Thesis

The When, Where and Why of Human Memory Retrieval

Qiong Zhang

June, 2019

Center for Neural Basis of Cognition
Machine Learning Departmant
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:
John R. Anderson, (Co-chair, CMU)
Robert E. Kass, (Co-chair, CMU)
Leila Wehbe (CMU)
Kenneth A. Norman, (Princeton)

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy jointly in Neural Computation and Machine Learning.

Abstract

Memory retrieval is fundamental in our daily experiences, whether it is to recognize a friend, to decide what to order from a menu or to navigate on the street. The process of memory retrieval, however, is latent and embedded among other cognitive processes such as perceptual encoding, decision making, and motor response. To track precisely *when* memory retrieval takes place, my research isolates individual cognitive processes from observed neural signal, by modeling the psychological activity in subjects' minds as a sequence of latent stages.

With precise timing of when memory retrieval occurs, I then examine *where* in the brain there is greater activity during the moment of memory retrieval. Developing a method that aligns neural recordings across subjects, and with better spatial resolution of an ECoG dataset, I provide a detailed mapping of the contributions of individual brain regions in a working memory task.

To further understand *why* memories are retrieved the way they are, I compare how well different cognitive mechanisms achieve the computational goal of a memory task. Principle of rationality posits that human cognition should adapt optimally to the task demands in the environment through learning and evolution. The more optimal cognitive mechanisms are more favorable to be used by human cognition. In a semantic fluency task, I demonstrate that an alternative memory search mechanism derived from reinforcement learning outperforms existing cognitive mechanisms both in their performance over simulations and in accounting for human behavioral data.

As a whole, my thesis work provides an integrated theory of human memory retrieval by uncovering its temporal dynamics, neural correlates, and underlying computational goal.

Acknowledgments

I am indebted to lots of people during the course my PhD. I cannot overstate how grateful I am to my advisor, John Anderson. I would not have found my course in studying human memory if I did not join his lab. In the past six years, he guided me not only on research projects but also on the way of doing research, which will have a lasting impact on my future academic career. I learned from him to associate rewards internally with the progressing of the work itself, how to get closer to the ground truth in science even in the absence of it, and the difficult balance to both follow our intuition and to abandon it in face of evidence. I also greatly thank my co-advisor Robert Kass, who genuinely cares about the personal well-being and growth of his students, and constantly guides us with a wealth of patience and experience. When I became over-focused on my projects, he reminded me of the importance in communicating research and always challenged me to consider the broader impact of the work. I also thank the rest of the thesis committee Leila Wehbe and Ken Norman for their valuable feedback and encouragement. Their questions and comments throughout the process were valuable to me in shaping the direction of the thesis.

The work presented in this document is the result of many fruitful collaborations. My foundations in associative memory were built up through four collaborative projects with Jelmer Borst. Marieke van Vugt has shaped my understanding in intracranial EEG data; it is also an inspiration to me the way she integrates her passion in Buddhism and dance into her research. I would also like to thank Matthew Walsh, for teaching me to collect EEG data, and getting me grounded in experimental work in addition to computational modeling. I also thank Ven Popov, with whom we push our discussions in class towards two exciting collaborative projects; and Marc Coutanche for his valuable guidance in these projects.

I would like to also thank Charles Kemp, Tom Mitchell, Lynn Reder, Valerie Ventura, Roy Maxion for their advice on teaching and research communication. Special thanks also go to Nicholas Cheadle and Clive Newstead in Eberly Teaching Center and Alex H at Global Communication Center for their constant help on my talks, classes, and writing.

I greatly appreciate the wonderful administrative support at CMU, Melissa Stupka at CNBC, Rebecca Finkel at Psychology, Diane Stidle at Machine Learning who always goes extra miles, and Judy van Rheenen at OIE for her advice and support during my most stressful visa times.

My PhD at CMU would not be complete without the help I have from outside CMU. I am extremely grateful for the wonderful people at Women in MathPsych, Pernille Hemmer, Marieke van Vugt, Leslie Blaha, Amy Criss, and Jennifer Trueblood, each and every of them has given me valuable career advice and influenced my decisions during job search. I also greatly thank Ken Norman, Mike Kahana, Tom Griffiths, Lynn Lohnas, Richard Shiffrin, Mike Jones, Jacob Feldman, the interaction with whom not only helped better frame my thesis in the past year, but also led to continuous discussions that is shaping my ongoing research trajectory.

A very special gratitude goes out to organizers and friends I met at MIT Brains, Minds & Machines

summer school. It was a turning point during my graduate school, when I truly see the importance of research interaction and being exposed to different trains of thought.

