomena concerning memory for item and associative information. Our first set of simulations concerns data on item and association recognition. Our second set of simulations concerns data on cued recall.

Simulations 1-4: Item and Associative Recognition

Simulation 1 examines the joint effects of recency and similarity on performance in a continuous recognition task based on the classic methods of Shepard and Teghtsoonian (1961) and Hockley (1982). Simulation 2 examines the contiguity effect of successive probes, fitting the data described in G. Schwartz, Howard, Jing, and Kahana (2005). Simulation 3 looks at the difference in forgetting rates of item recognition and associative recognition (Hockley, 1992). Simulation 4 examines the word frequency effect in item recognition (Lohnas & Kahana, 2013).

Experiment 1

To elucidate the joint effects of recency and similarity on recognition performance, we conducted a large online study using a continuous item recognition procedure.

Methods

Subjects. 657 subjects recruited from Amazon Mechanical Turk participated in the experiment. We excluded 109 subjects for not responding within the allotted time frame to more than 250 of the items on a list.

Procedure. Subjects completed a standard continuous recognition task (Shepard & Teghtsoonian, 1961). On each trial in this study, subjects observed a single word on the screen for a period of 1000 ms with an inter-trial interval of 1000 ms. Subjects were required to respond to each word by pressing one of the “a”, “s”, “d”, “f”, “h”, “j”, “k”, “l” keys, with each key representing both an old/new judgment and a confidence rating between 1 and 4 (Weidemann & Kahana, 2016). The items on each trial were drawn from a set of 300 words, each from a set of 25 semantic categories. This was the same wordpool used in a number of categorized free recall studies in our lab (Weidemann et al., 2019). The new items on the lists were constructed such that each word was either in a set of 2-6 items drawn from the same category or a block of 1-6 items, each drawn from a different category. 15 categories were selected to be part of the same category blocks and the other 10 categories were part of the different category blocks. After constructing the order of new word presentations, 280 old word presentations were placed randomly throughout the list (following the 20th new word presentation), with the old word drawn randomly from among the already presented new words. This procedure ensured variation in both temporal lag and semantic similarity among old words.

Analysis

Since the study-test lags approximately follow an exponential distribution, we partitioned study-test lags into five bins with respect to Euler's number e to ensure a fair count of items in each bin: Bin 1 (lag less than 8), Bin 2 (lag 8 to 20), Bin 3 (lag 21 to 54), Bin 4 (lag 55 to 148), Bin 5 (lag greater than 148)[2]. Replicating Hockley (1982), the hit rate decreased with the lag between study and test, as seen in Figure 1A. We defined items as having high semantic similarity when two or more same-category items occurred within the previous eight items; we defined low-similarity as having zero or one same-category item within the previous eight items. For statistical inference, we fitted a Linear Mixed Model with study-test lag bin, similarity level, and their interactions as fixed-effects and with a maximal random-effects structure that ensures non-singularity, following the procedure recommended by Bates, Kliegl, Vasishth, and Baayen (2018). The Satterthwaite-approximated ANOVA showed that both recency, as defined by study-test lag, and similarity have significant effects on hit rate (recency: $F(4, 804.82) = 85.72$, $p < .001$; similarity: $F(1, 246.34) = 5.71$, $p = .018$). The interaction between these factors is not significant ($F(4, 522.11) = 0.95$, $p = .435$).

A decreasing hit rate does not necessarily imply degraded memory. If, for example, subjects generally tend to become more conservative in responding with increasing study-test lag, you would see hit rates decline along with false alarm rates. We thus calculated local false alarm rates (Murdock & Kahana, 1993b) for each study-test lag bin and similarity level. To do this, we checked each old item to see if there is an adjacent new

[2] Items with lags less than e^2 belong to Bin 1, items with lags greater than e^2 while less than e^3 belong to Bin 2, and so on.

item. If so, we regarded this new item as in the same study-test lag bin as the old item[3]. The similarity level of this new item is determined by the number of same-category items within the previous eight items as described above (Figure 1C). For local false alarm rate, the ANOVA demonstrated significant effects of recency ($F(4, 1059.42) = 4.57$, $p = .001$) and of similarity ($F(1, 504.97) = 21.15$, $p < .001$), but the interaction between these factors did not cross the $p < 0.05$ threshold ($F(4, 1020.39) = 2.10$, $p = .079$). Taking Figure 1A and Figure 1C together, we see that greater recency improves recognition performance. However, higher similarity tends to elicit both higher hit rates and higher false alarm rates, raising the question of whether similarity impacts memory, decision criterion, or both.

Because we gathered confidence ratings in this experiment we could use these ratings to calculate A_z, an estimator of area under curve[4] (Stanislaw & Todorov, 1999). As shown in Figure 1E, the ANOVA found significant effects of recency and similarity as well as their interaction (recency: $F(4, 576.98) = 59.20$, $p < .001$; similarity: $F(1, 382.40) = 28.18$, $p < .001$; interaction: $F(4, 723.96) = 4.95$, $p < .001$). The Linear Mixed Model revealed that a high similarity level results in significantly lower A_z than a low similarity level for Bin 2 ($t(1608.26) = -2.36$, $p = .018$), Bin 3 ($t(1524.18) = -3.66$, $p < .001$), and Bin 4 ($t(1014.86) = -3.43$, $p < .001$), but not for Bin 1 ($t(1508.34) = -1.17$, $p = .244$) and Bin 5 ($t(236.06) = -0.19$, $p = .850$).

Simulation 1

We simulated Experiment 1 using CMR-IA and we programmed CMR-IA to perform a recognition test on the actual sequences of items used in the experiment. For

[3] A new item could be counted twice if it happens to be in the middle of two old items that belong to different study-test lag bins.

[4] The calculation of A_z is as follows. Say there are r possible ratings. We could have $r - 1$ different decision thresholds, thus obtaining $r - 1$ pairs of hit rate and false alarm rate. Then, we find the z scores of those hit rates and false alarm rates and fit a line in the z space, i.e., z-ROC curve. Finally, we define A_z using the intercept and the slope of the z-ROC curve,

$$A_z = \Phi\left[\frac{\text{intercept}}{\sqrt{1 + \text{slope}^2}}\right]$$

.

each item, the simulated subject first performed item recognition by comparing the current context with the context input brought by the item, following Equation A9. If the context similarity exceeds the item recognition threshold c_{recog}^{item}, the simulated subject would give a response of "old", otherwise "new". Then, the simulated subject encoded this item by updating the context and the M^{FC} matrix following Equation 2 ($\beta = \beta_{enc}$) and Equation 1. The context also drifted a bit between each trial, following Equation 2 ($\beta = \beta_{post}^{recall}$). We simulated 1000 sessions with CMR-IA.

Figure 1B, Figure 1D, and Figure 1E illustrate the data generated by CMR-IA. A_z is calculated by setting different thresholds on the context similarity, a notion comparable to the confidence in the real experiment. The model produces both recency and similarity effects, similar to those seen in the experiment. The model also predicts an interaction between recency and similarity, where a high similarity level tends to lower A_z only at large study-test lags. The contextual drift mechanism in CMR-IA produces the recency effects because increasing study-test lag reduces context similarity. To ensure reasonable performance at large study-test lags, we control the speed of contextual drift by setting both β_{enc} and β_{post}^{recall} to be low. For the similarity effects, preceding same-category items strengthen the representation of the current item in the context, which results in a higher context similarity and raises both the hit rate and false alarm rate. We also notice that a high similarity level tend to make the distribution of context similarity for both new and old items more left-skewed, thus making A_z smaller than that of a low similarity level. Parameters s_{FC} and γ_{FC} exert the greatest influence on this effect as both have an impact on the strength of semantic associations in M^{FC}. For the interaction between recency and similarity, since the distributions of context similarity at small study-test lags are already more left-skewed, a high similarity level has less impact on them than those at large study-test lags.

Simulation 2: Successive-probe Contiguity Effects

G. Schwartz et al. (2005) asked whether temporally defined associations play a role in item recognition. Using travel photographs as stimuli, they found that subjects were more likely to recognize a test item as old if the preceding old test item was studied in a temporally proximate position. On each list of their experiment, subjects ($N = 91$) studied 64 photographs and then performed an item recognition test on the same 64 studied photographs intermixed with 64 new photos drawn from the same pool. They compared responses to old photos following photos from adjacent list positions ($|\text{lag}| = 1$) with responses to old photos following photos from remote positions ($|\text{lag}| > 10$). Drawing out the ROC and z-ROC curve, they found that successive probes from adjacent positions exhibit better recognition performance than successive probes from remote positions (Figure 2A & Figure 2C). Each subject contributed data from six lists.

We simulated the same experimental paradigm to evaluate whether CMR-IA exhibits a successive-probe contiguity effect. For each simulation, we randomly sampled a stimulus sequence from the real experiment, though each stimulus here is not a real photograph but represented as a basis vector with a single element set to 1 and all other elements set to 0. To capture the possible semantic associations between photos, we initialized M^{FC} and M^{CF} as a random matrix with diagonal entries equal to 1. In the study phase, the simulated subject studied 64 items in sequence. The simulated subject encoded each item by updating the context and the M^{FC} matrix following Equation 2 ($\beta = \beta_{enc}$) and Equation 1. Between the study phase and the test phase, the context drifted following Equation 2 ($\beta = \beta_{distract}$). In the test phase, we presented 128 items in sequence to the simulated subject, where half of them were in the study phase and the rest were not. The simulated subject performed item recognition by comparing the current context with the context input retrieved by the probe item, following Equation 4. We treated context similarity as the confidence ratings reported by the simulated subject[5].

[5] We didn't introduce c_{recog}^{item} here, because we don't need a binary response for ROC analysis in this

After the recognition, the simulated subject updated the context with the context retrieved by the probe item, following Equation 2 ($\beta = \beta_{cue}$). Between each list, the context also drifted following Equation 2 ($\beta = \beta_{post}^{recall}$). we simulated 1000 sessions with CMR-IA.

To assess the effect of temporal contiguity on recognition performance, we conducted the same analysis as G. Schwartz et al. (2005). Figure 2B and Figure 2D illustrates the data generated by CMR-IA[6]. The model produces a successive-probe contiguity effect, replicating the findings of G. Schwartz et al. (2005). The effect arises from the contextual update after the judgment in CMR-IA, controlled by β_{cue}. This mechanism incorporates the context during the encoding of the test item into the current context, which makes the current context more similar to the context that is associated with the adjacently studied items, thus facilitating the recognition of these items. Parameter β_{enc} also exerts an influence on this effect since it determines how similar the associated contexts are among adjacent items. If we turn up β_{cue} or β_{enc}, the contiguity effect would be stronger.

Simulation 3: Differential Forgetting of Items and Associations

Hockley (1992) found that item information is more susceptible to decay than associative information. In his study, he used a variant of the continuous recognition paradigm. Each list comprised 160 trials with each trial consisting of a pair of words presented for study followed by a test probe which could either be a single word or a word pair: subjects made intact-rearranged judgments on pairs and old-new judgments on single words. For single-item recognition tests, half are old words and half are new words. For associative recognition tests, half are intact study pairs and half are rearranged pairs. The rearranged pairs are constructed by combining the left (or right) member of a study pair with the right (or left) member of an immediately preceding study pair. Test probes appeared at study-test lags of 2, 4, 6, 8, and 16 trials. Hockley (1992) demonstrated that d'

simulation and CMR-IA updates the context in the same way for “old” and “new” judgments.

[6] We constructed a “local” FAR separately for adjacent and remote lags by including data from new items that were contiguous with adjacent or remote old-item pairs, respectively, in a similar way as G. Schwartz et al. (2005).

for item recognition decreases more sharply than d' for association recognition (Figure 3A and Figure 3C).

We adopted the same experimental paradigm in CMR-IA to simulate Hockley's (1992) finding. While in the original experiment each subject completed five sessions on separate days and each session had two lists, here we simplified the procedure so that each simulation consists of only one list. We used the PEERS Word Pool (Kahana et al., 2024) instead of the Toronto Word Pool used by Hockley (1992). Each list consists of 160 trials, each consisting of one study phase and one test phase. In the study phase, the simulated subject encoded the word pair following Equation 5, Equation 2 ($\beta = \beta_{enc}$), and Equation 1. Then in the test phase, the simulated subject performed item recognition if the test presentation was an item or associative recognition if the test presentation was a pair. For item recognition, the simulated subject compared the current context with the context input brought by the item following Equation 4. If the context similarity exceeds the item recognition threshold c^{item}_{recog}, the simulated subject would give a response of "old", otherwise "new". For associative recognition, the simulated subject updated the current context using the first item following Equation 2 ($\beta = \beta_{cue}$), and then compared the current context with the context input brought by the second item following Equation A11. If the context similarity exceeds the associative recognition threshold c^{assoc}_{recog}, the simulated subject would give a response of "intact", otherwise "rearranged". After the judgment, the simulated subject updated the context with the test item or the second item in the test pair, following Equation 2 ($\beta = \beta_{cue}$). Between each trial, the context also drifted following Equation 2 ($\beta = \beta^{recall}_{post}$). We performed a total of 1000 simulation sessions.

Qualitatively replicated findings from Hockley (1992), CMR-IA demonstrated that item information is more susceptible to decay than associative information (Figure 3B and Figure 3D). This effect arises because of the different mechanisms for item recognition and associative recognition. For item recognition, the model directly compares the retrieved context with the current context, which is greatly influenced by the contextual drift in time (controlled by parameters β_{enc}, β_{cue}, and β_{post}^{recall}). In contrast, for associative recognition, the model first reinstates the context retrieved by one item into the current context (controlled by β_{cue}) and then compares the updated context with the retrieved context by the other item. This process is less influenced by the contextual drift in time.

Although CMR-IA accounts for the differential forgetting rate of items and associations, we should notice that, compared with the experimental data, the gap between d' for item and associative recognition is larger and the curvature of decay is less linear in CMR-IA. This might be because the learning of items and associations is controlled by the same parameters. Since a rearranged pair is constructed by blending a study pair with its preceding pair, distinguishing a rearranged pair from an intact pair requires a large β_{enc} to ensure an adequate contextual update. However, this, in turn, raises the already-good recognition performance of items when the study-test lag is small and makes the hit rate curve a steep exponential decay. It poses a challenge for CMR-IA to keep a balance among associative recognition performance, item recognition performance, and decaying curvature of memory.

Simulation 4: Word Frequency Effects in Recognition Memory

Studies of item recognition find that rare words are more easily recognized as targets and more easily rejected as lures than common words Lohnas and Kahana (2013)[7] (Figure 5A). An initial hypothesis is that the difference in semantic associations between rare words and common words gives rise to the word frequency effect. Indeed, we found that a word's average semantic associations with other words, $\mathbf{s}^{\text{mean}}$, is negatively correlated with its log word frequency ($p < .001$, $R^2_{adj} = 0.058$; Figure 4). This is consistent with a previous finding that rare words have more close neighbors in the semantic space (Monaco, Abbott, & Kahana, 2007).

[7] Lohnas and Kahana (2013) limited their analysis of the word frequency effect to a subset of 984 words drawn from the full 1638 item word pool used in the PEERS 1-3 studies. They did this because they could obtain imageability and concreteness measures for these words in the MRC database (Wilson, 1988). We followed their methods in our modeling work.

We tested this initial hypothesis by simulating a simple recognition paradigm in CMR-IA. In each simulation, a simulated subject studied a random list of 100 words and then was tested on a list of 200 words where half are old and the other half are new (Schulman, 1967). We chose the words randomly from the 984 words used by Lohnas and Kahana (2013). In the study phase, the simulated subject encoded each word by updating the context and the M^{FC} matrix following Equation 2 ($\beta = \beta_{enc}$) and Equation 1. Between the study phase and the test phase, the context drifted following Equation 2 ($\beta = \beta_{distract}$). In the test phase, the simulated subject performed item recognition for each item by comparing the current context with the context input brought by the item following Equation 4 and gave a response of "old" or "new" based on whether the context similarity exceeds the recognition threshold $c_{recog} = c^{item}_{recog}$. After the judgment, the simulated subject updated the context with this item, following Equation 2 ($\beta = \beta_{cue}$). We performed 1000 simulation sessions using particle swarm optimization to identify optimal parameter values. We binned the words into 10 bins according to their frequency and calculated the hit rate and false alarm rate for each bin of words in the same way as Lohnas and Kahana (2013). However, CMR-IA fails to produce the word frequency effect. Since common words tend to have weaker semantic associations with other words, the similarity between the associated context and the current context tends to be lower, so CMR-IA predicts a decrease in both hit rate and false alarm rate as the word frequency increases (Figure 5B). Moreover, we did not observe a better recognition performance for rare words than common words in CMR-IA. In the current simulation, the bin of the lowest word frequency has $d' = 1.89$ while the bin of the highest word frequency has $d' = 2.18$.

This is not surprising and aligns with prior theoretical work (e.g., Glanzer, Hilford, & Kim, 2004; Shiffrin & Steyvers, 1997). We therefore follow prior investigators by assuming that rare words attract greater attention than common words, leading to stronger associations in memory (Malmberg & Nelson, 2003; Criss & Malmberg, 2008). This is reasonable since people tend to pay more attention to novel items even in infancy

(Reynolds & Richards, 2005). We also assume that subjects make trial-by-trial adjustments to their recognition threshold, increasing their threshold for saying "yes" to rare words to strike a balance between increasing their hit rate and lowering their false alarm rate (Murdock, 1998; Benjamin, Diaz, & Wee, 2009; Monaco et al., 2007). Note that Monaco et al. (2007) were able to produce the word frequency effect using a Hopfield network with only the threshold adjustment mechanism but not the attention mechanism. Because of the difference in semantic associations between rare and common words, they observed that studied rare words are more proximal to strong attractors, thus naturally giving rise to a higher d'. However, we did not observe this phenomenon in our model, so we believe that raising attention for rare words is necessary to ensure a better recognition performance for rare words in CMR-IA.

