




Contents lists available at SciVerse ScienceDirect

Neuropsychologia

journal homepage: www.elsevier.com/locate/neuropsychologia



The neural dynamics of task context in free recall[☆]

Sean M. Polyn^{*}, James E. Kragel, Neal W. Morton, Joshua D. McCluey, Zachary D. Cohen

Department of Psychology, Vanderbilt University, United States

ARTICLE INFO

Article history:

Received 8 April 2011
Received in revised form 3 August 2011
Accepted 25 August 2011
Available online xxx

Keywords:

Organization
Computational modeling
Neural network
fMRI

ABSTRACT

Multivariate pattern analysis (MVPA) is a powerful tool for relating theories of cognitive function to the neural dynamics observed while people engage in cognitive tasks. Here, we use the Context Maintenance and Retrieval model of free recall (CMR; Polyn et al., 2009a) to interpret variability in the strength of task-specific patterns of distributed neural activity as participants study and recall lists of words. The CMR model describes how temporal and source-related (here, encoding task) information combine in a contextual representation that is responsible for guiding memory search. Each studied word in the free-recall paradigm is associated with one of two encoding tasks (size and animacy) that have distinct neural representations during encoding. We find evidence for the context retrieval hypothesis central to the CMR model: Task-specific patterns of neural activity are reactivated during memory search, as the participant recalls an item previously associated with a particular task. Furthermore, we find that the fidelity of these task representations during study is related to task-shifting, the serial position of the studied item, and variability in the magnitude of the recency effect across participants. The CMR model suggests that these effects may be related to a central parameter of the model that controls the rate that an internal contextual representation integrates information from the surrounding environment.

© 2011 Elsevier Ltd. All rights reserved.

1. Introduction

Multivariate pattern analysis (MVPA) techniques provide a natural framework for investigating the dynamics of topographic patterns of neural activity, which allow us to make contact with psychological theories that describe the dynamics of distributed representations of cognitive activity. These MVPA techniques, drawn from the domains of computer science and machine learning, allow one to classify a pattern as belonging to one of a number of categories (Duda, Hart, & Stork, 2001). Since Haxby et al. (2001) popularized these techniques by examining distributed patterns of activity in neuroimaging data using fMRI, researchers have used them to characterize the topography of the hemodynamic response in a wide variety of experimental paradigms (e.g., Haynes & Rees, 2005; Kamitani & Tong, 2005; Kriegeskorte, Goebel, & Bandettini, 2006; McDuff, Frankel, & Norman, 2009; Polyn, Natu, Cohen, & Norman, 2005).

However, these studies, while fueling theoretical debate on many levels (e.g., Cohen & Tong, 2001; Kanwisher & Yovel, 2006; Martin, 2007), have not utilized the formal computational machinery developed by cognitive researchers to characterize how the cognitive system might use distributed representations to carry out its most fundamental operations (e.g., Bower, 1972, chap. 5; Rumelhart & McClelland, 1986; Underwood, 1969). In particular, attribute-based theories of cognition, in which cognitive representations are modeled as vectors in a high-dimensional space, generally provide a wealth of predictions regarding the structure and dynamics of distributed patterns in the cognitive system. The attribute-based approach has been used in diverse domains within memory, including category learning (e.g., Nosofsky, 1986), recognition memory (e.g., Kahana & Sekuler, 2002; Murdock, 1982; Shiffrin & Steyvers, 1997), cued recall (e.g., Hintzman, 1984; Mensink & Raaijmakers, 1988), and free recall (e.g., Howard & Kahana, 2002a; Polyn, Norman, & Kahana, 2009a; Sederberg, Howard, & Kahana, 2008).

One class of attribute-based models has been quite successful in explaining the nuanced pattern of behavioral results observed in the free-recall paradigm: the *retrieved context* models of human memory (Howard & Kahana, 2002a; Howard, Wingfield, & Kahana, 2006; Polyn et al., 2009a; Sederberg et al., 2008). These models, described below, explain how an internal context representation is used to probe the contents of memory, and how information retrieved during memory search returns the context representation to a prior state, influencing the course of the subsequent memory

[☆] SMP and JEK wrote the manuscript, and JEK carried out the fMRI analyses. SMP and NWM developed the paradigm, with help from all authors, and all authors were involved in data collection. SMP and JDM carried out the simulations. ZDC and JDM carried out many of the behavioral analyses.

^{*} Corresponding author at: Vanderbilt University, PMB 407817, 2301 Vanderbilt Place, Nashville, TN 37240-7817, United States. Tel.: +1 615 322 2536; fax: +1 615 343 8449.

E-mail address: sean.polyn@vanderbilt.edu (S.M. Polyn).

URL: <http://memory.psy.vanderbilt.edu> (S.M. Polyn).

search. Relating these models to the neural systems that support memory search is a major goal of ongoing research (Polyn & Kahana, 2008).

Polyn et al. (2005) tested the core prediction of retrieved context models in a neuroimaging study of free recall in which the studied items were drawn from three neurally discriminable categories (celebrities, landmarks, and objects). They found that these category-related patterns reactivated during memory search, and that the rise in strength of a category-related pattern predicted upcoming recalls from that category. In terms of the MVPA approach, one can think of a set of categorized items, with neurally distinct representations, as a set of contrast dyes injected into the memory system as each item is studied. If one detects the neural representation of a particular category during the memory test, a trace of the dye is observed; the system has successfully revisited a past state. This is a powerful approach that has been used in a number of neurorecording studies of the human memory system. Researchers have found evidence for reinstatement of activity related to the modality of the studied material during recognition and cued recall (Nyberg, Habib, & Tulving, 2000; Wheeler, Petersen, & Buckner, 2000), the reinstatement of category-related information during cued recall (Kuhl, Rissman, Chun, & Wagner, 2011; Lewis-Peacock & Postle, 2008), and reinstatement of encoding task-related activity during a recognition task (Johnson & Rugg, 2007; Kahn, Davachi, & Wagner, 2004), as well as during a source memory task (McDuff et al., 2009; for a review see Danker & Anderson, 2010).

1.1. Computational models of retrieved context

Retrieved context models provide a framework with which to interpret the functional role of reactivated neural activity during memory search. A promising candidate framework is the Context Maintenance and Retrieval (CMR) model of free recall (Polyn et al., 2009a), which describes how the source context (reflecting the circumstances in which an item is encountered, as in Johnson, Hashtroudi, & Lindsay 1993) associated with studied items can be used to guide memory search for those items during free recall.

Retrieved context models of memory search, such as TCM (Howard & Kahana, 2002a), TCM-A (Sederberg et al., 2008), and CMR (Polyn et al., 2009a), consist of two interacting cognitive representations: a vector representation of the features of the study item, and a vector representation of the state of context. These two representations can influence one another through two sets of associative connections: item-to-context weights and context-to-item weights.¹

During the study period, the pattern of activity of the feature representation reflects the characteristics of the environment. This includes the physical characteristics of the study item, as well as the characteristics of the task cue that indicates to the participant which judgment task to use. Each task is represented by a distinct pattern of activity. As mentioned, these featural representations influence context through the item-to-context associations. The context elements have integrative machinery that allow them to maintain traces of their prior states: When new information related to the current state of the environment is projected along the item-to-context weights, the new state of context is a blend of the incoming information and the prior state of context.

This integrative machinery determines how much new information is pushed into the context representation as each item

is studied; a parameter of the model (β_{enc}) controls the rate of integration. In CMR, there are two subdivisions to the context representation: task context and temporal context.² The characteristics of the task cue influence the representation of task context and the characteristics of the studied item influence the representation of temporal context. The context integration parameter can be different for each of the context subdivisions: β_{enc}^{source} and β_{enc}^{temp} . A high value of β_{enc}^{source} will cause task context to update rapidly (i.e., when the model shifts from one task to another, the new task representation mostly pushes the old task representation out of task context). A high value of β_{enc}^{temp} will cause temporal context to update rapidly (i.e., when a new item is studied, the item representation mostly pushes the older item representations out of temporal context). Polyn et al. (2009a) assumed that these two parameters (β_{enc}^{source} and β_{enc}^{temp}) were independent, but did not assess the viability of a model variant in which the values for the two were related to one another (or controlled by a single parameter). Below, we explore the possibility that the values of β_{enc}^{source} and β_{enc}^{temp} are not independent in a particular person (i.e., if she has a high value for β_{enc}^{source} , she will also have a high value for β_{enc}^{temp}). This allows the model to account for an observed relationship between the discriminability of task representations (governed by task context), and the magnitude of the recency effect (governed by temporal context).