I am also very grateful for the internship opportunity at Facebook Oculus, where I got to appreciate how applications are not distractions but testbeds of theoretical work. I learned a great deal from my mentors James Hillis, Tanya Jonker and friends I met there: Lisa, Bas, Michael, Arun, Grace, Christof and Frank.

I am grateful for the stimulating discussions and peer support from my friends in neuroscience and psychology: Charles, Aria, Tina, Jayanth, Yuanning, Ven, Caitlin, Aryn, Jungaa, Jon, Dan, Jaeah, Roderick, Pierre, Elliot, Juliet, Robert; and from computer science and statistics: Ying, Yifei, Mariya, Lisa, Fan, Xun, Yuxiang, Su, Natalie, Zongge, Heejong, Yu, Lingxue, Mingda, Jining, Shashank, Sangwon, Bryan, Yufei. I would also like to thank Dimitry, Yan, Yuan, Koichiro, Michael, Yu, for the sleepless nights we were working together before course deadlines and the ping-pong breaks; Jisu, Kevin, Peter, Alex for being in a team of puzzle solving; Maya, Carmen and Santiago for bringing me back to dancing. Special thanks go to Aria, Charles and Marie who support me during the most down times in this journey, and my best friends Ye, Shujun, Menglu who are always a phone call away.

Last but not the least, I owe everything to my parents who support me unconditionally on anything I pursue, and my partner Nicolás who taught me to slow down and savor in this fast-moving world.

Contents

Chapter 1

Introduction

Memory retrieval is fundamental in our daily experiences, whether it is to recognize a friend, to decide what to order from a menu or to navigate on the street. The process of memory retrieval, however, is latent and embedded among other cognitive processes such as perceptual encoding, decision making, and motor response.

To track precisely when memory retrieval takes place, the challenge is to isolate the brief period of memory retrieval from other cognitive processes. Neural data such as EEG provides information in addition to behavioral data for tracking the cognitive processes. Conventional approaches in EEG literature, using event-related potentials (Luck, 2014), align the EEG signal to experimental events like stimulus onset or response commission and average the signal across trials. Under the assumption that there is no variability of event timing across multiple trials, the procedure of averaging can obscure the underlying neural signal (Luck, Woodman, & Vogel, 2000b). To account for the trial-to-trial variability, the first line of my thesis work provides a novel method that is capable of isolating individual cognitive processes by decomposing the task into a sequence of latent states (J. R. Anderson, Zhang, Borst, & Walsh, 2016a; Zhang, Walsh, & Anderson, 2017, 2018).

Once we know when memory retrieval takes place at the level of individual trials, we can locate better where in the brain it takes place around these timings. To achieve a better spatial resolution, my thesis work examines the ECoG activity, which gives finer spatial resolution in the cortical regions, and make it possible to examine subcortical activity such as that of medial temporal lobe (MTL), which plays an important role in (working) memory tasks (Ranganath, 2006; van Vugt, Schulze-Bonhage, Litt, Brandt, & Kahana,

2010). Combining the HSMM-MVPA method to recover the timing of memory retrieval, and a method to pool data across subject (Zhang, Walsh, & Anderson, 2017), a detailed mapping of brain regions are provided in a visual working memory task (Zhang, van Vugt, Borst, & Anderson, 2018).

Neural evidence of when and where memory retrieval takes place informs us about what the underlying cognitive mechanisms are. To further understand *why* memories are retrieved the way they are, I study a task of human memory search, where participants are asked to retrieve as many items as possible from a category in a fixed amount of time. Observed responses tend to be clustered semantically. A strategic search process under the marginal value theorem has been proposed, but it gives non-distinguishable patterns over human behavioral data compared with a non-strategic process. In the current work, I propose a rational analysis of the problem by examining what would be an optimal patch-switching policy under the framework of reinforcement learning. Built upon the random walk model and features of the local semantic patch, the resulted switching mechanism is more optimal than the marginal value theorem and better accounts for single-trial human behavioral data. These results provide theoretical justification of cognitive mechanisms used in human memory research, and shed light on how a rational account of the task can generate alternative hypotheses about human cognitive mechanism in the same task (Zhang & Anderson, 2018).

1.1 Temporal dynamics of memory retrieval

A longstanding interest in cognitive science is to identify the number and durations of different information processing stages involved in task performance (Donders, 1969). The challenge is to identify the number of stages, measure their durations, and understand how different experimental factors affect those durations. Sternberg (1969) proposed the additive factor method to deal with these challenges. The method entails the following assumptions: (1) time between a stimulus and response is occupied by a stream of successive processing stages, (2) each stage begins after the preceding stage ends, (3) different experimental factors affect the durations of different stages, and (4) the effects of different values of an experimental factor can be seen in overall RTs. While latency has the advan-

tage of an obvious relationship to the durations of underlying stages, it does not provide a direct measure of the individual stages but only of their cumulative time.

