Optimal policies for free recall

Qiong Zhang[1,2,3], Thomas L. Griffiths[5,6], Kenneth A. Norman[4,5]

Psychology Department[1], Computer Science Department[2], Center for Cognitive Science[3], Rutgers University–New Brunswick; Princeton Neuroscience Institute[4], Psychology Department[5], Computer Science Department[6], Princeton University

Author Note

This study was not preregistered. All datasets analyzed in this work are from publicly available datasets, which can be accessed from: http://memory.psych.upenn.edu/Data_Archive. All code and analysis can be accessed from: https://github.com/qiongzhang/rationalCMR. This work was supported by NSF DRL-2024587 awarded to K.A.N. and Q.Z., AFOSR FA9550-18-1-0077 awarded to T.L.G., and a CV Starr Fellowship awarded to Q.Z.. Correspondence concerning this article should be addressed to Qiong Zhang, Rutgers University–New Brunswick, NJ, 08854. E-mail: qiong.z@rutgers.edu

Abstract

There is rich structure in the order in which studied material is recalled in a free recall task (Howard and Kahana, 2002a). Extensive effort has been directed at understanding the processes and representations that give rise to this structure; however, it remains unclear why certain types of recall organization might be favored in the first place. We provide a rational analysis of the free recall task, deriving the optimal policy for recalling items under the internal representations and processes described by the Context Maintenance and Retrieval (CMR) model of memory search (Polyn et al., 2009a). Our model, which we call rational-CMR, shows that the optimal policy for free recall is to start from the beginning of the list and then sequentially recall forwards, providing a rational account of the primacy and forward asymmetry effects typically observed in free recall. In addition, when recall is not initiated from the beginning of list, it is optimal during recall transitions to minimize the amount of forward asymmetry. Predictions from the rational model are confirmed in human behavioral data: Top-performing human participants demonstrate a stronger tendency to initiate recall from the beginning of the list and carry forward recalls, and the amount of forward asymmetry in participants depends on whether they start recall from the beginning or end of the list. We discuss the resemblance of the optimal behavior in free recall to participants' behavior when applying mnemonic techniques such as the method of loci.

Keywords: free recall; memory search; rational analysis; computational modeling

Optimal policies for free recall

Psychologists have used the free recall task for decades to gain insight into the processes and representations underlying memory search (Murdock, 1960; Murdock, 1962; Roberts, 1972; Standing, 1973). In a free recall task, participants first study a list of items, and are later asked to freely recall as many items as they can from the list. One of the key benefits of the free recall task is that the data support a wide range of analyses. In addition to tracking overall performance (operationalized as the number of recalled items) studies have also looked at serial position effects (showing primacy – enhanced recall of items from the start of the list – and also recency – enhanced recall of items from the end of the list; Murdock, 1962) and contiguity effects (regularities in which items tend to be recalled together). Early studies of contiguity focused on semantic clustering, where studied items that are drawn from the same semantic category are recalled successively (Bousfield and Sedgewick, 1944; Bousfield, 1953; Cofer et al., 1966; Howard and Kahana, 2002b). The temporal contiguity effect is also a ubiquitous property of the recall sequences: Items studied in nearby serial positions tend to be recalled successively regardless of their degree of semantic association (Kahana, 1996). These temporal contiguity effects are bidirectional (participants show an enhanced probability of recalling items that preceded and followed the item at study) and typically show a forward asymmetry (i.e., forward transitions are more likely than backward transitions; Kahana, 1996; Howard and Kahana, 1999).

Several computational models have been developed to account for these regularities of free recall, including the Temporal Context Model (TCM; Howard and Kahana, 2002a) and its successor, the Context Maintenance and Retrieval model (CMR; Polyn et al., 2009a; Lohnas et al., 2015). We focus on CMR here, but our results should also generalize to TCM (Howard and Kahana, 2002a; Sederberg et al., 2008), given the high degree of similarity between the TCM-A model (Sederberg et al., 2008) and the variant of CMR considered here (see the Background section for details). CMR posits that a slowly changing internal context representation is associated with each of the studied items, and is

then used to guide memory search at recall. In CMR, recency arises because the time-of-test context at the beginning of the recall overlaps with the contexts associated with the most recent list items. Temporal contiguity arises because recalling one item leads to retrieval of its contextual state at study, which subsequently cues retrieval of adjacent items (see the *Background* section for detailed discussion of how CMR accounts for these and other findings).

The architectural assumptions of CMR are important in accounting for these regularities. Nevertheless, while CMR can be parameterized to show all of these regularities, it can also be parameterized so that these regularities do not appear. This leaves open the question of why people show these regularities in the first place. To our knowledge, our work is the first to address the latter question, by showing how these characteristics of free recall result from the cognitive system optimizing recall performance. This is achieved by applying rational analysis, which explains human behavior as an optimal solution to computational problems posed by the environment (Anderson and Milson, 1989; Anderson, 1990). While CMR (parameterized to fit human data) provides an explanation at Marr's algorithmic level (Marr, 2010), our work provides additional insight about the free recall task at the computational level. This distinction is in spirit similar to the contribution of the original rational analysis of human memory: Though recency and frequency effects had been known to memory researchers since the work of Ebbinghaus (Ebbinghaus, 1885), Anderson (1989, 1990) was the first to demonstrate that recency and frequency effects arise from human memory optimally adapting to the statistics of the environment. Similarly, models of categorization can capture human behavior in categorization tasks by making different assumptions about whether categories are represented by prototypes or exemplars (Medin and Schaffer, 1978; Reed, 1972; Nosofsky, 1984); a rational model of categorization can justify these assumptions by explaining how they are useful in achieving the computational goal underlying the categorization task (Anderson, 1990; Griffiths et al., 2007).

For our rational analysis, we define the goal of free recall to be correctly recalling as many items as possible. Our basic architectural assumptions about internal representations and processes during memory search are based upon the CMR model (Polyn et al., 2009a; Lohnas et al., 2015). That is, we treat the core architecture of the CMR model as a hard constraint on how memory works. This approach inherits from rational analysis the idea of explaining human behavior in terms of optimal solutions to problems posed by the environment (Anderson, 1990), but follows recent extensions of this approach to the case where that environment is partly specified by the cognitive processes and internal representations of the agent (cognitive-bounded rational analysis: Howes et al., 2009; see also resource-rational analysis: Lieder and Griffiths, 2020). In contrast to the model-fitting procedure typically used in CMR, which obtains a set of parameters to best fit human behavioral data, our model (which we call rational-CMR) obtains a set of parameters to maximize the total number of items recalled. The derived optimal behavior provides a rational account of why certain patterns in recall organization arise (i.e., because they lead to better task performance). The rational analysis also leads to testable predictions about the relationship between recall organization and overall performance (i.e., people whose behavior most closely approaches the optimal parameterization identified by rational-CMR are predicted to perform better).

In the remainder of the paper, we first give some background on the CMR model of human memory search, then we describe the steps to carry out a rational analysis on CMR. We show that – under the assumption that people can choose their entry point into the list – the optimal policy is to always start recall from the beginning of the list and then sequentially recall forwards, providing a rational account of primacy and forward asymmetry effects typically observed in free recall. To explain why people do not always start recall from the beginning of the list, we highlight the fact that – under CMR – people do not have the ability to simply choose their entry point into the list. Rather, the first item needs to be recalled based on cues available at the end of the list, which tend to

match end-of-list items more than start-of-list items; this results in participants sometimes initiating recall with end-of-list items (recency) instead of start-of-list items (primacy). We show that the optimal policy during recall is different when participants happen to initiate recall from the end of the list. In this case, the optimal policy after retrieving an item is to retrieve the state of temporal context associated with that item at study, and then to cue with this retrieved context; as described in the next section, retrieved temporal context is a symmetric cue that allows participants to move backwards as well as forward in time – this backward movement is especially useful when participants start from the end of the list. We also discuss how the optimal parameters to use at encoding differ depending on whether – at recall – participants end up initiating recall from the start of the list (primacy) or the end of the list (recency); given that it is not knowable at encoding whether participants will succeed in getting back to the start of the list at recall, the optimal strategy at encoding is for participants to interpolate between parameters that yield optimal performance given primacy vs. recency.

After describing the model's predictions, we compare the theoretical predictions from the optimal policy with human behavioral data, finding that top-performing human participants are using this policy more often than the rest of the participants. Finally, we discuss the implications of the optimal policy beyond free recall task and its connection to mnemonic techniques; among other things, we argue that common mnemonic techniques like the method of loci help participants to get back to the start of the list, thereby giving them access to the benefits (noted above) of initiating recall at the start of the list.

Background: Context Maintenance and Retrieval Model (CMR)

It has been long recognized that associations are made not only among items but also between items and context (Estes, 1955; Bower, 1967; Anderson and Bower, 1972). Building on the classic stimulus-sampling theory developed by Estes (1955), Bower's (1967) model posits a context that slowly fluctuates in a context space, and binds with any

to-be-remembered external or internal experiences (see also random context models; Murdock, 1997). As a consequence, experiences that are encoded close together in time also share similar context vectors. This accounts for the recency effect in free recall, where items at the end of the list are better recalled, since they are close in time and share similar context with when the recall starts.

These earlier formulations of context models provide a sound theoretical footing for more recent computational models such as TCM and CMR (Howard and Kahana, 2002a and Polyn et al., 2009a; Lohnas et al., 2015). In addition to having a slowly drifting context, these more recent models also allow for the retrieval of previously encountered contextual states, which can in turn serves as cues for subsequent recalls (Howard and Kahana, 2002a). The state of context at time t, denoted as c_t, follows the process

$$c_t = \rho c_{t-1} + \beta c^{IN} \tag{1}$$

where c^{IN} represents the retrieved context induced by an encountered experience, β is a parameter that determines the rate of contextual drifting ($0 \leq \beta \leq 1$), and ρ is set to a scalar that ensures $||c_t|| = 1$. When an item is presented during the encoding phase, the retrieved context c^{IN} corresponds to the pre-experimental context of the item – a representation of the item's features based on all of the encounters with the item prior to the start of the experiment. This can be expressed as:

$$c^{IN} = M_{pre}^{FC} f_t \tag{2}$$

where M_{pre}^{FC} encodes the contexts that were associated with different items prior to the experiment, and f_t is a binary vector that is all zeros except at the position that represents the presented item at time t; therefore, $M_{pre}^{FC} f_t$ is the context previously associated with the presented item. In addition to a fixed M_{pre}^{FC} that represents item-context associations prior to the experiment, there are also M_{exp}^{FC} and M_{exp}^{CF} that capture the item-to-context associations and context-to-item associations acquired during the experiment. These matrices are initially set to zero, and are incremented during the encoding process based on

the Hebbian outer-product learning rule $\Delta M^{FC}_{exp} = \Delta M^{CF}_{exp} = f_t c^T_{t-1}$ to instantiate the associative learning that takes place between a presented item f_t and the current context c_{t-1}. The overall effect of having context drift towards the retrieved context of presented items, as described in Equations (1) and (2), together with associative learning, is that each item is embedded at a location in the context space corresponding to the representations of recently-encountered items. With each new presented item, the context drifts slightly towards the pre-experimental context associated with that item, so that items presented close in time to each other are linked to similar contextual states. We followed the order-of-operations consistent with Howard and Kahana (2002a), Sederberg, Howard, and Kahana (2008), Lohnas et al. (2015), and Rouhani et al. (2020), which is to associate an item with the context first, before the context vector drifts towards the item's pre-experimental context.

During recall, when memories of these items are needed again, the memory search process is driven by the current state of the context representation c_t. The support for recalling each item depends on how much the current context matches the items' study context. The starting context at recall is close in time to (and thus similar to) the end-of-list context during encoding; this gives rise to better recalls for items studied at the end of the list (i.e., the recency effect). As recall continues, context evolves under the same process as it did during the encoding phase, $c_t = \rho c_{t-1} + \beta c^{IN}$ as in Equation (1), but with the retrieved context c^{IN} introduced differently. The key difference is that items are encountered for the first time during the experiment in the encoding stage, whereas – in the recall phase – items are encountered for the second time when they are recalled. The first time that an item is encountered during its presentation at $t = i$, there is no association between this item and the experimental context yet; therefore, the retrieved context c^{IN} consists of solely the pre-experimental context associated with this item, expressed in $c^{IN}_{enc} = M^{FC}_{pre} f_i$; however, when the same item is encountered a second time when it is recalled at $t = j$, the retrieved context c^{IN}_{rec} can come from both the

pre-experimental context associated with the item $M_{pre}^{FC} f_j$ and the experimental context associated with the item $M_{exp}^{FC} f_j$ which has been acquired through Hebbian learning during the encoding phase. The extent of retrieving the pre-experimental context versus retrieving the experimental context is regulated by the parameter $\gamma_{fc} \in [0, 1]$,

$$c_{rec}^{IN} = (1 - \gamma_{fc}) M_{pre}^{FC} f_j + \gamma_{fc} M_{exp}^{FC} f_j = (1 - \gamma_{fc}) c_{enc}^{IN} + \gamma_{fc} c_{i-1}. \tag{3}$$

The value of γ_{fc} has important implications for the recall transition patterns. When $\gamma_{fc} = 0$, the retrieved context entirely consists of pre-experimental context associated with the item $M_{pre}^{FC} f_j$, which is identical to the retrieved context c_{enc}^{IN} when the item was first encountered during encoding. When $\gamma_{fc} = 1$, the retrieved context entirely consists of experimental context associated with the item $M_{exp}^{FC} f_j$, which is essentially the context c_{i-1} that was associated with the item at $t = i$ during encoding. Both c_{enc}^{IN} and c_{i-1} are part of the original study context, as illustrated in Figure 1, so reinstating them into the current context has the effect of mentally "jumping back in time". As a consequence of this "jumping back in time", items that were studied close in time to the just-recalled item have a higher chance of being recalled next, because of the similarity between their study context and the current context. This contributes to the temporal contiguity effect commonly observed in free recall.