When an item is studied, an episodic memory is formed by linking the item features to the currently active pattern of contextual activity. This is accomplished by altering the item-to-context and the context-to-item weights in the network using a Hebbian learning rule (which increases the associative strength between co-active elements). This episodic association serves two major purposes. The associations formed on the context-to-item weights allow the context representation to serve as a retrieval cue: If a particular context representation is reactivated, it can now be used to revive the pattern of item features that co-occurred with that particular context state. The associations formed on the item-to-context weights allow the model to engage in contextual retrieval: If a particular item representation is reactivated, it can now be used to revive the pattern of contextual activity that prevailed when it was originally studied. This is the heart of the retrieved context framework; when a participant remembers a particular item, the memory system retrieves not only the features of that item, but also the context representation that prevailed when that item was originally encountered. This process, thought to be unique to human memory retrieval, has been described as mental time travel (Tulving, 2002).

These dynamics allow retrieved context models to explain the organizational phenomena observed during free recall. There are three major forms of organization: temporal, semantic, and source (Polyn et al., 2009a). We will describe temporal and source organization here (for a discussion of semantic organization, refer to Polyn, Erlichman, & Kahana, 2011; Polyn et al., 2009a). Temporal organization is exemplified by the *contiguity effect* (Kahana, Howard, & Polyn, 2008), the tendency for successively recalled items to come from neighboring serial positions on the study list. Temporal organization arises because the temporal context representation changes slowly over the course of the list. Thus, neighboring items are associated with similar states of the temporal context representation. When a particular item is recalled, the system retrieves the contextual state that prevailed when that item was originally studied. This retrieved context is a good cue for the items in neighboring list positions to the just-recalled item, leading to temporal organization.

¹ Here, we describe the dynamics of CMR, although nearly all of the non-source-related machinery is equivalent to that of the other models (TCM and TCM-A).

² The term *temporal* is not meant in terms of clock time, but rather in terms of list position; context is updated each time a new item is studied.

During memory search, items associated with the same source context also tend to be recalled successively. Recall organization is influenced both by external source characteristics such as modality, typeface, and gender of a speaker (Hintzman, Block, & Inskip, 1972; Murdock & Walker, 1969) and internal source characteristics such as encoding task (Polyn et al., 2009a; Polyn, Norman, & Kahana, 2009b). According to CMR, retrieved context contains information about the task used to process the studied item. This task information is used to guide the next recall attempt, so successively recalled items tend to be associated with the same study task. Thus, CMR predicts that task-related patterns of neural activity should be reactivated during memory search, to the extent that the task-related patterns reflect the operation of task context. The integrative machinery of the context representation leads to predictions about variability in the fidelity of task representations during the study period as well. Specifically, when the participant shifts from one task to another, the model predicts that residual activation related to the previous task will remain in the context representation. The predictions of CMR regarding the fidelity of neural task representations will be examined more thoroughly in a simulation experiment described below.

With this framework in mind, we turn to the results of a neuroimaging study of free recall, in which participants shifted between two encoding tasks (a size judgment and an animacy judgment) within-list. We find evidence for distinct neural representations of encoding tasks during the study period, and find that the fidelity of these representations varies both within participant and across participants in reliable ways. These task-specific patterns of neural activity are reactivated during recall, allowing a classifier to identify the encoding task that was used to study the remembered item. The results are discussed in terms of the retrieved context modeling framework.

2. Methods

2.1. Participants

Participants were 20 (12 female) native English speakers between 18 and 35 years of age. Consent was obtained in accordance with procedures approved by the Vanderbilt University Institutional Review Board. Participants were paid \$20/h for their participation, with up to an additional \$10 earned dependent upon performance in the task.

2.2. Experimental procedure

Participants were run in a variant of the free-recall paradigm, in which they studied a sequence of 12 lists (spanning two separate sessions), each of which contained 24 study items. Stimuli were presented with a computer running *PyEPL* (Geller, Schleifer, Sederberg, Jacobs, & Kahana, 2007). Each list was directly followed by either a free-recall test or a source recognition test. Each session included 3 free recall lists and 3 source recognition lists; the conditions were pseudo-randomly ordered within each session. Here, we focus primarily on the free-recall trials. Two different encoding tasks were used in the experiment. The *size task* involved judging whether an item was bigger or smaller than a shoebox, while the *animacy task* involved deciding whether an item was living or nonliving. Prior to stimulus presentation, a cue appeared on the screen for 0.7 s in order to indicate which task would be performed next. The size task cue was a line drawing of a shoebox; the animacy task cue was a line drawing of a heart. The line drawings were of similar size and contrast. After presentation of the task cue, there was a fixation interval of 0.3 ± 0.1 s, followed by presentation of the study item for 2.5 s, during which the participant made her judgment. After each stimulus presentation, an inter-stimulus fixation interval followed for between 0.5 and 5 s. The order of tasks and the durations of the inter-stimulus intervals were determined pseudo-randomly to optimize estimation of the hemodynamic response to each task (Dale, 1999), with the constraints that a task was not repeated more than 6 consecutive times, and that half of the items on a given list were studied with each of the two tasks.

After the final study item was presented on a given list, participants were given 75 s to freely recall as many items as they could remember from the most recent list, in any order. Vocal responses were recorded using a scanner-safe voice recording system (Resonance Technologies, Inc.). Scanner noise was removed from audio recordings using custom software (Cusack, Cumming, Bor, Norris, & Lyzenga, 2005)

and the *Noise Removal* tool in Audacity, and vocalized responses were scored offline using *PyParse* (Solway, Geller, Sederberg, & Kahana, 2010) and *Penn TotalRecall*.³

Each of the 20 participants performed 4 sessions of the free-recall paradigm. The first two sessions were performed in the laboratory while we recorded scalp electroencephalographic (EEG) activity, not reported here. These two sessions allowed participants to become familiar with the task, and allowed us to obtain stable estimates of behavioral measures, such as reaction time and response consistency. The two subsequent sessions were performed in the 3T Phillips MRI scanner at the Vanderbilt University Institute of Imaging Science (VUIIS). The 20 participants who were tested at VUIIS are a subset of 55 participants who were tested in the laboratory using this paradigm. Below, we report behavioral results from the full 55 participant study. The neuroimaging analyses in this report focus on the two experimental sessions in the scanner. The only paradigmatic difference between the laboratory sessions and the sessions carried out at VUIIS is the appearance of same-task trials (trials in which all items were studied using the same encoding task) in the laboratory sessions. These same-task trials were used to calculate a baseline for our measure of source organization, described below.

2.3. Word pool and construction of the study lists

Words were drawn from a word pool of 812 items. This is a subset of words from the USF free-association study (Nelson, McEvoy, & Schreiber, 2004) which were used in the Word Association Spaces study of the semantic meanings of words (Steyvers, Shiffrin, & Nelson, 2004). This subset was chosen based on the appropriateness of the two encoding tasks for judging the word, and based on the consistency of responses made by a separate set of 42 participants making the same semantic judgments on the words. If greater than 75% of the participants in this independent study made the same judgment on a particular word (e.g., *big* for *truck*), the word was eligible for inclusion.

A set of 24 words were randomly sampled without replacement from a larger pool to create a given study list. The task order (and inter-stimulus intervals) for a given list was determined by software designed to optimize estimation of the hemodynamic response (as mentioned above; Dale, 1999). Lists were inspected to ensure that the number of items associated with each response (i.e., *big*, *small*, *living*, *nonliving*) was balanced: if more than 70% of the items judged with a particular task were associated with the same response, the list was recreated. We also inspected WAS semantic similarity scores for the list items: if any pair of items had a similarity score ≥ 0.55 (measured using cosine distance), the list was recreated.