Multiple experimental factors can affect the duration and interacts at the same stage (McClelland, 1979). In addition, one experimental factor can affect the durations of multiple stages, and its impact on the duration of each stage cannot be easily identified. Worse yet, if a factor increases the duration of one stage while decreasing the duration of another, different levels of the factor may produce near-equivalent overall RTs. The cumulative nature of RTs limits conclusions that can be made about the effects of an experimental factor on the durations of individual processing stages.

To evaluate the consequence of such limitations, we examine a case where an experimental factor is hypothesized to impact two processing stages in an associative recognition task in opposing ways. The factor (i.e., probe similarity) is the degree of match between the three words of a triple presented to participants (i.e., probe) and the words in a previously studied triple. Participants were asked to decide if they had studied the triple before. Our theory of associative recognition predicts that similarity of the probe to a studied triple will decrease the duration of a retrieval stage while increasing the duration of a comparison stage. Consequently, the theory predicts small effects of similarity on overall RTs despite its larger effects on the durations of individual processing stages.

To overcome such limitations, we develope a non RT-based method that uses neuroimaging data gathered using EEG. The method involves applying hidden semi-Markov models and multivariate pattern analysis (HSMM- MVPA) to the EEG data to identify latent processing stages (J. R. Anderson & Fincham, 2014; J. R. Anderson et al., 2016a; J. R. Anderson, Zhang, Borst, & Walsh, 2016b). This is achieved by explicitly modeling the trial-to-trial variability of ERP components that would otherwise be distorted or lost in the average waveforms. The captured ERPs signify changes in the information processing and mark the transition from one processing stage to the next processing stage. We apply this HSMM-MVPA method to test our theory of how probe similarity affects associative recognition. In doing so, we advance understanding of associative recognition and demonstrate the utility of using the HSMM-MVPA method.

1.2 Neural correlates of memory retrieval

With EEG data in the previous study, we can map HSMM-MVPA stages to individual cognitive processes and identify when memory retrieval takes place, but poor spatial resolution in EEG limits our ability to determine how different brain regions are engaged during these periods. To achieve a better spatial resolution, in the next study, we look into the ECoG activity during the periods of interest. ECoG recordings from epileptic patients not only give finer spatial resolution in the cortical regions, but also make it possible to examine subcortical activity such as that of medial temporal lobe (MTL), which plays an important role in (working) memory tasks (Ranganath, 2006; van Vugt et al., 2010). Recent studies show that MTL is not uniquely involved in long-term memory, but also critical to short-term memory even when the retention period is as short as 2-10s (Hannula, Tranel, & Cohen, 2006; Holdstock, Gutnikov, Gaffan, & Mayes, 2000; Holdstock, Shaw, & Aggleton, 1995; Owen, Sahakian, Semple, Polkey, & Robbins, 1995; van Vugt et al., 2010). To further evaluate the role of MTL in short-term memory, we focus on a visual working memory task from a published dataset (van Vugt, Sekuler, Wilson, & Kahana, 2013). In each trial, participants first studied a list of faces, then, after a short delay, they were cued with a probe face and asked to judge if it had been among the just-studied faces (i.e., a Sternberg task).

Information can be preserved in working memory across a short delay without active maintenance (Lewis-Peacock, Drysdale, Oberauer, & Postle, 2012; Owen et al., 1995). Our interest is in the process by which this information is later retrieved. There are two questions asked in the current study. First, what is the time course of this retrieval process; second, what are the neural correlates during the identified time course of this retrieval process. The answer to the first question can be addressed similarly to the previous EEG study applying the HSMM-MVPA method. Note that knowing when memory retrieval takes place can lead to better identification of where in the brain it takes place. This is because there is trial-to-trial variability of the time courses of the retrieval process. Previous studies have focused on time windows of fixed length locked to the motor response as a way to locate when memory retrieval takes place. This approximation can attenuate effects that can be otherwise uncovered if we know exactly when memory retrieval takes place on a single trial basis.

There are two additional challenges when applying the HSMM-MVPA method to ECoG data, compared with the previous study when applying the same method to EEG data. First, the number of subjects in ECoG data is small given the limited opportunities in collecting such data among patient population. This has an impact on the analysis power when identifying how stage durations vary across different experimental conditions. This analysis is critical because it provides evidence on how latent stages identified in the HSMM-MVPA method map to specific cognitive processes including memory retrieval. To overcome the limitation of having too few subjects in the ECoG dataset, current study also analyzes a regular scalp EEG dataset of the same visual working memory experiment. Second, recording sites in ECoG data vary across subjects in both numbers and locations, we need a method of subject alignment before pooling data from different subjects. Multi-set canonical correlation analysis (M-CCA) is used for this purpose to transform electrode activity in each subject to a common neural representational space, where the inter-subject correlations of the transformed data are maximized across subjects. We have previously demonstrated the reliability of M-CCA in aligning subjects in MEG data (Zhang, Walsh, & Anderson, 2017). In the current work, we apply M-CCA to align both ECoG and EEG data from individual subjects. This allows us to compare the obtained common dimensions and the HSMM-MVPA results across two experiments in two different measurement modalities.