Though both contribute to temporal contiguity, reinstating c_{enc}^{IN} and reinstating c_{i-1} bias forward recalls and backward recalls differently. This is illustrated in Figure 1A: During encoding, at $t = i$, the drifting part of the study context ρc_{i-1} (in orange) is similar to contexts both before and after $t = i$, and has the chance to be associated with items that come before and after $t = i$. Therefore, reinstating c_{i-1} gives rise to both forward recalls and backward recalls. However, since an item's pre-experimental context is incorporated into temporal context after the item is presented, the retrieved context c_{enc}^{IN} (in blue) does not share any similarity with contexts before $t = i$, and only has an opportunity to be associated with items that come after. Therefore, reinstating c_{enc}^{IN} gives rise to only forward recalls, which accounts for the forward asymmetry commonly seen in

free recall experiments. Figure 1B illustrates the effect of γ_{fc} on the recall transition patterns. When $\gamma_{fc} = 0$ only c_{enc}^{IN} is reinstated, which leads to asymmetry in forward and backward recalls; this asymmetry can be seen in the conditional response probability curve, which is computed by dividing the number of times a transition of that lag is actually made by the number of times it could have been made (Kahana, 1996). When $\gamma_{fc} = 1$ only c_{i-1} is reinstated, which leads to symmetric forward and backward recalls. Reinstating c_{enc}^{IN} and c_{i-1} both contribute to temporal contiguity; for the purpose of distinguishing the two in the rest of the paper, we refer to the asymmetric part of the temporal contiguity as "asymmetric contiguity", which is also the forward asymmetry effect, and the symmetric part as "symmetric contiguity". In practice, experimental participants adopt an intermediate value of γ_{fc}, demonstrating a combination of both effects.

At this point, we have described the key properties of CMR (Howard and Kahana, 2002a; Polyn et al., 2009a, Lohnas et al., 2015), which closely resembles the TCM-A model (Sederberg et al., 2008) in terms of the features relevant for this rational analysis. To be able to fully simulate the behavioral patterns, the model also needs to be equipped with a retrieval rule and a stopping rule. Here, we use the softmax function as the retrieval rule $p_i = \frac{e^{ks_i}}{\sum_j e^{ks_j}}$, with s_i as the support to retrieve item i and parameter k governing the amount of noise during the retrieval (Howard and Kahana, 2002a, Sederberg et al., 2008). The stopping probability is a function of summed support for the already recalled items s_r and the not-yet-recalled items s_{nr}, expressed as $p_{stop} = e^{-\epsilon_d s_{nr}/s_r}$ (Kragel et al., 2015). Importantly, the key model properties reviewed above are capable of explaining recency and temporal contiguity, but they do not account for the primacy effect, where early items in the list are better recalled than items in the middle of the list. This has been captured by assuming that there is increased attention during recall to beginning-of-list items, with the support to each item i scaled by $\phi_i = \phi_s e^{-\phi_d(i-1)} + 1$ (Polyn et al., 2009a, Lohnas et al., 2015). This introduces an additional two parameters ϕ_s and ϕ_d into the model.

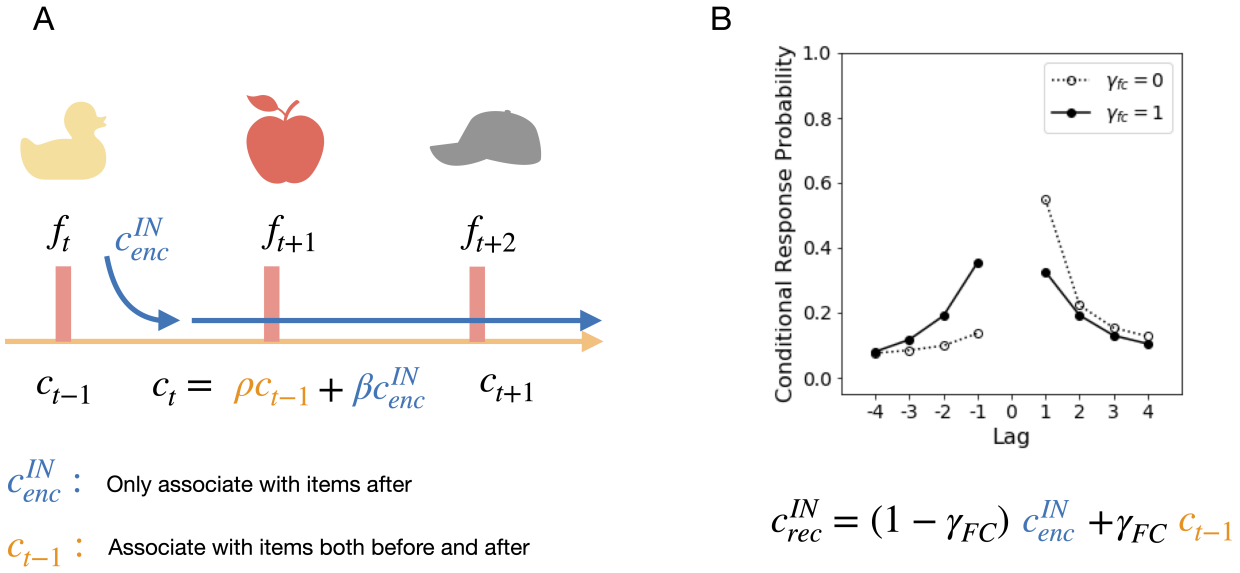

Figure 1. Illustration of how the value of γ_{fc} regulates the extent of forward recalls versus backward recalls. (A) Context can be updated in two ways at encoding: using the pre-experimental context c^{IN}_{enc} introduced by the just-encoded item f_t, and using the context drifting from the previous time step c_{t-1}. From the figure, it is clear that only the items that come after c^{IN}_{enc} have a chance to be associated with it. Therefore when c^{IN}_{enc} is reinstated at recall, it only gives rise to forward recalls. In contrast, the drifting part of the context is associated with both items before time t and items after time t, therefore giving rise to both forward and backward recalls. (B) The value of γ_{fc} regulates the amount of retrieved context from c^{IN}_{enc} versus c_{t-1} when item f_t is encountered again during recall. When $\gamma_{fc} = 0$, only c^{IN}_{enc} is reinstated, which leads to forward asymmetry (i.e., asymmetric contiguity) as seen from the conditional response probability. When $\gamma_{fc} = 1$, only c_{i-1} is reinstated, which leads to symmetric forward and backward recalls (i.e., symmetric contiguity). Parameters used in the simulation were obtained from the CMR fit in Figure 2.

A rational analysis of the free recall task

In this section, we derive an optimal solution to the free recall task via rational analysis. We adopt the steps laid out in Anderson (1990):

1. Precisely specify the goals of the cognitive system.
2. Develop a formal model of the environment to which the system is adapted.
3. Make minimal assumptions about computational limitations.
4. Derive the optimal behavior function, given items 1 through 3
5. Examine the empirical evidence to see whether the predictions of the behavior function are confirmed

Anderson (1990) argues that human memory is adapted to the statistical patterns in the physical environment; in our analysis, we commit to the architectural assumptions in CMR, which has been shown to capture human behavior in memory search. We specify the corresponding components for each step with respect to the free recall task and the CMR model, and organize the rest of the manuscript based on these steps.

First, we specify the goals of the cognitive system. In the task of free recall, the goal is to recall as many items as possible from the studied list, within the time constraints imposed by the recall task (e.g., participants might be told that they have to complete recall within 2 minutes), while at the same time minimizing recall errors such as those from extra-list intrusions. There is more than one way to combine these goals. However, for the purpose of the current analysis, we set the goal as maximizing the number of items correctly recalled from the list. This definition of the goal is consistent with the instructions participants commonly receive in free recall experiments, where they are told to recall as many items as possible from the list without penalizing them for recall errors, or rewarding them for their speed.

Next, we build a formal model of the environment based on CMR, making minimal additional assumptions about computational limitations under this environment. In contrast to previous studies of rational analysis, the environment here incorporates not only the local task environment of free recall, but also incorporates the architectural assumptions from CMR. Our approach thus closely resembles cognitive-bounded rational analysis (Howes et al., 2009) and resource-rational analysis (Lieder and Griffiths, 2020), which identifies optimal cognitive processes given a utility function that is a result of both constraints imposed by their cognitive architecture and the local task environment. CMR describes how the cognitive system navigates a context space to encode and later recollect a list of items. The space of possible behaviors during memory search is fully specified in the architectural assumptions of CMR and its parameters. For example, one of the main assumptions is how the cognitive system updates the current context in a drifting process based on a just-encoded or recalled item, which has been shown to capture human behavioral patterns in empirical data (Howard and Kahana, 2002a; Polyn et al., 2009a; Lohnas et al., 2015). Specifying an environment without these assumptions, such as one that allows traversing freely to any position in the context space regardless of previous context, can easily achieve the goals in memory search but does not provide a realistic model of how human memory functions.

For the current rational analysis, we divide the set of CMR parameters into two distinct categories. Some parameters, including k in the retrieval rule and ϵ_d in the stopping rule, correspond to the core architectural constraints of the memory system. We assume that these constraints that describe the noise level in the memory retrieval process (k) and the extent that one can still continue retrieval given memory strengths of different items (ϵ_d) are not themselves subject to optimization. The rest of the parameters in CMR specify the space of possible encoding and recall behavior during memory search. The goal of the current work is to examine how these parameters can be optimized by the cognitive system, while keeping the core architectural constraints of the memory system fixed.

Parameters that fall into the second category include β_{enc} and β_{rec}, which describe the amount of drifting during encoding and recall, γ_{fc}, which describes the proportion of pre-experimental and experimental associations used to update context, and c_{start} (a newly introduced parameter to replace ϕ_s and ϕ_d; see text below), which describes the amount of primacy. Our rational analysis consists of fixing the first set of parameters to some reasonable values to capture the architectural constraints during the retrieval process, and then optimizing the parameter values in the second set that describe the space of encoding and recall policies to achieve maximum free recall performance. Essentially, our analysis will provide a rational account of the following three issues:

1. What amounts of drifting during encoding (β_{rec}) and during recall (β_{rec}) lead to optimal performance?

2. What is the ideal balance between retrieving pre-experimental context versus experimental context through γ_{fc}?

3. How does regulating the amount of primacy affect overall free recall performance?

By searching the space of these encoding and recall policies through these parameters, the cognitive system has the opportunity to find the optimal solution in the current task. Since these components are optimized jointly, not separately, there will be an opportunity to identify the subspace of encoding and recall policies where different components interact to maximize overall recall.

As noted above, the standard way of addressing primacy in CMR is to provide a boost to early-list items at recall via the parameters ϕ_s and ϕ_d; while this parameterization provides a way of descriptively fitting the primacy effect, it does not correspond to a mechanistic claim about how recall works, and thus we do not consider it one of the "constraints on the recall process" provided by CMR. To provide a mechanism for primacy, Kragel, Morton, and Polyn (2015) argue (following Laming, 1999) that primacy arises because participants partially reinstate the start-of-list context when they initiate recall.

We instantiate this in our model by replacing ϕ_s and ϕ_d with a different parameter c_{start} that specifies which study context (by serial position) to reactivate at the beginning of recall. When $c_{start} = 0$, the starting context at recall is set to the beginning-of-list context that was active before the presentation of the first item; When $c_{start} = 10$ (given that the list length is 10), the starting context at recall is set to the end-of-list context that was active after the last item was presented. This parameterization assumes that participants have the ability to reinstate the context associated with any serial position at the outset of recall; this is of course a major simplification – just because participants want to recall starting at a particular point does not mean they will succeed. We will address this issue (i.e., what if participants do not succeed at reinstating the context associated with their desired starting point) later in the paper.

Before we can proceed to obtain the optimal encoding and recall policy in the free recall task, we need to fix the parameters that characterize the cognitive constraints during the retrieval process. To obtain a set of reasonable values for these parameters, we fit CMR to an immediate free recall dataset (Kahana et al., 2002). We use Bayesian optimization to search the space of CMR parameters, in order to minimize the normalized root-mean-square error between the CMR simulations and the data. Results from the obtained CMR model can be found in Figure 2. It captures the probability of recall by serial position as shown in the serial position curve, the probability of the first recall by serial position, and the conditional response probability, computed by dividing the number of times a transition of that lag is actually made by the number of times it could have been made. More details on the CMR fitting procedure and the dataset used can be found in the Appendix.

Deriving the optimal behavior function

Next, having fixed the parameters that specify hard constraints on the retrieval process, ϵ_d and k, based on the model fits shown in Figure 2, we ran a second round of

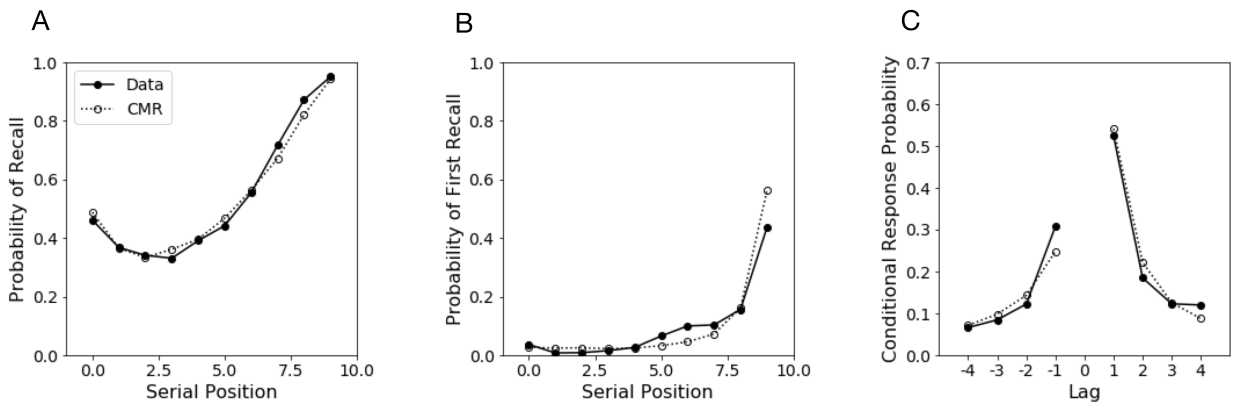

Figure 2. Behavioral patterns in the dataset in Kahana (2002) and the CMR fit to these patterns. From left to right are the serial position curve, probability of the first recall, and conditional response probability. CMR parameters obtained using Bayesian optimization: $\beta_{enc} = 0.79$, $\beta_{rec} = 0.49$, $\gamma_{fc} = 0.40$, $\phi_s = 4.66$, $\phi_d = 2.73$, $k = 5.38$ and $\epsilon_d = 2.72$.