Specifically, we used $\mathbf{s}^{\text{mean}}$ as an indicator of word frequency in CMR-IA[8]. When updating the M^{FC} matrix, we incorporated an attention mechanism by adding an attention coefficient ψ into Equation 1 so that rare words form a stronger association with the context. When making the old-new judgments, the threshold c_{recog} is not fixed but related to the frequency of the test word according to Equation A10 (see Appendix A for more details). We performed another 1000 simulation sessions and did the same analysis as above. This time, CMR-IA predicts an increase in false alarm rate and a decrease in hit rate as the word frequency increases (Figure 5C), replicating the word frequency effect. The recognition threshold c_{recog}, determined by the parameters c_s and c_{recog}^{item}, is crucial for decreasing the false alarm rate for rare words. Especially, c_s determines the sensitivity of how the recognition threshold would be influenced by word frequency, thus controlling the slope of the false alarm rate. A positive c_s which decreases the false alarm rate for rare words also decreases their hit rate. To ensure a higher hit rate for rare words, the attention coefficient ψ, determined by the parameters ψ_s and ψ_c, plays an important role. Especially,

[8] This avoids introducing extra input information to our model, compared with using an explicit representation of word frequency.

ψ_s determines how sensitive the attention coefficient is to the word frequency, thus controlling the slope of the hit rate.

Simulations 5-8: Recency, Similarity, and Contiguity in Cued Recall

As our work aimed to develop a CMR-based model of both item and associative information, we now turn to the analysis of data from cued recall of paired associated as this paradigm depends explicitly on the retrieval of associations between the cue and the target. Simulation 5 examines serial position effects in memory for paired associates. For this simulation, we have chosen a classic dataset from Murdock (1967). Simulation 6 examines CMR-IA's prediction regarding data on associative symmetry, fitting data described in Kahana (1993) and Kahana (2002). Simulation 7 examines data on prior-list and intra-list intrusions in cued recall (Davis, Geller, Rizzuto, & Kahana, 2008), focusing on the effects of recency and contiguity. Simulation 8 looks at similarity effects and their role in producing intra-list intrusions focusing on data from a study of memory for name-face pairs (Pantelis, van Vugt, Sekuler, Wilson, & Kahana, 2008). In a subsequent section, we evaluate CMR-IA's ability to account for the interactions between item and associative information, as seen in data from successive testing procedures (Kahana, 2000).

Simulation 5: Serial Position Effects in Cued Recall

Studies of short-term memory for associations often present subjects with a short series of paired items (typically words) and then test each of the studied pairs in a random order (e.g., Tulving & Arbuckle, 1963). Such experiments typically show very modest serial position effects overall. However, pronounced recency appears when considering the first tested pair. For example, Figure 6A shows a large recency effect in a paired-associate experiment in which subjects ($N = 16$) attempted to recall a single probed pair following the study of a six-pair list (Murdock, 1967).

Here we asked whether CMR-IA can match the serial position effect seen in Murdock's (1967) experiment. Using the PEERS Word Pool (Kahana et al., 2024), we simulated Murdock's experimental procedure with CMR-IA. Each simulation session consisted of 78 lists (as in the original experiment). Each list comprised both a study phase and a test phase. In the study phase, the simulated subject studied six word pairs in sequence, encoding each pair by updating the context as well as the M^{FC} and M^{CF} matrix following Equation 5, Equation 2 ($\beta = \beta_{enc}$), and Equation 1. Between the study phase and the test phase, context drifted following Equation 2 ($\beta = \beta_{distract}$). In the test phase, we randomly chose a studied pair and presented the first item as a cue. The simulated subject performed cued recall by first updating the current context with the cue, following Equation 2 ($\beta = \beta_{cue}$), and then trying to retrieve an item from memory, following Equation 3 and Equation A14. Following successful recall, the simulated subject updated the context with the retrieved item following Equation 2 ($\beta = \beta_{rec}$). Between each list, the context also drifted following Equation 2 ($\beta = \beta_{post}^{recall}$). We simulated a total of 1000 sessions.

We replicated the recency effect in cued recall in CMR-IA (Figure 6B). The contextual drift during the study-test lag, which is controlled by β_{enc} and $\beta_{distract}$, determines such an effect. The associated context of more recently studied pairs has a larger overlap with the current context, resulting in a better recall performance.

Simulation 6: Associative Symmetry

One of the classic questions concerning associations relates to their directionality. According to Ebbinghaus (1885/1913), studying a sequence of items leads to stronger forward than backward associations. Indeed, the concept of directional associations goes back to Herbart's (Herbart, 1834) theory of associative learning. According to this view, representations of the first and the second members of the pair associate via a directional bond that is stronger in the forward than the backward direction. An opposing view, however, emerged in Gestalt theories of associations. Kohler (1940) summarized the Gestalt position, which saw associative learning as creating a new memory composed of the elements of each to-be-associated item. This new memory does not possess any forward-directionality but rather combines the two constituents into a new "whole". Asch and Ebenholtz (1962) argued that empirical findings on associative recall more closely align with this Gestalt hypothesis, citing early work on backward associations by Murdock (1956). For a summary of the empirical literature on associative symmetry see Kahana (2002).

To examine the directionality of associative recall we modified Simulation 5 by probing memory in either the forward or backward direction[9]. We preassigned each simulated pair to be tested in either the forward direction (present item A_i as a cue to recall item B_i) or in the backward direction (present item B_i as a cue to recall item A_i) and tested them in a reversed order. Figure 7B shows CMR-IA's predicted probabilities of forward vs. backward recall; for comparison, Figure 7B shows empirical data reported by Kahana (1993). CMR-IA predicts equivalent forward and backward recall, consistent with experimental data. The reason that CMR-IA predicts symmetry for pairs even though it predicts forward asymmetry for free recall of lists is that both items in a given pair become associated with a single representation of context. In contrast, when simulating lists of discretely presented items, CMR updates the context following each encoding event.

Kahana (2002) suggested a stronger test of the associative symmetry hypothesis. Specifically, if subjects encode forward and backward associations holistically, then successively testing these associations should exhibit similar correlations when the repeated tests involve congruent cues (e.g., cueing with A for recall of B, or vice versa, on both tests) vs. incongruent cues (cueing with A for B on the first test, and B for A on the second test; or vice versa). Here we examined CMR-IA's predictions under similar conditions. Each simulation session consists of six lists. In the study phase of each list, we presented 12 word pairs randomly sampled from the PEERS Word Pool (Kahana et al., 2024) to the simulated subject. The simulated subject studied each pair once by updating

[9] There are other differences between this simulation and Simulation 5. To be as comparable as the experiment from Kahana (1993), in this simulation, the simulated subject is tested on all studied pairs in the reverse order of presentation. Also, each simulation session consists of 42 lists. Other details remain the same.

the context as well as the M^{FC} and M^{CF} matrix following Equation 5, Equation 2 ($\beta = \beta_{enc}$), and Equation 1. Then, all the studied pairs are tested twice, once in each of the two test phases, designated as Test 1 and Test 2. On Test 1, half of the studied pairs are cued in the forward direction ($A-?$), and the other half are cued in the backward direction ($? - B$). On Test 2, half of the studied pairs are tested in the same direction as Test 1, while the other half are tested in the reverse direction. We used the term *congruent* to refer to the two same-direction conditions (forward-forward and backward-backward) and the term *incongruent* to refer to the two different-direction conditions (forward-backward and backward-forward). The simulated subject performed cued recall by first updating the current context with the cue, following Equation 2 ($\beta = \beta_{cue}$), and then trying to retrieve an item from memory, following Equation 3 and Equation A14. Following successful recall, the simulated subject updated the context with the retrieved item following Equation 2 ($\beta = \beta_{rec}$). Here, we introduced output encoding, because we believe it is important for successive tests. Following successful recall, the simulated subject also encoded the cue and the retrieved item as a pair by updating M^{FC} and M^{CF} following Equation 5 and Equation 1. Context also drifted between phases following Equation 2 ($\beta = \beta_{distract}$) and between lists ($\beta = \beta^{post}_{recall}$). We performed the simulation for 1000 sessions. In line with Kahana (2002), CMR-IA predicts a similar correlation between congruent cues and incongruent cues (Table 2). This, again, is due to the fact that CMR-IA encodes both items in a pair symmetrically by associating them with a single representation of context.

The present simulation of cued recall requires that we take a stand on the process of updating context during the test phase of the experiment. In line with a large body of empirical evidence (see Carpenter, Pan, & Butler, 2022, for a review), we allow the model to learn during tests, so that the association between the cue and the retrieved item would be learned by updating M^{FC} and M^{CF}.

Simulation 7: Prior-list and Intra-list Intrusions in Cued Recall

Intrusions in cued recall tend to come from the target list more often than from prior lists (we refer to these two types of intrusions as intra-list intrusions, ILIs, and prior-list intrusions, PLIs). PLIs exhibit a recency effect insofar as they tend to come from recent lists more often than from remote lists (Davis et al., 2008). ILIs exhibit a contiguity effect insofar as they tend to come from positions near that of the just tested pair, an effect also present in serial recall (Solway, Murdock, & Kahana, 2012) and probed recall tasks (Kahana & Caplan, 2002). Figure 8A and 8C illustrate the conditional-response probability for ILIs and the recency effect for PLIs.

Here we ask whether CMR-IA exhibits these basic patterns. We simulated the same experimental paradigm as Davis et al. (2008) to assess PLIs and ILIs in CMR-IA. In each session, we presented 16 lists to the simulated subject, where each list comprised 12 word pairs sampled from the PEERS Word Pool (Kahana et al., 2024). Following the list presentation we tested the simulated subject on eight of the studied pairs. In the study phase of each list, the simulated subject encoded each pair by updating the context as well as the M^{FC} and M^{CF} matrix following Equation 5, Equation 2 ($\beta = \beta_{enc}$), and Equation 1. Between study and test phases context drifted according to Equation 2 ($\beta = \beta_{distract}$). In the test phase, we randomly chose eight studied pairs, where half are probed in the forward direction and the rest are probed in the backward direction. For each test cue, the simulated subject performed cued recall by first updating the current context with the cue, following Equation 2 ($\beta = \beta_{cue}$), and then trying to retrieve an item from memory, following Equation 3 and Equation A14. Following successful recall, the simulated subject updated the context with the retrieved item, following Equation 2 ($\beta = \beta_{rec}$). We performed the simulation 1000 times in CMR-IA.

CMR-IA reproduced the PLI recency effect (Figure 8B) and the ILI contiguity effect (Figure 8D). The rate of contextual drift, determined by the various β parameters, determines the degree to which PLIs from more recent lists will win the retrieval competition (as compared with PLIs from more remote lists). This is because a word pair studied in a more recent list would be associated with a context that is more similar to the current context, thus more likely to be retrieved. For ILIs, the rate of contextual drift determines the similarity among the contexts associated with neighboring studied pairs. As such, higher rates of contextual drift will lead to a strong ILI contiguity effect.

Simulation 8: Similarity Effects

Simulation 8 examines similarity effects and their role in producing intra-list intrusions (ILIs). Whereas most prior studies using verbal materials manipulated similarity in a categorical manner (Osgood, 1949), we sought to test CMR-IA using a dataset that varied similarity parametrically. For this purpose, we simulated data from a study of memory for name-face associations, where the researchers synthesized faces using a perceptual model that provided a metric space for computing inter-item similarity (Pantelis et al., 2008). In their experiment, Pantelis et al. (2008) asked subjects ($N = 25$) to learn a set of eight name-face pairs across 10 repeated study-test trials. They then examined both the probability of a correct recall and of an ILI as a function of the similarities among the faces.

We simulated an experimental paradigm to evaluate the effect of similarity on cued recall in CMR-IA. As in Pantelis et al. (2008), simulated subjects learned lists of eight randomly chosen face-name pairs drawn from a pool of 16 faces and 16 names. Rather than constructing correlated feature vectors for each name and face, we used orthogonal feature vectors and simulated inter-item similarity by setting the value of the pre-experimental associations represented by M^{FC}_{pre} and M^{CF}_{pre}. We used the coordinates of the 16 faces in the multidimentional scaling (MDS) space provided by Pantelis et al. (2008) to calculate the cosine similarity between faces as their pre-experimental associations. We assumed that there are no pre-existing associations among names and between names and faces. Our simulation deviated from Pantelis et al. (2008) in that we did not attempt to simulate multi-trial learning effects; rather, we just simulated the encoding and retrieval of a once-presented list of paired associates. During the encoding phase, context updated according Equation 5 and Equation 2 ($\beta = \beta_{enc}$); following each study pair, we updated M^{FC} and M^{CF} matrices according to Equation 1. Between the study phase and the test phase, context drifted following Equation 2 ($\beta = \beta_{distract}$). In the test phase, we probed the simulated subject with faces from each of the eight studied pairs to determine which name the model would recall (if any). For each test face, the simulated subject performed cued recall by first updating the current context with the face item, following Equation 2 ($\beta = \beta_{cue}$), and then trying to retrieve a name from memory, following Equation 3 and Equation A14[10]. Following successful recall, the simulated subject updated the context with the retrieved item following Equation 2 ($\beta = \beta_{rec}$). We performed the simulation 10,000 times.

Pantelis et al. (2008) first asked whether cueing memory with a face that has many neighbors (i.e., similar faces) among the set of studied name-face pairs resulted in poorer memory for the associated name. Defining neighbors as faces that fall within a radius of three units of the probe face in the MDS space, they found that recall probability decreased monotonically with the number of neighboring faces in the list (the “neighborhood” effect, see Figure 9A). They also asked whether ILIs are more likely to come from neighboring faces. Picking out ILIs, they calculated the Euclidean distance between the cue face and the face belonging to the incorrectly recalled name in the MDS space. They found that the conditional probability of making an ILI decreased monotonically with the distance of the two faces(Figure 9C). These results align with the predictions of CMR-IA (Figure 9B and Figure 9D). In CMR-IA, pre-experimental associations in the context-to-item weight matrix, controlled by the parameter s_{FC}, represent the similarity among faces[11]. When performing the cued recall, before retrieving a name, the simulated subject updated the

[10] We prohibited CMR-IA from recalling faces by only letting names enter the leaky accumulator in this simulation.

[11] s_{CF} does not have any influence in this simulation mathematically, because we confined pre-experimental associations among faces and prohibited faces from entering the leaky accumulator.

current context with the test face first, so the updated current context would have strong representations for the neighbors of the test face. The more neighbors it has, the more likely confusions will arise during name recall. Also, the nearer one neighbor is (i.e., closer distance in face space), the more likely its paired name will be wrongly retrieved.

4. Successive Tests of Item and Associative Information

Here we consider the relation between item and associative information as measured using the successive testing technique (Kahana, 2000). This method enables researchers to measure the correlation between successful retrieval of a given item under different conditions. In Tulving's classic procedure (Flexser & Tulving, 1978), subjects studied a list of $A - B$ word pairs and were then tested successively, first by item recognition and then by cued recall. In the item-recognition test, subjects were shown list items (B items from each of the studied pairs) intermixed with non-list items. Subjects responded "yes" to items if they remembered seeing them in the study list. In the cued-recall test, subjects attempted to recall the B items when given the A items as cues. In this manner, memory for each of the B items was tested twice — first by recognition and then later by recall. Flexser and Tulving observed a moderate degree of dependency between item recognition and cued recall at the level of individual subject-items. Expressed using Yule's $\mathcal{Q}$, a measure of association for 2×2 contingency tables, one finds values ranging from 0.45 to 0.65 across a wide range of experiments (Kahana et al., 2005)[12].

Kahana et al. (2005) analyzed the correlation between recognition and recall in four classes of linear distributed memory models that differed in mathematical operation they used to form and retrieve associations and in the process they used to represent and evaluate item information. In their formulation, however, none of the models possessed a representation of context, a factor critical to any episodic memory task (Kahana, 2020). Kahana et al. (2005) further showed that observed correlation between recognition and

[12] Yule's $\mathcal{Q}$ behaves much like an ordinary Pearson correlation, ranging from -1.0 (perfect negative correlation) to +1.0 (perfect positive correlation).

recall not only reflects the inherent model-derived correlation, but also reflects the effects of variability during encoding and retrieval, and the effect of the first test on the second test. Variability during encoding will influence both item-specific and relational information, thereby boosting the observed correlation. Output encoding can, under some assumptions, increase the correlation because information stored at test may differentially affect recognized and non-recognized items, and thereby increase the probability of recall. On the other hand, varibility in the criterion for successful recognition or recall can reduce the correlation between tasks. This is because the two tests of a given pair are widely separated in time and the retrieval conditions therefore introduce variance that is unique to each test.

Here we consider the more general question of contingency between successive tests of item and associative information. Comparisons of recognition and recall can be simplified by separating differences due to the response stage from differences in the information being tested. This can be achieved by making two comparisons, first between item recognition and cued recall and then between associative recognition and cued recall. Differences between item recognition and cued recall reflect both differences in information and differences in the response stage. Differences between associative recognition and cued recall should primarily reflect differences in the response stage. To the extent that the differences between item recognition and cued recall exceed those between associative recognition and cued recall, these can be attributed to differences in the type of information tested in the two tasks.

In a study comparing dependency relations among different memory tests, Wallace (1978) found a high degree of dependency between pair recognition and cued recall. In his pair recognition task, subjects distinguished intact study pairs from new pairs. This is not strictly an associative-recognition task because subjects could use both item-specific and associative information to distinguish the pairs (in the associative recognition paradigm, subjects distinguish intact pairs from rearranged pairs).