1.3 A rational analysis of memory retrieval during semantic search

A central paradigm to study human memory search is the semantic fluency task, where participants are asked to retrieve as many items as possible from a category (e.g. animals) in a fixed period. It is observed that responses tend to be clustered semantically (e.g. "cat"' follows "dog") (Hills, Jones, & Todd, 2012). Evidence over human behavioral data shows that marginal value theorem accounts for how our minds decide to switch from one cluster/patch to the next (Hills et al., 2012). Under this mechanism, people make strategic decisions to search the semantic space, similar to how animals optimally forage in a patchy spatial environment: one forages locally in one food patch, then switches to a new

patch when the resources in the current patch are depleted. It was observed that participants leave a patch in memory search when current rate of finding items is near the average rate for the entire task (Hills et al., 2012), consistent with what the marginal value theorem predicts in optimal foraging (Charnov, 1976). Recent work, however, demonstrated that similar behavioral patterns (i.e. that are consistent with the marginal value theorem) can e merge using a random walk simulation on a semantic network generated by human word-association experiments (Abbott, Austerweil, & Griffiths, 2015). The random walk model moves from one patch to another stochastically, without basing its decision on the information about the current patch. The fact that a strategic switching mechanism and a non-strategic switching mechanism predict similar temporal profiles around the switches poses challenges in understanding the exact cognition mechanism used by humans.

The goal of the current work is to further examine whether humans use the marginal value theorem during memory search. To provide more evidence in the comparison across alternative mechanisms, we consider the abstract computational problem posed by searching a semantic memory network and explore what would be an optimal strategy in this task (i.e. which strategy can generate the most items in a semantic fluency task). This approach is based on the principle of rationality, which explains human behavior as an optimal solution to the computational problems posed by our environment (J. Anderson, 1990; J. R. Anderson & Milson, 1989); see also bounded rationality in (Simon, 1978) and ecological rationality in (Todd & Gigerenzer, 2007)). Examining which cognitive mechanism better solves the computational problem gives additional justification on why it should be used by humans.

The principle of rationality can be applied not only to compare existing cognitive mechanisms, i.e. the marginal value theorem versus the stochastic random walk, but also to propose new hypothesis of alternative mechanisms. This is especially the case when existing mechanisms are not optimal for the given task. The marginal value theorem decides whether to switch by comparing an instantaneous reward to an overall average rewards; it is optimal under a set of conditions, including the assumption that local but not global resources are depleted during the course of the search (Charnov, 1976). This assumption does not hold in the semantic fluency task, as humans can re-enter the same patches, and they are depleted over time as more items are recalled. To perform well in this task re-

quires a decision policy that takes into account such dynamics in the environment. We propose an alternative mechanism based on reinforcement learning that directly optimizes performance under this specific task environment. In addition to comparing performance of different mechanisms in simulations, we also test how well they account patterns in the human behavioral data.

As a whole, my research addresses key aspects of the memory retrieval process by uncovering its temporal dynamics, neural correlates, and underlying computational goal. In the remaining thesis, they are organized as Chapter 2, Chapter 3 and Chapter 4 respectively. The last chapter consists of work proposed during the thesis proposal.

Chapter 2

Temporal dynamics of memory retrieval in associative memory

In this study, we investigated the information processing stages underlying associative recognition. We recorded EEG data while participants performed a task that involved deciding whether a probe word triple matched any previously studied triple. We varied the similarity between probes and studied triples. According to a model of associative recognition developed in the Adaptive Control of Thought-Rational cognitive architecture, probe similarity affects the duration of the retrieval stage: Retrieval is fastest when the probe is similar to a studied triple. This effect may be obscured, however, by the duration of the comparison stage, which is fastest when the probe is not similar to the retrieved triple. Owing to the opposing effects of probe similarity on retrieval and comparison, overall RTs provide little information about each stage's duration. As such, we evaluated the model using a novel approach that decomposes the EEG signal into a sequence of latent states and provides information about the durations of the underlying information processing stages. The approach uses a hidden semi-Markov model to identify brief sinusoidal peaks (called bumps) that mark the onsets of distinct cognitive stages. The analysis confirmed that probe type has opposite effects on retrieval and comparison stages.