Bayesian optimization on the remaining parameters, which determine encoding and retrieval policies. The goal of this second round of Bayesian optimization was to implement our rational analysis: That is, we set out to identify the configurations of these parameters (i.e., the encoding and retrieval policies) that maximize recall performance, under the constraints imposed by the CMR architecture and the other (fixed) parameters. We refer to the resulting model as rational-CMR. The experiment in the simulations was set up based on the stimuli and trial structures in Kahana et al. (2002). More details on obtaining the rational-CMR parameters can be found in the Appendix. This study was not preregistered. All datasets analyzed in this work are from publicly available datasets, which can be accessed from: http://memory.psych.upenn.edu/Data_Archive. All code and analysis can be accessed from: https://github.com/qiongzhang/rationalCMR.

Results from rational-CMR can be found in Figure 3. Figure 3A shows that the optimal policy is one that always starts recalling from the beginning of the list according to the probability of first recall, and then sequentially recalls in the forward direction according to the conditional response probability ($c_{start} = 0$, $\beta_{enc} \approx 1$, $\beta_{rec} \approx 1$, $\gamma_{fc} \approx 0$;

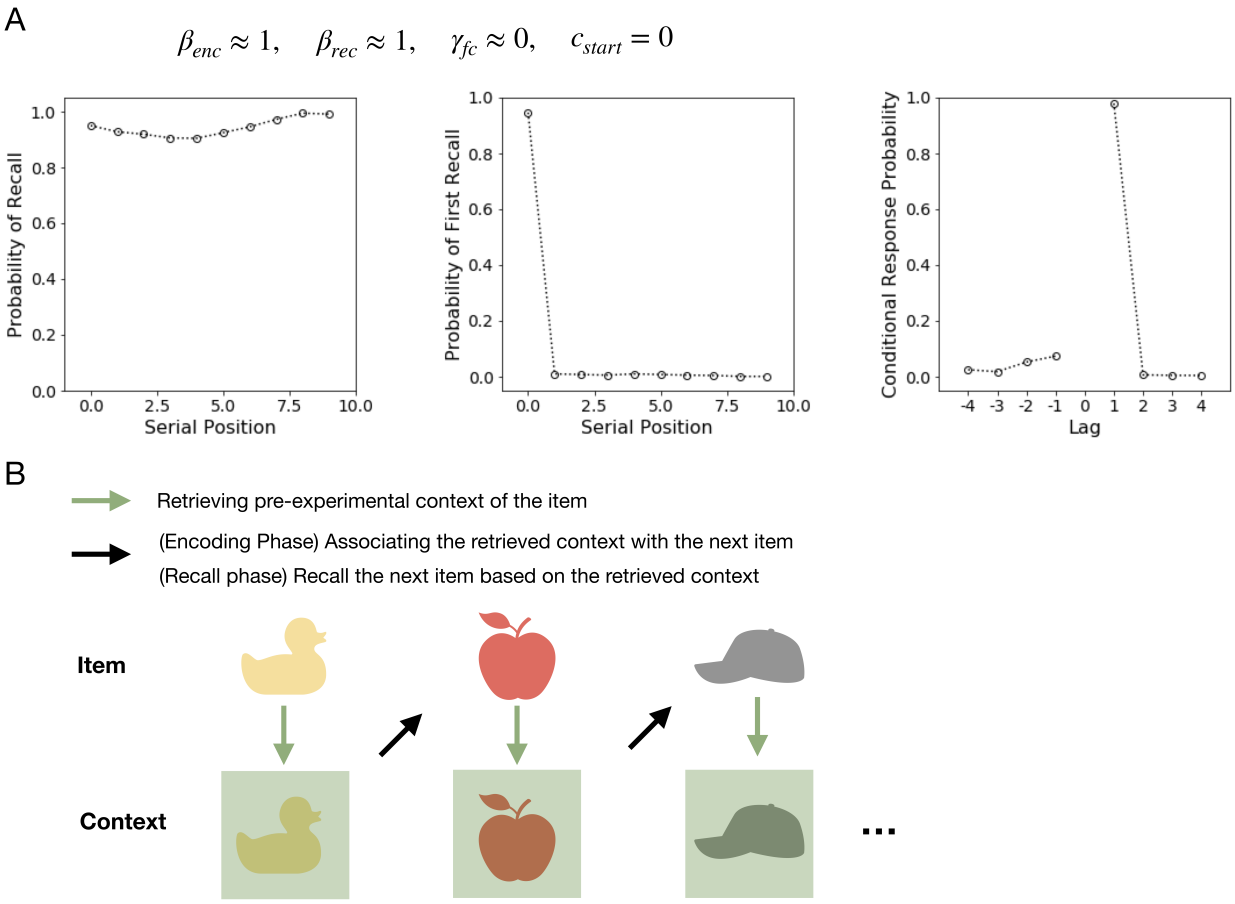

Figure 3. A rational account of primacy and forward asymmetry. (A) Proportion of items correctly recalled, optimized over encoding and recall policy. The optimal policy always starts by recalling from the beginning of the list with $c_{start} = 0$, and then sequentially recalls forwards. The optimal parameters are: $\beta^*_{enc} = 0.99$, $\beta^*_{rec} = 0.99$, $\gamma^*_{fc} = 0.001$. The fixed parameters are: $k = 5.38$, $\epsilon_d = 2.72$. (B) This optimal encoding and recall policy can be understood as a strategy that chains items through the context. With $\beta_{enc} = 1$, $\beta_{rec} = 1$, and $\gamma_{fc} = 0$, when an item is presented (at encoding) or retrieved (at test), context always drifts all the way to that item's pre-experimental context. For example, at encoding, when a duck is studied, context drifts fully to the pre-experimental context for duck, which is then associated with the next item (apple). At test, when duck is retrieved, context fully drifts (once again) to the pre-experimental context for duck, which triggers recall of apple. This way, the same chain is traversed during encoding and recall.

there are no alternative policies with equal level of recall performance, as the optimal policy identified here consistently arises over multiple runs of Bayesian optimization with differently initialized parameters). This result provides a rational account of why it is beneficial to have primacy ($c_{start} = 0$) and forward asymmetry ($\gamma_{fc} \approx 0$) during free recall.

The optimal policy corresponds to a strategy that chains items through context (Figure 3B). With $\beta_{enc} = 1$ (the optimal value identified by rational-CMR), when an item is studied (call it item n), the model drifts context all the way to the pre-experimental context of that item; once the next item (item $n + 1$) is presented, the just-retrieved context from item n is associated with item $n + 1$. This process iterates until the encoding reaches the end of the list. During the recall phase, the optimal starting context identified by rational-CMR is the context at the beginning of the encoding list ($c_{start} = 0$). Since the start-of-list context was associated with the first item on the study list, reinstating the start-of-list context makes it possible to recall the first item reliably. From this point on, the optimal recall policy is to only retrieve the pre-experimental context but not the experimental context of a just-recalled item (i.e., setting $\gamma_{fc} = 0$), and then have the current context drift all the way to the just-retrieved context (i.e., setting $\beta_{rec} = 1$). This way, retrieval follows the same chain that was traversed during encoding: Retrieving item n's pre-experimental context triggers recall of item $n + 1$ via the associative link that was forged at study, leading to retrieval of item $n + 1$'s pre-experimental context, which cues recall of item $n + 2$, and so on. We explored how sensitive the optimal policy is to the specific variant of CMR that was used. A similar optimal policy was obtained when we allowed intrusions from other lists in the simulations, in a different variant of CMR (CMR2; Lohnas et al., 2015). We also found that the optimal policy was qualitatively identical when we used an alternative primacy mechanism: the attention-based account instantiated in TCM and CMR (Polyn et al., 2009b, Lohnas et al., 2015). Full results of these analyses can be found in the Appendix. When obtaining the optimal policy, we assumed that the fixed parameters that characterize the cognitive constraints during the retrieval process are

shared across all participants. We observe that splitting all participants into two performance groups when fitting CMR, i.e. assuming that better-performing subjects and worse-performing subjects are different in their cognitive constraints, does not change our main conclusions about the optimal policy (see more details in the Appendix).

The derived optimal policy is non-trivial for several reasons. First, it is not obvious that the optimal policy is to serially recall in a task where one can recall items in any order; second, even if one figures out intuitively that chaining (in this case, indirectly through context) is a useful strategy because it minimizes the odds that an item will be inadvertently skipped, it is not obvious that one should carry out forward chaining rather than backward chaining, given that the active context at the start of recall is the end-of-list context. We will come back to the comparison between forward chaining and backward chaining in the General Discussion. Interestingly, the optimal policy resembles the stereotyped walk that characterizes the method of loci, a mnemonic technique commonly adopted for memorizing a long list of information (Yates, 1966). In the method of loci, one mentally walks through a pre-defined route, associating each item with loci along this route during encoding; when it comes to recall, one simply mentally walks through the same route and uses loci along the route as memory cues to subsequently retrieve each item. We will come back to the comparison with mnemonic techniques in the General Discussion.

There are two features of the optimal policy that might appear non-intuitive and require closer examination. First, under the strategy to drift all the way to each encoded item ($\beta_{enc} = 1$), each item is associated with the item presented right before it but not any other items. This appears to contradict the compound cuing effect commonly observed in free recall experiments (Lohnas and Kahana, 2014), where multiple prior items combine to form a compound cue for the next response. To obtain the compound cuing effect in CMR, β_{enc} must be set to a value less than one. Here we show the conditions when it is optimal to have $\beta_{enc} < 1$ instead of $\beta_{enc} = 1$. With increasing noise during encoding (Figure 4B) or during recall (Figure 4C), the optimal value of β_{enc}, expressed as β^*_{enc}, decreases.

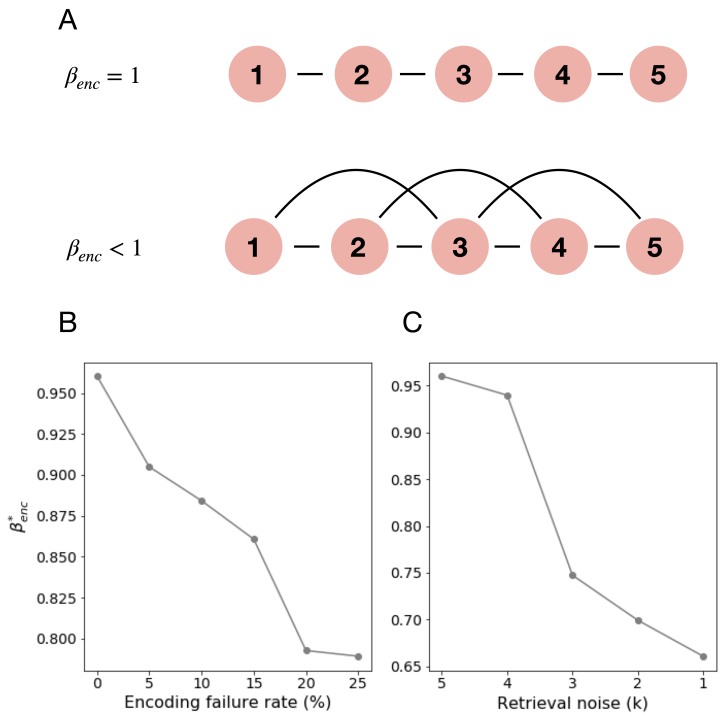

Figure 4. The effect of encoding and retrieval noise on the optimal policy. (A) When $\beta_{enc} = 1$ in the optimal policy, each item is only associated with the item before and after (chained through context); when $\beta_{enc} < 1$, items are also connected with items not immediately adjacent. (B) Encoding noise is increased by introducing the possibility that, on a percentage of trials (horizontal-axis), an item fails to be associated with the current context. The optimal drift rate during encoding β^*_{enc} decreases when encoding noise increases. The optimal drift rate during recall β^*_{rec} also decreases accordingly from $\beta_{rec} = 1$ to $\beta_{rec} = 0.90$ at the highest level of encoding noise. Optimal γ_{fc} stays at 0. The fixed parameters in the simulation are: $k = 5.38$, $\epsilon_d = 2.72$. (C) Retrieval noise is increased through the parameter k in the softmax retrieval rule. Smaller k values correspond to larger retrieval noise. β^*_{enc} decreases when retrieval noise increases. β^*_{rec} also decreases accordingly from $\beta_{rec} = 1$ to $\beta_{rec} = 0.56$ at the highest level of retrieval noise. Optimal γ_{fc} stays at 0. For this simulation, $\epsilon_d = 2.72$.

Intuitively this makes sense – when there is more noise in the system, "chaining" one item to the next through context is suboptimal, because of the high likelihood that the chain will be broken. Though having β_{enc} values smaller than 1 can cause occasional skipping of an item during recall, which is not ideal, it is more robust to such an interruption: When an item is associated with other nearby items in addition to the immediate adjacent item, occasionally failing to encode or recall the next item in the chain due to noise will not prevent other items from being recalled, as shown in Figure 4A. In contrast, when $\beta_{enc} = 1$, failing to encode or recall the next item in the chain will prevent all of the successive items from being recalled. In the simulations, encoding noise is implemented as the percentage of time when there is a failure to associate an item with the current context; the recall noise is controlled by the scale parameter k in the softmax retrieval rule during retrieval.

The second feature in the optimal policy that merits closer examination is the optimal starting context. It is optimal to initiate recall with the beginning-of-list context during study. If humans are indeed rational, one might wonder why participants often start recalling items from the end of the list. This can be explained based on the idea that participants start recall with the end-of-list context active in their mind; their ability to adopt the optimal policy is constrained by how much they are able to suppress the end-of-list context and reactivate the beginning-of-list context.

To account for this cognitive constraint in the model, we derive the optimal policy under the scenario when one is not able to reactivate the beginning-of-list context. Figure 5A shows the proportion of items recalled, as a function of the starting study context, assuming (for illustrative purposes) that participants can initiate recall at any serial position. We are most interested in two scenarios: One is when the participant is able to reactivate the beginning-of-list context ($c_{start} = 0$), and the other is when the participant is occupied with the end-of-list context ($c_{start} = 10$). When starting recall from the beginning of the list, the optimal value of β_{enc} is large (Figure 5B), for reasons described in the previous section; by contrast, if starting recall from the end of the list, the optimal value of

β_{enc} is small (Figure 5C) – in this case, the model can no longer rely on the strategy of chaining items through context in a forward manner, and instead the model makes items accessible at retrieval by placing them close by in the context space during encoding. The challenge here is that – at encoding – the participant does not know whether they will succeed or fail at initiating recall from the start of the list. To address this challenge, an optimal participant would need to choose the single value of β_{enc} that yields the highest expected value of recall performance, factoring in the relative probabilities of the two outcomes (i.e., initiating recall from the start vs. the end of the list). Put another way: While a participant cannot anticipate whether they will succeed at initiating recall from the start of the list on a particular trial, it is reasonable to think that the participant would know how successful they are *on average*, and they can use this information when deciding how to parameterize encoding.