2.4. fMRI data acquisition

Functional images were collected using an EPI pulse sequence (TR = 2000 ms, TE = 30 ms, voxel size = 3.0 mm \times 3.0 mm \times 3.6 mm, flip angle = 75°, FOV = 192 mm). During the functional EPI scans, 30 oblique slices were collected over the whole brain, oriented parallel to the AC-PC plane. Whole brain MP-RAGE structural scans were collected (TR = 2500 ms, TE = 4.38 ms, voxel size = 1.0 mm \times 1.0 mm \times 1.0 mm, flip angle = 8°, FOV = 256 mm).

2.5. fMRI data analysis

2.5.1. Preprocessing

The first four volumes of each functional run were removed to allow for equilibration of scanner signal. Preprocessing of fMRI data was performed using routines as implemented by the SPM8 software package. All volumes from the twelve functional runs were realigned to the first functional volume of the first run, correcting for head motion. A mean functional image was generated from the realigned time-series and co-registered to the T1-weighted whole brain anatomical scan. These T1-weighted anatomical images were then segmented into gray matter, white matter, and cerebrospinal fluid, and normalized to a template in Montreal Neurological Institute (MNI) stereotactic space using the unified segmentation approach as implemented in the *New Segment* tool in SPM8 (Ashburner & Friston, 2005). Images were resampled to 3-mm isotropic voxels and spatially smoothed with a 8-mm FWHM Gaussian kernel.

2.5.2. MVPA methods

Pattern classification analyses were implemented in MATLAB using the Princeton MVPA Toolbox⁴ (Norman, Polyn, Detre, & Haxby, 2006). All fMRI time-series data were detrended, and z-scored within run. Classification was carried out using penalized logistic regression, using L_2 regularization. Regularization prevents overfitting by punishing large weights during classifier training (Duda et al., 2001). For all classification analyses, a penalty parameter of 50 was used.

2.5.2.1. The study period. In order to train a classifier to discriminate between task representations, time-points from the study period were labeled as belonging to

³ This software package is available at <http://memory.psych.upenn.edu/TotalRecall>.

⁴ This software package is available at <http://www.pni.princeton.edu/mvpa>.

either the size or animacy task, based on expected neural responses. For each session, the stimulus onsets were convolved with a canonical hemodynamic response function and thresholded to capture the peak of the response, creating binary labels for each of the two tasks. Time-points in which both or neither tasks were present were removed from the analysis. The classifier was trained on a portion of the fMRI data from a particular participant, and was tested on a separate subset from the same participant. Discriminability of the neural task representations was assessed during the study period using a 12-fold cross-validation scheme: The classifier was trained on all but one of the functional runs collected during study, and tested on the excluded run. This process was iterated twelve times, until every run served as the test run. Overall classifier performance was defined as the average performance across the twelve iterations.

A feature selection process was used to remove uninformative voxels from the classification. For each participant, a univariate ANOVA was run on each training set from the study period, testing the extent to which the signal of each voxel varied according to study task. Voxels were thresholded at $p < 0.05$. The mean number of task-sensitive voxels across participants and cross-validation folds was 6177, with a range from 1833 to 13,093 voxels. This feature selection process, coupled with L_2 regularization in the classifier, served to improve classification accuracy by removing noise sources from the data.

2.5.2.2. The recall period. In order to examine task-related activity during the retrieval portion of the experiment, the classifier was trained on all of the encoding period data, and then tested on the recall period data. Classifier performance during the recall period was evaluated by constructing an event-related average of classifier performance for the 4 s preceding and 10 s following each recall event (as in Fig. 2). Classifier performance was calculated in terms of the proportion of times the classifier correctly identified the encoding task associated with the recalled item. To prevent contamination from activity related to previous recalls, recall events were included in the analysis only if no recalls occurred in the preceding 10 s time period. After this inclusion criterion was applied, 15% of the total recall events remained for the event-related analysis (the range of events remaining for individual subjects was 3–35%). Many events were excluded from the initial portion of the recall period, during which participants tended to quickly recall a number of studied items in succession. This criterion also excludes items that were recalled quickly following another item; quick inter-response times have been associated with a strong temporal, semantic, or source relationship between the two items (e.g., Howard & Kahana, 2002b; Kahana, 1996; Polyn et al., 2009a). The classifier was trained on a set of voxels identified using a univariate ANOVA run on the study period data with study task identity as a factor (as above). Voxels were thresholded at $p < 0.05$. The mean number of included voxels from this procedure was 5225, with a range of 2924 to 11,400 voxels across subjects.

2.5.2.3. Classifier importance maps. Classifier importance maps were generated to determine the influence of each voxel on the classifier's estimate of each task, as in previous MVPA studies (McDuff et al., 2009; Rissman, Greely, & Wagner, 2010). After training, the classifier weights were inspected. The weight connecting a given voxel to a particular output unit is multiplied by the voxel's mean activity for the task associated with that output unit. In the case that the weight and the mean activity for a given task were both positive, a positive importance weight was assigned. If the weight and the mean activity were both negative, a negative importance weight was assigned. Voxels with different signs for the weight and activation values were assigned an importance value of zero. The group importance map was constructed by averaging individual importance maps from the study phase cross-validation procedure, across subjects. This method differs from the one used by Polyn et al. (2005), which only examines net positive or negative contributions to output units of the classifier, and is described in detail by McDuff et al. (2009, Supplementary Materials).

3. Results

3.1. Behavioral performance

During the study period, participants made a judgment on each of the 24 items. Performance was similar for the two encoding tasks on a number of measures, though participants were slightly faster and more accurate for the animacy judgments. Mean response time for the size task was 1.20 s (SEM 0.020), and was 1.13 s (SEM 0.022) for the animacy task; this difference of about 70 ms was significant across the full set of participants ($t(54) = 6.5$; $p < 0.001$).

We calculated a measure of response accuracy which quantified whether the responses made by the participant were generally consistent with the responses made by an independent set of participants making size and animacy judgments on the same words (see Section 2). Mean response accuracy was 0.90 (SEM 0.01); for

size judgments accuracy was 0.88 (SEM 0.01), and for animacy judgments response accuracy was 0.92 (SEM 0.01). This difference was significant ($t(54) = 2.84$; $p < 0.01$), but likely had more to do with there being slightly more ambiguity regarding the correct answer for a size judgment as compared to an animacy judgment, than a difference in how engaged the participants were for each task. Response accuracy was stable across list positions (for serial positions 1–4: 0.90 [SEM 0.01]; for serial positions 5–20: 0.90 [SEM 0.01]; for serial positions 21–24: 0.89 [SEM 0.01]). Finally, participants recalled a similar proportion of items studied with each task. Mean proportion recalled for the size task was 0.38 (SEM 0.02), and was 0.38 (SEM 0.01) for the animacy task; this difference was not significant ($t(54) = 0.21$; $p > 0.5$). For the analyses in the manuscript, we collapse results across the two encoding tasks.

A standard serial position curve was observed when probability of recall was calculated as a function of serial position. To summarize these results, we divided the studied items into three serial position groups (primary: 1–4, mid-list: 5–20, and terminal: 21–24), and separately calculated the proportion of items recalled from each of these groups. Proportion recalled was 0.42 (SEM 0.02) for the primary items, 0.33 (SEM 0.02) for the mid-list items, and 0.53 (SEM 0.01) for the terminal items. Both primary and terminal items were better recalled than the mid-list items (primary vs. mid-list, $t(54) = 5.98$, $p < 0.001$; terminal vs. mid-list, $t(54) = 14.6$, $p < 0.001$).