2.1 Background

A longstanding interest in cognitive science is to identify the number and durations of different information processing stages involved in task performance (Donders, 1969, translation). The challenge is to identify the number of stages, measure their durations, and understand how different experimental factors affect those durations. Sternberg (1969) proposed the additive-factor method to deal with these challenges. The method entails the following assumptions: (1) time between a stimulus and response is occupied by a stream of successive processing stages; (2) each stage begins after the preceding stage ends; (3) different experimental factors affect the durations of different stages; and (4) the effects of different values of an experimental factor can be seen in overall reaction times. A limitation of the additive-factor method is that if one experimental factor affects the durations of multiple stages, its impact on the duration of each stage cannot be easily identified. Worse yet, if a factor increases the duration of one stage while decreasing the duration of another, different levels of the factor may produce near-equivalent overall reaction times. The cumulative nature of reaction times limits conclusions that can be made about the effects of an experimental factor on the durations of individual processing stages.

In this paper, we examine a case where an experimental factor is hypothesized to impact two processing stages in an associative recognition task in opposing ways. The factor (i.e. probe similarity) is the degree of match between the three words of a triple presented to participants (i.e. probe), and the words in a previously studied triple. Participants were asked to decide if they had studied the triple before. As described below, our theory of associative recognition predicts that similarity of the probe to a studied triple will decrease the duration of a retrieval stage while increasing the duration of a comparison stage. Consequently, the theory predicts small effects of similarity on overall reaction times despite its larger effects on the durations of individual processing stages.

To overcome limitations of overall RT, we apply a non RT-based method that uses neuroimaging data gathered using electroencephalography (EEG). The method involves applying hidden-semi Markov models and multivariate pattern analysis (HSMM-MVPA) to the EEG data to identify latent processing stages (Anderson, Zhang, Walsh, & Borst, 2016; Anderson & Fincham, 2014a,b). Information about the number and durations of processing stages based on the EEG can be used to evaluate the predictions of an existing

theory, or to guide the development of a new theory. We use the HSMM-MVPA method to test our theory of how probe similarity affects associative recognition. In doing so, we advance understanding of associative recognition and demonstrate the utility of using the HSMM-MVPA method.

2.1.1 Associative Recognition

Associative recognition involves judging whether two or more items were previously encountered together. For example, participants in our experiment decided whether three words in a triple had been studied together. According to one class of recall-to-reject models, such judgments would be made by retrieving a studied item from memory, and comparing it to the probe to determine whether they match (Anderson & Reder, 1999; Malmberg, 2008; Rotello & Heit, 2000; Rotello, Macmillan, & Van Tassel, 2000). One example of a recall-to-reject model is the process-level account based on Adaptive Control of Thought Rational (ACT-R; Anderson, 2007). The ACT-R model consists of four general processing stages (Anderson et al., 2016; Borst, Schneider, Walsh, & Anderson; 2013; Schneider & Anderson, 2012): encoding the word in the probe, retrieving a related memory, comparing the retrieved item to the probe, and responding.

According to the ACT-R's theory of declarative memory (Anderson, 2007), the time to retrieve a memory is an inverse function of its activation A_i:

$$T_i = Fe^{-A_i} \tag{2.1}$$

where F is a latency scaling parameter. The activation of an item is a sum of the item's inherent strength and its strength of associations with items present in the current context:

$$A_i = B_i + \sum_{j \in C} W_j S_{ji} \tag{2.2}$$

where B_i is its base-level activation, C is the context defined as the set of retrieval cues,W_j is the attentional weight assigned to each cue j, and S_{ji} is the strength of the

association between each cue j and item i. During associative recognition, words in the probe act as cues for retrieving a studied associate. The strength of association S_{ji} can be understood as the probability that cue j predicts item i. S_{ji} is expressed in ACT-R's associative strength equation:

$$S_{ji} = S - ln(fan_j) \tag{2.3}$$

S is a cue's maximum associative strength, and fan_j is the number of associates of cue j. In our experiment, we manipulated fan by varying the number of triples in which certain words appeared; the more triples a word appeared in, the less effective the word was as a retrieval cue. Many behavioral experiments have confirmed that RTs become longer as the number of associates, or fan, of a probe increases (e.g., Pirolli & Anderson, 1985; Anderson, 1974; for reviews, see Anderson, 2007; Anderson & Reder, 1999).