To capture this, Figure 5D shows how overall recall accuracy varies as a function of β_{enc} and also the average proportion of trials where the participant succeeds at initiating recall from the start (vs. end) of the list. To obtain the results shown in Figure 5D, we fixed the values of these parameters (β_{enc} and the probability of initiating recall from the start of the list), and then selected values of the parameters that control retrieval policy (γ_{fc} and β_{rec}) in order to optimize recall performance. Importantly, at test, participants have the option of parameterizing recall differently depending on whether they were successful (on that trial) at initiating recall from the beginning of the list. Accordingly, the model was allowed to select different values of γ_{fc} and β_{rec} on trials where recall was initiated from the start vs. the end of the list. Figure 5D shows that, as the probability of initiating recall from the start of the list increases, the optimal value of β_{enc} increases.

So far, we have provided a rational account of primacy and forward asymmetry. As discussed earlier, forward asymmetry arises from updating context at retrieval using the pre-experimental context associated with the retrieved item ($\gamma_{fc} = 0$). Next, we identify the conditions where it is optimal to update context at retrieval using the context that was

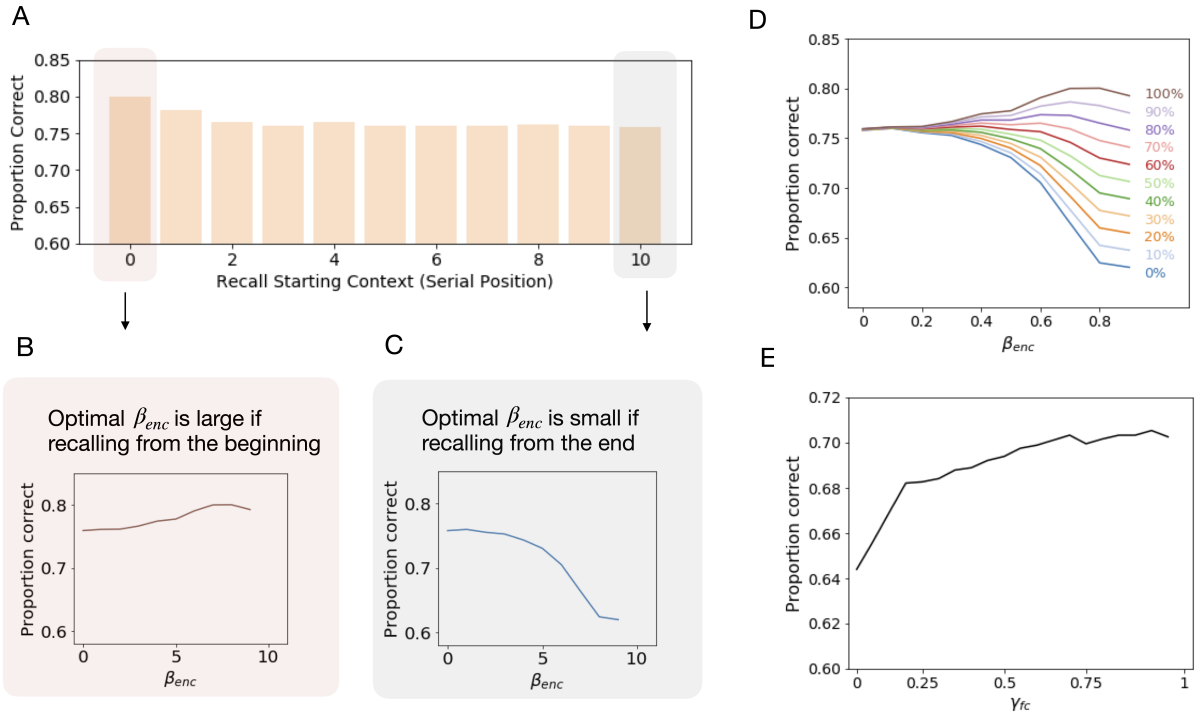

Figure 5. Optimal drift rate at encoding varies according to how recall is initiated. (A) Proportion of items correctly recalled, optimized over encoding and recall policy, assuming (for illustrative purposes) that participants can start recall by reinstating the context associated with an arbitrary serial position in the list. (B) The optimal drift rate at encoding β^*_{enc} is large if starting recall from the beginning of list. (C) β^*_{enc} is small if starting recall from the end of the list. (D) β^*_{enc} is constrained by how much one is able to initiate recall with the beginning-of-list context. The plot shows the effect of β_{enc} on overall recall accuracy, as a function of the proportion of trials in which participants succeed at starting recall from the beginning of the list (each colored line corresponds to a different proportion), under noisy retrieval conditions ($k = 3$). A total of five parameters are optimized in the simulations: β_{enc} for all trials, and different β_{rec} and γ_{fc} for trials that start from the beginning or start from the end; the fixed parameters are $k = 3$, $\epsilon_d = 2.72$. (E) On trials where the starting recall context is the end-of-list context, it is optimal to use symmetric contiguity (i.e., large γ_{fc}) to recall backwards. The plot here assumes that the probability of initiating recall with the beginning-of-list context is 80%, in which case the optimal parameters for trials with end-of-list context are $\gamma^*_{fc} = 0.93$, $\beta^*_{enc} = 0.6$ and $\beta^*_{rec} = 0.81$. The other fixed parameters are $k = 3$, $\epsilon_d = 2.72$.

associated with the retrieved item at study ($\gamma_{fc} = 1$), which gives rise to symmetric contiguity effects (i.e., backward recalls in addition to forward recalls). Figure 5E shows that, on trials when one is not able to jump back to the beginning of the list (i.e., the starting context is the end-of-list context), it is optimal to use both backward recalls and forward recalls by setting $\gamma_{fc} = 1$, providing a rational account of the symmetric contiguity effect. This pattern holds when β_{enc} is smaller than 0.8: When β_{enc} is too large, the context at encoding $c_t = \rho c_{t-1} + \beta_{enc} c^{IN}$ is dominated by the pre-experimental context induced by newly encoded items $\beta_{enc} c^{IN}$, leaving a very weak amount of experimental context c_{t-1} to generate backward recalls, even if one can fully reinstate this experimental context with $\gamma_{fc} = 1$ ($\beta_{enc} = 0.6$ in the specific example given in Figure 5E – this is the optimal β_{enc} value for when the proportion of start-of-list trials is 80%).

Testing the predictions of the optimal behavior function on empirical data

We have learned from rational-CMR that the optimal behavior function in free recall is to start recalling from the beginning of the list and then sequentially recall forward, relying on the forward asymmetry induced by updating context with the recalled item's pre-experimental context ($\gamma_{fc} = 0$). In addition, when recall cannot be initiated at the beginning-of-list context, it is optimal to minimize the amount of forward asymmetry and instead rely on symmetric contiguity, induced by updating context with the context that was associated with the recalled item at study ($\gamma_{fc} = 1$). These observations serve as predictions for human behavioral data. Here, we test these predictions at the level of individuals, using a large free recall database: the Penn Electrophysiology of Encoding and Retrieval Study (PEERS). The present analysis is based on the behavioral data from the 171 participants (age range: 18-30) in Experiment 1 of the PEERS dataset, which used an immediate free recall task. More details of the PEERS dataset can be found in previous studies (Lohnas et al., 2015; Healey and Kahana, 2016; Healey et al., 2018). The present analysis also incorporates the behavioral data from the 61 participants in Howard and

Kahana (2008) with an immediate free recall condition and a delayed free recall condition. The criteria for including participants and sessions into the current analysis can be found in the Appendix.

Prediction 1: Better-performing participants will demonstrate a stronger tendency to initiate recall from the beginning of the list and recall in a forward direction. According to rational-CMR, participants who perform better in the free recall task should demonstrate a stronger tendency to initiate recall from the beginning of the list, and they should make more forward recalls than backward recalls. Figure 6A-C summarizes the serial position curve, the probability of first recall, and the conditional response probability in the PEERS dataset. Participants were divided into two groups based on their average performance in the experiment: the better-performing participants are from the top 10% (grey) and the worse-performing participants are from the remaining 90% (orange). Dividing the participants at this percentile highlights the behavior of the participants who completed the free recall task with high performance (on average 85% correct recall) versus the worse-performing participants (on average 59% correct recall). Figure 6B shows that, consistent with the model's predictions, better-performing participants showed a higher tendency to start the recall process from the first item compared with the rest of the participants(Mann–Whitney $U = 940$, $n_1 = 17$, $n_2 = 154$, $p = 0.02$, two-tailed), and they also showed a lower tendency to start the recall process from the last item compared with the rest of the participants (Mann–Whitney $U = 719.5$, $n_1 = 17$, $n_2 = 154$, $p < 0.01$, two-tailed). In addition, Figure 6C shows that better-performing participants were more likely to carry out forward recalls, compared to the rest of the participants. To further examine this effect, we summarized forward asymmetry as a single value – the difference between the summed probability of four forward lags and the summed probability of four backward lags from the conditional response probability curve. Figure 6D shows that the Spearman's correlation between the forward asymmetry and free recall performance was $\rho = 0.25$ among the 171 individuals

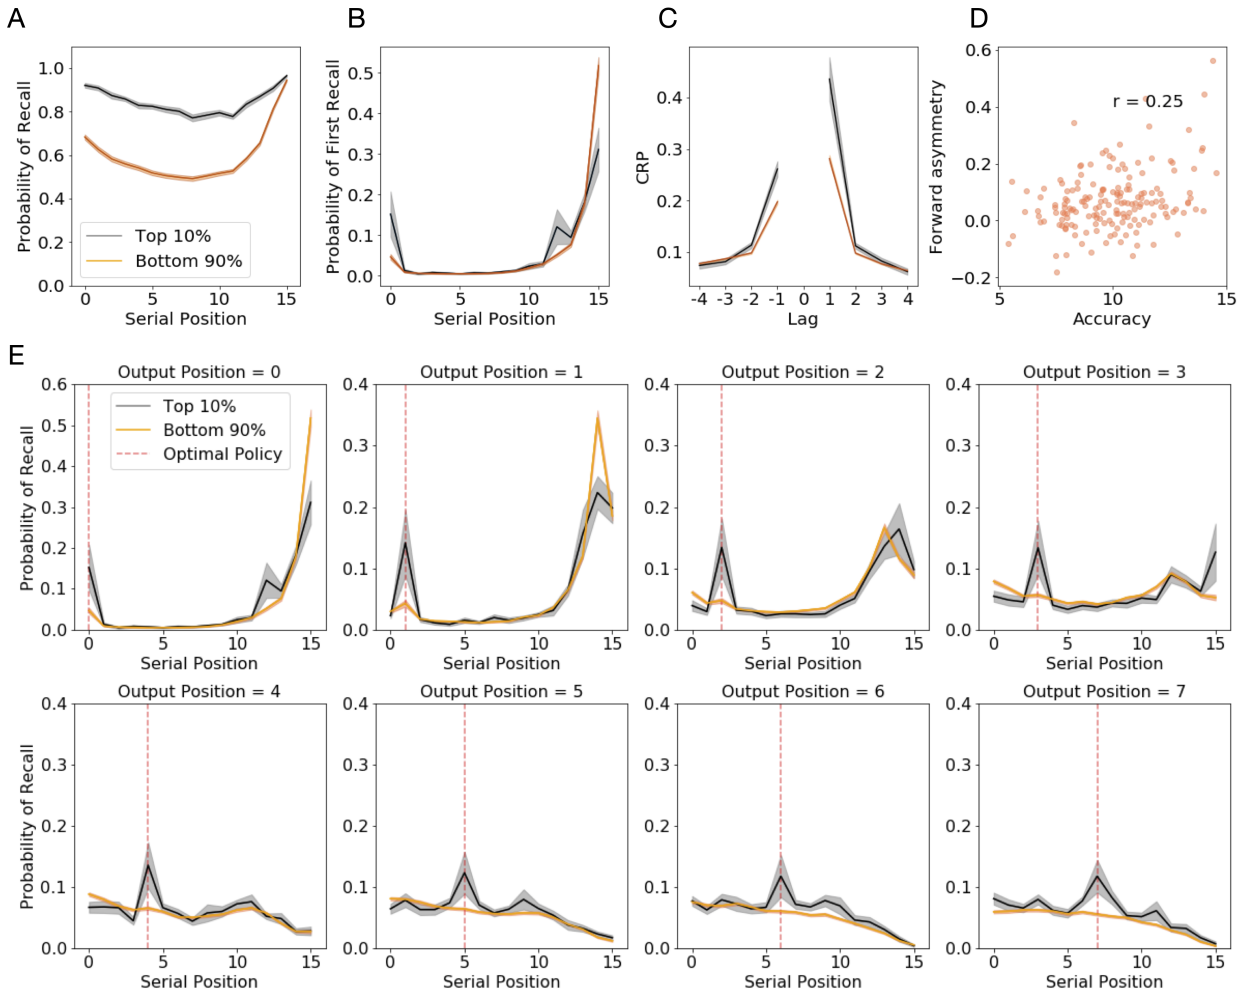

Figure 6. Better-performing participants demonstrate a stronger tendency to initiate recall from the beginning of the list and carry out forward recalls. (A-C) Behavioral patterns of the top 10% of participants (grey) in terms of performance in PEERS dataset compared with the rest of the participants (orange): serial position curve, probability of first recall, conditional response probability. The first three recalls are removed when computing the conditional response probability curve as they are heavily influenced by the recency effect. (D) Correlation between the forward recall asymmetry and participants' recall performance. Forward recall asymmetry is calculated as the difference between the summed probability of four forward lags and the summed probability of four backward lags from the conditional response probability. (E) Serial position curves for each output position for the top 10% and the bottom 90% of performers. The vertical dashed lines represent the optimal policy.

($p < 0.001$). Importantly, this correlation between forward asymmetry and performance is not a trivial consequence of better-performing individuals exhibiting higher primacy. When participants start recall from the beginning of the list, this naturally leads to a larger number of available forward recalls than backward recalls; however, it does not necessarily lead to forward asymmetry, since our measure of forward asymmetry (calculated from the conditional response probability) controls for the availability of recall transitions (i.e., the conditional probability of a forward transition at a particular lag is the observed probability of that transition, divided by the number of times when that transition was possible).