Many distributed memory models, including the four analyzed by Kahana et al.

(2005), assume that associative recognition and cued recall operate on the same underlying representation. As such, these models predict that the informational correlation would be very high. Similarly, the increased overlap of information would lead to a greater possible effect of output encoding. That is, an associative recognition test may boost subsequent cued recall more than an item recognition test. This will tend to exaggerate the predicted contingency beyond that derived from the basic model equations theory. Experiment 2 and the associated modeling examine this issue in detail.

To more fully evaluate the relation between recognition and recall we conducted another experiment that compared successive recognition tests where the first test is either item or associative information, the second test is either item or associative recognition, and the items tested in the first and second test are either the same or different. By introducing conditions of total informational overlap and no informational overlap, we can estimate the effects of variability in encoding and retrieval, as we explain below.

Experiment 2

Experiment 2 compared contingency relations between three types of recognition tasks and a subsequent cued recall test for the same items: In the associative recognition task, subjects judged probe pairs as intact or rearranged; in the item recognition task, subjects judged individual items as old or new; in the pair recognition task, subjects judged pairs of items as old or new. Although our central question concerned the difference in correlation between recognition and recall under these three conditions, we also examined the hypothesis that asymmetrical associations may underlie the recognition-failure of recallable words (Rabinowitz, Mandler, & Barsalou, 1977). According to this account, the backward association from B to A is significantly weaker than the forward association, and therefore, when tested on recognition of B alone, subjects often fail to recall A, resulting in an incomplete match to the stored information. As expressed by Rabinowitz et al. (1977, p. 661) “Retrieval must be asymmetric in order for recognition failure to occur, and it is expected to be asymmetric given the nature of the instructions (study the pair so that you can recall the TBR [to-be-remembered item] given the LC [context item] as a cue).” We did not expect to find support for this hypothesis because of the extensive body of literature showing symmetric recall of paired associates under a wide range of experimental conditions (see Kahana, 2002, for a review). Nonetheless, none of these studies had specifically considered symmetry in the recognition-failure paradigm, so it seemed worth exploring.

We tested the asymmetry hypothesis by manipulating the order of the pairs in each condition of this experiment. In the cued recall phase, subjects were either shown A and asked to recall B, or they were shown B and asked to recall A. Order effects were also tested in the three recognition tests. In pair and associative recognition, subjects were either shown $A - B$ or $B - A$, and in item recognition subjects were either tested on the A item or the B item. It was hypothesized that symmetry would be obtained in cued recall and item recognition but not in associative or pair recognition. The hypothesized violations of symmetry in the pair and associative recognition tests were motivated by the belief that all of the information about the stimulus (including perceptual features) is encoded at study and enters the memory comparison at test. When the full pair is shown at test and it is reversed in order from the presentation at study, the mismatching perceptual information should lower performance on the recognition test. This should not happen in recall or item recognition because only a single item is provided as the test cue.

Methods

Subjects

Ninety undergraduate students at the University of Toronto participated for extra-credit in an introductory psychology course. All subjects were native English speakers.

Procedure

We randomly assigned subjects to each of three groups. Group one performed a standard yes-no item recognition test. Group two performed the pair recognition test used by Wallace (1978). Group three performed an associative recognition test. Twenty-one subjects failed to recall at least 10% of the list items and were excluded from the data analyses.

Subjects studied a list of 48 word-pairs and then took part in successive recognition and cued-recall tests. During the study phase, word pairs, randomly selected from the Toronto word pool, were presented visually at a rate of one pair every two seconds. We did not test subjects on the last eight word pairs on the study list, designating these as recency buffers.

During the test phase, subjects in group one (item recognition) made yes/no judgments on old and new single items. Subjects in group two (pair recognition) made yes/no judgments on intact and new pairs. Subjects in group three (associative recognition) made yes/no judgments on intact (old) and rearranged (new) word pairs. Intact pairs matched those shown during study, whereas rearranged pairs combined items from two different study pairs to form a new pair. Half of the test probes were tested in the forward order ($A - B$) and half in the backward order ($B - A$). If a subject did not respond within eight seconds, the computer moved along to the next item.

After completing the recognition-memory test, we administered a cued-recall test on the studied pairs. In the associative recognition group, we only tested intact pairs during the cued-recall phase. Half of the items were tested with forward recall; the other half were tested with backward recall. Subjects had 12 seconds to type their responses on a computer keyboard. The entire study-test process repeated multiple times. On each repetition, a new random sample of words was selected from the Toronto word pool.

Results

Before turning to our primary questions concerning the dependency relations between successive recognition and recall tests, we comment briefly on the differences in overall recognition and recall performance. Performance on cued recall differed across the test groups, $F(2, 66) = 11.0, MSe = 0.025, p < 0.001$ (see Table 3). Tukey HSD post-hoc comparisons revealed significant differences between associative and pair recognition groups ($p < 0.05$), and between item and associative recognition groups ($p < 0.01$). The difference between item and pair recognition groups just missed our threshold for statistical significance ($p = 0.06$). These differences likely reflected the differing numbers of items intervening between original study and cued recall for each of the three groups. For the associative recognition group, 40 pairs were tested in the recognition phase, for the pair recognition group, 80 pairs were tested in the recognition phase, and for the item recognition group 80 individual items were tested in the recognition phase. Recognition performance, as measured by d', also varied across the three groups, $F(2, 66) = 13.1, MSe = 0.3, p < 0.01$. Performance in the pair-recognition group was significantly higher than either the associative recognition or the item recognition groups ($p < 0.01$). In contrast, no significant difference in performance emerged between the item and associative recognition groups.

Dependency Relations

Both the associative and pair recognition tests yielded higher recognition-recall dependencies than the item recognition test (see Table 3). $\mathcal{Q}$ varied significantly as a function of the type of recognition test, $F(2,66) = 10.6, MSe = 0.029, p < 0.001$. Tukey HSD post-hoc comparisons revealed that recognition-recall dependencies were significantly higher when associative recognition ($p < 0.001$) or pair recognition ($p < 0.05$) were tested than when item recognition was tested. The difference between the associative and pair recognition exhibited a trend suggestive of a potential difference between these conditions ($p = 0.08$).

Correlations between Recognition and Recall Performance across Subjects

Along with examining the item by item dependencies, one can consider the inter-subject correlations between recognition and recall performance. The correlation between associative recognition and cued recall ($r^2 = 0.80$) was higher than the correlation between pair recognition and cued recall ($r^2 = 0.735$) and the correlation between item recognition and cued recall ($r^2 = 0.33$). The item-recognition vs. cued-recall correlation accords well with classic data reported by Underwood, Boruch, and Malmi (1978). All correlations were statistically significant.

Associative Symmetry

Although experimental conditions were designed to promote sequential processing of the pairs[13], we failed to observe any significant difference between forward and backward cued recall; the mean difference score (forward recall - backward recall) was 0.007 (± 0.020 [95% CI]). Unlike cued recall, we observed significantly better recognition of intact probes than of reversed probes in the associative recognition task ($t = 4.9, df = 37, p < 0.001$). If the magnitude of asymmetry found in associative recognition is taken as an estimate of the asymmetry that might be expected in cued recall (if the null hypothesis were false), then

[13] Subjects read the word pairs aloud from left to right as soon as they appeared. During the recall phase, the probe word appeared on the same side of the screen as it had appeared during the study phase.

the present data has a power of 0.93 for rejecting the null hypothesis of symmetry in cued recall. As a general rule, successful cued recall of meaningful pairs does not depend on the order of study (Kahana, 2002).

Our findings demonstrate a clear dissociation between associative recognition and cued recall. This dissociation contrasts with the high dependency relations and inter-subject correlations between associative recognition and cued recall. Nonetheless, one can explain the asymmetry effect by assuming that subjects base their recognition judgments partially on the perceptual familiarity of the pair. If one reverses the order of the pair, perceptual familiarity can not be used in making the recognition judgment. Memory models can account for this effect if item vectors include source information on the perceptual attributes of an item (e.g., Kahana, 2002; Polyn et al., 2009).

Manipulating recall direction allowed us to test the hypothesis that failed backward recall produces recognition-failure of recallable words (e.g., Rabinowitz et al., 1977). The present data argue against this explanation. Because typical levels of recognition failure were coupled with near-perfect symmetry between forward and backward recall, we do not find any evidence for the Rabinowitz et al. (1977) hypothesis.

Experiment 3

Experiment 3 examined the effect of informational overlap on the correlation between successive memory tasks. We presented each study pair six times to reduce encoding variability resulting from attentional fluctuations. We included comparisons of both identical probes (on Tests 1 and 2) and unique probes to allow us to directly estimate the effects of variability during encoding and retrieval on the Test 1 – Test 2 correlation.

Method

Subjects

Sixty-three undergraduate students at Brandeis University participated for extra credit in an introductory psychology course. All subjects were native English speakers.

Procedure

A single trial consisted of a study phase followed by two successive recognition test phases (termed Test 1 and Test 2). Each study list was composed of 25 unique word pairs (denoted $A - B$) drawn at random from the noun subset of the Toronto noun pool (Friendly, Franklin, Hoffman, & Rubin, 1982). Across the study list, each $A - B$ pair appeared six times subject to the constraint that at least one distinct pair intervened between repeated pairs. Each pair appeared individually in the center of a computer screen for 2s. Subjects studied these pairs for a later recognition test. After studying a given list and before the successive recognition tests, subjects made same-different discrimination judgments on 15 pairs of patterns. This distractor task served to attenuate recency effects. Each pattern was constructed by randomly filling elements in a 4×4 square matrix.

Successive recognition tests followed the distractor task. The cues given in these tests reflected seven different within-subject conditions. Each of the 25 word pairs in the study list and an additional 5 pairs that were not in the study list were assigned to a condition, which pre-determined the cues given in Test 1 and Test 2. The non-studied pairs serve as lures that were tested in both Test 1 and Test 2. Five target conditions represent the following cues at Test 1 and Test 2: Six "Same Item" cues (A or B at Test 1 and the same item at Test 2); six "Different Item" cues (A or B at Test 1 and the other item at Test 2); three "Intact Pair" cues ($A - B$ at Test 1 and $A - B$ at Test 2); nine "Item/Pair" cues (A or B at Test 1, and the pair $A - B$ at Test 2); six "Pair/Item" cues (the pair $A - B$ at Test 1 and A or B at Test 2)[14]. In addition to testing the 25 studied pairs and the five lures that appeared in both Test 1 and Test 2, there were 30 additional lures presented at each test. Of these 30 additional lures, half were single items and half were pairs.

In the recognition test phases, we asked subjects "Was the word on the list?" or "Was the word pair on the list?" If the probe item is a word, then either A, B, or X

[14] The five non-studied pairs were randomly assigned to one of the target conditions along with the 25 studied pairs. The following analysis only included the non-studied pairs that were tested with "Same Item" cues and "Intact Pair" cues. The result is represented by the "Repeated Lure" condition in Table 4.

(where X is a lure not from the study list) appeared in the center of the screen. If the probe item is a word pair, then either $A - B$ or $X - Y$ (where X and Y denote lures) are presented in the center of the screen. Subjects responded “yes” or “no” by pressing the left and right control keys. For half of the subjects, the left key indicated a “yes” response; for the other half, the right key indicated a “yes” response.

The pattern-matching distractor task began after each of the recognition tests. The entire experiment consisted of three trials (each trial having a study list followed by two successive recognition tests as described above). We designated the first trial as practice and excluded these data from subsequent analyses. The stimuli for each trial consisted of a different random sample taken from the Toronto noun pool, with no words repeated across trials.

Results

Table 4 reports overall accuracy across conditions for each of the successive memory tests. For the first test, the proportion of "yes" responses (hit rate) depends primarily on whether the cue is a single item or a pair. For single-item cues, performance hovers around 0.81, whereas for pair cues, performance hovers around 0.90. This difference likely reflects the benefits of having a richer cue to compare with memory.

For the second test, the hit rate for single items not present on the first test drops precipitously (from 0.82 to 0.68), whereas the hit rate for items present on the first test increases slightly (from 0.81 to 0.82). For the second test, the hit rate for pairs present on the first test increases slightly (from 0.90 to 0.92). This pattern strongly suggests an effect of output encoding.

Dependency Relations

The $\mathcal{Q}$ values varied widely across the different Test 1 – Test 2 conditions. If associative recognition and cued recall are more alike than they are different (e.g., Nobel & Shiffrin, 2001), the item recognition — associative recognition condition should yield data comparable to the values obtained in Experiment 2. Consistent with this position, $\mathcal{Q}$ for item recognition followed by associative recognition was 0.64 in this experiment, only slightly higher than the values obtained in Experiment 2 for comparisons of item recognition and cued recall.

When the order of item and associative tests reversed, with the associative test given first, the correlation does not change significantly ($\mathcal{Q} = 0.59$). Testing the same information on Test 1 and Test 2 yielded very high correlations ($\mathcal{Q} = 0.94$ for repeated associative recognition tests, and $\mathcal{Q} = 0.86$ for repeated item recognition tests); with values exceeding the correlation between associative recognition and cued recall.

Subjects exhibited a significant positive contingency between successive recognition judgments made on the first and second members of a studied pair, even those the two tests are widely separated in time ($\mathcal{Q} = 0.26$). Such a finding could be explained in one of two ways: 1) the correlation is induced by shared variability, across pairs, in the quality of encoding both the A_i and the B_i items. 2) when A_i is presented as a cue in Test 1, subjects implicitly recall the $A_i - B_i$ pair, thereby strengthening the representation of B_i in memory (output encoding).

5. Modeling the Dependency between Items and Associations

Here, we evaluate CMR-IA's fit to data from Experiments 2 and 3. For both experiments, we created simulations to match the procedures used with actual subjects. We describe the methods and results for each experiment below.

Simulating Experiment 2

In the study phase, we presented a list of 48 word pairs randomly selected from the PEERS Word Pool (Kahana et al., 2024) to the simulated subject. The simulated subject encoded each pair by updating the context as well as the M^{FC} and M^{CF} matrix following Equation 5, Equation 2 ($\beta = \beta_{enc}$), and Equation 1. For item recognition, the simulated subject compared the current context with the context input brought by the item following Equation 4. If the context similarity exceeds the item recognition threshold c_{recog}^{item}, the simulated subject would give a response of "old", otherwise "new". For associative recognition and pair recognition, the simulated subject updated the current context using the first item following Equation 2 ($\beta = \beta_{cue}$) and then compared the current context with the context input brought by the second item following Equation A11. If the context similarity exceeds the associative recognition threshold c_{recog}^{assoc}, the simulated subject would give a response of "intact", otherwise "new/rearranged". After the recognition judgment, the simulated subject updated the context with the test item or the second item in the test pair, following Equation 2 ($\beta = \beta_{cue}$). For iitems/pairs judged as old, the simulated subject updated the M^{FC} and M^{CF} matrix following Equation 5 and Equation 1, simulating output encoding. For cued recall, the simulated subject first updated the current context with the cue, following Equation 2 ($\beta = \beta_{cue}$), and then trying to retrieve an item from memory, following Equation 3 and Equation A14. Following successful recall, the simulated subject updated the context with the retrieved item according to Equation 2 ($\beta = \beta_{rec}$). The simulated subject also encoded the cue and the retrieved item as a pair following Equation 5 and Equation 1. Following Lohnas et al. (2015), $\beta_{distract}$ determined the context change between phases (study and test, or successive tests) and β^{recall}_{post} determined the context change between study-test-test trials. We performed 100 simulation sessions for each group using particle swarm optimization to identify optimal parameter values (all groups shared the same set of model parameters).

Table 3 provides simulation results alongside the experimental data. Performance on cued recall differs significantly across simulation groups ($F(2, 297) = 1173.05$, $p < .001$). Tukey HSD post-hoc comparisons reveal significant differences between all three groups (ps $< .001$). Recognition performance, as measured by d', also varies across three groups ($F(2, 297) = 2911.04$, $p < .001$). Tukey HSD post-hoc comparisons reveal significant differences between all three groups (ps $< .001$). The dependency relations, as measured by $\mathcal{Q}$, vary according to the types of recognition test ($F(2, 297) = 87.02, p < .001$). Tukey HSD post-hoc comparisons reveal significant differences between all three groups (ps $< .001$). Although we report statistics on the model, we recognize that in modeling the group data rather than individual subjects (as is commonly done), one cannot easily relate the statistical findings between the model and the data.

For item recognition, CMR-IA directly compares the retrieved context with the current context. In contrast, for both associative recognition and cued recall, the cue item or the first item in the test pair reinstates its associated context, CMR-IA then uses the updated context to recall or recognize the other item. The congruence in how the model uses the stored association accounts for the high dependency between associative recognition and cued recall. Such differences between item recognition and associative recognition also interact with the process of output encoding. CMR-IA only encodes the recognized probe for item recognition, which helps little in retrieving the other item. However, for associative recognition, the model encodes the recognized pair, thus strengthening the association between the two items in the pair. Nevertheless, CMR-IA does not have a different mechanism for pair recognition compared with associative recognition, which might account for the somewhat poorer fit to the pair recognition data.

We also examine the hypothesis of associative symmetry in CMR-IA. Consistent with our Gestalt approach in modeling the encoding of word pairs, we find no significant difference between forward and backward in cued recall ($p = .394$). However, our model fails to produce the empirically-observed asymmetry in associative recognition tests. In CMR-IA an intact pair and a reversed pair would not be distinguishable for the simulated subject. A revision of the model that encodes the location of the words on the screen in the item vectors could overcome this limitation.