During associative recognition, ACT-R rejects foils (i.e. non-studied associates) by retrieving the closest matching studied associate, and determining that it does not perfectly match the probe. The retrieved associates in our experiment share one or more words with the probe. The time to retrieve the closest non-matching associate to a foil will be longer than the time to retrieve the matching associate to a target (i.e. a studied associate). This is because activation from the words in the foil spreads to different memories, whereas activation from the words in a target converges on a single memory. In other words, the number of sources spreading activation to the retrieved associate is greater for targets than for foils. This prediction, though straightforward, is difficult to test. The durations of other, non-retrieval related processes may also vary, obscuring the effects of number of sources of spreading activation. In particular, the duration of the subsequent comparison stage in the retrieve-to-reject model may be longer when the retrieved item is similar to the probe.

Neuroimaging studies provide additional evidence for ACT-R's activation and associative strength equations. In two such studies (Danker, Gunn, & Anderson, 2008; Sohn, et al., 2005), participants memorized word triples. They were then shown a set of probes and asked to decide whether they had studied the probes. The associative fans of words that made up the triples varied (Fan 1, 2, and 3). Consistent with behavioral studies, RTs

increased with associative fan. Additionally, the fMRI BOLD (Blood-oxygen-level dependent) response in the left prefrontal cortex, which is postulate to reflect ACT-R's retrieval module, increased with associative fan. This supports the idea that the most demanding retrievals engaged the left prefrontal cortex for the longest, producing the greatest BOLD response. Danker (2010) further found that foils produced greater LIPFC BOLD response, consistent with the proposal that they have the longest retrievals. While these results are consistent with the ACT-R theory of associative recognition, they cannot be strongly tied to a retrieval stage because of fMRI's low temporal resolution. EEG, unlike fMRI, has millisecond resolution, making it an attractive alternative for studying brief cognitive processes like those involved in associate recognition. EEG experiments of recognition memory have mainly focused on two ERP components: an early frontocentral negativity called the FN400 (Curran, 2000; Walsh, Paynter, Zhang, & Reder, 2016), and a later posterior positivity called the parietal old/new effect (Curran, 2000; Duzel, Yonelinas, Mangus, Heinze, & Tulving, 1997). These components have typically been interpreted in the context of dual-process theories of recollection memory (Diana, Reder, Arndt, & Park, 2006; Rugg & Curran, 2007; Yonelinas, 2002). The FN400 is thought to reflect a familiarity process that provides information about whether an item has been seen before, but does not involve retrieval of contextual information (Curran, 2000). The parietal old/new effect corresponds to a recollection process and does involve retrieval of contextual information. Given their role in recognition memory, one would think that these components would be informative with respect to models of associative recognition. This is especially true of the parietal old/new effect, which involves the retrieval of associative information. However, the conventional approach to isolating ERP components requires aligning the EEG signal to experiment events like stimulus onset or response commission, and averaging the signal across trials to create ERPs. However, differences in the onset latencies of components between trials and conditions can obscure or create the illusion of differences in ERP component amplitudes (Luck, 2005). This is especially problematic of components that are only weakly locked to overt experiment events.

2.1.2 HSMM-MVPA applied to an EEG data of Associative Recognition Task

To cope with the trial-by-trial variability in the onset latencies of ERP components, we developed a novel method that involves applying hidden semi-Markov models and multivariate pattern analysis (HSMM-MVPA) to EEG data (Anderson et al., 2016). Transitions from one cognitive stage to the next are signified by the onsets of bumps that summate with ongoing sinusoidal noise in the EEG signal. The bumps have finite durations, amplitudes, and topographical distributions. The postulation that processing stages are signaled by such bumps is consistent with theories of ERP generation in EEG data (Makeig et al., 2002; Yeung, Bogacz, Holroyd, & Cohen, 2004). Using HSMM-MVPA, it is possible to recover the number, timing, and topographical distributions of bumps that maximize the likelihood of the EEG data.

We previously applied HSMM-MVPA to an EEG study of associative recognition that manipulated fan and probe type (Anderson et al., 2016). Participants saw targets made up of two words previously studied together and foils made up of two words previously studied in separate pairs. Figure 2.1 shows the swimlane representation of the ACT-R retrieve-to-reject model for targets and foils in that experiment, with module activities initiated by production rules. To elaborate how the EEG signal is mapped to different processing stages within the framework of the ACT-R theory, a production evokes a change in neural processing, which produces a phasic response characterized by a bump in the EEG signal. The HSMM-MVPA method identifies both the distinctive scalp profiles and the variable latencies of such bumps in the single-trial EEG. Two early bumps marked the encoding of the word pair, and a third bump marked the retrieval onset. Following a variable period, a fourth bump marked the completion of retrieval and the start of comparison. A final, fifth bump marked completion of comparing the retrieved and target word pairs with a behavioral response. Several of these bumps related to ERP components evoked in standard recognition memory paradigms. In particular, during the time surrounding the fifth bump, voltages were more positive over the posterior scalp for targets than foils, corresponding to the parietal old/new effect.