To visualize in more detail how serial positions change by output position during recall, Figure 6E plots serial position curve for each output position, averaged across the top 10% and the bottom 90% of performers. The optimal policy of starting recall from the beginning of the list and then sequentially recalling forwards corresponds to the dashed vertical line for each subplot, where the first recall is at the first serial position, and the second recall is at the second serial position, etc. Consistent with our model predictions, Figure 6E shows that recall transitions among the top 10% of performers are more aligned with the optimal policy than recall transitions among the bottom 90% of performers.

One might worry that higher-performing participants will always show greater consistency with the optimal policy, regardless of whether our rational-CMR theory is correct; however, this is not the case. A simple counterexample to this would be if the participants all follow random policies that demonstrate some variability, but none of these policies are close to the optimal policy. In this case, even if we looked at the high-performing participants, those participants would be following strategies that differ from the optimal policy (see Appendix for a simulation that demonstrates this).

Prediction 2: Better-performing participants will demonstrate a stronger effect of forward asymmetry when recall is initiated from the beginning of the list and a stronger effect of symmetric contiguity when recall is initiated from the end of the list. Participants are not able to activate the beginning-of-list context on

all trials. According to rational-CMR, patterns of recall transitions (operationalized using conditional response probability) should depend on whether recall can be initiated at the beginning of-list context. Specifically, the model predicts that it is optimal to maximize forward asymmetry in recall transitions if the recall is initiated from the beginning-of-list context; conversely, it is optimal to minimize forward asymmetry and maximize the ability to carry out backward recalls through the symmetric contiguity effect if the recall is initiated from the end-of-list context. In Figure 7, trials are divided into primacy trials (A-C) and recency trials (D-F) based on patterns of recall initiation. In human data, there is no "ground truth" way of measuring the starting context at recall. Consequently, we sorted trials into primacy trials and recency trials based on behavioral patterns. The most obvious way to define primacy trials would be to select trials where recall started from the first item in the study list (Ward et al., 2010). However, applying this criterion would leave us with fewer than 7 out of 96 trials for more than half of the 171 participants and zero primacy trials for 25% of all participants (who must therefore be removed from the analysis). To improve the power of our analysis, we relaxed the definition of the primacy trials to encompass all trials where the first item at study was retrieved during the first four recalls (we nevertheless examined the results under both definitions of primacy trials and observed that the major conclusions of the analysis did not change). This is based on the observation that sometimes participants can recall from the beginning of the list – they just cannot do that right from the first recall. In the PEERS free recall dataset, the first item during study is retrieved the most frequently at the fourth position, as the first three recalls are dominated by the recency effect. Next, we had to operationally define recency trials. Here, we wanted to select trials where performance was unlikely to have been influenced by the beginning-of-list context at any point during the recall period; to meet this criterion, we defined recency trials as those without recalls of the first item.

Figure 7A plots the conditional response probability for the top 10% and the bottom 90% of performers in the human data, where the analysis was limited to primacy trials and

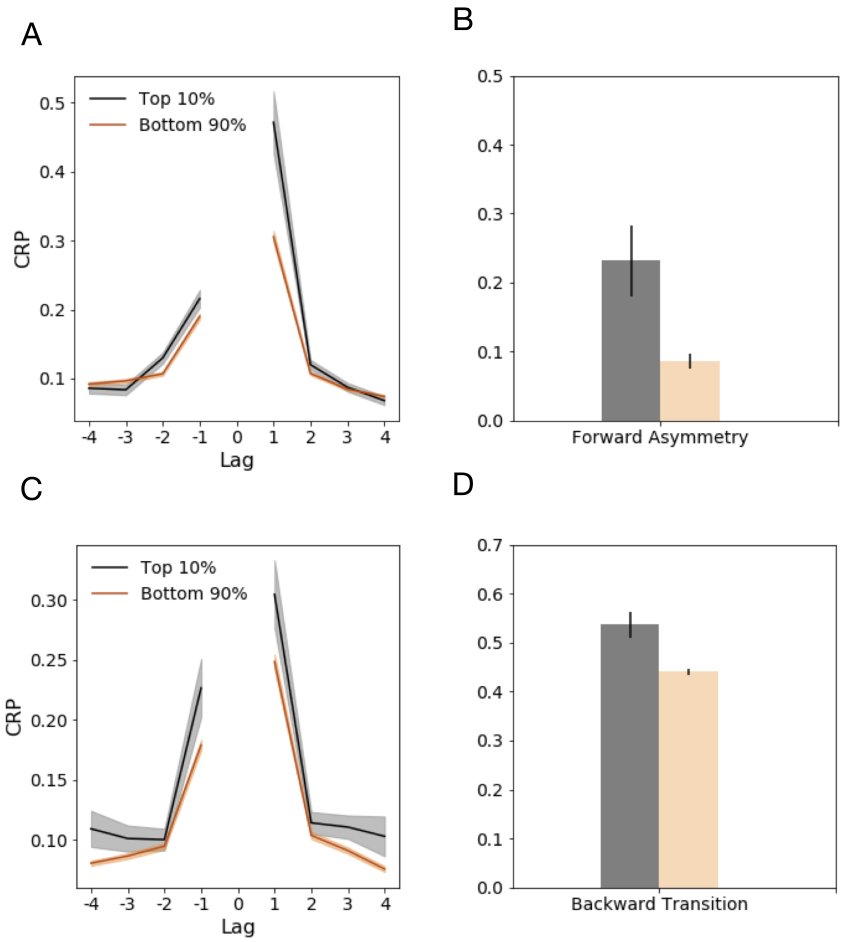

Figure 7. Better-performing participants demonstrate a stronger effect of forward asymmetry than other participants on primacy trials (A-B), and a stronger effect of symmetric contiguity than other participants on recency trials, as characterized by backward transitions (C-D); see text for definitions of primacy and recency trials. The behavioral score of forward asymmetry in B is calculated as the difference between the summed probability of four forward lags and the summed probability of four backward lags from the conditional response probability of each participants (N=171). The behavioral score of backward transitions in D is calculated as the summed probability of four backward lags from the conditional response probability curve of each participant (N=171). In A and B, "Top 10%" refers to participants whose performance during primacy trials is in the top 10^{th} percentile of the population. In C and D, "Top 10%" refers to participants whose performance during recency trials is in the top 10^{th} percentile of the population. Error bars indicate the SEM.

"top 10%" and "bottom 90%" were computed based on performance on these primacy trials. We found that the better performers showed higher levels of forward asymmetry. To formally test this, we calculated a behavioral forward asymmetry score (plotted in Figure 7B separately for the top 10% and the bottom 90% of performers), operationalized as the difference between the summed probability of four forward lags and the summed probability of four backward lags from the conditional response probability curve in Figure 7A; this score was higher in the top 10% of performers (Mann–Whitney $U = 807$, $n_1 = 17$, $n_2 = 154$, $p < 0.01$, two-tailed) than in rest of the participants.

Figure 7C plots the conditional response probability for the top 10% and the bottom 90% of performers in the human data, where the analysis was limited to recency trials and "top 10%" and "bottom 90%" were computed based on performance on these recency trials. We found that the better performers showed a stronger effect of symmetric contiguity. As discussed earlier, CMR posits that forward transitions (as captured in the conditional response probability curve) are contributed by both forward asymmetry and symmetric contiguity, whereas backward transitions are uniquely contributed by symmetric contiguity; we therefore calculated the behavioral influence of symmetric contiguity by computing a backward transition score (plotted in Figure 7D separately for the top 10% and the bottom 90% of performers), operationalized as the summed probability of four backward lags from the conditional response probability curve in Figure 7C. The score was higher in the top 10% of performers (Mann–Whitney $U = 641$, $n_1 = 17$, $n_2 = 154$, $p < 0.001$, two-tailed). Though the forward asymmetry in the top 10% of performers is still sizable and similar with the bottom 90% of performers (Mann–Whitney $U = 1235$, $n_1 = 17$, $n_2 = 17$, $p = 0.35$, two-tailed), there is decreased forward asymmetry in the top 10% of performers during recency trials compared with that in primacy trials (Mann–Whitney $U = 1235$, $n_1 = 154$, $n_2 = 154$, $p = 0.03$, two-tailed) but this is not observed in the bottom 90% of performers (Mann–Whitney $U = 1235$, $n_1 = 17$, $n_2 = 154$, $p = 0.30$, two-tailed). We also carried out the same analysis using an alternative way of scoring forward asymmetry, calculated as the

ratio, instead of the difference, between the summed probability of four forward lags and the summed probability of four backward lags from the conditional response probability of each participant. This does not change the main conclusions.

Analysis of the PEERS dataset provided evidence for the model prediction that it is optimal to maximize forward asymmetry when the recall is initiated from the beginning-of-list context, and maximize backward recall when the recall is initiated from the end-of-list context. However, demonstrating that recall transitions are correlated with how recall is initiated is not the same as demonstrating that recall transitions depend on the how recall is initiated, where the accessibility of the beginning-of-list context (relative to the end-of-list context) is experimentally manipulated. The latter is a stronger test of the model predictions. To investigate this latter scenario, we analyzed publicly available datasets from Howard and Kahana (1999), where the recall phase of the free recall experiments was conducted either immediately after studying the list, or with a delayed period of distractor activities. Participants could not predict during encoding whether it would be an immediate or delayed condition; they were warned that sometimes there would be a math test between the end of the list and the signal to recall and sometimes there would not. Adding a period of delay after studying the list helps with disengaging from the end-of-list context (Murdock, 1962), and leads to increased accessibility of the beginning-of-list context during recall, as can been from Figure 8B, where there is a higher tendency to start the recall process from the first item in the delayed free recall condition than the immediate free recall condition (Mann–Whitney $U = 185.500$, $n_1 = 65$, $n_2 = 62$, $p = 0.04$, two-tailed).

Having confirmed that recall initiation is affected by the manipulation of immediate vs. delayed recall in this dataset, the next step is to examine whether optimal recall transitions vary in the immediate vs. delayed recall conditions. Consistent with the model predictions, there is stronger forward asymmetry in the delayed free recall condition than in the immediate free recall condition (Figure 8C; Mann–Whitney $U = 1188.00$, $n_1 = 65$,

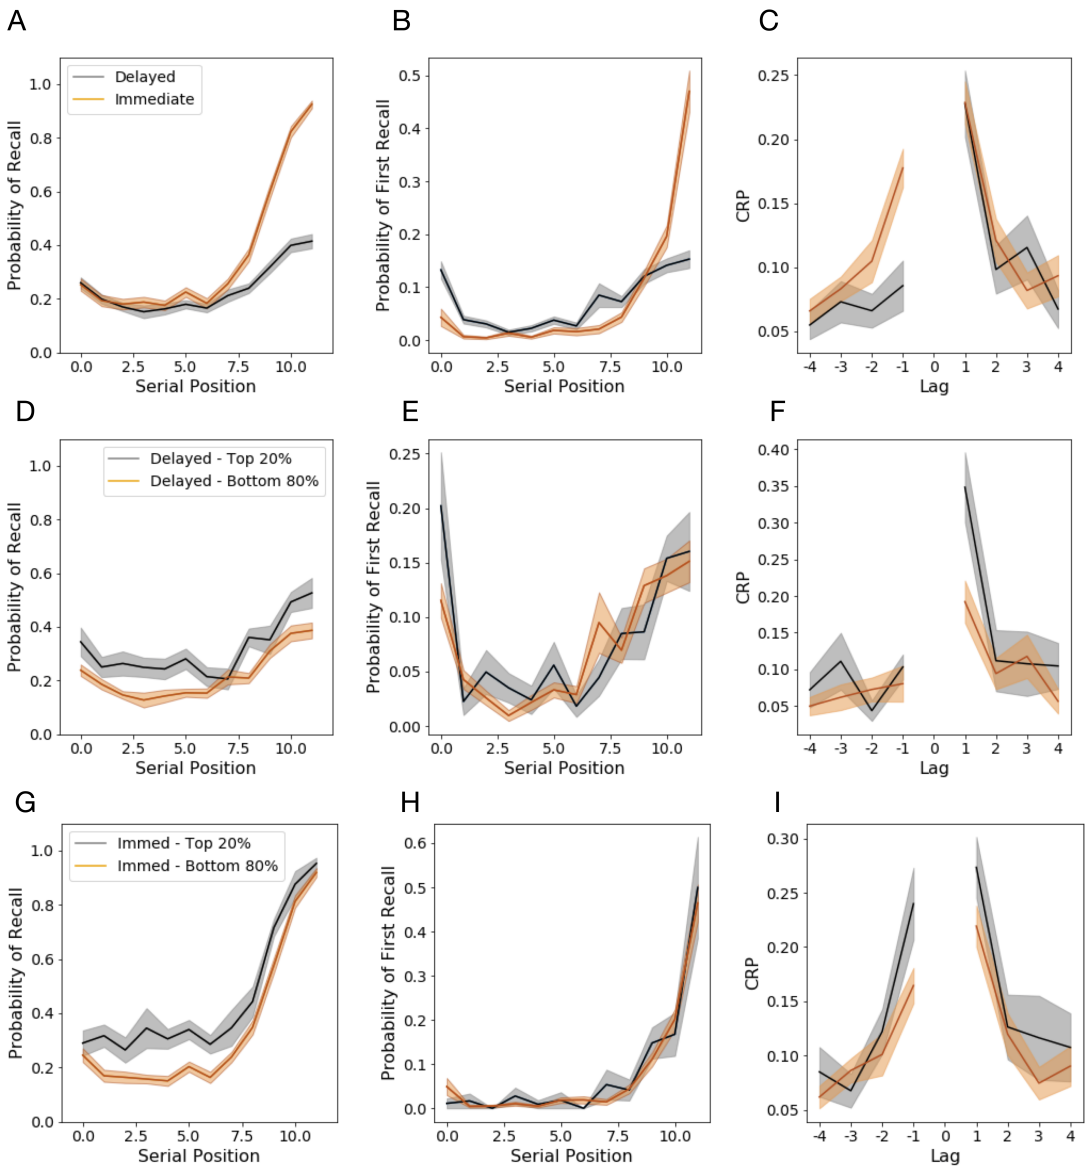

Figure 8. Behavioral patterns in the dataset of Howard and Kahana (1999) for the delayed free-recall condition and the immediate free-recall condition (A-C), the delayed condition split by performance level of participants (D-F), and the immediate condition split by performance level of participants (G-I). From left to right, the three columns show the serial position curve, probability of the first recall, and conditional response probability.