Simulating Experiment 3

We replicated the same experimental paradigm to simulate Experiment 3 in CMR-IA. In the study phase, we presented a list consisting of 25 unique word pairs randomly selected from the PEERS Word Pool (Kahana et al., 2024) to the simulated subject. Each pair appeared six times in a random order. The simulated subject encoded each pair by updating the context as well as the M^{FC} and M^{CF} matrix following Equation 5, Equation 2 ($\beta = \beta_{enc}$), and Equation 1. In the test phase, we let the simulated subject complete two recognition tests successively. For item recognition, the simulated subject compared the current context with the context input brought by the item following Equation 4. If the context similarity exceeds a threshold c determined by Equation A10, the simulated subject would give a response of “old”, otherwise “new”. For associative recognition, the simulated subject updated the current context using the first item following Equation 2 ($\beta = \beta_{cue}$) and then compared the current context with the context input brought by the second item following Equation A11. If the context similarity exceeds a threshold c determined by Equation A12, the simulated subject would give a response of “old”, otherwise “new”. Note that we introduced a new parameter σ for randomness in the item and associative recognition thresholds to model retrieval variability. After the judgment, the simulated subject updated the context with the test item or the second item in the test pair, following Equation 2 ($\beta = \beta_{cue}$). We introduced output encoding only for items/pairs judged as old. In these cases, the simulated subject updated the M^{FC} and M^{CF} matrix following Equation 5 and Equation 1. Following Lohnas et al. (2015), $\beta_{distract}$ determined the context change between phases (study and test, or successive tests), and β_{post}^{recall} determined the context change between study-test-test trials. We repeated the study-test-test trials three times in one session while the first was not included in the analysis, as in the real experiment. We performed 300 simulation sessions using particle swarm optimization to identify optimal parameter values.

Table 4 provides simulation results alongside the experimental data. As in the data, CMR-IA exhibited high correlations between repeated item recognition tests ($\mathcal{Q} = 0.81$) and repeated associative recognition tests ($\mathcal{Q} = 0.81$). The reinstatement of context during Test 1 ensures the comparable recognition performance during Test 2. The mechanism of output encoding also contributes to the slight increase in recognition performance when the same items/pairs are tested twice.

When testing different information, $\mathcal{Q}$ for item recognition followed by associative recognition is 0.61, and $\mathcal{Q}$ for associative recognition followed by item recognition is 0.64. CMR-IA predicts a decrease in the correlation (as compared with repeated tests of the same information) owing to the mechanistic differences between item and associative recognition. As stated in the previous section, the model uses the stored association differently for item and associative recognition. In addition, the process of output encoding during item recognition does not help to strengthen the association between the two items within a pair.

Repeatedly testing different items in a word pair reveals the weakest dependency relation ($\mathcal{Q} = 0.27$) as compared with all other conditions. This is because the mnemonic information probed in Test 2 does not benefit from output encoding during Test 1.

In the repeated lure condition, a recognition lure from Test 1 would re-appear on

Test 2. Even though neither item appeared in the study list, successive tests of these lures exhibited very strong dependencies. CMR-IA captures this result, demonstrating moderately high correlations between repeated lure responses similar to those seen in the data (see Table 4). This positive correlation arises because CMR-IA encodes those lures that are wrongly judged as old in Test 1 thus raising the probability of judging them as old again in Test 2. Retrieval variability (controlled by the parameter σ) moderates this correlation, bringing it in line with the levels observed in the data.

6. General Discussion

We propose CMR-IA as a unified theory of memory for items and associations. Most prior developments of retrieved-context theory sought to explain recall dynamics in free recall. Recent papers have extended these models to account for data on item recognition (Healey & Kahana, 2016) and serial recall (Lohnas, 2025). The present paper fills a key gap by extending this class of models to data on cued recall and associative recognition and comparing memory for associations with memory for items, including novel comparisons of successive tests. Simultaneously accounting for data on the relations between item and associative memory has long constituted a key challenge for computational models of human memory (Murdock, 1992). We adopted the Gestalt approach to associative learning by associating the two constituent items in a pair with the same context. This is realized by adding the two feature vectors together, updating the context with the combined feature vector, and storing the association in the associative matrix. Our approach ensures the symmetry of the two items within a pair. An alternative approach is to treat the pair as a list of two items. However, this approach implies an update in context and makes the two items associated with different contexts, thereby producing a forward asymmetry similar to that seen in free recall (Kahana, 1996; Healey, Long, & Kahana, 2019).

We designed a mechanism to exploit the stored associative information and demonstrate that it could let our model perform associative recognition and cued recall. To perform associative recognition, one of the items in the test pair reinstates its associated context into the current context and the model compares the updated context with the associated context of the other item. To perform cued recall, the cue item reinstates its associated context into the current context and the model uses the updated context to retrieve the target. In contrast to item recognition and free recall where the model directly uses the current context to perform recognition or recall, the additional step of reinstatement in both the processes of associative recognition and cued recall enables CMR-IA to bring the stored associative information into play. With such a mechanism, we showed that CMR-IA successfully simulated several benchmark phenomena in item and associative memory.

Item and Associative Recognition

Our first set of simulations aimed to evaluate whether CMR-IA could provide a good qualitative (and in some cases quantitative) account of select phenomena concerning item and associative recognition.

For our first simulation, we sought to evaluate whether the model could account for the joint effects of recency and similarity on item recognition. As we could not find an adequate archival dataset whose data were still available for secondary analyses we conducted a novel experiment. Subjects performed a continuous recognition task where we presented an item twice with different lags and grouped same-category items into blocks during the presentation. As expected, data from this experiment show that subjects more easily recognized recent items and those with fewer similar prior items. CMR-IA captured these effects because drifting context produces recency and because pre-experimental associations between items (our model of similarity) increase false alarms to new items.

G. Schwartz et al. (2005) documented a contiguity effect in item recognition (see, also Osth et al., 2018; Malmberg & Annis, 2012). They found that when successively testing two old items on a recognition memory test, subjects were more likely to make successive high-confidence hits when the items were studied at nearby list positions. Context reinstatement in CMR-IA produces this contiguity effect as a test item would retrieve a context that is similar to the context associated with the items studied nearby, thus facilitating the recognition of these items.

Hockley (1992) found that items and associations have different forgetting rate. Specifically, he reported that the recognition performance for items drops more quickly than associations when the study-test lag increases. In CMR-IA, the difference in mechanism for item recognition and associative recognition explains this phenomenon. An additional step of reinstatement in the process of associative recognition enables CMR-IA to make more use of the stored associative information and depend less on the current temporal context. Thus, the performance of associative recognition is less influenced by the study-test lag than the performance of item recognition.

Word frequency effects are well-documented effects in item recognition that rare words are more easily recognized as targets and more easily rejected as lures (Schulman, 1967; Lohnas & Kahana, 2013). When simulating the word frequency effect in CMR-IA, we find that elevated attention for rare words plus adjustments to the recognition threshold is necessary for a successful simulation. The increased attention for rare words in CMR-IA ensures that rare words are better encoded during the study phase, thus they have a higher hit rate. Also, CMR-IA assigns higher recognition thresholds to rare words during recognition, leading to a lower false alarm rate.

Cued Recall

We next asked whether CMR-IA could account for a set of key results concerning paired-associate (cued) recall. We began by asking whether our model could account for the classic serial position effect in cued recall. Specifically, the last couple of studied word pairs enjoy substantially higher recall accuracy than other word pairs, especially when tested immediately after study (Murdock, 1967). As in free recall, CMR-IA generates this recency effect due to the evolution of temporal context: contexts associated with final list items will bear greater overlap with time-of-test context, especially on an immediate recall test.

We next considered the phenomenon of associative symmetry (Kahana, 2002).

Although subjects in free recall exhibit a marked forward asymmetry in their recall transitions, this asymmetry disappears when testing paired associates. Specifically, after studying a list of pairs in the order $A_i - B_i$ subjects will be able to recall A_i given B_i just as easily as they can recall B_i given A_i. In CMR-IA, this effect arises naturally under our approach to modeling associative learning. As A_i and B_i share the same context, cueing with either member of the pair will retrieve the same context which serves as an equally effect cue for the other item.

Recall errors can provide valuable information about the nature of memory encoding and retrieval. Specifically, in the case of cued recall, subjects will occasionally commit intrusions, which at times come from associations learned on prior lists (prior-list intrusions, PLIs) or from other pairs within the same list (intra-list intrusions, ILIs). PLIs exhibit a marked recency effect, being more likely to come from recent than from remote lists. ILIs tend to come from pairs studied near the test pair, demonstrating a contiguity effect (Davis et al., 2008). In CMR-IA, contextual drift accounts for the recency effects of PLIs, and the similarity in context among proximally studied pairs accounts for the contiguity effects of ILIs.

Similarity can have a significant effect on the nature of both correct responses and errors in cued recall tasks. Specifically, items with more neighbors (i.e., items with similarity above a threshold) during the study tend to have a lower probability of correct recall, and ILIs are more likely to come from similar items (Pantelis et al., 2008). CMR-IA captured this effect because pre-experiment associations lead more similar items to have more similar associated contexts, resulting in a higher probability of incorrect retrievals.

Successive Tests of Item and Associative Information

To further explore the relation between item and associative information, we report two new experiments and fit CMR-IA to the results. The first experiment concerns successive tests of recognition and cued recall, showing that associative recognition has a higher correlation with cued recall than item association. The second experiment concerns successive tests of recognition, demonstrating that a larger informational overlap results in a higher correlation between successive memory tasks.

We fit CMR-IA quantitatively to the results of the two experiments and it turns out that CMR-IA could fit the data well. Therefore, the discrepancy between item and associative information is explained by the difference in the retrieval process in CMR-IA. Item information corresponds to the similarity between the current context and the retrieved context by the item. However, associative information depends less on the current context. Rather, it depends more on the association between the two items, represented by the similarity between the two retrieved contexts by the two items. Moreover, output encoding and retrieval variability ensure that CMR-IA captures the full picture of the successive tests, such as the moderate dependency between repeated lures.

Relation to Other Models

Murdock (1982)'s Theory of Distributed Associative Memory (TODAM) provided one of the first attempts at a theory of item and associative memory. As described in the Introduction, TODAM represented each studied item as an N-dimensional vector and assumed that learning involved the accumulation of these memories by vector addition. TODAM assumed that learning a novel pair of items involved the discrete (vector) convolution of the constituent items, and that summing the convolution along with the items enabled the storage of both items and associations in a common memory vector. TODAM simulated item and associative recognition as a similarity operation: If the dot-product between the probe item or association and the memory vector exceeded a threshold, the model would endorse the probe item as old; otherwise, the model would consider the probe item as new. Cued recall involved computing the inverse convolution of the probe with the memory vector to retrieve an approximate representation of the target. A hypothesized "deblurring" stage would convert the retrieved representation into a target memory, again using a similarity operation to determine the best match. In this manner, TODAM could simulate item recognition, associative recognition, and cued recall in a common framework. Concurrent with Murdock's TODAM, Metcalfe (Metcalfe-Eich, 1982) proposed a similar approach but used auto-convolution rather than direct summation as the basis for storing item information.

Subsequent years saw a gold rush of theory development. Gillund and Shiffrin (1984) extended their Search of Associative Memory model of free recall to account for data on item and associative information. Humphreys et al. (1989) developed a tensor model that allowed for three-way associations between context, items, and associations. Although list context did not directly play a role in Murdock and Metcalfe's original implementations, most subsequent applications included a representation of context and in several cases, an autocorrelated representation of context (Murdock, 1997; Mensink & Raaijmakers, 1988). Following Murdock, most models assumed distinct storage of item and associative information (Cox, 2024; Osth & Dennis, 2015) and distinct retrieval operations for recognition and recall (Shiffrin & Steyvers, 1997; Osth & Dennis, 2015).

CMR-IA parts from most other theories by assuming a common learning rule for items and associations. Whether learning two neighboring items in a free recall task or learning a paired associate in a cued recall task, the model associates each item with a representation of context. However, in learning a pair of items, the model associates the same context with both items. In recognition and cued recall, the vector representing the cue item retrieves the history of previously associated context vectors. This retrieval operation, which sits at the heart of the model, drives both recognition and recall decisions. The similarity between retrieved context and current context drives recognition decisions. Because CMR-IA does not directly associate the constituent items of a paired associate, the model cannot use a direct similarity operation for associative recognition. Rather, the first item processed by the subject retrieves its associated context, which in turn updates the current context before comparing the retrieved context of the second item and the current context.

Limitations and Future Directions

To the extent that item and associative memory tasks share fundamental properties with tasks already modeled using CMR2 and its variants, our model's success should not surprise anyone. Many fundamental principles of memory, including recency, similarity, contiguity, associative interference, and context-dependency, appear conserved across a wide range of tasks (Kahana, Diamond, & Aka, in press). The data that will offer the strongest model tests and thus highlight potential limitations will come from dissociations among these tasks. For example, Simulation 3 showed that whereas CMR-IA naturally produced the differential forgetting rates of item and associative information, the model struggled to simultaneously produce the appropriate levels of performance in both tasks. Other dissociations between item and associative tasks come from studies of response times (Nobel & Shiffrin, 2001) and dynamic retrieval has played a key role in recent models of item and associative memory, enabling these models to account for both accuracy and RT distributions (Diller, Nobel, & Shiffrin, 2001; Cox, 2024). Our model, however, lacks a mechanism to account for reaction time in recognition memory tasks and thus cannot answer for these data.

The predicted differences between item and associative memory stem from the central assumption in our model: namely, the associative process that binds two items being encoded as a pair. We adopted the strong position that paired items associate with a single shared context vector. This allowed us to account for symmetric associations within pairs (Kahana, 2000) even though the associative process that links neighboring items in a free or serial recall experiment is inherently forward biased (Howard & Kahana, 2002). Our model did not provide a natural means of explaining the observed asymmetry in associative recognition.

More broadly, we acknowledge that in learning consistent pairings of items subjects would likely create new conjunctive features. Compound words, such as ice cream, pancake, and laptop, naturally illustrate the creation of such conjunctive features and, ultimately, new units in memory. CMR-IA does not explain how new feature vectors emerge through the history of shared experience. We see this as a fundamental challenge for future work.

Our work adhered to the common practice of fitting group average data rather than data from individual subjects. Thus, we only use one set of parameters to simulate each effect, mimicking the situation of an average subject performing many trials. This approach could not provide insights into individual differences in item and associative memory, including the correlations across subjects reported in our novel experimental data. A better approach would be to fit different sets of parameters for different subjects, thus accounting for the variability in responses. For example, Healey and Kahana (2016) fit individual young and old subjects in free recall tasks and showed that age-related memory change attributes to attention sustainability, context retrieval ability, intrusion screening ability, and retrieval noise. Cohen and Kahana (2022) and Lohnas (in press-b) also demonstrate the utility of individual subject fits in evaluating the predictions of memory theories.

Conclusions

The present work shows how a context-maintenance and retrieval model (CMR-IA) can explain a wide range of facts about item and associative memory. Our model captures associative effects by assuming that each item in a pair binds with an identical representation of context and context drifts between pairs and lists according to the standard CMR context evolution equation (Lohnas et al., 2015; Cohen & Kahana, 2022). The shared context among the items thus serves as the glue binding them together, symmetrically. As in free recall, context dynamics generate recency and contiguity at multiple time scales. Learning during recall combines with the core model mechanisms to explain the correlations between successive tests. Pre-experimental associations allow the model to account for similarity and interference effects. These findings demonstrate that retrieved-context theory readily extends from self-cued to experimenter-cued memory tasks.

References

Asch, S. E., & Ebenholtz, S. M. (1962). The principle of associative symmetry. *Proceedings of the American Philosophical Society*, *106*, 135–163.

Bahrick, H. P. (1970). Two-phase model for prompted recall. *Psychological Review*, *77*(3), 215–222. doi: 10.1037/h0029099

Bates, D., Kliegl, R., Vasishth, S., & Baayen, H. (2018). Parsimonious mixed models. *arXiv*(arXiv:1506.04967). doi: 10.48550/arXiv.1506.04967

Benjamin, A. S., Diaz, M., & Wee, S. (2009). Signal detection with criterion noise: applications to recognition memory. *Psychological review*, *116*(1), 84. doi: 10.1037/a0014351

Bower, G. H. (1972). Stimulus-sampling theory of encoding variability. In A. W. Melton & E. Martin (Eds.), *Coding processes in human memory* (pp. 85–121). New York: John Wiley and Sons.

Caplan, J. B. (2005). Associative isolation: Unifying associative and list memory. *Journal of Mathematical Psychology*, *49*, 383–402.

Carpenter, S. K., Pan, S. C., & Butler, A. C. (2022). The science of effective learning with spacing and retrieval practice. *Nature Reviews Psychology*, *1*(9), 496–511.

Carr, H. A. (1931). The laws of association. *Psychological Review*, *38*, 212-228.

Clark, S. E. (1992). Word frequency effects in associative and item recognition. *Memory & Cognition*, *20(3)*, 231–243.

Cohen, R. T., & Kahana, M. J. (2022). A memory based theory of emotional disorders. *Psychological Review*, *129*(4), 742-776. doi: 10.1037/rev0000334

Cox, G. E. (2024). Dynamic retrieval of events and associations from memory: An integrated account of item and associative recognition. *Psychological Review*.

Cox, G. E., & Shiffrin, R. M. (in press). Oxford handbook of human memory. In M. J. Kahana & A. D. Wagner (Eds.), (2nd ed., chap. Computational Models of Event Memory). Oxford, U. K.: Oxford University Press.

Criss, A. H., & Malmberg, K. J. (2008). Evidence in favor of the early-phase elevated-attention hypothesis: The effects of letter frequency and object frequency. *Journal of Memory and Language*, *59*, 331-345.