2.1.3 Motivation and Overview of Current Experiment

Our model of associative recognition places the effect of probe type in the retrieval stage. However, the results from other studies suggest that the effects of probe type may extend to the comparison stage in a different and conflicting way. The more similar the retrieved item is to the probe, the longer the comparison process will be. Studies of perceptual decision making consistently show strong effects of foil similarity in response latency (for a review, see Farell, 1985). Foils that share more features with targets are rejected more slowly. Likewise, King and Anderson (1976) found an effect of foil similarity on response latencies in associative recognition.

According to ACT-R's activation equations, overlapping foils will result in greater activation for the retrieved memory and, consequently, shorter retrieval times. The finding that responses were slower for overlapping foils indicates that the effect of probe similarity extends beyond the retrieval stage. We suspect that similarity also influences the comparison stage. Interestingly, and problematically for measures of overall RT, the similarity of the probe to a studied item in memory may reduce the duration of the retrieval stage and increase the duration of the comparison stage. If we can isolate these stage durations using the HSMM-MVPA method, we can test two strong predictions about the effects of probe similarity:

i. Reflecting the effect of activation of the retrieved memory, the duration of the retrieval stage should decrease with the degree of match (i.e. probe similarity) between the probe and the studied pairs.

ii. Reflecting the number of comparisons needed to detect a mismatch between the retrieved memory and the probe, the duration of the comparison stage should increase with the degree of match between the probe and the studied pairs.

To test these predictions, we conducted an experiment with a study phase and a test phase. In the study phase, participants learned 24 word triples. In half of the triples, both the person and location had one associate (Fan 1), and in half of the triples, both the person and location had two associates (Fan 2). All verbs had three associates. In the test phase, during which EEG was recorded, participants completed an associative recognition task where they distinguished between word triples they had studied (targets) and those they had not (dissimilar foils, similar 1 foils, and similar 2 foils). By using word triples

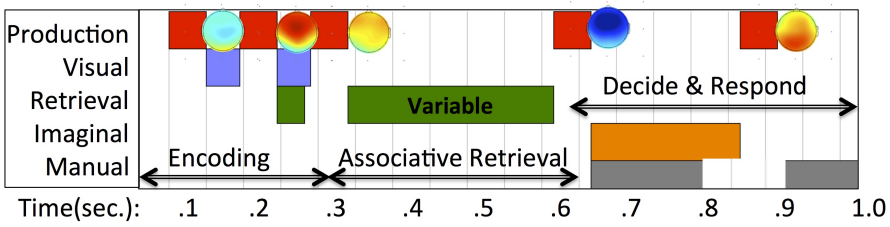

Figure 2.1: Swimlane representation of the ACT-R retrieve-to-reject model for targets and foils, with the scalp profiles identified from the HSMM-MVPA method.

instead of word pairs, we could examine the effects of multiple levels of probe similarity on the durations of the retrieval and comparison stages with the use of the HSMM-MVPA (Anderson et al., 2016).

2.2 Methods

2.2.1 Participants

20 individuals from the Carnegie Mellon University community participated in a single three-hour session for monetary compensation (14 men and 6 women, ages range from 19 to 35). All but one were right-handed. None reported a history of neurological impairment. One subject with less than 30% artifact-free EEG recording was excluded, leaving a total of 19 subjects.

2.2.2 Materials

Participants memorized word triples comprised of a person, a verb, and a location, shortened as "P-V-L" (e.g. Musician – Walk – Factory). The triples were created from lists of 18 people, 8 verbs and 18 locations (Appendix I.). Word length was 4 to 8 letters for people, 4 to 5 letters for verbs, and 4 to 7 letters for locations. Each participant memorized a list of 24 randomly generated word triples during the study phase of the experiment.

Foils were created by combining a person, verb, and location that had been studied in different triples from one another. Words were only swapped with other words that had the same number of associates in order to preserve the fan manipulation (e.g., Fan 2 foils

were created using people and locations from other studied Fan 2 triples). There were three types of foils (Table 1): (1) dissimiliar foils – None of the words in the triple had been studied together; (2) similar 1 foils – P-L had been studied together, but P-V and P-L had not; and (3) similar 2 foils – P-V and V-L had been studied together, but P-L had not. In total, participants were tested on 72 word triples (24 targets, 24 dissimiliar foils, 12 similar 1 foils, and 12 similar 2 foils). Half were Fan 1 and half were Fan 2 (Table 1). Foils were constructed such that every Fan 1 Person and Location appeared equally often, every Fan 2 Person and Location appeared equally often, and every Verb appeared equally often. Appendix II contains an example of 24 studies triples (targets) and a set of foils.