Participants showed reliable temporal and source organization of the studied material. The amount of temporal organization was measured using *temporal factor*, a percentile-based ranking measure (see Polyn et al., 2009a, for details). This measure ranges from 0 to 1; a score of 1 would indicate perfect temporal organization (all items recited in perfect forward or backward serial order), and a score of 0.5 would indicate no evidence for temporal organization. Mean temporal factor across the set of 55 participants was 0.59 (SEM 0.0085), which was significantly greater than a chance score of 0.5 ($t(54) = 10.6$; $p < 0.001$). To quantify source organization, we tabulated the proportion of recall transitions between same-task items, and compared this to the proportion of same-task transitions that one would expect by chance if the task labels were meaningless (this same technique was used by Polyn et al. (2009a), and is described in more detail there). The mean source factor across participants was 0.52 (SEM 0.010), which was significantly greater than the chance score of 0.48 (SEM 0.024) calculated from relabeled same-task trials ($t(54) = 4.12$; $p < 0.001$). The source organization score is modest in size (perfect source organization given this recall performance would be close to 0.9), but statistically quite reliable.

3.2. Distinct neural task representations during study

A logistic regression-based classification analysis identified distinct patterns of neural activity for the two encoding tasks (*size* and *animacy*) during study. The analysis is carried out on the individual participant level, allowing the classifier to characterize each participant's idiosyncratic pattern of task-related neural activity, and then identify the presence of these task-specific patterns on a left-out testing set. The mean proportion of correctly identified items was 0.61 (0.01 SEM); this was significantly above chance ($t(19) = 10.79$; $p < 0.001$). Classifier performance for the size items was 0.62 (SEM 0.01), and mean performance for the animacy items was 0.61 (SEM 0.01). Classifier performance was significantly above chance both for size items ($t(19) = 11.35$; $p < 0.0001$) and animacy items ($t(19) = 11.17$; $p < 0.0001$), and there was no difference in accuracy between the two tasks ($t(19) = 1.43$; $p > 0.05$). Classifier performance was reliably above chance for the 20 participants, with a range of mean performance of 0.53–0.70. We examined the classifier performance as a function of subsequent memory status of the studied item (whether or not it would be later recalled). We




Fig. 1. Average classification importance map projected onto the cortical surface. Warm colors denote regions associated with voxels that were assigned large weights towards making a *size* decision regarding encoding task identity. No regions were identified as important for making an *animacy* decision at the group level, despite the presence of *animacy*-related regions at the level of individual subjects. Voxels with an average weight of less than 0.0005 were removed from the map for illustrative purposes.

found that subsequent memory status did not significantly affect classifier performance. Mean performance for recalled items was 0.62 (0.01 SEM); mean performance for unrecalled items was 0.60 (0.01 SEM), a non-significant difference ($t(19)=2.03$; $p=0.056$).

3.3. Neuroanatomical regions informing the classifier

Fig. 1 depicts the set of voxels that most reliably influenced the classifier's decision across individual subjects. Despite idiosyncrasies in each participant's task patterns, a number of brain regions showed reliable task-related neural activity. Voxels whose activity is informative regarding encoding task are assigned large weights by the classifier. We averaged these weights across participants and projected them onto a model of the cortical surface (Fig. 1). A number of regions were identified as reliably more active for size judgments, including occipital and limbic sub-regions. While individual subjects showed reliable animacy-related signal, the anatomical location of this signal was quite varied across participants and as such does not appear on the group-level map. The classifier, being trained separately on each individual participant, was able to take advantage of these within-participant regularities.

3.4. Task-specific representations are reinstated during recall

At the heart of the Context Maintenance and Retrieval (CMR) model is a contextual representation that contains information about recently experienced events, and which returns to prior states during memory search (Polyn et al., 2009a). When studied items are associated with distinct source characteristics, the context retrieval mechanism will cause source context to be reactivated when an item is retrieved. This prediction was confirmed: Task-specific patterns of neural activity, identified during the study period, were reactivated during memory search. Fig. 2 shows an event-related average of classifier performance relative to the onset of the recall event. For each time bin, we calculated the proportion of events for which the classifier correctly identified the task identity of the remembered item. Task-related reactivation is concurrent with the vocalization of the recalled item, and the task pattern is sustained for several seconds after the recall event. The neural response depicted on this graph is uncorrected for lag in hemodynamic response (usually estimated as 4–6 s to peak response), raising the possibility that the onset of task-specific reactivation preceded the vocalization of the recalled item.




Fig. 2. Task-specific neural representations are reinstated during recall. This plot shows an event-related average of classifier performance centered on the vocalization time of each remembered item. Each point on the x-axis represents a whole-brain scan (each lasting 2 s). The y-axis shows the proportion of the time that the classifier correctly identified the encoding task associated with the item recalled at time 0. Points marked with a star show performance that is significantly greater than chance performance of 0.5 ($p < 0.005$, uncorrected); points marked with a plus are significant at $p < 0.05$, uncorrected. Error bars represent standard error of the mean.

3.5. Task context and neural patterns during study

The integrative mechanism supporting the context representation causes recently active task representations to be maintained, which means that when the participant shifts between tasks, the context representation contains residual activation related to the previous task. We compared classifier performance for task-shift items (preceded by the other task) and task-repeat items (preceded by the same task), and found no reliable difference in task discriminability for these two classes of items ($t(19)=0.70$; $p > 0.4$). However, we found a reliable decline in classification performance as a function of serial position of the studied item, using the same serial position bins that we used to characterize the probability of recall above. As shown in Fig. 3, classifier performance declined steadily over the course of the list (though classifier performance was well above chance for each of the three groups). Performance on the terminal positions was significantly lower than performance on the primary positions ($t(19)=3.2$; $p < 0.004$) and the mid-list positions ($t(19)=3.1$; $p < 0.006$). A follow-up analysis examined performance across free-recall trials within a session; we found that performance was stable across trials (Trial 1 mean and SEM: 0.62 [0.01], Trial 2: 0.61 [0.02], Trial 3: 0.61 [0.02], all differences not significant), suggesting that this is a within-list decline in discriminability of task representations that resets for successive lists.

We ran a simulation experiment with CMR to examine how task-shifting and serial position affect task discriminability in the model. This simulation confirmed that task representations get less discriminable across serial positions, but also revealed a large effect of task shifting on task discriminability, as well as an interaction between the two effects. For this simulation, we used the best-fit parameters reported in Polyn et al. (2009a), a recent simulation study of an experiment quite similar to this (Table 1). For the simulations we quantified task discriminability as the mean Euclidean distance between task-A and task-B representations (in the model) as a function of serial position group.⁵ Fig. 5a shows a decline in task

⁵ The task representations that the Euclidean distance calculation is being performed on are 2-dimensional. A city-block distance metric reveals the same pattern of results.

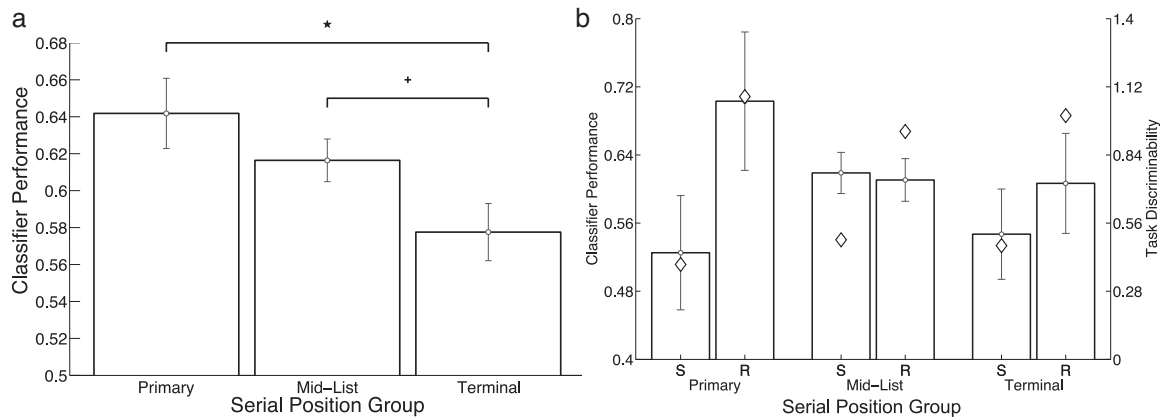


Fig. 3. (a) Discriminability of the neural task representations declines as a function of serial position. The study list was divided into three serial position groups: primary (positions 1–4), mid-list (positions 5–20), and terminal (positions 21–24). While classification performance for all three groups is significantly above chance, there is a significant decline in classifier performance between the primary and terminal positions, and the mid-list and terminal positions. Stars and pluses denote differences significant at $p < 0.005$ and $p < 0.05$, respectively. (b) The interaction of serial position group and shift (S)/repeat (R) status. Bars represent classifier performance on neural data, and diamonds represent discriminability of task representations in the model. Error bars are 95% confidence intervals based on within-subject error (Loftus & Masson, 1994).

discriminability across serial position groups. The rate of this decline is sensitive to a parameter β_{enc} (discussed above) that controls the rate at which task information is integrated into the context representation. This decline in task discriminability was reliably seen for a wide range of values of β_{enc} (i.e., $0.3 \leq \beta_{enc} \leq 0.9$, a range in which the model could produce reasonable behavioral results).