Davis, O. C., Geller, A. S., Rizzuto, D. S., & Kahana, M. J. (2008). Temporal associative processes revealed by intrusions in paired-associate recall. *Psychonomic Bulletin & Review*, *15*(1), 64-69.

Diller, D. E., Nobel, P. A., & Shiffrin, R. M. (2001). An ARC–REM model for accuracy and response time in recognition and recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*(2), 414–435. Retrieved from `http://dx.doi.org/10.1037/0278-7393.27.2.414` doi: 10.1037/0278-7393.27.2.414

Ebbinghaus, H. (1885/1913). *Memory: A contribution to experimental psychology.* New York: Teachers College, Columbia University.

Estes, W. K. (1955). Statistical theory of spontaneous recovery and regression. *Psychological Review*, *62*, 145-154. doi: 10.1037/h0048509

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

Friendly, M., Franklin, P. E., Hoffman, D., & Rubin, D. C. (1982). The Toronto Word Pool: Norms for imagery, concreteness, orthographic variables, and grammatical usage for 1,080 words. *Behavior Research Methods and Instrumentation*, *14*, 375–399. doi: 10.3758/BF03203275

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

Glanzer, M., Hilford, A., & Kim, K. (2004). Six regularities of source recognition. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *30(6)*, 1176–1195.

Glenberg, A. M., & Bradley, M. M. (1979). Mental contiguity. *Journal of Experiment Psychology: Human Learning and Memory*, *5*(2), 88-97.

Glenberg, A. M., & Swanson, N. G. (1986). A temporal distinctiveness theory of recency and modality effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *12*(1), 3-15. doi: 10.1037/0278-7393.12.1.3

Gregg, V. (1976). Word frequency, recognition and recall. In J. Brown (Ed.), *Recall and recognition.* Oxford, England: John Wiley and Sons.

Healey, M. K., & Kahana, M. J. (2014). Is memory search governed by universal principles or idiosyncratic strategies? *Journal of Experimental Psychology: General*, *143*(2), 575–596. doi: 10.1037/a0033715

Healey, M. K., & Kahana, M. J. (2016). A four-component model of age-related memory change. *Psychological Review*, *123*(1), 23–69. doi: 10.1037/rev0000015

Healey, M. K., Long, N. M., & Kahana, M. J. (2019). Contiguity in episodic memory. *Psychonomic Bulletin & Review*, *26*(3), 699–720. doi: 10.3758/s13423-018-1537-3

Herbart, J. F. (1834). Lehrbuch zur Psychologie. In K. Kehrbach (Ed.), *Sämtliche Werke* (2nd ed., Vol. 4, p. 376-377). Königsberg: August Wilhelm Unzer.

Hockley, W. E. (1982). Retrieval processes in continuous recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *8*, 497-512.

Hockley, W. E. (1992). Item versus associative information: Further comparisons of forgetting rates. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 1321-1330.

Hockley, W. E., & Murdock, B. B. (1987). A decision model for accuracy and response latency in recognition memory. *Psychological Review*, *94*, 341-358.

Holdstock, J. S., Mayes, A. R., Roberts, N., Cezayirli, E., Isaac, C. L., O'Reilly, R. C., & Norman, K. A. (2002). Under what conditions is recognition spared relative to recall after selective hippocampal damage in humans. *Hippocampus*, *12*(3), 341–351. doi: 10.1002/hipo.10011

Hollingsworth, H. L. (1928). *Psychology: Its facts and principles.* D. Appleton and Company.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, *79*, 2554–2558.

Howard, M. W., & Kahana, M. J. (1999). Contextual variability and serial position effects in free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*(4), 923–941. doi: 10.1037/0278-7393.25.4.923

Howard, M. W., & Kahana, M. J. (2002). A distributed representation of temporal context. *Journal of Mathematical Psychology*, *46*(3), 269–299. doi: 10.1006/jmps.2001.1388

Humphreys, M. S. (1978). Item and relational information: A case for context independent retrieval. *Journal of Verbal Learning and Verbal Behavior*, *17*, 175-187.

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.

Kahana, M. J. (1993). Interactions between item, associative, and serial order information. *Unpublished doctoral dissertation.*

Kahana, M. J. (1996). Associative retrieval processes in free recall. *Memory & Cognition*, *24*(1), 103–109. doi: 10.3758/BF03197276

Kahana, M. J. (2000). Contingency analyses of memory. In E. Tulving & F. I. M. Craik (Eds.), *Oxford handbook of human memory* (pp. 323–384). England: Oxford Press.

Kahana, M. J. (2002). Associative symmetry and memory theory. *Memory & Cognition*, *30*, 823–840.

Kahana, M. J. (2012). *Foundations of human memory.* New York, NY: Oxford University Press.

Kahana, M. J. (2020). Computational models of memory search. *Annual Review of Psychology*, *71*(1), 107–138. doi: 10.1146/annurev-psych-010418-103358

Kahana, M. J., & Caplan, J. B. (2002). Associative asymmetry in probed recall of serial lists. *Memory & Cognition*, *30*(6), 841-849.

Kahana, M. J., Diamond, N. B., & Aka, A. (in press). Laws of human memory. In M. J. Kahana & A. D. Wagner (Eds.), *Oxford handbook of human memory* (Vol. 1). Oxford University Press.

Kahana, M. J., Lohnas, L. J., Healey, M. K., Aka, A., Broitman, A. W., Crutchley, P., ... Weidemann, C. T. (2024). The penn electrophysiology of encoding and retrieval study. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. doi: 10.1037/xlm0001319

Kahana, M. J., Rizzuto, D. S., & Schneider, A. (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.

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

Kohler, W. (1940). *Dynamics in psychology.* New York: Liveright.

Logan, G. D. (2021). Serial order in perception, memory, and action. *Psychological Review*, *128*(1), 1-44.

Logan, G. D., & Cox, G. E. (2023). Serial order depends on item-dependent and item-independent contexts. *Psychological Review*. doi: 10.1037/rev0000422

Lohnas, L. J. (in press-a). Oxford handbook of human memory. In M. J. Kahana & A. D. Wagner (Eds.), (2nd ed., chap. Free Recall and Memory Search). Oxford, U. K.: Oxford University Press.

Lohnas, L. J. (in press-b). A retrieved context model of serial recall and free recall. *Computational Brain & Behavior*.

Lohnas, L. J., & Kahana, M. J. (2013). Parametric effects of word frequency in memory for mixed frequency lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*(6), 1943–1946. doi: 10.1037/a0033669

Lohnas, L. J., Polyn, S. M., & Kahana, M. J. (2015). Expanding the scope of memory search: Modeling intralist and interlist effects in free recall. *Psychological Review*, *122*(2), 337–363. doi: 10.1037/a0039036

MacLeod, C. M., & Kampe, K. (1996). Word frequency effects on recall, recognition, and word fragment completion tests. *Journal of Experimental Psychology*, *22*(1), 132-142.

Malmberg, K. J., & Annis, J. (2012). On the relationship between memory and perception: sequential dependencies in recognition memory testing. *Journal of Experimental Psychology: General*, *141*(2), 233. doi: 10.1037/a0025277

Malmberg, K. J., & Nelson, T. O. (2003). The word frequency effect for recognition memory and the elevated-attention hypothesis. *Memory & Cognition*, *31*, 35-43.

Manning, J. R. (in press). Oxford handbook of human memory. In M. J. Kahana & A. D. Wagner (Eds.), (2nd ed., chap. Context Reinstatement). Oxford, U. K.: Oxford University Press.

McGeoch, J. A. (1932). Forgetting and the law of disuse. *Psychological Review*, *39*, 352-70. doi: 10.1037/h0069819

Mensink, G.-J. M., & Raaijmakers, J. G. W. (1988). A model for interference and forgetting. *Psychological Review*, *95*, 434-55. doi: 10.1037/0033-295X.95.4.434

Metcalfe, J. (1985). Levels of processing, encoding specificity, elaboration, and CHARM. *Psychological Review*, *92*, 1-38.

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

Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arxiv preprint arXiv:1301.3781*. doi: 10.48550/arXiv.1301.3781

Monaco, J. D., Abbott, L. F., & Kahana, M. J. (2007). Lexico-semantic structure and the word-frequency effect in recognition memory. *Learning and Memory*, *14*(2), 204–13.

Murdock, B. B. (1956). Backward learning in paired associates. *Journal of Experimental Psychology*, *51*, 213-215.

Murdock, B. B. (1967). Recent developments in short-term memory. *British Journal of Psychology*, *58*(3-4), 421-433. doi: 10.1111/j.2044-8295.1967.tb01099.x

Murdock, B. B. (1974). *Human memory: Theory and data.* Potomac, MD: Lawrence Erlbaum and Associates.

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

Murdock, B. B. (1992). Item and associative information in a distributed memory model. *Journal of Mathematical Psychology*, *36*, 68-99.

Murdock, B. B. (1997). Context and mediators in a theory of distributed associative memory (TODAM2). *Psychological Review*, *104*, 839-862. doi: 10.1037/0033-295X.104.4.839

Murdock, B. B. (1998). The mirror effect and attention-likelihood theory: A reflective analysis. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24(2)*, 524-534.

Murdock, B. B., & Kahana, M. J. (1993a). Analysis of the list strength effect. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *19*(3), 689-697.

Murdock, B. B., & Kahana, M. J. (1993b). List-strength and list-length effects: Reply to Shiffrin, Ratcliff, Murnane, and Nobel. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *19*(6), 1450-1453.

Murdock, B. B., & Lamon, M. (1988). The replacement effect: Repeating some items while replacing others. *Memory & Cognition*, *16*, 91-101.

Nobel, P. A., & Shiffrin, R. M. (2001). Retrieval processes in recognition and cued recall. *Journal of Experimental Psychology: Learning Memory, and Cognition*, *27*, 384-413.

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

Osth, A. F., & Dennis, S. (2015). Sources of interference in item and associative recognition memory. *Psychological Review*, *122*(2), 260—311. doi: 10.1037/a0038692

Osth, A. F., & Dennis, S. (in press). Oxford handbook of human memory. In M. J. Kahana & A. D. Wagner (Eds.), (2nd ed., chap. Global matching models of recognition memory). Oxford, U. K.: Oxford University Press.

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. doi: 10.1016/j.cogpsych.2018.04.002

Pantelis, P. C., van Vugt, M. K., Sekuler, R., Wilson, H. R., & Kahana, M. J. (2008). Why are some people's names easier to learn than others? The effects of similarity on memory for face-name associations. *Memory & Cognition*, *36*(6), 1182–1195.

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. doi: 10.1037/a0014420

Rabinowitz, J. C., Mandler, G., & Barsalou, L. W. (1977). Recognition failure: Another case of retrieval failure. *Journal of Verbal Learning and Verbal Behavior*, *16*, 639-663.

Reynolds, G. D., & Richards, J. E. (2005). Familiarization, attention, and recognition memory in infancy; an event-related potential and cortical source localization study. *Developmental Psychology*, *41(4)*, 598-615.

Robinson, E. S. (1932). *Association theory to-day; an essay in systematic psychology.* New York: The Century Co.

Schulman, A. I. (1967). Word length and rarity in recognition memory. *Psychonomic Science*, *9*(4), 211–212. doi: 10.3758/BF03330834

Schwartz, G., Howard, M. W., Jing, B., & Kahana, M. J. (2005). Shadows of the past: Temporal retrieval effects in recognition memory. *Psychological Science*, *16*(11), 898-904. doi: 10.1111/j.1467-9280.2005.01634.x

Schwartz, R. M., & Humphreys, M. S. (1974). Recognition and recall as a function of instructional manipulations of organization. *Journal of Experimental Psychology*, *102*(3), 517–519. doi: 10.1037/h0035901

Sederberg, P. B., Howard, M. W., & Kahana, M. J. (2008). A context-based theory of recency and contiguity in free recall. *Psychological Review*, *115*(4), 893–912. doi: 10.1037/a0013396

Shepard, R. N., & Teghtsoonian, M. (1961). Retention of information under conditions approaching a steady state. *Journal of Experimental Psychology*, *62*, 302-309.

Shiffrin, R. M., & Steyvers, M. (1997). A model for recognition memory: REM—retrieving effectively from memory. *Psychonomic Bulletin and Review*, *4*, 145. doi: 10.3758/BF03209391

Siegel, L. L., & Kahana, M. J. (2014). A retrieved context account of spacing and repetition effects in free recall. *Journal of Experimental Psychology: Learning Memory and Cognition*, *40*(3), 755-764. doi: 10.1037/a0035585

Solway, A., Murdock, B. B., & Kahana, M. J. (2012). Positional and temporal clustering in serial order memory. *Memory & Cognition*, *40*(2), 177–190. doi: 10.3758/s13421-011-0142-8

Stanislaw, H., & Todorov, N. (1999). Calculation of signal detection theory measures. *Behavior research methods, instruments, & computers*, *31*(1), 137–149. doi: 10.3758/bf03207704

Tulving, E. (1968). Theoretical issues in free recall. In T. R. Dixon & D. L. Horton (Eds.), *Verbal behavior and general behavior theory* (p. 2-36). Englewood Cliffs, N. J.: Prentice-Hall.

Tulving, E., & Arbuckle, T. Y. (1963). Sources of intratrial interference in immediate recall of paired associates. *Journal of Verbal Learning and Verbal Behavior*, *1*, 321-334.

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

Underwood, B. J. (1945). The effect of successive interpolations and proactive inhibition. *Psychological Monographs*, *59*, v + 33. doi: 10.1037/h0093547

Underwood, B. J., Boruch, R. F., & Malmi, R. A. (1978). Composition of episodic memory. *Journal of Experimental Psychology: General*, *107*, 393-419.

Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, *108*(3), 550–592. doi: 10.1037/0033-295x.108.3.550

Wachter, J. A., & Kahana, M. J. (2024, 10). A retrieved-context theory of financial decisions. *The Quarterly Journal of Economics*, *139*(2), 1095-1147. doi: 10.1093/qje/qjad050

Wallace, W. (1978). Recognition failure of recallable words and recognizable words. *Journal of Experimental Psychology: Human Learning and Memory*, *4*, 441-452.

Weidemann, C. T., & Kahana, M. J. (2016). Assessing recognition memory using confidence ratings and response times. *Royal Society Open Science*, *3*, 150670. doi: 10.1098/rsos.150670

Weidemann, C. T., Kragel, J. E., Lega, B. C., Worrell, G. A., Sperling, M. R., Sharan, A. D., … Kahana, M. J. (2019). Neural activity reveals interactions between episodic and semantic memory systems during retrieval. *Journal of Experimental Psychology: General*, *148*(1), 1–12. doi: 10.1037/xge0000480

Wilson, M. D. (1988). The MRC psycholinguistic database: Machine readable dictionary, version 2. *Behavior Research Methods, Instruments & Computers*, *20*(1), 6-11.

Yang, J., Zhao, P., Zhu, Z., Mecklinger, A., Fang, Z., & Li, H. (2013). Memory asymmetry of forward and backward associations in recognition tasks. [doi:10.1037/a0028875]. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*(1), 253–269. doi: 10.1037/a0028875

Appendix A

Formal Description of CMR-IA

Structure and Initialization

Inheriting from CMR2 (Lohnas et al., 2015; Healey & Kahana, 2016), CMR-IA has two representational layers, each defined as a vector: the feature layer F and the context layer C. The i^{th} item or pair presented to the model activates its associated features in F and is denoted by the vector $\mathbf{f}_i$. This in turn updates the state of context to the vector $\mathbf{c}_i$, as described below. Each of these vectors is initialized to 0.

The layers interact through two associative matrices: M^{FC}, which stores the strengths of the feature-to-context associations, and M^{CF}, which stores the strength of the context-to-feature associations. Each association matrix is a weighted sum of a pre-experimental component (M^{FC}_{pre} and M^{CF}_{pre}) that reflects longstanding semantic relationships and an experimental component (M^{FC}_{exp} and M^{CF}_{exp}) that reflects new learning occurring during the experiment. At the start of each experimental session, the experimental associations are initialized to 0. The semantic associations in M^{FC}_{pre} and M^{CF}_{pre} are defined using Word2Vec, which is a method of estimating the representations of words in vector space, or word "embeddings" (Mikolov, Chen, Corrado, & Dean, 2013). The cosine of the angle between two words' vector representations serves as the measure of semantic similarity. Thus, each element (a, b) in M^{FC}_{pre} and M^{CF}_{pre} is determined by taking the cosine similarity value between words a and b (though here we define an item's similarity with itself as 0).

Item Presentation

As in CMR2, we assume that each item has a localist representation for simplicity (Lohnas et al., 2015; Healey & Kahana, 2016). The first item presented in each simulated session is not a list item, but rather a distraction item used to initialize the state of the context vector to have nonzero elements. Each $\mathbf{f}_i$ retrieves the context states to which those features have previously been associated as the input to C:

$$\mathbf{c}_i^{\mathbf{IN}} = \frac{M^{FC}\mathbf{f}_i}{\|M^{FC}\mathbf{f}_i\|} \tag{A1}$$

This retrieved context, $\mathbf{c}_i^{\mathrm{IN}}$, which is normalized to have a length of 1 so that it can be directly compared to other states of context, is incorporated into the context representation by adding it to the current context vector $\mathbf{c}_{i-1}$:

$$\mathbf{c}_i = \rho_i \mathbf{c}_{i-1} + \beta \mathbf{c}_i^{\mathbf{IN}} \tag{A2}$$

Note that Equation A1 and Equation A2 are equivalent to Equation 2 in the main text. In this way, context is a recency-weighted sum of past context states. During item encoding, $\beta = \beta_{enc}$ determines how much new information ($\mathbf{c}_i^{\mathrm{IN}}$) is added into context with each studied item. ρ_i weakens the previous state of context such that $\|\mathbf{c}_i\| = 1$:

$$\rho_i = \sqrt{1 + \beta^2[(\mathbf{c}_{i-1} \cdot \mathbf{c}_i^{\mathbf{IN}})^2 - 1]} - \beta(\mathbf{c}_{i-1} \cdot \mathbf{c}_i^{\mathbf{IN}}) \tag{A3}$$

Pair Presentation

In simulations of cued recall and associative recognition, pairs of items are presented. Suppose the two items in a pair are represented as $\mathbf{f}_{i_1}$ and $\mathbf{f}_{i_2}$, respectively. The model adds them together to form a unified vector:

$$\mathbf{f}_i = \mathbf{f}_{i_1} + \mathbf{f}_{i_2} \tag{A4}$$

The model then follows Equation A1 and Equation A2 to update the context.