$n_2 = 62$, $p = 0.020$, two-tailed), where there is higher accessibility of the beginning-of-list context. Moreover, we can divide participants into two groups based on their average performance in the experiment: the better-performing participants are from the top 20% (grey) and the worse-performing participants are from the remaining 80% (orange). We used a 20/80 split here (instead of dividing the participants into the top 10% and remaining 90% as in the PEERS dataset) to make sure that there were at least 10 participants in each group. Figure 8F shows that, in the delayed free recall condition (when it is easier to access the beginning-of-list context), better-performing participants demonstrate numerically (but not significantly) stronger forward asymmetry compared to the worse-performing participants (Mann–Whitney $U = 198.50$, $n_1 = 12$, $n_2 = 41$, $p = 0.159$, two-tailed; however, at the +1 lag of the conditional response probability curve, the difference between better-performing participants and worse-performing participants is significant: Mann–Whitney $U = 187.50$, $n_1 = 12$, $n_2 = 41$, $p = 0.007$, two-tailed); conversely, in the immediate free recall condition (when the end-of-list context is relatively more accessible) better-performing participants demonstrate stronger backward transitions compared to the worse-performing participants (Mann–Whitney $U = 144.00$, $n_1 = 10$, $n_2 = 48$, $p = 0.025$, two-tailed). Taken together, these results provide converging support for the model's predictions about how recall initiation will affect recall policy.

General Discussion

Much of the work in the free recall literature has sought to understand the processes and representations that give rise to different patterns of recall organization; however, it remains unclear why certain types of recall organization arise in the first place. Our work provides a rational analysis of the free recall task, deriving the optimal policy under the internal representations and processes of memory search described by CMR. Our model, which we call rational-CMR, demonstrates that the optimal policy in free recall is to start recall from the beginning of the list and then sequentially recall forwards, providing a

rational account of primacy and forward asymmetry effects typically observed in free recall. In addition, the rational model also makes a novel prediction that the optimal policy in recall transitions should depend on recall initiation. These predictions from the rational model were confirmed in human behavioral data, where top human participants demonstrated stronger primacy and forward asymmetry than the rest of the participants, and the amount of forward asymmetry in participants depended on the starting context at recall. We now turn to the broader implications of these results.

Why is the optimal policy optimal?

We show that the optimal policy in free recall is to start recall from the beginning of the list and then sequentially recall forwards, which corresponds to a strategy that chains items through context (Figure 3B). Intuitively, chaining (in this case, indirectly through context) is a useful strategy because it minimizes the odds that an item will be inadvertently skipped. However, if the chained structure in the optimal policy is the key factor for effective recalls, one might ask why is it necessary to start recall from the beginning and then recall forwards, rather than starting recall from the end and then recalling backwards. The latter would be more convenient, because one already has access to the end-of-list context at the start of recall, whereas re-activating the beginning-of-list context requires an extra retrieval step. To understand this better, we need to review the encoding mechanism in CMR in Figure 1. In CMR, there is a reliable way to proceed forward (by reinstating c_{enc}^{IN}) but there is not a reliable way to proceed backward – the only way to promote backward transitions is to reinstate the drifting contextual state associated with the item at study, which is equally likely to propel recall forward and backward. Since there is a reliable way to propel recall forward but no reliable way to propel recall backward, proceeding forward yields better results (assuming that the participant is able to succeed in initiating recall from the start of the list).

A rational account of the primacy effect

In the free recall task, words presented at the beginning and at the end of the list are better remembered than those presented in the middle of the list. This U-shape in the serial position curve of free recall is one of the most robust findings in cognitive psychology (Murdock, 1962). There exist a number of theories that account for primacy. Rehearsal-based explanations of the primacy effect posit that increased rehearsal opportunities for early items in the list lead to increased encoding in the short-term memory buffer (Atkinson and Shiffrin, 1968; Murdock and Metcalfe, 1978; Rundus, 1971; Rundus and Atkinson, 1970; Brodie and Murdock, 1977). Even if no short-term buffer is presumed, the repetition, the distribution, and the recency of rehearsals can make primacy items more accessible at test (Tan and Ward, 2000). Despite the success of the rehearsal account in explaining primacy, there are empirical studies showing that primacy is still found when rehearsal is minimized, such as when a five-item list is presented in less than one second (Neath and Crowder, 1996), when there is semantic judgement required encoding each word (Howard and Kahana, 1999), or when there is an interval of distractor activity between every two adjacent words that are presented (Polyn et al., 2009b). Other accounts of primacy do not depend on the rehearsal process. For example, some accounts posit that the first items are more distinctive than the other items (Murdock, 1960; Neath, 1993), although this has been critiqued on the grounds that distinctiveness is an effect rather than an explanation (Hunt and Lamb, 2001). Other accounts posit that primacy is caused by contextual change at the beginning of the list, induced by either a prediction error (Rouhani et al., 2020) or novelty (Davelaar, 2013). CMR captures the primacy effect by assuming a gradient by serial position, without explicitly committing to a particular mechanism of primacy.

In contrast to the debate in the literature on the mechanisms that give rise to primacy, our work provides a rational account of why primacy is important to have in the first place. In a rational analysis of the free recall task based on CMR simulations, we

showed that initiating recall from the beginning of the list gives better performance than initiating recall from other positions of the list. This model prediction was tested using a large free recall dataset. Consistent with the model's predictions, the better-performing participants in the dataset showed a higher probability of initiating recall from the beginning of the list and a lower probability of initiating recall from the end of the list, compared with those of the worse-performing participants. Notably, the first recall in the better-performing participants did not always consist of the first item in the study list, but rather was dominated by the last few items. This can be explained in terms of the idea that, at the start of recall, the end-of-list context is active; to activate the beginning-of-list context, participants need to suppress the end-of-list context and retrieve the beginning-of-list context – both of these processes are prone to failure, and the difficulty of activating the first part of the list can be modulated by situational factors (e.g., the presence of a distinctive event at the start of the list). Thus, the challenging nature of replacing the end-of-list context with the start-of-list context serves as a constraint on optimal behavior (Sahakyan and Kelley, 2002). The second part of the model simulations take into account this constraint, by deriving what the optimal behavior is when one is only able to retrieve the beginning-of-list of context on a proportion of trials. Implications of these results are discussed in the next section.

There has been other work that has considered how participants vary how they initiate recall as a function of task instructions. Tan et al. (2016) found that the tendency to start recall from the first item occurred only when participants were required to recall as many items as possible. When participants were asked to recall only 1 or 2 items, they tended to initiate recall with end-of-list items (Tan et al., 2016). These findings show that different retrieval strategies exist for searching for different numbers of items, and that participants have some control over these retrieval strategies (Ward and Tan, 2019). Our analysis provides a rational explanation for why participants would adopt different retrieval strategies under different task demands. When the goal of the memory task is to retrieve

as many items as possible, recalling from the beginning meets the goal better than recalling from the end, as shown in the optimal policy. When the goal of the memory task is to retrieve only 1 or 2 items, recalling from anywhere in the list can meet the goal equally well; under this scenario, participants do not have to expend the effort to reinstate the beginning-of-list context – it is optimal to start recalling from the end-of-list context, which is already active at the start of recall.

A rational account of forward asymmetry

During memory search, forward transitions are more likely than backward transitions for small absolute values of lag (Kahana, 1996). This effect has been shown in immediate free recall (Kahana, 1996), as well as delayed free recall (Howard and Kahana, 1999). Unlike primacy, for which CMR does not commit to a particular mechanism, CMR can account for the mechanism of forward asymmetry. However, CMR does not address why there is forward asymmetry. It is possible under the space of CMR parameterization to turn off the asymmetry entirely ($\gamma_{fc} = 1$). Our work provides a rational account of why forward asymmetry is an important feature in free recall. By making a connection between forward asymmetry and overall task performance, we were able to account for the individual differences in free recall performance observed in a large free recall database. Specifically, there was a significant correlation between the amount of forward asymmetry and overall performance at the level of individuals.

In addition to explaining the forward asymmetry, rational-CMR can also explain the presence of backward transitions. Our results show that optimal recall transitions depend on recall initiation; forward asymmetry is only important when recalls are initiated from the beginning of the list. When participants are not able to initiate recall from the beginning of the list, it is optimal to rely on symmetric contiguity to recall backwards instead of forward asymmetry. This model prediction was further tested over a large free recall database and a dataset containing both immediate and delayed free recall conditions.

Consistent with the model's predictions, better-performing participants demonstrated a stronger effect of forward asymmetry when recall was initiated from the beginning of the list and a stronger effect of the symmetric contiguity when recall was initiated from the end of the list.

Both forward asymmetry and symmetric contiguity contribute to temporal contiguity effects commonly observed in free recall (Kahana, 1996), and it has been shown that participants' temporal contiguity effects correlate highly with their recall performance (Sederberg et al., 2010). Results from rational-CMR are consistent with this empirical finding, while also providing additional insight into which type of temporal contiguity is most useful for performance, conditioned on how recall is initiated.

Our results also connect to several empirical results in the literature that examined the relationship between primacy/recency and memory performance (Dalezman, 1976). Dalezman (1976) instructed participants to recall either from the beginning, the end or the middle of a list before they recalled the rest of the items. Though the amount of primacy and recency can be reliably altered by the recall instruction, the overall performance largely stays unchanged. Why is it the case that altering recall initiation does not affect the overall performance? Our results emphasize that one needs to coordinate both recall initiation and recall transitions to achieve optimal recall. Though the optimal policy is to demonstrate primacy and forward asymmetry, having primacy alone, as induced from one of the recall conditions in Dalezman (1976), does not guarantee improvement in memory performance.

Other work by Unsworth et al. (2011) has looked at individual differences in primacy/recency and memory performance. Participants were clustered into three groups according to their serial position curves: a high-primacy low-recency group, a low-primacy high-recency group, and a high-primacy high-recency group; Unsworth et al. (2011) found that the last group demonstrated the highest level of memory performance. In the Appendix, we include a "proof-of-concept" simulation showing one way that our model can explain this pattern of results. We assume that all participants have a similar mix of trials

where they start recall from the beginning vs. the end of the list, but they vary in how well they can optimize their recall strategy as a function of how they initiate recall. In our simulation, we have one group that pursues the optimal policy (adjusting γ_{fc} based on how recall is initiated), another group that uses the "optimal-for-beginning-of-list" γ_{fc} regardless of how recall is initiated, and a third group that uses the "optimal-of-end-of-list" γ_{fc} parameters regardless of how recall is initiated. With these assumptions, our model qualitatively replicates the key results from Unsworth et al. (2011) – the "optimal policy" group has the best overall performance and demonstrates a high level of performance for both primacy and recency items compared with other groups, whereas the other two groups demonstrate high levels of performance for primacy or recency items (but not both) and lower overall levels of performance. Details of the simulation can be found in the Appendix.

Connection to mnemonic techniques

Rational-CMR may also help to explain the efficacy of mnemonic techniques such as the method of loci. The method of loci was first documented in Roman and Greek rhetorical treatises, where orators adopted this technique to memorize key points in a long speech (Yates, 1966). Practicing this method first involves familiarizing oneself with an environment, such as a street or a building, and mentally constructing an ordered route between well-defined locations (loci). When given a list of items to remember, the method of loci is characterized by a stereotyped walk during both encoding and recall: During encoding, one mentally walks through the pre-defined route, associating each item with loci along this route; during recall, one simply imagines oneself walking through the same familiar route, using loci along the route as memory cues to subsequently retrieve each item (Yates, 1966).

The stereotyped walk used in the method of loci resembles the optimal policy derived in rational-CMR, which is to start recall from the beginning of the list and sequentially recall forwards. We propose that one of the reasons why the method of loci works so well is

that it enforces the optimal policy of memory search through explicit instructions and pre-training. During free recall, without using mnemonic techniques, there are two potential reasons why it is hard for individual participants to adopt the exact optimal policy derived in the rational-CMR model. First, participants start recall with the end-of-list context in their mind. Their ability to adopt the optimal policy is constrained by how much they are able to drop the end-of-list context and reactivate the beginning-of-list context. Second, participants have limited encoding time to fully process and store an item. Their ability to start recall from the beginning of the list and recall sequentially forwards is constrained by the fidelity of their encoding. Using the method of loci, one can enforce the usage of optimal policy during free recall by directing the recall process to always start from the beginning; also, the method of loci leverages a pre-learned spatial representation so that information can be more easily encoded in a sequence on the fly during the encoding stage.

Our results also shed light on the recent empirical finding that two other mnemonics – the peg method and the temporal encoding method – achieve comparable improvements in recall performance compared to the method of loci (Bouffard et al., 2018; Caplan et al., 2019). Both methods emphasize a temporal feature as the scaffold. The peg method works by pre-memorizing a list of "peg words"; during encoding, items to study are associated with this list of words in the order that they are presented. The temporal encoding method works similarly to the peg method. Instead of pre-memorizing a list of words, participants pre-define a chronological timeline using a list of their most vivid memories and later associate items to study with these memories. Despite their differences with the method of loci (in the exact instructions that are used, and their use of a non-spatial "scaffold" instead of spatial knowledge) these techniques all share an encoding and recall policy that resembles the optimal policy in free recall, which may help to explain why they achieve comparable improvements in recall performance.

Limitations and future work

Committing to the architectural assumptions of CMR, the current work derived the optimal policy and demonstrated that the derived optimal policy can explain performance observed in human data. It is possible that there is additional variance in memory performance that can be explained by altering the architectural assumptions of CMR. In future work, we will explore the consequences on the optimal policy if specific assumptions of CMR are altered (for example, if we allowed parameters to vary dynamically as the recall process unfolds, instead of being held constant). For simplicity, the current work assumes that the semantic representations of words are independent from each other. We acknowledge that semantic associations between words drive a significant amount of the recall dynamics. Typically, in categorized-list free recall experiments, with semantically-related items located at different locations of the list, the semantic structure can overshadow the effect from temporal context, with semantic context driving recalls in different directions than the temporal context. Therefore, we might not be able to arrive at a stable solution for optimal policy under the influence of semantic associations. The first step of the rational analysis is to isolate the two and focus on discovering optimal policy without effect from semantic associations. Future work will examine the two components jointly.