In Fig. 3b we examine the relationship between task-shifting and serial position more closely, for both neural data and simulated data. We divide each serial position group into shift and repeat items. An ANOVA examining classifier performance in terms of the two factors (shift/repeat status and serial position group) revealed a main effect of shift/repeat status ($F(1, 19) = 11.49; p < 0.005$), no main effect of serial position ($F(2, 19) = 1.21; p < 0.3$), and an interaction between the two factors ($F(2, 19) = 6.49; p < 0.005$). A similar analysis on the simulated data revealed a main effect of shift/repeat status ($F(1, 39) = 476; p < 0.005$), a main effect of serial position ($F(2, 39) = 61; p < 0.005$), and an interaction between

the two factors ($F(2, 39) = 1775; p < 0.005$). In the neural data, task-repeat items enjoy a great advantage in discriminability relative to task-shift items in the first group of serial positions ($t(19) = 3.61; p < 0.005$), but this advantage fades for the middle group of serial positions ($t(19) = -1.45; p > 0.1$). A small difference between task-shift and task-repeat items is observed for the final serial position group, but is of marginal significance considered on its own ($t(19) = 1.40; p > 0.1$).

In contrast to the pattern of neural results, the model predicts that some difference between task-shift and task-repeat items will be observed for all serial position groups. However, the model does predict that the difference in discriminability between task-shift and task-repeat items will reduce as the list progresses. The discriminability of the task context representations changes across list positions because the model begins the list with a pure representation of the task context associated with the first list item, but once the model shifts between encoding tasks, the task context representation becomes a blend of information related to each of the two task contexts (due to the integrative properties of context). According to the model, task-shift items should be less discriminable (relative to task-repeat items) in early serial positions, since they are being integrated into a representation that is purely representing the other task. In later serial positions, the context representation is already a blend of the two tasks (owing to several task shifts) and the difference in discriminability between task-shift and task-repeat items will be reduced.

While the absence of a shift/repeat effect for mid-list serial positions could be related to a lack of statistical power, other neurorecording studies of free-recall provide important guidance. In particular, Sederberg et al. (2006) showed that the neural response to a studied item changes as a function of list position, perhaps owing to across-list changes in the effectiveness of encoding (also Wiswede, Russeler, & Munte, 2007). The possibility that model dynamics change with respect to list position may provide insight into an aspect of the neural data not captured by the model: the steady decline in task discriminability across all serial positions. The model predicts that task discriminability will level off at mid-list serial positions (Fig. 5a). A follow-up simulation provides one possible mechanism by which the model could capture this decline: A variable β_{enc} parameter that decreases steadily across the list (representing perhaps a participant able to extract less information from each successive studied item). Fig. 5b shows this within-list decline in task discriminability. We return to this modified version of the model in the discussion.

Table 1
 Best-fit parameters of the Context Maintenance and Retrieval model.

| Parameter | Value | Description |
|------------------------|-------|--|
| β_{enc}^{temp} | 0.776 | Contextual integration rate for temporal information during study |
| β_{rec}^{temp} | 0.510 | Contextual integration rate for temporal information during recall |
| β_{enc}^{source} | 0.588 | Contextual integration rate for source information during study |
| β_{rec}^{source} | 0.588 | Contextual integration rate for source information during recall |
| L_{sw}^{CF} | 0.129 | Association strength between source-context and item features |
| d | 0.767 | Controls disruption of temporal context due to task-shift |
| γ^{FC} | 0.898 | Relative weight of expt. to pre-expt. context |
| s | 2.78 | Strength of semantic associations between studied items |
| κ | 0.111 | decay of activation in the decision competition |
| λ | 0.338 | Lateral inhibition in the decision competition |
| η | 0.159 | Magnitude of noise in the decision competition |
| τ | 0.174 | Time-constant for the decision competition |
| ϕ_s | 1.07 | Added associative strength for primary list items |
| ϕ_d | 0.981 | Decay of increased associative strength for primary items |

Note: The best-fit CMR parameters from Polyn et al. (2009a), used in the current simulations.

3.6. Classification performance correlates with the magnitude of the recency effect

Finding that task-specific patterns of neural activity are reinstated during the recall period raises the possibility that these neural patterns are functionally relevant to the memory search process itself. As such, the fidelity of these representations during the study period may allow us to predict how participants will behave during the memory test. We examined whether individual differences in the ability of the classifier to discriminate the two tasks during encoding were correlated with individual differences in performance at the behavioral level. We conducted a correlation analysis in which the classifier performance across participants was compared to a number of behavioral measures. We included measures of proportion of items recalled, temporal factor (as described above), source factor (as described above), the primacy effect (proportion of primary items recalled as compared to mid-list items), and the recency effect (the probability that a participant begins his or her recall sequence with the final item from the study list). We found a strong and significant positive relationship between the magnitude of the recency effect and the ability of the classifier to discriminate items studied using each of the two encoding tasks (Fig. 4a; $r = 0.63$, $p < 0.01$, Bonferroni corrected). The other comparisons were not significant. It is unclear, upon first consideration, why there should be a relationship between the fidelity of neural task representations and the magnitude of the behavioral recency effect. However, the CMR model provides a possible explanation.

As described in the introduction, CMR proposes that memory search is guided by an internal context representation which contains information about the temporal context and the task context of the studied material. Fig. 5d shows the effect of altering β_{enc}^{temp} on the recency effect (measured here as the probability that a participant begins recall with the final item from the study list). As β_{enc}^{temp} increases (the *high integration* simulation), the recency effect gets sharper (i.e., the simulated participant is more likely to begin recall with the final studied item). A high value for β_{enc}^{temp} means that when new information is projected from the feature representation to context (along the item-to-context associations), this incoming information forces out information relating to prior studied items (if β_{enc}^{temp} was set to 1, the new information would completely overwrite whatever information had been in context). The recall process is competitive: all items from the study list are fighting with one another to be the first item recalled. An item's effectiveness in the recall competition is directly related to how well the item's featural representation is associated with the current contextual cue. As reviewed above, a high value for β_{enc}^{temp} means that at the beginning of the recall period, the context representation is highly associated with the final item (much more so than with earlier items). Thus, the terminal list item competes more effectively when integration rate is high.

Examining the task representations in the model, we find that with higher values of β_{enc}^{source} , the task representations associated with the set of task-A and task-B items are more distinct from one another. This is shown in Fig. 5c, where we present the mean Euclidean distance between task context states for the set of A items, and for the set of B items. As the integration rate increases, the mean distance between the contextual states for the A items and the B items increases. This is simply because participants with a higher integration rate are able to more quickly clear out their context representations when a new task comes along. For participants with a low integration rate, task context is more of a blend between task A and task B representations. Comparing Fig. 5c and d, we see that there are two consequences to altering integration rate in the model: A change in the magnitude of the recency effect, and a change in the fidelity of task representations. Fig. 4b

shows this relationship graphically; altering β_{enc} reveals a strong correlation between recency magnitude and task discriminability ($r = 0.98$; $p < 0.0007$).