Forming Associations Between Items and Context

As each new item or pair is presented, new experimental associations are formed between the feature representation and the current state of context. These associations are formed according to a Hebbian outer-product learning rule:

$$\begin{aligned}\Delta M^{FC}_{exp} &= \mathbf{c}_i \mathbf{f}_i^\top \\ \Delta M^{CF}_{exp} &= \mathbf{f}_i \mathbf{c}_i^\top\end{aligned} \tag{A5}$$

We choose this association definition based on the CMR model developed by Polyn et al. (2009) rather than the method presented in CMR2 (Lohnas et al., 2015; Healey & Kahana, 2016), as our definition is more plausible considering form associations between the two items in a pair. When studying a pair, $\mathbf{c}_{i-1}$ will not contain information on either of the two items, thus associating $\mathbf{f}_i$ and $\mathbf{c}_{i-1}$ will not form new associations between the two items in M^{FC}_{exp} and M^{CF}_{exp}. In contrast, our definition ensures new associations are learned between the two items so that one can retrieve the other in the following recognition or recall process.

The newly formed experimental associations are incorporated with pre-experimental semantic associations. The balance between new and existing associations is controlled by parameters γ_{FC} and γ_{CF}:

$$\begin{aligned}M^{FC} &= (1-\gamma_{FC})(I + s_{FC} M^{FC}_{pre}) + \gamma_{FC} \psi_i M^{FC}_{exp} \\ M^{CF} &= (1-\gamma_{CF})(I + s_{CF} M^{CF}_{pre}) + \gamma_{CF} \phi_i M^{CF}_{exp}\end{aligned} \tag{A6}$$

where s_{FC} and s_{CF} are scaling parameters that control the influence of pre-experimental semantic associations. I is an identity matrix the same size as M^{CF}_{pre} and M^{FC}_{pre}. Effectively, this means that the on-diagonal terms are not multiplied by the s parameter. This allows the s parameter to scale semantic relations between pairs of different items while having no effect on auto-associations.

ϕ_i simulates increased attention to beginning-of-list items, producing a primacy effect by scaling the magnitude of context-to-feature associations:

$$\phi_i = \phi_s e^{-\phi_d (i-1)} + 1 \tag{A7}$$

where ϕ_s is a model parameter that scales the overall level of primacy, and the model parameter ϕ_d scales the degree to which primacy decays for each list item presented after the first item.

ψ_i simulates increased attention to rare words by scaling the magnitude of feature-to-context associations. Instead of adding a direct representation of word frequency into the model, we use a word's average pre-experimental semantic associations with other words ($\mathbf{s}_i^{\text{mean}}$) as the indicator of word frequency, since we have found a negative linear relationship between log word frequency and $\mathbf{s}_i^{\text{mean}}$ (see Simulation 4). We assume a linear relationship between the increased attention and $\mathbf{s}_i^{\text{mean}}$:

$$\psi_i = \psi_s \mathbf{s}_i^{\text{mean}} + \psi_c \tag{A8}$$

where ψ_s is a model parameter for slope and ψ_c is a model parameter for intercept. ψ_s is above 0 so that a rare word (with a high $\mathbf{s}_i^{\text{mean}}$) corresponds to a high ψ_i. Note that this mechanism is only necessary for simulating the word frequency effect in Simulation 4. For other simulations, we set $\psi_s = 0$ and $\psi_c = 1$ so that $\psi_i = 1$ for all items, thus being consistent with previous CMR models.

The Item Recognition Process

During item recognition, the test probe, represented as $\mathbf{f}_i$, is presented and it retrieves its associated context $\mathbf{c}_i^{\mathbf{IN}}$ following Equation A1. The current state of context $\mathbf{c}_t$ is used to compare with $\mathbf{c}_i^{\mathbf{IN}}$:

$$\mathbf{c}_i^{\mathbf{IN}} \cdot \mathbf{c}_t \tag{A9}$$

For old items, the match between the retrieved context and the current context will depend on how much contextual drift has occurred between the original presentation and the recognition event. If the comparison returns a value that is above a threshold c_{recog_i}, the probe is judged as an old item, otherwise, it is rejected as a new item. Participants might make trial-by-trail adjustments to the recognition threshold by setting a high threshold for saying "yes" to a rare word, as they may want to strike a balance between increasing their hit rate and lowering their false alarm rate. We assume a linear relationship between c_{recog_i} and $\mathbf{s}_i^{\text{mean}}$:

$$c_{recog_i} = c_s \mathbf{s}_i^{\text{mean}} + c_{thresh}^{item} + \epsilon \tag{A10}$$

where c_s is a model parameter for slope and c_{recog}^{item} is a model parameter for intercept. c_s is below 0 so that a rare word (with a high $\mathbf{s}_i^{\text{mean}}$) corresponds to a high c_{recog_i}. Such a relationship between the recognition threshold and the word frequency is only necessary for simulating the word frequency effect in Simulation 4. Representing retrieval variability, ϵ is a random value drawn from a uniform distribution with an interval of $[-\sigma, \sigma]$, where σ is a model parameter. This mechanism is only necessary for simulating Experiment 3. For other simulations, we set $c_s = 0$ and $\sigma = 0$ so that $c_{recog_i} \equiv c_{recog}^{item}$ for all items, thus being consistent with previous CMR models.

After the judgment, no matter what the judgment is, the context updates following Equation A2, with a drift rate parameter for the cue, β_{cue}. If output encoding is on, the model also updates M^{FC} and M^{CF} for judged-as-old items following Equation A5.

The Associative Recognition Process

During associative recognition, a pair of test probes, represented as $\mathbf{f}_{i_1}$ and $\mathbf{f}_{i_2}$, is presented. The model first updates the current state of context with $\mathbf{f}_{i_1}$ following Equation A1 and Equation A2 ($\beta = \beta_{cue}$). Then, the updated context vector $\mathbf{c}_t'$ is used to compare with the context $\mathbf{c}_{i_2}^{\mathbf{IN}}$ retrieved by $\mathbf{f}_{i_2}$ following Equation A1:

$$\mathbf{c}_{i_2}^{\mathbf{IN}} \cdot \mathbf{c}_t' \tag{A11}$$

Because two items in a studied pair are associated with the same context that prevailed when they were originally presented, the retrieved contexts of these two items would be similar. Thus, the updated context (incorporating the retrieved context of the first item) should have a high similarity value with the retrieved context of the second item for an old (intact) pair, but a low similarity value for a new (rearranged) pair. If the comparison returns a value that is above a threshold c_{recog}:

$$c_{recog} = c_{recog}^{assoc} + \epsilon \tag{A12}$$

then the test pair is judged as an old (intact) pair, otherwise, it is rejected as a new (rearranged) pair. c_{recog}^{assoc} is a model parameter for baseline associative recognition threshold and ϵ is a random value drawn from a uniform distribution with an interval of $[-\sigma, \sigma]$ representing retrieval variability. Again, this mechanism of retrieval variability is only necessary for simulating Experiment 3. For other simulations, we set $\sigma = 0$ so that $c_{recog} \equiv c_{recog}^{assoc}$ for all pairs.

After the judgment, no matter what the judgment is, $\mathbf{c}_{i_2}^{\mathbf{IN}}$ updates the state of context following Equation A2 with β_{cue}. If output encoding is on, the model also updates M^{FC} and M^{CF} for judged-as-old pairs following Equation A4 and Equation A5.

The Cued Recall Process

During cued recall, the cue item, represented as $\mathbf{f}_i$, is presented and it updates the current state of context following Equation A1 and Equation A2 ($\beta = \beta_{cue}$). Then, the updated contextual state $\mathbf{c}'_t$ is used for retrieval via the M^{CF} associations, similar to the free recall process in CMR2 (Lohnas et al., 2015; Healey & Kahana, 2016):

$$\mathbf{f}_t^{\mathbf{IN}} = M^{CF}\mathbf{c}'_t \tag{A13}$$

The resulting $\mathbf{f}_t^{\mathbf{IN}}$ gives the degree of support, or activation, for each item in the model's vocabulary. Items with low activation values are unlikely to be recalled by the retrieval process described below, and considering them as candidates for retrieval is extremely computationally expensive. Therefore, only a certain number of items[15] with the highest activation values are assigned to a vector, $\mathbf{a}$, of retrieval candidates. The values in $\mathbf{a}$ are then used as the initial input for a set of competitive accumulators, one for each candidate, according to the leaky competitive accumulator model (Usher & McClelland, 2001):

$$\begin{aligned} \mathbf{x}_n &= (1 - \tau\kappa - \tau\lambda\mathbf{N})\mathbf{x}_{n-1} + \tau\mathbf{a} + \epsilon \\ \mathbf{x}_n &\rightarrow \max(\mathbf{x}_n, 0) \end{aligned} \tag{A14}$$

where $\mathbf{x}_n$ is a vector with one element for each retrieval candidate in $\mathbf{a}$. When the retrieval competition starts, all elements are set to zero (i.e., $\mathbf{x}_0 = 0$) and the activation for each item given in a is used as its starting line in the race to threshold. τ is a fixed time constant. κ is a parameter that determines the decay rate for item activations. λ is the lateral inhibition parameter, scaling the strength of an inhibitory matrix $\mathbf{N}$ that subtracts each item's activation from all of the others except itself. ϵ is a random vector whose elements are drawn from a normal distribution with mean 0 and SD as a model parameter η. The second line of Equation A14 means that the accumulating elements cannot take on negative values. $\mathbf{x}_n$ continues to be updated until one of the activation values exceeds its threshold or until the recall period ends.

Similar to CMR2, CMR-IA dynamically sets the retrieval threshold of each item as the recall period progresses to allow items that were recalled earlier in the period to participate in, but not dominate, current retrieval competitions (Lohnas et al., 2015; Healey & Kahana, 2016). Specifically, at the beginning of the recall period, each item, i, has a threshold of $\Theta_i = 1$. If item i is retrieved, its threshold is incremented by a value ω

[15] Typically this number is four times the number of items in a study list. The two exceptions are that in Simulation 8 this number is 16 as only 16 items occur in one simulation session and that in Simulating Experiment 2 this number is 176, which is two times the number of items in the study list, as the study list is too long.

and then gradually returns to 1 with subsequent recalls:

$$\Theta_i = 1 + \omega\alpha^j \tag{A15}$$

where j is the number of subsequent retrievals, α is a parameter between 0 and 1; the larger the value of α, the more intervening retrievals are needed before an already-recalled item is likely to be retrieved again. Using this same mechanism, CMR-IA prevents the cue item from being retrieved by raising the threshold of the cue item following Equation A15 immediately after the presentation of the cue.

The first word to accumulate enough activation to cross its threshold wins the retrieval competition. The winner's representation is reactivated on F, allowing the model to retrieve the contextual state associated with the item. Context is updated with the same mechanism used during the study period, following Equation A2 with β_{rec}. If output encoding is on, the model also encodes the cue and the retrieved item as a pair by updating M^{FC} and M^{CF} following Equation A4 and Equation A5.

As in CMR2, before the item is actually outputted by the model, however, it undergoes a post-retrieval editing phase, consistent with the observation that people often report thinking of items that they do not overtly recall during free recall experiments. Editing is accomplished by comparing the context representation retrieved by the candidate item with the currently active context representation, similar to Equation A9. If the comparison returns a value that is beneath a threshold parameter, c_{thresh}, the item is filtered out.

Shift Between Different Phases

It is possible that the temporal context shifts when switching between different phases in an experiment. Some experiments even exaggerate this shift by adding a distraction task between studying and testing. CMR-IA simulates this shift in temporal context by activating a unique "disruption" item on the feature layer and allowing this item to update context using Equation A2, with a drift rate parameter $\beta_{distract}$ when shifting from a study phase to a test phase and another drift rate parameter β_{post}^{recall} when ending a test phase.

Category	Parameter	1	2	3	4	5	6a	6b	7	8	Exp 2	Exp 3
CMR Basic	β_{enc}	0.060	0.400	0.200	0.550	0.600	0.700	0.954	0.600	0.600	0.406	0.500
	β_{rec}					0.500	0.500	0.520	0.100	0.500	0.037	
	β_{cue}		0.400	0.500	0.100	0.600	0.550	0.854	0.600	0.800	0.872	0.381
	$\beta_{distract}$		0.900		0.200	0.100	0.010	0.831	0.100	0.100	0.978	0.112
	β_{post}^{recall}	0.050	0.900	0.300		0.990	0.990	0.372	0.100		0.979	0.608
	γ_{FC}	0.130	0.050	0.650	0.140	0.800	0.280	0.402	0.550	0.600	0.202	0.142
	γ_{CF}					0.500	0.280	0.542	0.700	0.600	0.802	
Semantics	s_{FC}	0.150	0.600	0.400	0.100	0.100	0.100	0.051	0.150	0.500	0.130	0.424
	s_{CF}					0.100	0.100	0.149	0.150		0.015	
Primacy	ϕ_s					0.400	0.400	2.175	0.250	1.000	4.907	
	ϕ_d					1.450	1.300	0.731	0.500	0.600	4.763	
Accumulator	κ					0.080	0.040	0.036	0.040	0.080	0.297	
	λ					0.050	0.020	0.030	0.010	0.040	0.022	
	η					0.020	0.010	0.010	0.020	0.010	0.0002	
	ω					5.000	3.000	6.158	5.000	2.000	9.709	
	α					0.500	0.900	0.958	1.000	0.500	0.806	
Criteria	c_{thresh}					0.010	0.010	0.302	0.200	0.450	0.165	
	c_{recog}^{item}	0.684		0.762	-0.036						0.222	0.698
	c_{recog}^{assoc}			0.851							0.307	0.741
	c_s				2.200							
	σ											0.005
Attention	ψ_s				40.000							
	ψ_c				-3.000							

Table 1

Best-fit parameters of CMR-IA.

		Congruent				Incongruent		
			Test 1				Test 1	
			+	-			+	-
Data	Test 2	+	0.319	0.012	Test 2	+	0.293	0.122
		-	0.006	0.663		-	0.049	0.537
		Yule's $\mathcal{Q}$ = 0.94				Yule's $\mathcal{Q}$ = 0.96		
			Test 1				Test 1	
			+	-			+	-
CMR-IA	Test 2	+	0.307	0.020	Test 2	+	0.308	0.113
		-	0.025	0.647		-	0.026	0.552
		Yule's $\mathcal{Q}$ = 0.98				Yule's $\mathcal{Q}$ = 0.94		

Table 2

Contingency tables for congruent and incongruent successive tests in Kahana (2002) and in CMR-IA.

Condition	Data					CMR-IA				
	P(Rc)	HR	FAR	d'	$\mathcal{Q}$	P(Rc)	HR	FAR	d'	$\mathcal{Q}$
Item	0.19	0.67	0.15	1.56	0.57	0.11	0.62	0.22	1.05	0.58
	(0.01)	(0.02)	(0.02)	(0.07)	(0.05)	(0.001)	(0.004)	(0.004)	(0.015)	(0.013)
Pair	0.30	0.80	0.12	2.20	0.71	0.32	0.78	0.01	2.81	0.72
	(0.03)	(0.02)	(0.01)	(0.11)	(0.04)	(0.005)	(0.004)	(0.001)	(0.019)	(0.011)
Associative	0.42	0.72	0.22	1.46	0.81	0.41	0.77	0.22	1.43	0.81
	(0.04)	(0.03)	(0.02)	(0.13)	(0.02)	(0.006)	(0.004)	(0.005)	(0.018)	(0.011)

Table 3

Descriptive statistics for Experiment 2 and its simulation in CMR-IA (numbers in parentheses are standard errors of the mean). Cue types for the three test conditions were as follows: Item: A vs. X, Pair: $A - B$ vs. $X - Y$, Associative: $A - B$ vs. $A - B_r$. Note: $\mathcal{Q}$ was determined separately for each subject, and 0.5 was added to each of the cells in the contingency table prior to calculating the $\mathcal{Q}$ value.