Conclusion

The discovery of the primacy effect and the recency effect in free recall dates back to the times of Hermann Ebbinghaus (Ebbinghaus, 1885). Since then, psychologists have used the free recall task to gain insight into processes and representations underlying memory search (Murdock, 1960; Murdock, 1962; Roberts, 1972; Standing, 1973). The field has primarily focused on identifying an array of behavioral patterns that are consistently shown across free recall experiments, and building computational models to capture these behavioral patterns. To our knowledge, our work is the first to examine why humans

demonstrate these recall patterns in the first place. Under the architectural assumptions about internal representations and processes during memory search instantiated by the CMR model, we showed that recall patterns including primacy, forward asymmetry, compound cuing, and symmetric contiguity can arise naturally from optimizing the overall recall performance. In particular, we showed that, although participants are free to recall items in any order (hence "free" recall), the optimal policy is one where participants recall items in the same order that items are studied. We connected these results with the literature on mnemonic techniques, and proposed that mnemonic techniques support superior memory by enforcing the optimal policy of memory search through explicit instructions and pre-training. These findings also offer tantalizing insights into how to further improve human memory through the intelligent choice of memory cues.

References

Anderson, J. R. (1990). *The adaptive character of thought*. Psychology Press.

Anderson, J. R., & Bower, G. H. (1972). Recognition and retrieval processes in free recall. *Psychological review*, *79*(2), 97–123.

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

Atkinson, R., & Shiffrin, R. (1968). Human memory: A proposed system and its control processes. *Psychology of Learning and Motivation*, *2*(4), 89–195. https://doi.org/10.1016/s0079-7421(08)60422-3

Bouffard, N., Stokes, J., Kramer, H. J., & Ekstrom, A. D. (2018). Temporal encoding strategies result in boosts to final free recall performance comparable to spatial ones. *Memory & Cognition*, *46*(1), 17–31.

Bousfield, W. A. (1953). The occurrence of clustering in the recall of randomly arranged associates. *The Journal of General Psychology*, *49*(2), 229–240.

Bousfield, W. A., & Sedgewick, C. H. W. (1944). An analysis of sequences of restricted associative responses. *The Journal of General Psychology*, *30*(2), 149–165.

Bower, G. (1967). A multicomponent theory of the memory trace. *Psychology of Learning and Motivation*, *1*, 229–325.

Brodie, D. A., & Murdock, B. B. (1977). Effect of presentation time on nominal and functional serial-position curves of free recall. *Journal of Verbal Learning and Verbal Behavior*, *16*(2), 185–200. https://doi.org/10.1016/s0022-5371(77)80046-7

Caplan, J. B., Legge, E. L., Cheng, B., & Madan, C. R. (2019). Effectiveness of the method of loci is only minimally related to factors that should influence imagined navigation. *Quarterly Journal of Experimental Psychology*, *72*(10), 2541–2553. https://doi.org/10.1177/1747021819858041

Cofer, C. N., Bruce, D. R., & Reicher, G. M. (1966). Clustering in free recall as a function of certain methodological variations. *Journal of Experimental Psychology*, *71*(6), 858–866.

Dalezman, J. J. (1976). Effects of output order on immediate, delayed, and final recall performance. *Journal of Experimental Psychology: Human Learning and Memory*, *2*(5), 597.

Davelaar, E. J. (2013). A novelty-induced change in episodic (NICE) context account of primacy effects in free recall. *Psychology*, *04*(09), 695–703. https://doi.org/10.4236/psych.2013.49099

Ebbinghaus, H. (1885). Memory: A contribution to experimental psychology (translated by ha ruger & ce bussenues, 1913). *New York: Teachers College, Columbia University.*

Estes, W. K. (1955). Statistical theory of spontaneous recovery and regression. *Psychological Review*, *62*(3), 145–154.

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 & Instrumentation*, *14*(4), 375–399.

Griffiths, T., Canini, K., Sanborn, A., & Navarro, D. (2007). Unifying rational models of categorization via the hierarchical dirichlet process.

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

Healey, M. K., Long, N. M., & Kahana, M. J. (2018). Contiguity in episodic memory. *Psychonomic Bulletin & Review*, *26*(3), 699–720. https://doi.org/10.3758/s13423-018-1537-3

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. https://doi.org/10.1037/0278-7393.25.4.923

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

Howard, M. W., & Kahana, M. J. (2002b). When does semantic similarity help episodic retrieval? *Journal of Memory and Language*, *46*(1), 85–98.

Howes, A., Lewis, R. L., & Vera, A. (2009). Rational adaptation under task and processing constraints: Implications for testing theories of cognition and action. *Psychological review*, *116*(4), 717.

Hunt, R. R., & Lamb, C. A. (2001). What causes the isolation effect? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*(6), 1359–66.

Kahana, M. J. (1996). Associative retrieval processes in free recall. *Memory & Cognition*, *24*(1), 103–109. https://doi.org/10.3758/bf03197276

Kahana, M. J., Howard, M. H., Zaromb, F., & Wingfield, A. (2002). Age dissociates recency and lag recency effects in free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*(3), 530–540.

Kragel, J. E., Morton, N. W., & Polyn, S. M. (2015). Neural activity in the medial temporal lobe reveals the fidelity of mental time travel. *Journal of Neuroscience*, *35*(7), 2914–2926.

Laming, D. (1999). Testing the idea of distinct storage mechanisms in memory. *International Journal of Psychology*, *34*(5-6), 419–426.

Lieder, F., & Griffiths, T. L. (2020). Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, *43*.

Lohnas, L. J., & Kahana, M. J. (2014). Compound cuing in free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *40*(1), 12–24. https://doi.org/10.1037/a0033698

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.

Marr, D. (2010). *Vision: A computational investigation into the human representation and processing of visual information*. MIT press.

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

Miller, J. F., Weidemann, C. T., & Kahana, M. J. (2012). Recall termination in free recall. *Memory & cognition*, *40*(4), 540–550.

Mockus, J., Tiesis, V., & Zilinskas, A. (1978). The application of Bayesian methods for seeking the extremum. *Towards Global Optimization*, *2*(117-129), 2.

Murdock, B. B. (1960). The distinctiveness of stimuli. *Psychological Review*, *67*(1), 16.

Murdock, B. B. (1962). The serial position effect of free recall. *Journal of Experimental Psychology*, *64*(5), 482–488. https://doi.org/10.1037/h0045106

Murdock, B. B. (1997). Context and mediators in a theory of distributed associative memory (TODAM2). *Psychological Review*, *104*(4), 839.

Murdock, B. B., & Metcalfe, J. (1978). Controlled rehearsal in single-trial free recall. *Journal of Verbal Learning and Verbal Behavior*, *17*(3), 309–324.

Neath, I. (1993). Distinctiveness and serial position effects in recognition. *Memory & Cognition*, *21*(5), 689–698.

Neath, I., & Crowder, R. G. (1996). Distinctiveness and very short-term serial position effects. *Memory*, *4*(3), 225–242.

Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *10*, 104–114.

Polyn, S. M., Norman, K. A., & Kahana, M. J. (2009a). A context maintenance and retrieval model of organizational processes in free recall. *Psychological Review*, *116*(1), 129–156. https://doi.org/10.1037/a0014420

Polyn, S. M., Norman, K. A., & Kahana, M. J. (2009b). Task context and organization in free recall. *Neuropsychologia*, *47*(11), 2158–2163. https://doi.org/10.1016/j.neuropsychologia.2009.02.013

Reed, S. K. (1972). Pattern recognition and categorization. *Cognitive Psychology*, *3*, 382–407.

Roberts, W. A. (1972). Free recall of word lists varying in length and rate of presentation: A test of total-time hypotheses. *Journal of Experimental Psychology*, *92*(3), 365.

Rouhani, N., Norman, K. A., Niv, Y., & Bornstein, A. M. (2020). Reward prediction errors create event boundaries in memory. *Cognition*, *203*, 104269. https://doi.org/10.1016/j.cognition.2020.104269

Rundus, D. (1971). Analysis of rehearsal processes in free recall. *Journal of Experimental Psychology*, *89*(1), 63–77. https://doi.org/10.1037/h0031185

Rundus, D., & Atkinson, R. C. (1970). Rehearsal processes in free recall: A procedure for direct observation. *Journal of Verbal Learning and Verbal Behavior*, *9*(1), 99–105. https://doi.org/10.1016/s0022-5371(70)80015-9

Sahakyan, L., & Kelley, C. M. (2002). A contextual change account of the directed forgetting effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*(6), 1064.

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. https://doi.org/10.1037/a0013396

Sederberg, P. B., Miller, J. F., Howard, M. W., & Kahana, M. J. (2010). The temporal contiguity effect predicts episodic memory performance. *Memory & Cognition*, *38*(6), 689–699. https://doi.org/10.3758/mc.38.6.689

Snoek, J., Larochelle, H., & Adams, R. P. (2012). Practical bayesian optimization of machine learning algorithms. In F. Pereira, C. J. C. Burges, L. Bottou, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems 25*

(pp. 2951–2959). Curran Associates, Inc. http://papers.nips.cc/paper/4522-practical-bayesian-optimization-of-machine-learning-algorithms.pdf

Standing, L. (1973). Learning 10000 pictures. *Quarterly Journal of Experimental Psychology*, *25*(2), 207–222. https://doi.org/10.1080/14640747308400340

Tan, L., & Ward, G. (2000). A recency-based account of the primacy effect in free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*(6), 1589–1625. https://doi.org/10.1037/0278-7393.26.6.1589

Tan, L., Ward, G., Paulauskaite, L., & Markou, M. (2016). Beginning at the beginning: Recall order and the number of words to be recalled. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *42*(8), 1282.

Ward, G., & Tan, L. (2019). Control processes in short-term storage: Retrieval strategies in immediate recall depend upon the number of words to be recalled. *Memory & cognition*, *47*(4), 658–682.

Ward, G., Tan, L., & Grenfell-Essam, R. (2010). Examining the relationship between free recall and immediate serial recall: The effects of list length and output order. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *36*(5), 1207.

Yates, F. A. (1966). *Art of memory*. University of Chicago Press.

Appendix

Additional CMR simulations

The effect of CMR variants on the optimal policy. We explored how sensitive the optimal policy is to the specific variant of CMR that was used. For simplicity, the variant of CMR we used only recalls from the current list and does not incorporate intrusions from other lists in the experiment. It is possible that intrusions could reduce the efficacy of the chaining policy in the optimal strategy (by "breaking the chain"). To capture intrusion errors, we repeated the analysis on CMR2 where memories accumulate both within and across lists (Lohnas et al., 2015), enabling the model to account for recall of both current list items ("correct recalls") and prior list items ("intrusions"). In addition to intrusions, we also explored the effect of an alternative primacy mechanism: the attention-based account instantiated in TCM and CMR (Polyn et al., 2009b, Lohnas et al., 2015), which posits that there is increased attention during recall to the early list items, thereby increasing the encoding strengths of these items. This will provide a constant pull to the beginning of the list even after recall initiation, which can potentially affect the success of the chaining strategy.

To incorporate these effects, we fit CMR2 to the same dataset (Kahana, 2002); CMR2 successfully captures the amount of intrusions observed in empirical recalls (6.1% of total recalls in the data and 6.3% in the model). As with CMR, CMR2 parameters can be divided into two distinct categories: parameters that correspond to core architectural constraints of the memory system, and parameters that specify the space of possible policies that can be pursued (given these constraints) during encoding and recall. CMR2 adds some additional parameters to the former ("core architectural constraints") group: parameters that characterize the retrieval process ($\kappa = 0.31$, $\lambda = 0.13$), and parameters that prevent intrusions and recall repetitions ($c_{thresh} = 0.073$, $\omega = 11.89$, $\alpha = 0.68$) correspond to the core architectural constraints of the memory system. Differing from our previous analysis, we also considered the attention-based primacy mechanism to be part of

the constraints of the memory system ($\phi_d = 0.99$, $\phi_s = 1.81$). The latter ("policy specifying") group of parameters is the same for CMR2 and CMR; the fit values for these parameters were $\beta_{rec} = 0.63$, $\beta_{enc} = 0.52$, and $\gamma_{fc} = 0.43$. Repeating the rational analysis as before, we fix the first set of parameters to capture the architectural constraints during the retrieval process, and then optimize the parameter values in the second set that describe the space of encoding and recall policies to achieve maximum free recall performance.

We observed a similar optimal policy with large $\beta_{enc} = 0.94$, large $\beta_{rec} = 0.88$ and small $\gamma_{fc} = 0.17$ for CMR2 (which incorporates intrusions and attention-based primacy), compared with $\beta_{enc} = 0.99$, $\beta_{rec} = 0.99$, and $\gamma_{fc} = 0.001$ for CMR. There are also some differences, with $\beta_{rec} < 1$ and $\gamma_{fc} > 0$ in CMR2. Allowing intrusions does not qualitatively change the optimal policy, for two reasons: First, and most importantly, intrusions are rare; second, using a β_{rec} value that is < 1 (but still large) implements a kind of "soft chaining" that allows for the chain to sometimes be continued even after an intrusion. Specifically, when $\beta_{rec} < 1$, the current context is not entirely dominated by the retrieved context from the intrusion – there is still some context from the item recalled before the intrusion, which can sometimes cue recall of the correct next item (although, more often, the intrusion results in stoppage of recall, mirroring experiments results from Miller et al., 2012). We also found that there is no influence on the optimal policy from adding the attention-based primacy effect. Even though there is stronger encoding of the early list items, we do not observe a constant pull to the beginning of the list in recall patterns. The strengthening of early list items is negligible compared with the memory support to recall the next item in the chaining strategy.

The effect of alternative assumptions for the cognitive constraints on the optimal policy. To identify the optimal encoding and recall policy in the free recall task, we need to fix the parameters that characterize the cognitive constraints during the retrieval process. To obtain a set of reasonable values for these parameters, we fit CMR to all participants in an immediate free recall dataset (Kahana et al., 2002). To explore the

consequence that participants at high and low levels of performance demonstrate different cognitive constraints to start with, we fit CMR separately to the top half of participants and the bottom half of participants in Kahana (2002). The obtained fixed parameters are $\epsilon_d = 3.01$, $k = 5.47$ and $\epsilon_d = 2.11$, $k = 4.41$, respectively. In other words, the top half of participants are less likely to terminate recall, and have a lower noise level during retrieval, compared with the bottom half of participants.

We repeated the same analysis to optimize recall performance on the two groups of participants, given the differences in their cognitive constraints. We observe that splitting all participants into two performance groups when fitting CMR, i.e. assuming that better-performing participants and worse-performing participants are different in their cognitive constraints, does not change our main conclusions about the optimal policy.