In Polyn et al. (2009a), the two context subdivisions were controlled by separate parameters, β_{enc}^{temp} and β_{enc}^{source} . However, the observed relationship between task discriminability and recency suggests that perhaps the two β_{enc} sub-parameters are yoked to one another: if a participant has high integration of temporal information, they also have high integration of task information. We return to this point later, in a broader discussion of individual differences.

The question arises as to why, of the 5 behavioral measures we examined (overall recall, temporal organization, source organization, recency effect, and primacy effect), only recency was related to the discriminability of neural task representations. If one varies β_{enc} in CMR, each of these 5 behavioral measures varies somewhat. However, the recency effect, measured as the likelihood that one's very first response is from the final serial position of the list, is unique among the 5 as being the one that is least sensitive to other CMR parameters that control recall dynamics. For example, consider the model parameter controlling contextual retrieval: β_{rec} . Varying this parameter causes large variability in each of these behavioral measures, except for the recency effect, which is completely unaffected by changes in β_{rec} . This is because contextual retrieval is a process that occurs after each item is recalled, and since the recency effect (as we measure it here) is based only on the first recall, it is unaffected. In other words, the relationship between recency and task discriminability may arise because recency is selectively sensitive to individual differences in contextual integration rate during encoding. In future work, we hope to obtain estimates of both β_{enc} and β_{rec} (as well as other model parameters) for each individual, which will allow us to assess the joint influence of these parameters on a wide range of behavioral measures.

4. Discussion

In the case of free recall of categorized materials (as in Polyn et al., 2005), it is difficult to determine whether reinstated category-specific neural activity reflects the category-specific characteristics of the retrieved materials themselves, or the contextual cue used to probe for those materials. However, in the current paradigm, the same study items could be presented in the context of either of the encoding tasks. The observation of reactivation of source-specific neural patterns is direct evidence in support of the *retrieved context hypothesis*: When a particular item is remembered, the memory system also retrieves the broader contextual details associated with the study event. The other analyses in this report characterize the dynamics of these source-specific neural patterns during the study period. We show that the discriminability of these task representations is related to both serial position and task-shifting, and that individual differences in the fidelity of these patterns (as measured by classifier performance during the study period) are related to the magnitude of the recency effect (as measured by the tendency to begin recall with the final list item).

4.1. The neural reactivation of source representations during memory search

Previous studies have shown reactivation of source-related patterns of neural activity during recognition (Johnson & Rugg, 2007), and during source recall (McDuff et al., 2009), but this is the first report (to our knowledge) of source-related reactivation during free recall. We seed the list with items associated with two well-matched but neurally discriminable study tasks. From a levels-of-processing standpoint, these two tasks are of equivalent depth,




Fig. 4. Examining the relationship between the magnitude of the recency effect and classifier performance on task identity during the study period. (a) A scatterplot of the relationship between these two variables; dashed line indicates the best-fit line of a regression analysis ($r=0.61$; $p<0.05$). (b) A simulation analysis using CMR examining the relationship between the recency effect and the discriminability of task representations (mean Euclidean distance between task representations for task-A and task-B items). Each point represents a simulation experiment sweeping the β_{enc} parameter from low to high values, while holding the other parameters at the best-fit levels reported in Table 1. Regression analysis revealed a similar correlation between the two measures ($r=0.98$; $p<0.007$) for the simulation.

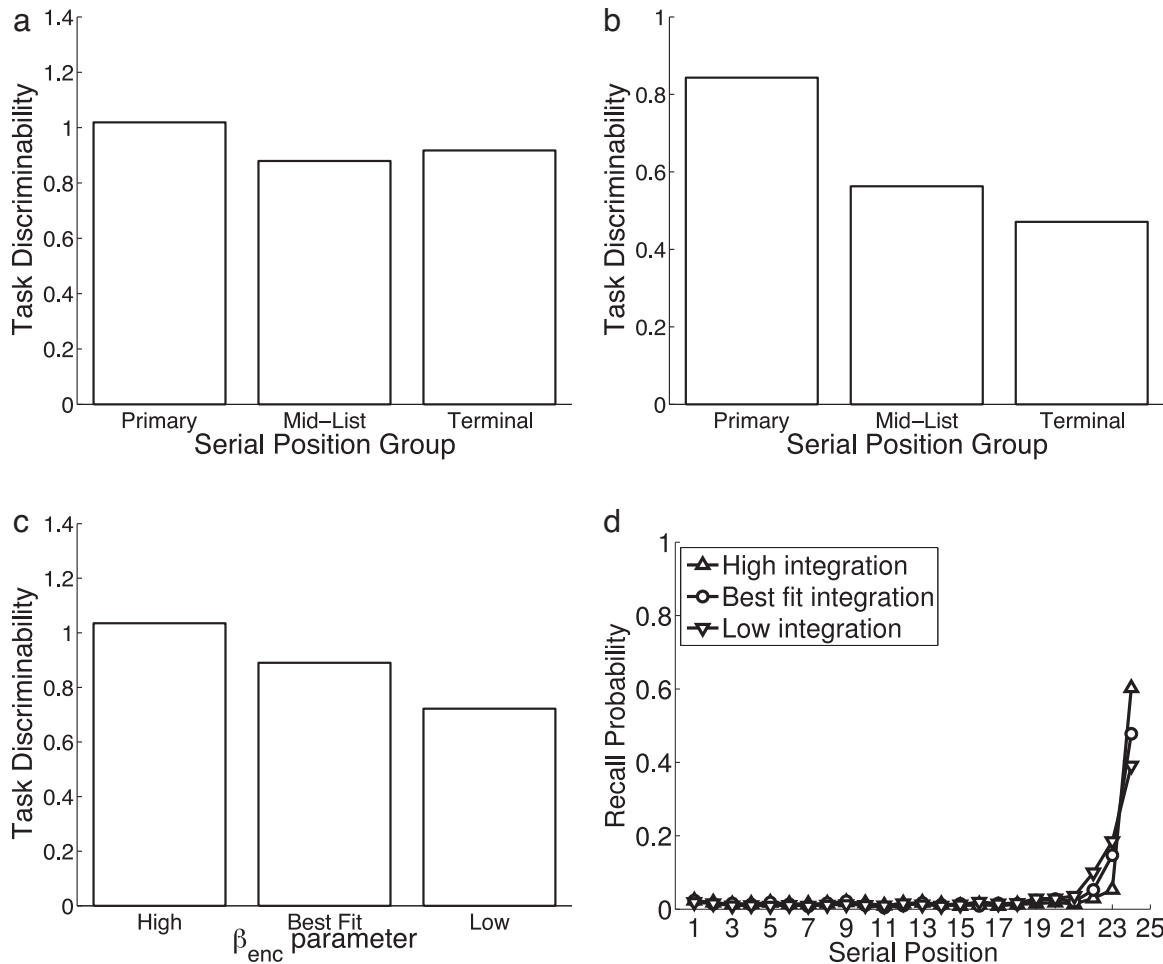


Fig. 5. Simulation analyses using CMR. (a) Simulation examining the decline in task discriminability with serial position group. (b) If β_{enc} declines steadily across serial positions (0.02 per position) then task discriminability decreases across all three serial position groups. (c) Simulations examining the internal representation of task context in CMR for three levels of the context integration parameter β_{enc} . Average Euclidean distance between the two task context states decreases as β_{enc} decreases. (d) Simulation showing how the magnitude of the recency effect (as measured by the probability of first recall by serial position) is also influenced by β_{enc} . The Best fit parameter set corresponds to the parameters reported in Polyn et al. (2009a) for a free-recall paradigm very similar to the one reported here. For High integration and Low integration simulations we manipulated the β_{enc} parameter (see text for details).

given the equivalent memorability of the items associated with each task (Craig & Lockhart, 1972).

We found that patterns of neural activity sensitive to the identity of the encoding tasks were reactivated while participants searched their memories for the studied material. A major open question regards the functional nature of this reactivated activity. One might suggest that we are observing reactivation of a task representation used by the executive system to ensure that a stimulus is processed appropriately (as in Cohen, Dunbar, & McClelland, 1990). It is also possible that this neural activity represents a reactivation of the cognitive procedures used while making the judgment itself (Kolers & Roediger, 1984). Future work can approach this problem both through further development of the computational model (to account for the dynamics of task performance as well as the dynamics of memory search) as well as further investigations of the neural signals underlying these effects (to determine other possible points of contact between these representations and behavioral performance).