Condition	**Test 1 Cue**	P("yes")		**Test 2 Cue**	P("yes")		Yule's Q	
		Data	CMR-IA		Data	CMR-IA	Data	CMR-IA
Different Item	A	0.82	0.78	B	0.68	0.77	0.26	0.27
	B	(.020)	(.008)	A	(.030)	(.007)	(0.10)	(.028)
Item/Pair	A	0.82	0.78	A-B	0.85	0.88	0.64	0.61
	B	(.016)	(.006)	A-B	(.020)	(.005)	(0.12)	(.020)
Pair/Item	A-B	0.91	0.88	A	0.85	0.83	0.59	0.64
	A-B	(.018)	(.006)	B	(.021)	(.007)	(0.10)	(.018)
Same Item	A	0.81	0.79	A	0.82	0.82	0.86	0.81
	B	(.017)	(.008)	B	(.017)	(.007)	(0.03)	(.012)
Intact Pair	A-B	0.90	0.88	A-B	0.92	0.91	0.94	0.81
		(.022)	(.009)		(.019)	(.007)	(0.02)	(.008)
Repeated Lure	X	0.07	0.10	X	0.15	0.13	0.54	0.61
	X-Y	(.014)	(.011)	X-Y	(.018)	(.012)	(0.12)	(.018)
Non-repeated Lure	C	0.07	0.09	D	0.06	0.09	—-	—-
	E-F	(.009)	(.002)	G-H	(.009)	(.002)		

Table 4

Accuracy and Yules Q values for successive recognition tests in Experiment 3 and its simulation in CMR-IA (numbers in parentheses are standard errors of the mean). Note: Q was determined separately for each subject, and 0.5 *was added to each of the cells in the contingency table prior to calculating the Q value.*

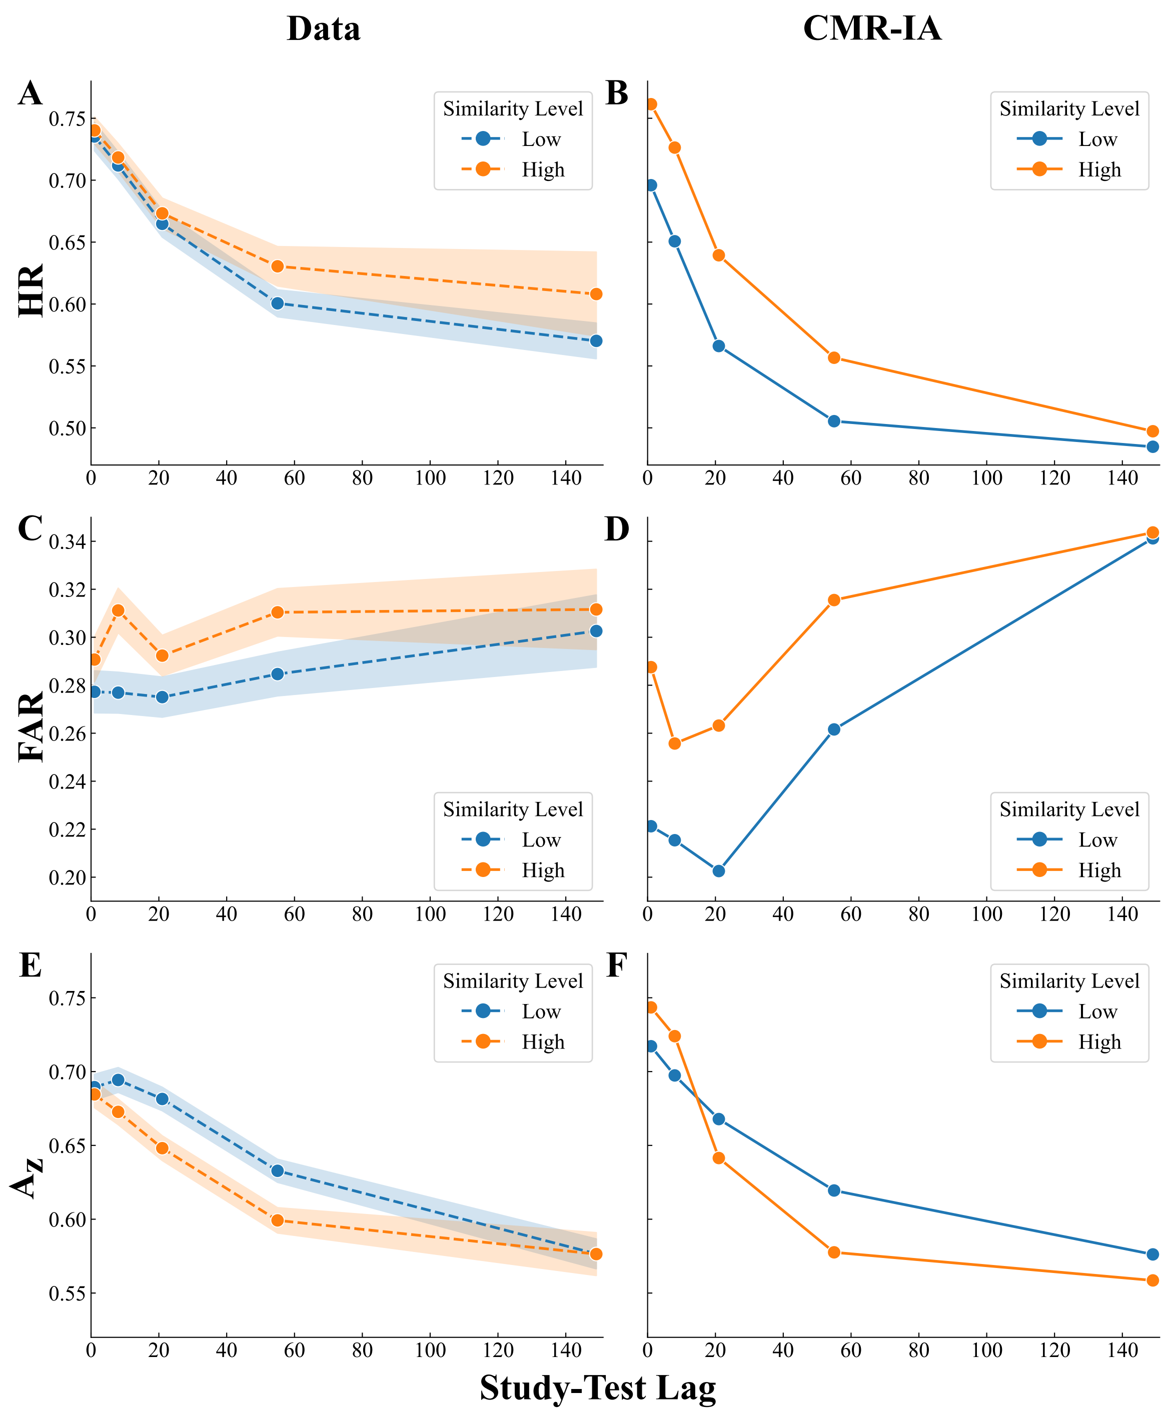

Figure 1

Joint effects of recency and similarity in continuous recognition. High similarity level indicates the presence of at least two same-category items within the previous eight items; low similarity level indicates the presence of zero or one same-category items within the previous eight items. A, C, and E. Hit rate (HR), false alarm rate (FAR), and A_z as a joint function of similarity and recency in Experiment 1 (error bar indicates standard error of mean). B, D, and F. CMR-IA predictions for Experiment 1.

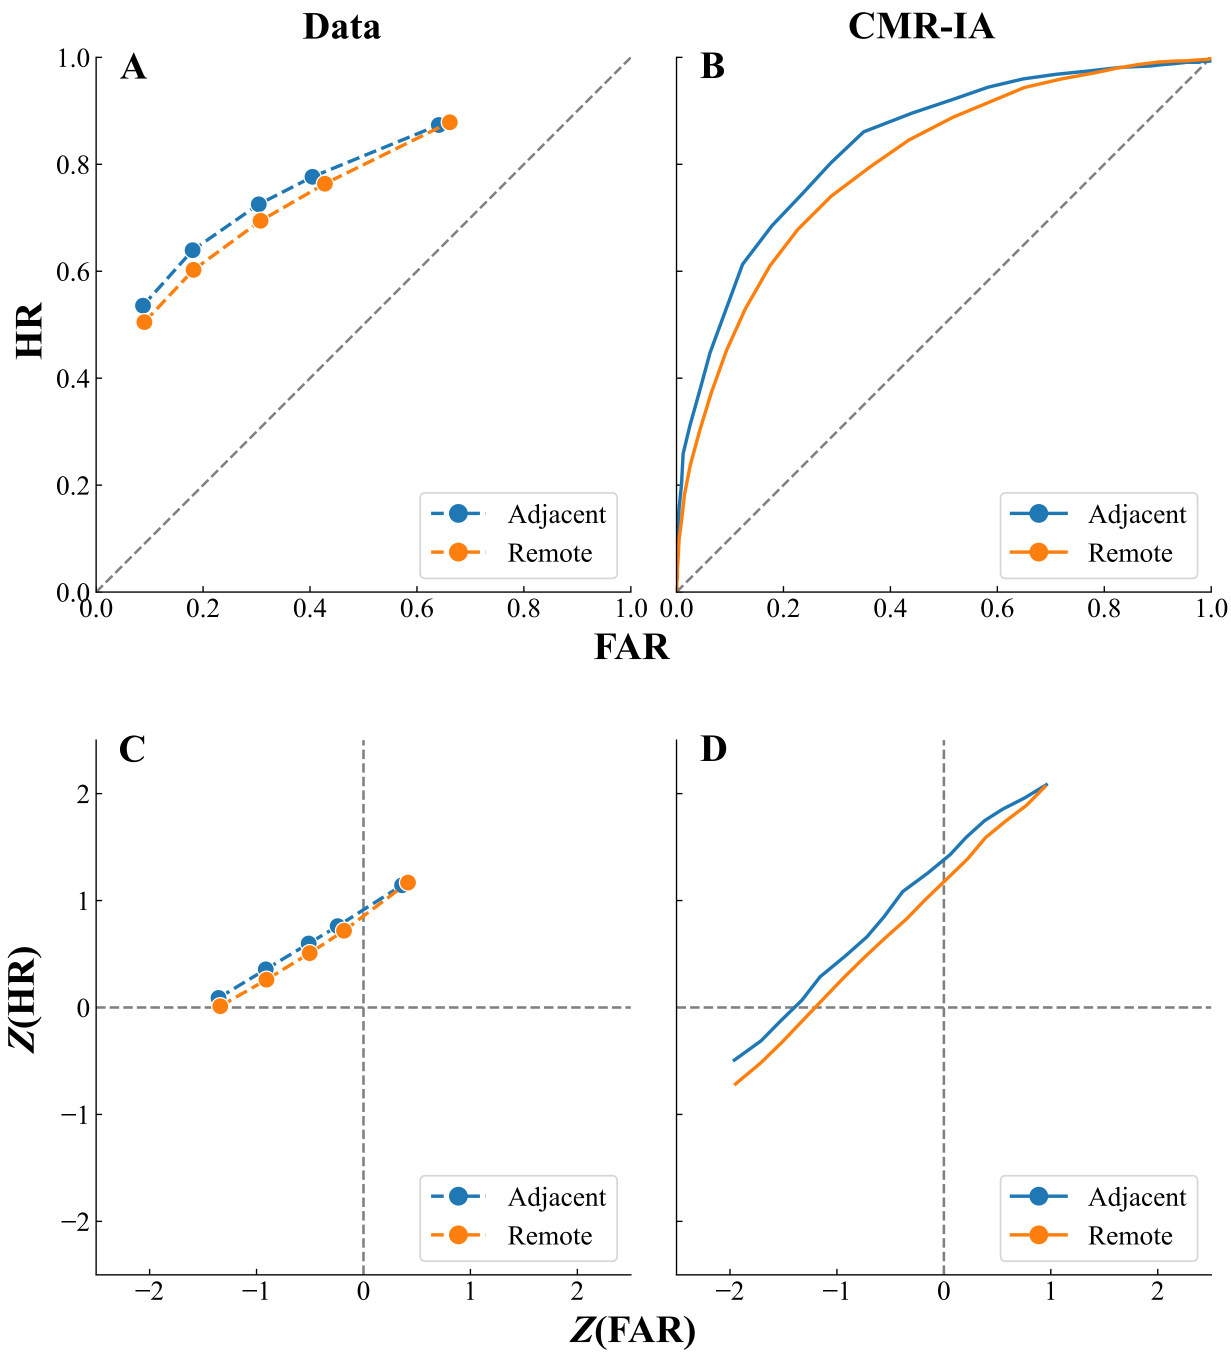

Figure 2

Successive-probe contiguity effects. A and C. The ROC curve and the z-ROC curve for adjacent and remote test probes (data from G. Schwartz et al., 2005). B and D. The ROC curve and the z-ROC curve for adjacent and remote test probes in CMR-IA.

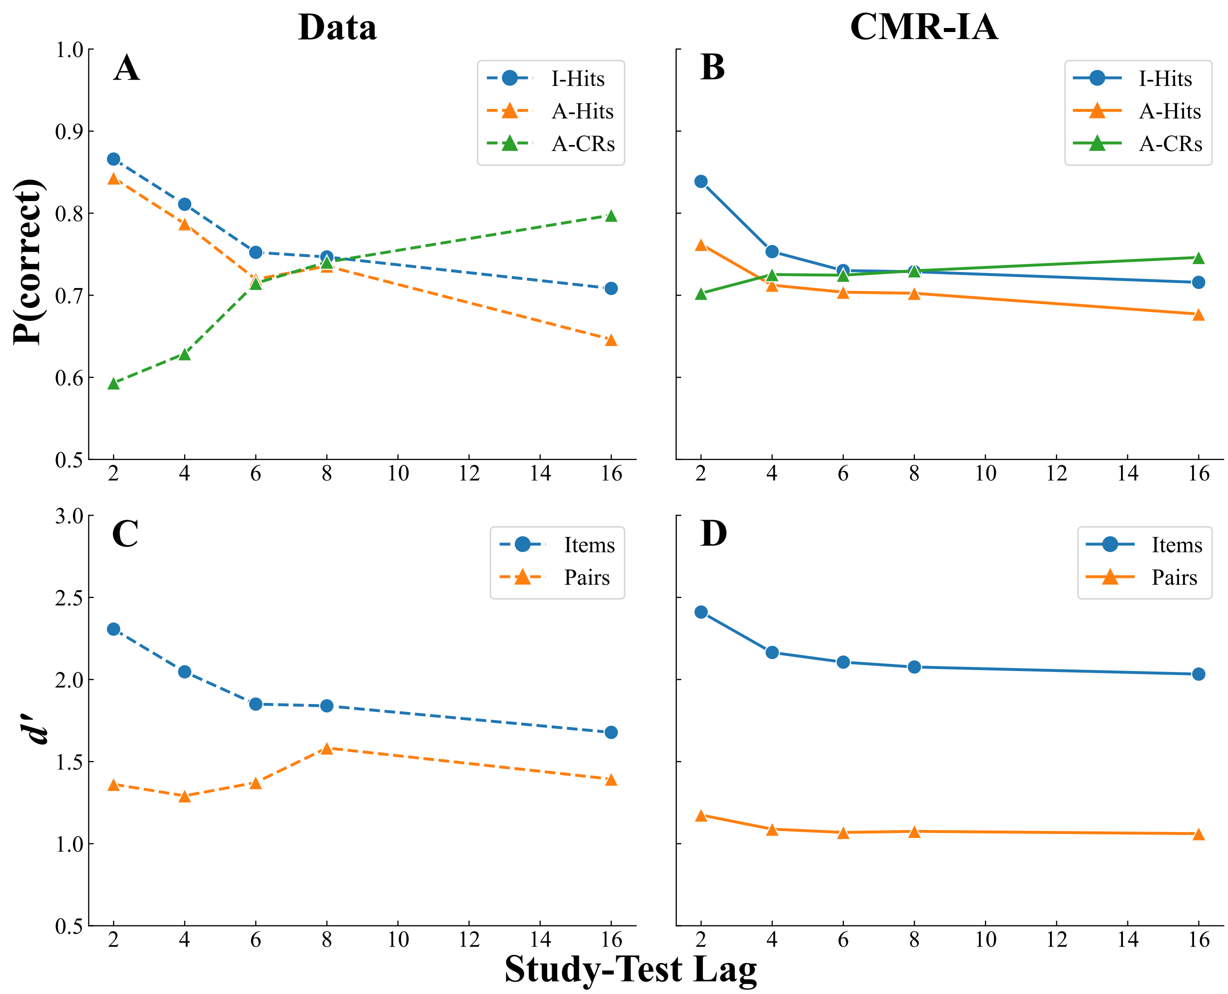

Figure 3

Differential forgetting rate for item and associative recognition. A and B. Hit (Hits) and correct rejection (CRs) rates for associative recognition (A) and hit rates for item recognition (I). C and D. d′ values for item (items) and associative (pairs) recognition as a function of test lag. A and C are data from Hockley (1992), B and D are simulation results in CMR-IA.

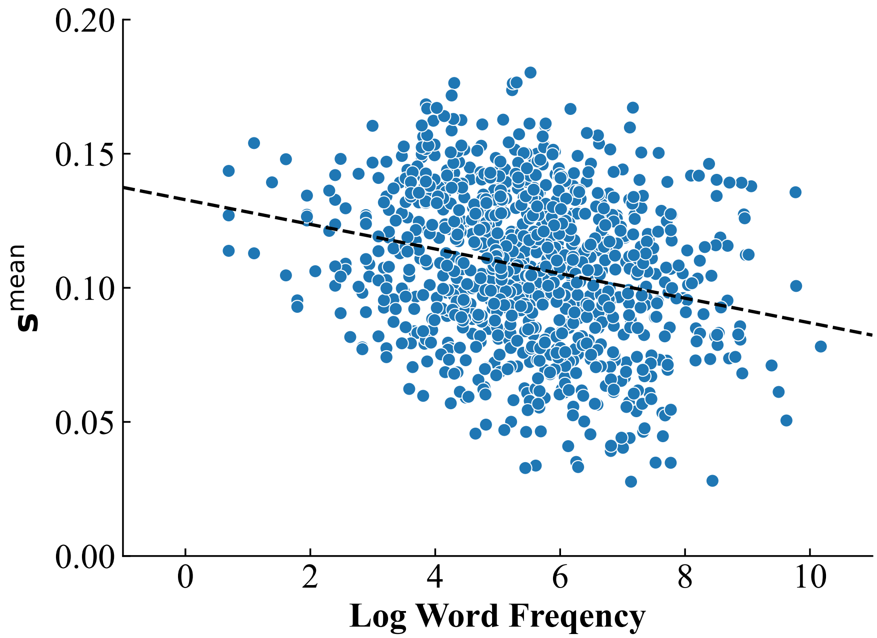

Figure 4
Relationship between log word frequency and average semantic associations for the words used by Lohnas and Kahana (2013).