For the top half of participants, when we allow the model to decide where to start recall, it is optimal to recall from the beginning-of-list context, with $\gamma_{fc} = 0.007$, $\beta_{enc} = 0.99$, and $\beta_{rec} = 0.98$. When recalling from the end-of-list context, under large $\beta_{enc} = 0.6$, the optimal parameters for $\gamma_{fc} = 0.98$ and $\beta_{rec} = 0.98$. For the bottom half of participants, when we allow the model decide where to start recall, it is optimal to recall from the beginning-of-list context, with $\gamma_{fc} = 0.004$, $\beta_{enc} = 1.00$, and $\beta_{rec} = 0.98$. When recalling from the end-of-list context, under large $\beta_{enc} = 0.6$, the optimal parameters are $\gamma_{fc} = 0.99$ and $\beta_{rec} = 1.00$.

Is there always greater consistency with the optimal policy among higher performing participants? Our finding that higher-performing participants show greater consistency with the optimal policy is not trivially true. A simple counterexample to this would be if the participants all follow random policies that demonstrate some variability, but none of these policies are close to the optimal policy. To demonstrate this, we simulated a population of participants who follow random policies in the part of the parameter space far from the optimal policy (for each simulated participant, their γ_{fc}, β_{enc}, β_{rec}, and the proportion of primacy trials are uniformly sampled from 0.3 to 0.6.) and

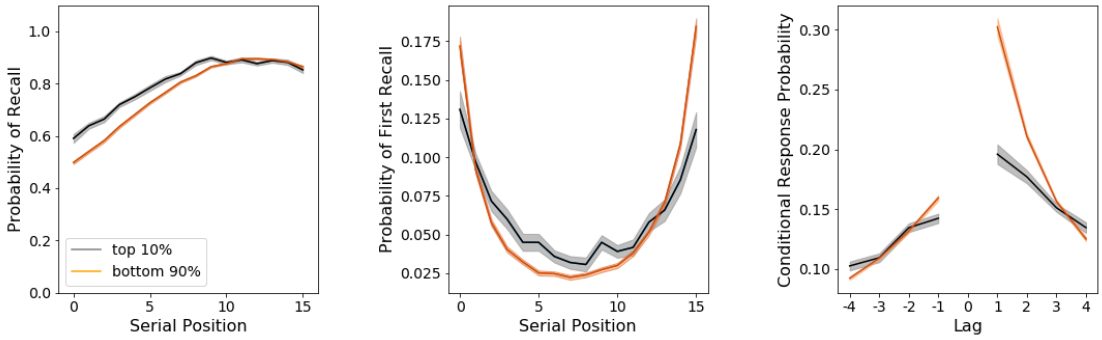

Figure A.1. Behavioral patterns for the better-performing and worse-performing simulated participants. 171 participants are simulated following random policies that are inconsistent with the optimal policy, with their γ_{fc}, β_{enc}, β_{rec}, and the proportion of primacy trials uniformly sampled from 0.3 to 0.6. "Top 10%" refers to simulated participants whose performance is in the top 10^{th} percentile of the population. From left to right of the three columns are the serial position curve, probability of the first recall, and conditional response probability. The fixed parameters in the simulations are $\epsilon_d = 4.60$, $k = 5.09$.

subjected them to the same analysis as Figure 6. If we looked at the high-performing participants in Figure A.1, they follow strategies (lower primacy and lower forward symmetry) that differ from the optimal policy (higher primacy and higher forward asymmetry). These simulation results do not look similar to the patterns we observe with human participants in Figure 6.

Simulations of qualitative patterns in Unsworth et al. (2011). We ran a "proof of concept" simulation to show that our model can account for the qualitative patterns shown in Unsworth et al. (2011). In Unsworth et al. (2011), participants were clustered into three groups according to their serial position curves: a high-primacy low-recency group, a low-primacy high-recency group, and a high-primacy high-recency group, with the last group demonstrating the highest memory performance. We simulate three groups of 50 participants. The three groups have a similar proportion of trials in which they start recalling from the beginning versus from the end of a list ($p = 0.5$), but

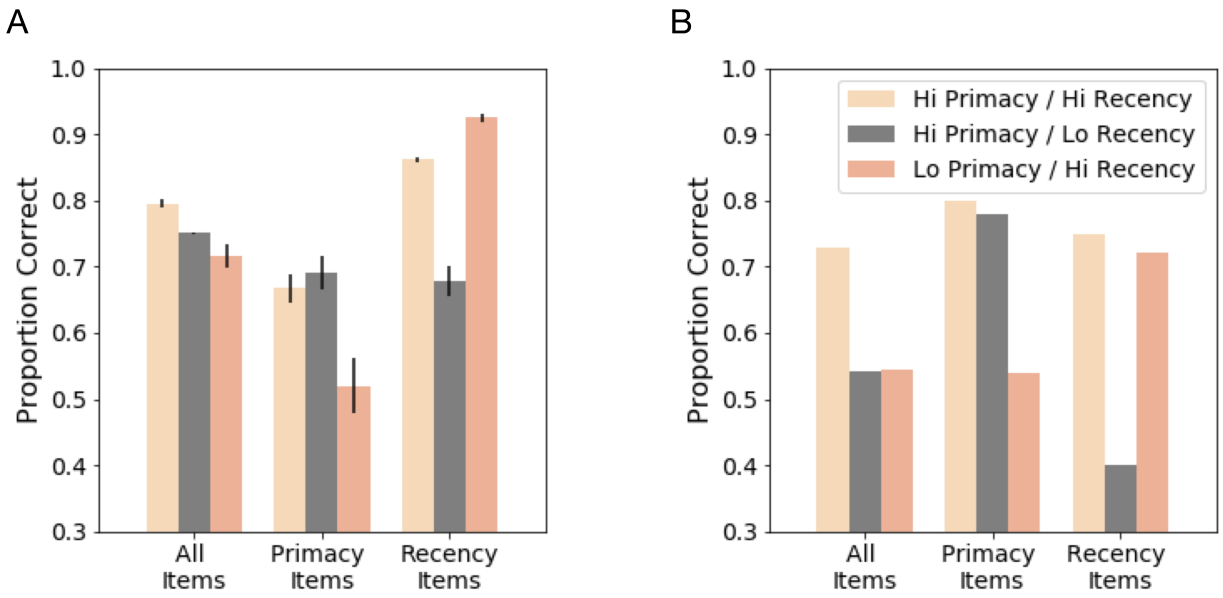

Figure A.2. Proportion correct for all items, primacy items (the first three items only) and recency items (the last three items only) for each group ($n = 50$), comparing current simulations (A) and Unsworth et al. (2011) (B). For our simulations, "Hi Primacy / Hi Recency" represents the high-primacy high-recency group that changes γ_{fc} based on where recall starts; "Hi Primacy / Lo Recency" represents the high-primacy low-recency group that fails to have optimal transitions when recall starts from the beginning (fixed $\gamma_{fc} = 1$); "Hi Recency / Lo Primacy" represents the low-primacy high-recency group that fails to have optimal transitions when recall starts from the end (fixed $\gamma_{fc} = 0$). The fixed parameters in the simulations are $\beta_{enc} = 0.6$, $\beta_{rec} = 1.0$, $\epsilon_d = 4.60$, $k = 5.09$. For the data shown in B, group labels (e.g., "Hi Primacy / Hi Recency") represent the groups identified in the clustering analysis conducted by Unsworth et al. (2011). Error bars represent one standard error of the mean. Values in B were estimated from Figure 2 and Table 1 in Unsworth et al. (2011).

they differ in how optimal their recall transitions are. The best-performing group is able to optimally adjust the amount of forward asymmetry given the recall initiation on a given trial ($\gamma_{fc} = 0$ when recall is initiated from the start of the list and $\gamma_{fc} = 1$ when recall is initiated from the end of the list). In contrast, the high-primacy low-recency group fails to increase forward asymmetry during the primacy trials ($\gamma_{fc} = 1$ in all cases), and the low-primacy high-recency group fails to increase backward transitions during the recency trials ($\gamma_{fc} = 0$ in all cases). Figure A.2A plots the proportion correct for all items, primacy items (the first three items only) and recency items (the last three items only) for each group; these results closely resemble the qualitative patterns from Unsworth et al. (2011), as shown in Figure A.2B: the "optimal policy" group has the best overall performance and demonstrates a high level of performance for both primacy and recency items compared with other groups, whereas the other two groups demonstrate high levels of performance for primacy or recency items only (but not both), and show lower levels of performance overall.

Parameter estimation in CMR and rational-CMR

CMR. There are three sets of behavioral patterns and a total of $N = 2l + 8$ data points we are interested in modeling: the serial position curve ($S = \{s_1, s_1, ..., s_l\}$), the probability of first recall curve ($R = \{r_1, r_1, ..., r_l\}$), and the conditional response probability curve with -4/+4 lag ($C = \{c_1, c_1, ..., c_8\}$). l is the list length of the free recall data. The goal in fitting CMR to a human behavioral dataset is to obtain a set of parameter values θ that minimize the difference between the simulated values s, r, c as a function of θ, and empirical values s', f', c' of these data points, expressed in the normalized root mean square error (NRMSE):

$$f(x) = \frac{\sqrt{\frac{\sum_{i=1}^{l}(s_i - s_i')^2}{l}}}{s'_{max} - s'_{min}} + \frac{\sqrt{\frac{\sum_{i=1}^{l}(r_i - r_i')^2}{l}}}{r'_{max} - r'_{min}} + \frac{\sqrt{\frac{\sum_{i=1}^{8}(c_i - c_i')^2}{l}}}{c'_{max} - c'_{min}}$$

where the normalization in each component rescales each set of behavioral patterns so that they are comparable across each other. Given some domain of the parameter space Θ, the

goal is to find the set of parameters θ that minimizes the error function $f(\theta)$:

$$\theta^* = \text{argmin}_{\theta \in \Theta} f(\theta).$$

The exact functional form for f is not available, but rather f arises as a complex function of the free recall behavior simulated based on mechanisms specified in CMR. Applying a grid search to the parameter space would not be feasible since it requires a total of m^{10} evaluations assuming there are m possible values for each parameter. Each evaluation is expensive, as it requires simulating a new dataset and calculating behavioral patterns s, r, c from it.

We apply Bayesian optimization to estimate the parameters in CMR (Mockus et al., 1978), as it has recently become popular for training expensive machine-learning models whose behavior depend in a complicated way on their parameters such as convolutional neural networks (Snoek et al., 2012). Bayesian optimization models f as a Gaussian process, with a prior $p(f) = \text{GP}(f; \mu, K)$. Given existing evaluations of f in $D = (\Theta, f)$, the posterior of f can be written as $p(f|D) = GP(f; \mu_{f|D}, K_{f|D})$. Then, a new set of parameters θ_{new} to evaluate can be proposed by estimating how desirable evaluating f at θ_{new} is expected to be for the minimization problem, measured as the probability of improvement function $E(u(\theta)|\theta, D)$, where $u(\theta) = 1$ if $f(\theta) \leq \min f$ or 0 if $f(\theta) > \min f$.

Parameters values reported in the current work are based on 200 iterations of Bayesian optimization, after 500 random initializations, which sums up to a total of 700 evaluations of f.

Rational-CMR. Instead of fitting the model to minimize the discrepancy with the human behavioral data, the goal in rational-CMR is to find the set of parameter values θ that maximize task performance f:

$$\theta^* = \text{argmax}_{\theta \in \Theta} f(\theta).$$

where task performance f is a function of θ, and can be obtained from the simulated dataset as the average proportion of items correctly recalled. θ is then obtained using the

same Bayesian optimization procedure that was described in the previous section.

Kahana et al. (2002) dataset

Fitting of CMR and rational-CMR is based on free recall data from 31 university undergraduates (age range: 17-21) in Experiment 1 from Kahana et al. (2002). Participants studied 33 lists of words in an immediate free recall experiment (the first three lists were practice and were not included in the analyses). Lists were composed of 10 common, two-syllable nouns chosen at random and without replacement from the Toronto Noun Pool (Friendly et al., 1982). At the start of each trial, a fixation cross appeared for 1400 ms, followed by a 100 ms blank inter-stimulus interval. Then, the computer screen displayed each list item for 1400 ms, followed by a 100 ms blank inter-stimulus interval. During list presentation, participants were instructed to say each word aloud. Immediately following the list, participants were given 30s to recall list items. It was made clear that they need not attempt to recall the items in the order of presentation. Each session lasted around 1.5h.

Penn Electrophysiology of Encoding and Retrieval Study

We tested the predictions from rational-CMR using the Penn Electrophysiology of Encoding and Retrieval Study (PEERS), a large database characterizing the electrophysiological correlates of memory encoding and retrieval; see full details of the experiment in Lohnas et al. (2015), Healey and Kahana (2016) and Healey et al. (2018). The present analysis is based on the behavioral data from the 171 young participants (age range: 18-30) who completed Experiment 1 of the PEERS dataset. Participants performed an immediate free recall experiment. Each participant completed one practice session and six subsequent experimental sessions, each with 16 lists. Only the experimental sessions are included in the analysis. Each study list was composed of 16 words, and was followed by an immediate free recall test. Words were either presented concurrently with a task cue (size judgement or animacy judgment), or with no encoding task. There were three conditions of

study lists: no-task lists, single-task lists and task-shift lists. List and task order were counterbalanced across sessions and participants. Each item was presented on the screen for 3000 ms, followed by a jittered inter-stimulus interval uniformly drawn from 800-1200 ms. If the word was associated with an encoding task, participants indicated their response via a keypress. Once the presentation of the last item was completed, after a jittered delay of 1200-1400 ms, the participant was given 75 seconds to recall any items from the just-presented list. Some sessions were randomly selected for final free recall task and a recognition task after the immediate free recall task. Only the immediate free recall task was considered in the analysis.

Howard and Kahana (1999) dataset

We tested the predictions from rational-CMR using the dataset from Howard and Kahana (1999). 61 participants studied lists of words for a subsequent free-recall test. Lists were composed of 12 common, two-syllable nouns chosen at random and without replacement from the Toronto Noun Pool (Friendly et al., 1982), and were presented visually at a rate of 1 word per second. During list presentation, participants were required to perform a semantic orienting task on the presented words. The participants were to press the left control key if they judged the word to be concrete and the right control key if they judged it to be abstract. In an immediate condition, the free-recall test was given immediately after list presentation. In a delayed condition, participants performed an arithmetic distractor task for 10 s before recall. Participants were given 45s to recall all items in the list in any order. Participants could not predict during encoding whether it would be an immediate or delayed condition; they were warned that sometimes there would be a math test between the end of the list and the signal to recall and sometimes there would not.