4.2. The dynamics of source representations during the study period

The discriminability of task-related neural representations was affected by task-shifting and the serial position of the studied item, but unaffected by whether the studied item would later be recalled. This suggests that variability in classification performance was not necessarily related to lapses in attention on the part of the participant; the encoding task properly oriented the participant to the studied item. In line with this observation, we found that response accuracy did not change over the course of the list. Even though responses remained accurate, classifier performance declined significantly over the course of the list. Performance did not continually decline across multiple lists, rather, task discriminability seemed to reset itself with each new study list.

Our simulations using the CMR model provide explanations for the difference in discriminability between task-shift and task-repeat items, and the decline in discriminability with increasing serial position. As the list progresses, the dynamics of the context representation cause task representations to become less distinct. When the first item is studied, the representation for whichever encoding task is used is integrated into task context. Since this is the first list item, we assume there is no information related to the other task in the context representation, resulting in a relatively pure representation of the encoding task. As the list progresses, and the participant is asked to shift back and forth between the tasks, the context representation becomes blended; the integration process causes task context to represent a mix of the two task representations.

Further work is needed to understand some of the differences between the simulated data and the neural data. For example, our simulations suggest (Fig. 3b, diamonds) that the difference in discriminability between shift and repeat items (with repeat items having more discriminable task representations) should be maintained across all serial positions, whereas in the neural data this difference is only reliably exhibited for primary list positions. However, the model does predict that mid-list shift and repeat items will show less of a difference in task discriminability than primary shift and repeat items. Further work is needed to determine whether this difference between simulation and data is a true qualitative difference, or is simply due to a low signal-to-noise ratio in the neural data.

The simulations also suggest (Fig. 5a) that the decline in discriminability with serial position should asymptote by mid-list, whereas the neural data shows a continual decline in task discriminability with serial position (Fig. 3a). The literature suggests a number of ways to interpret a steadily declining ability to integrate

information about successively studied items. The first comes from the scalp electroencephalographic (EEG) literature. Sederberg et al. (2006) suggest, based on changes in oscillatory dynamics across the study list, that an initial burst of focused encoding (attending exclusively to the item on-screen) gives way to divided encoding (dividing cognitive resources between the current item and the many items that have been recently presented). This is related to the idea that participants sometimes think back to previous items as the list progresses (Brodie & Murdock, 1977; Rundus, 1971; Tan & Ward, 2000). An ever-increasing tendency to engage in rehearsal of previous items would cause these items to retrieve their associated task context, which would make the task context representation more of a blend of the two tasks. By this hypothesis, individual differences in rehearsal strategies could be related to the decline in task performance over serial positions. This rehearsal or divided-encoding hypothesis also provides a satisfactory account of why task discriminability resets between lists: If task discriminability declines because participants are thinking back to earlier list items, when they start a new list, there are very few items to think back to. The divided-encoding hypothesis can be instantiated in the model by steadily decreasing the context integration parameter over the course of the list. This represents a declining ability to effectively extract information from a stimulus presented in the environment, and allows the model to capture the continual decline in task discriminability across serial positions.

Characterizing rehearsal is particularly challenging given the covert nature of the process. Pattern classification techniques may provide some leverage in this domain. If a list was constructed such that participants studied a number of items drawn from one class, followed by a number of items from a second class, one could examine whether neural patterns related to the first set of items are detected in the latter part of the list (potentially indexing the degree of covert rehearsal). The CMR model has not been extended to capture the behavioral dynamics associated with rehearsal. However, such an approach may be straightforward, as a number of studies suggest that the dynamics of rehearsal mirror the dynamics of recall itself (Laming, 2006; Murdock & Metcalfe, 1978). Allowing the model to engage in rehearsal may simply involve activating memory search during the inter-item intervals as the list progresses.

4.3. Using CMR to explore individual differences at a neural and behavioral level

An advantage of having a formal computational framework to describe human memory is that the framework can be tuned to explain the differences in memory performance seen between individuals. In examining how the performance of the classifier (applied to the study period) was related to various behavioral measures, we found a strong correlation between the ability of the classifier to discriminate the two encoding tasks, and the magnitude of the recency effect. Simulation analyses showed that a single model parameter (β_{enc}) controlling the rate of context integration during encoding can explain the relationship between these two phenomena. The simulations described by Polyn et al. (2009a) did not consider the possibility that the two sub-parameters related to contextual integration during study (β_{enc}^{temp} and β_{enc}^{source}) were linked to one another, in part because this prior work examined aggregate behavior, averaged across participants. In this study, the observed relationship between the recency effect and neural task discriminability motivated examining a variant of CMR in which contextual integration for temporal information and for source information is related.

The recency effect, as measured by the tendency to start the recall sequence with the terminal list item, has been shown to correlate with working memory capacity (e.g., Unsworth & Engle,

2007; Fig. 5). Participants with a low working memory span have a sharp recency effect (a high probability of initiating recall with the terminal list item), and participants with a high working memory span have a more shallow recency effect. This provides an interesting way to interface the CMR model with the broader literature on individual differences in memory. By this story, the disadvantage exhibited by low-span individuals in memory tests is related to a high context integration rate; the context representation at any given time is focused on a rather narrow temporal interval, representing only the quite recent past. This hypothesis seems quite complementary to ongoing work characterizing individual differences in free recall and working memory (Unsworth & Engle, 2007).

Acknowledgments

The authors would like to thank Kenneth Norman, Michael Kahana, Per Sederberg, Tom Palmeri, and Gordon Logan for valuable discussions, Kristen McCabe for assistance in data collection and processing, Evan Stein for assistance with model evaluation, and the Kahana lab for general discussions and support of the behavioral analysis software.