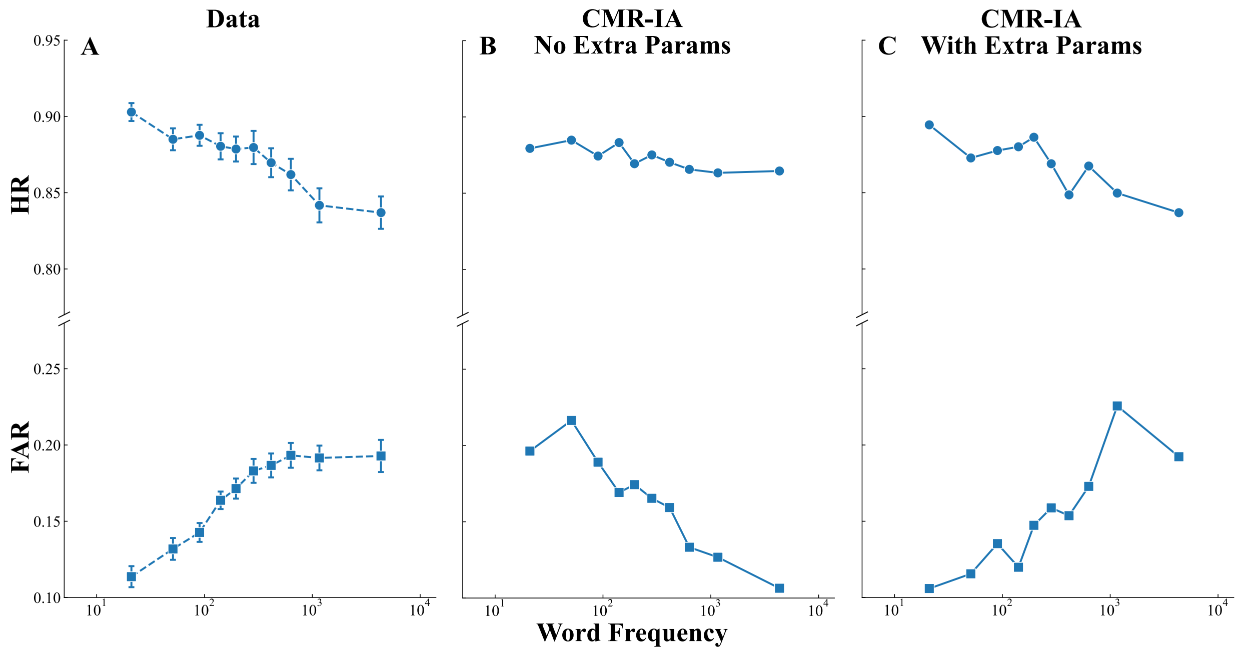

Figure 5

Word frequency effects. A. Hit rate (HR) and false alarm rate (FAR) as a function of word frequency (data from Lohnas and Kahana, 2013). B. CMR-IA fails to produce the word frequency effect without introducing the attention and recognition threshold adjustment mechanisms. C. Effect of word frequency on HR and FAR in CMR-IA after introducing the attention and recognition threshold adjustment mechanisms.

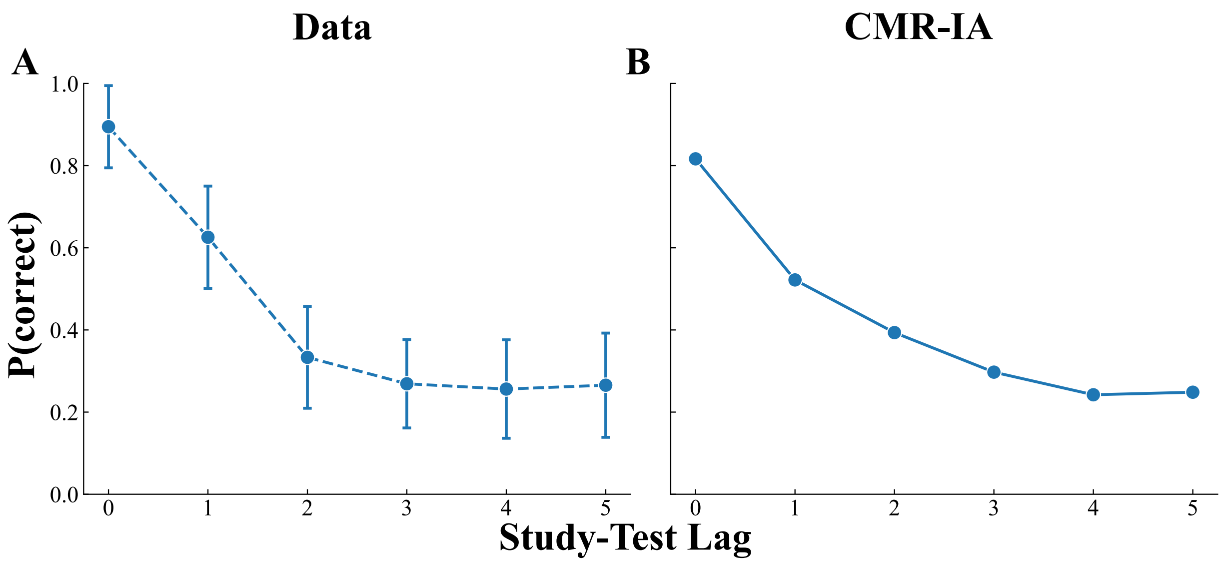

Figure 6

Retention curve in cued recall. A. Data from Murdock (1967). B. Simulation in CMR-IA.

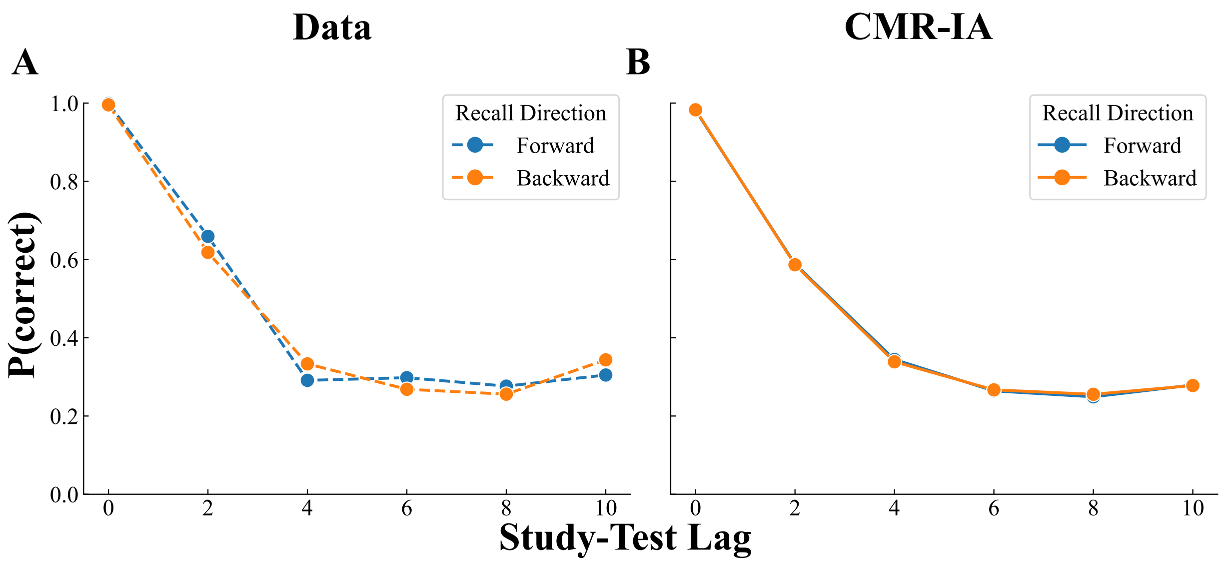

Figure 7

Forward and backward cued recall. A. Experimental data comes from Kahana's (1993) replication of an early study by Murdock (1965) involving the study and test of six word pairs. B. Simulation in CMR-IA.

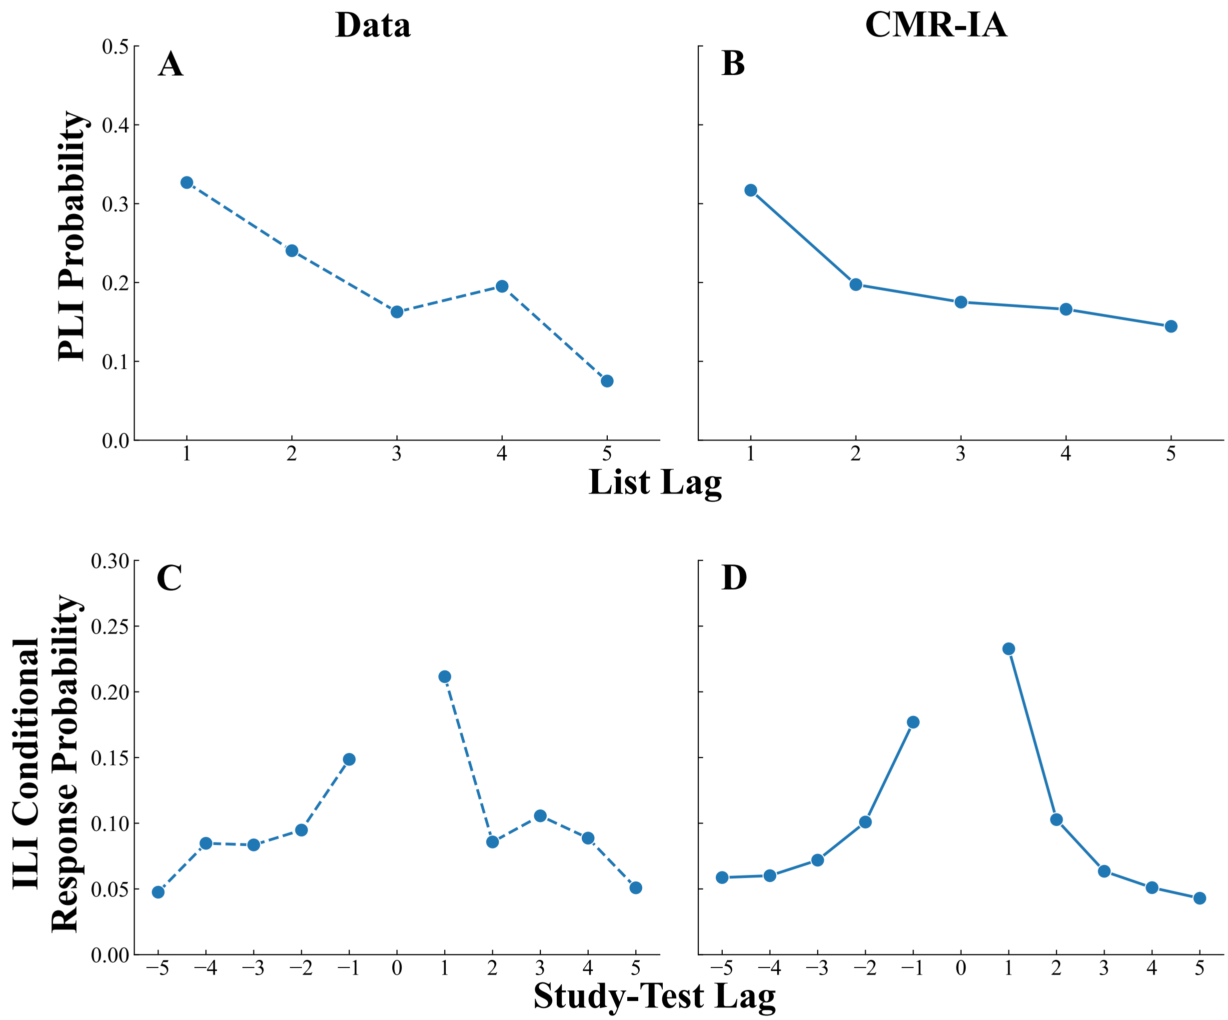

Figure 8

Prior-list intrusions (PLIs) and intra-list intrusions (ILIs) in cued recall. A. The PLI recency effect. When subjects make a PLI in cued recall, their incorrect recalls tend to come from recent lists. Data from Davis et al. (2008). B. Simulation of the PLI recency effect in CMR-IA. C. The ILI contiguity effect. ILIs occur when a subject incorrectly recalls an item paired with a different cue item on the target list, and these show a contiguity effect such that intra-list intrusions tend to come from pairs studied at nearby positions with the list. Data from Davis et al. (2008). Following study of a list of N pairs $(A_1 - B_1, A_2 - B_2, \ldots, A_N - B_N)$*, an intra-list intrusion would occur when a subject recalls* B_{i+lag} *in response to the cue* A_i *(lag indicates the number of pairs rather than lists). D. Simulation of the ILI contiguity effect in CMR-IA.*

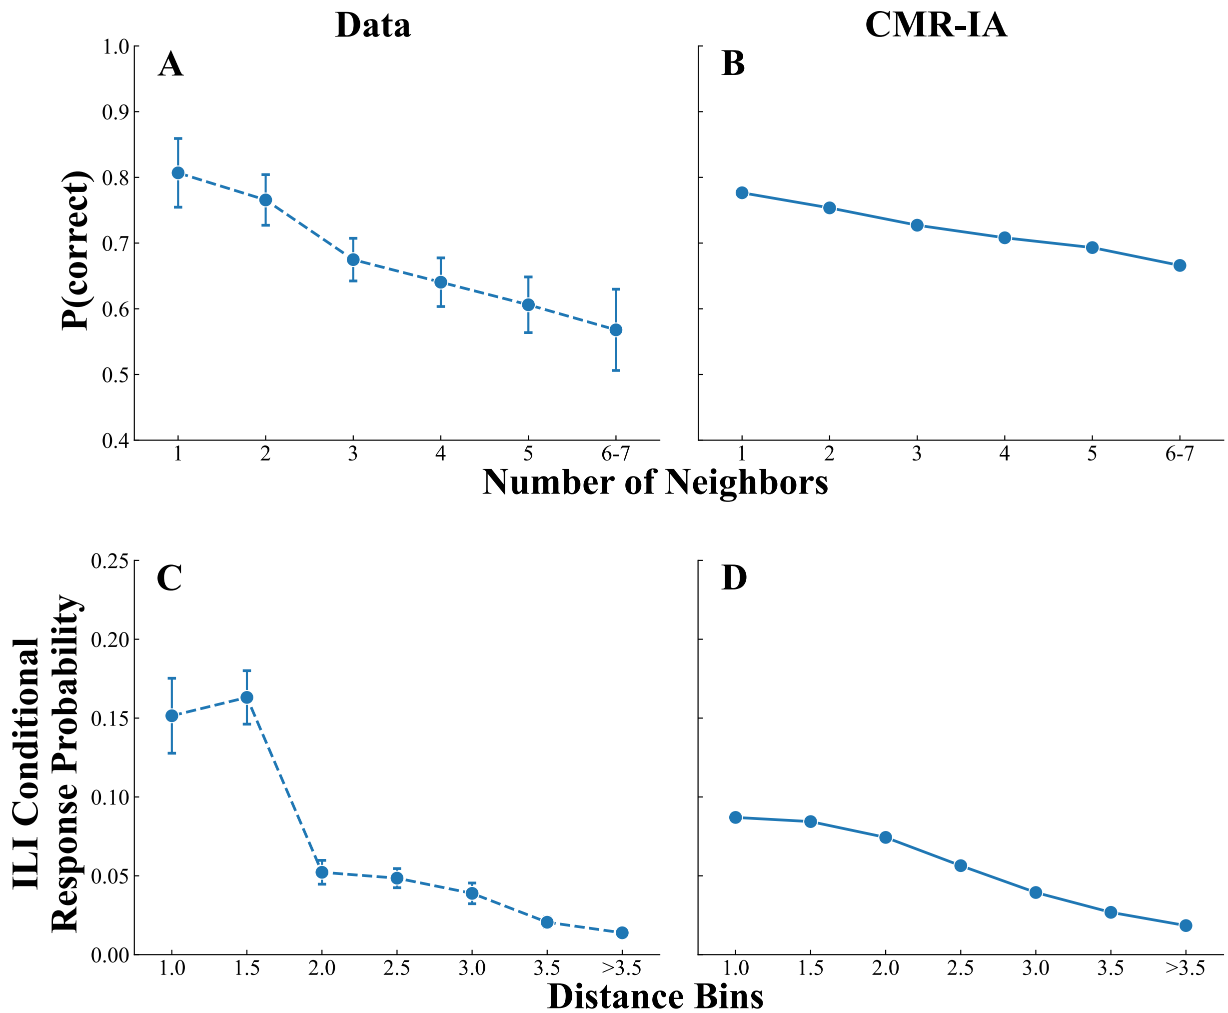

Figure 9

Similarity effects in cued recall. A. The "neighborhood" effect indicates how the probability of correct recall decreases as a function of the number of neighboring faces (Data from Pantelis et al., 2008). B. Simulation of the "neighborhood" effect in CMR-IA. C. The conditional probability of making an intra-list intrusion (ILI) decreases as a function of the Euclidean distance between the face corresponding to the recalled name and the cue face. D. Simulation of the ILI-distance effect in CMR-IA.