References

- Ashburner, J. & Friston, K. J. (2005). Unified segmentation. *NeuroImage*, 26, 839–851.
- Bower, G. H. (1972). Stimulus-sampling theory of encoding variability. In A. W. Melton, & E. Martin (Eds.), *Coding processes in human memory* (pp. 85–121). New York: John Wiley and Sons.
- Brodie, D. A. & Murdock, B. B. (1977). Effects of presentation time on nominal and functional serial position curves in free recall. *Journal of Verbal Learning and Verbal Behavior*, 16, 185–200.
- Cohen, J. D., Dunbar, K. & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing model of the Stroop effect. *Psychological Review*, 97(3), 332–361.
- Cohen, J. D. & Tong, F. (2001). The face of controversy. *Science*, 293(5539), 2405–2407.
- Craik, F. I. M. & Lockhart, R. S. (1972). Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior*, 11, 671–684.
- Cusack, R., Cumming, N., Bor, D., Norris, D. & Lyzenga, J. (2005). Automated post-hoc noise cancellation tool for audio recordings acquired in an MRI scanner. *Human Brain Mapping*, 24(4), 299–304.
- Dale, A. M. (1999). Optimal experimental design for event-related fMRI. *Human Brain Mapping*, 8, 109–114.
- Danker, J. F. & Anderson, J. R. (2010). The ghosts of brain states past: Remembering reactivates the brain regions engaged during encoding. *Psychological Bulletin*, 136(1), 87–102.
- Duda, R. O., Hart, P. E. & Stork, D. G. (2001). *Pattern classification* (2nd ed.). New York: Wiley.
- Geller, A. S., Schleifer, I. K., Sederberg, P. B., Jacobs, J. & Kahana, M. J. (2007). PyEPL: A cross-platform experiment-programming library. *Behavior Research Methods*, 39(4), 950–958.
- Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L. & Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, 293, 2425–2429.
- Haynes, J. D. & Rees, G. (2005). Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nature Neuroscience*, 8(5), 686–691.
- Hintzman, D. L. (1984). MINERVA 2: A simulation of human memory. *Behavioral Research Methods, Instrumentation, and Computers*, 26, 96–101.
- Hintzman, D. L., Block, R. A. & Inskoop, N. R. (1972). Memory for mode of input. *Journal of Verbal Learning and Verbal Behavior*, 11, 741–749.
- Howard, M. W. & Kahana, M. J. (2002a). A distributed representation of temporal context. *Journal of Mathematical Psychology*, 46, 269–299.
- Howard, M. W. & Kahana, M. J. (2002b). When does semantic similarity help episodic retrieval? *Journal of Memory and Language*, 46, 85–98.
- Howard, M. W., Wingfield, A. & Kahana, M. J. (2006). Modeling the age-related associative deficit in self-initiated recall. *Psychonomic Bulletin & Review*, 13(3), 439–445.
- Johnson, J. D. & Rugg, M. D. (2007). Recollection and the reinstatement of encoding-related cortical activity. *Cerebral Cortex*, 17, 2507–2515.
- Johnson, M. K., Hashtroudi, S. & Lindsay, D. S. (1993). Source monitoring. *Psychological Bulletin*, 114(1), 3–28.
- Kahana, M. J. (1996). Associative retrieval processes in free recall. *Memory & Cognition*, 24, 103–109.
- Kahana, M. J., Howard, M. W. & Polyn, S. M. (2008). Associative retrieval processes in episodic memory. In H. L. Roediger III (Ed.), *Cognitive psychology of memory. Vol. 2 of Learning and memory: A comprehensive reference, 4 vols (J. Byrne, Ed.)* (pp. 467–490). Oxford: Elsevier.
- Kahana, M. J. & Sekuler, R. (2002). Recognizing spatial patterns: A noisy exemplar approach. *Vision Research*, 42, 2177–2192.
- Kahn, I., Davachi, L. & Wagner, A. D. (2004). Functional-neuroanatomic correlates of recollection: Implications for models of recognition memory. *Journal of Neuroscience*, 24(17), 4172–4180.
- Kamitani, Y. & Tong, F. (2005). Decoding the visual and subjective contents of the human brain. *Nature Neuroscience*, 8, 679–685.
- Kanwisher, N. & Yovel, G. (2006). The fusiform face area: A cortical region specialized for the perception of faces. *Philosophical Transactions of the Royal Society of London B*, 361, 2109–2128.
- Kolers, P. A. & Roediger, H. L. (1984). Procedures of mind. *Journal of Verbal Learning and Verbal Behavior*, 23, 425–449.
- Kriegeskorte, N., Goebel, R. & Bandettini, P. (2006). Information-based functional brain mapping. *Proceedings of the National Academy of Sciences of the United States of America*, 103(10), 3863–3868.
- Kuhl, B. A., Rissman, J., Chun, M. M. & Wagner, A. D. (2011). Fidelity of neural reactivation reveals competition between memories. *Proceedings of the National Academy of Sciences of the United States of America*, 108(14), 5903–5908.
- Laming, D. L. (2006). Predicting free recalls. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32(5), 1146–1163.
- Lewis-Peacock, J. A. & Postle, B. R. (2008). Temporary activation of long-term memory supports working memory. *Journal of Neuroscience*, 28(35), 8765–8771.
- Loftus, G. R. & Masson, M. E. J. (1994). Using confidence intervals in within-subject designs. *Psychonomic Bulletin & Review*, 1, 476–490.
- Martin, A. (2007). The representation of object concepts in the brain. *Annual Review of Psychology*, 58, 25–45.
- McDuff, S. G. R., Frankel, H. C. & Norman, K. A. (2009). Multivoxel pattern analysis reveals increased memory targeting and reduced use of retrieved details during single-agenda source monitoring. *The Journal of Neuroscience*, 29(2), 508–516.
- Mensink, G.-J. M. & Raaijmakers, J. G. W. (1988). A model for interference and forgetting. *Psychological Review*, 95, 434–455.
- Murdock, B. B. (1982). A theory for the storage and retrieval of item and associative information. *Psychological Review*, 89, 609–626.
- Murdock, B. B. & Metcalfe, J. (1978). Controlled rehearsal in single-trial free recall. *Journal of Verbal Learning and Verbal Behavior*, 17, 309–324.
- Murdock, B. B. & Walker, K. D. (1969). Modality effects in free recall. *Journal of Verbal Learning and Verbal Behavior*, 8, 665–676.
- Nelson, D. L., McEvoy, C. L. & Schreiber, T. A. (2004). The University of South Florida free association, rhyme, and word fragment norms. *Behavior Research Methods, Instruments and Computers*, 36(3), 402–407.
- Norman, K. A., Polyn, S. M., Detre, G. J. & Haxby, J. V. (2006). Beyond mind-reading: Multi-voxel pattern analysis of fMRI data. *Trends in Cognitive Sciences*, 10(9), 424–430.
- Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, 115, 39–57.
- Nyberg, L., Habib, R. & Tulving, E. (2000). Reactivation of encoding-related brain activity during memory retrieval. *Proceedings of the National Academy of Sciences of the United States of America*, 97, 11120–11124.
- Polyn, S. M., Erlichman, G. & Kahana, M. J. (2011). Semantic cuing and the scale-insensitivity of recency and contiguity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*.
- Polyn, S. M. & Kahana, M. J. (2008). Memory search and the neural representation of context. *Trends in Cognitive Sciences*, 12, 24–30.
- Polyn, S. M., Natu, V. S., Cohen, J. D. & Norman, K. A. (2005). Category-specific cortical activity precedes retrieval during memory search. *Science*, 310, 1963–1966.
- 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.
- Polyn, S. M., Norman, K. A. & Kahana, M. J. (2009b). Task context and organization in free recall. *Neuropsychologia*, 47, 2158–2163.
- Rissman, J., Greely, H. T. & Wagner, A. D. (2010). Detecting individual memories through the neural decoding of memory states and past experience. *Proceedings of the National Academy of Sciences of the United States of America*, 107(21), 9849–9854. <http://www.pnas.org/content/107/21/9849.abstract>
- Rumelhart, D. E., McClelland, J. L., & the PDP Research Group. (1986). *Parallel distributed processing*. MIT Press.
- Rundus, D. (1971). An analysis of rehearsal processes in free recall. *Journal of Experimental Psychology*, 89, 63–77.
- Sederberg, P. B., Gauthier, L. V., Terushkin, V., Miller, J. F., Barnathan, J. A. & Kahana, M. J. (2006). Oscillatory correlates of the primacy effect in episodic memory. *NeuroImage*, 32(3), 1422–1431.
- 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.
- Shiffrin, R. M. & Steyvers, M. (1997). A model for recognition memory: REM—Retrieving effectively from memory. *Psychonomic Bulletin and Review*, 4, 145–166.
- Solway, A., Geller, A. S., Sederberg, P. B. & Kahana, M. J. (2010). Pyparse: A semiautomated system for scoring spoken recall data. *Behavior Research Methods*, 42(1), 141–147.
- Steyvers, M., Shiffrin, R. M. & Nelson, D. L. (2004). Word association spaces for predicting semantic similarity effects in episodic memory. In A. F. Healy (Ed.), *Cognitive psychology and its applications: Festschrift in Honor of Lyle Bourne, Walter Kintsch, and Thomas Landauer* (pp. 237–249). Washington, DC: American Psychological Association.
- 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, 1589–1626.

- Tulving, E. (2002). Episodic memory: From mind to brain. *Annual Review of Psychology*, 53, 1–25.
- Underwood, B. J. (1969). Attributes of memory. *Psychological Review*, 76(6), 559–573.
- Unsworth, N. & Engle, R. W. (2007). The nature of individual differences in working memory capacity: Active maintenance in primary memory and controlled search from secondary memory. *Psychological Review*, 114(1), 104–132.
- Wheeler, M. E., Petersen, S. E. & Buckner, R. L. (2000). Memory's echo: Vivid remembering reactivates sensory-specific cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 97, 11125–11129.
- Wiswede, D., Russeler, J. & Munte, T. F. (2007). Serial position effects in free memory recall—An ERP-study. *Biological Psychology*, 75(2), 185–193.