In addition to manipulating probe type and associative fan, we varied the vertical ordering of the words on the screen during the test phase. Half of the triples appeared in the studied order (P-V-L), and half appeared in a shuffled order (i.e. P-L-V; L-P-V; L-V-P; V-P-L; or V-L-P). Participants were instructed to decide whether the words had been studied together regardless of the ordering.

Table 1. Probe Types, Item Matches, and Probe Counts for Fan 1 and Fan 2 Triples

Probe Type	*Person-Verb Match?*	*Person-Location Match?*	*Verb-Location Match?*	*Fan 1*	*Fan 2*
Target	Yes	Yes	Yes	12	12
Dissimilar	No	No	No	12	12
Similar 1	No	Yes	No	6	6
Similar 2	Yes	No	Yes	6	6

2.2.3 Procedure

The experiment began with a study phase where participants learned a list of 24 word triples. When a triple was presented for the first time, it appeared at the center of the screen for 8000 ms. Participants were instructed to read and memorize the triple. After all triples appeared once, participants advanced into the drop-out portion of the study phase. In each trial, two of the three words in a triple appeared and participants were required to recall and type the omitted word. There was no time limit to respond. If the response was incorrect, the correct answer was displayed for 4000 ms, and if the answer was correct, the word CORRECT appeared for 4000 ms. If a triple elicited an error, it was repeated again after all other triples had appeared. A block of trials ended when all triples had elicited a

correct response. The drop-out portion consisted of three blocks of trials. A different word from each triple was omitted during each block. In the subsequent test phase, participants completed an associative recognition task where they distinguished between word triples they had studied and those they had not. Each trial began with a centrally presented fixation cross for a variable duration (sampled uniformly from 400 to 600 ms). Next, a probe word triple appeared vertically on the screen. Participants responded with a key press to indicate whether or not the triple had been studied. To respond "yes", they pressed the J key with the right index finger, and to respond "no", they pressed the K key with the right middle finger.

Participants were instructed to respond as quickly and accurately as possible. This was reinforced with a bonus point system. 100 points were deducted for every incorrect response, and 100-50*Time points were awarded for every correct response (if negative, the participant was given 0). Time was in seconds. After the participant responded, feedback appeared on the screen for 1000 ms. If the answer was incorrect, the word INCORRECT appeared along with the correct answer and the amount of points deducted. If the answer was correct, the word CORRECT appeared along with the amount of points awarded. Participants completed a total of 12 blocks of 72 trials. Each of the 72 word triples appeared once per test block.

2.2.4 EEG recording and analysis

Stimuli appeared on a 60 Hz LCD monitor set 60 cm from participants. The EEG was recorded from 128 Ag-AgCl sintered electrodes (10-20 system) using a Biosemi Active II System (BioSemi, Amsterdam, Netherlands). The EEG was re-referenced online to the combined common mode sense (CMS) and driven right leg (DRL) circuit. Electrodes were also placed on the right and left mastoids. Scalp recordings were algebraically re-referenced offline to the average of the right and left mastoids. The EEG and EOG signals were filtered with a bandpass of .1 to 70.0 Hz and were digitized at 512 Hz. The EEG recording was decomposed into independent components using the EEGLAB FastICA algorithm (Delorme & Makeig, 2004). Components associated with eye blinks were automatically identified and projected out of the EEG recording. Epochs of 1100 ms (including a 100 ms baseline) were then extracted from the continuous recording and corrected over

the prestimulus interval. Epochs containing voltages above +100 μV or below -100 μV were excluded ($< 4\%$ epochs).

EEG data were analyzed from trials with correct responses. Data were averaged across contiguous electrodes to create four regions: a left anterior/superior (LAS) region (FFC3h, F3, F1, FFC5h, FFC1h, FC3, and FC1), a right anterior/superior (RAS) region (FFC2h, F2, F4, FFC4h, FFC6h, FC2, and FC4), a left posterior/superior (LPS) region (CP3, CP1, CPP5h, CPP3h, CPP1h, P3, and P1), and a right posterior/superior (RPS) region (CP2, CP4, CPP2h, CPP4h, CPP6h, P2, and P4). We analyzed data from these regions during an early time window (300 to 500 ms) corresponding to the FN400, and a later time window (600 to 1000 ms) corresponding to the parietal old/new effect. Data from each time window were entered into a 4 (probe type) x 2 (fan) x 2 (laterality) x 2 (anterior/superior) repeated measures ANOVAs. For all analyses involving probe type (the only factor with more than two levels), we adjusted the p values using the Greenhouse-Geisser correction. We also analyzed response-locked waveforms which were baseline corrected using the 100 ms pre-stimulus interval before the triple appeared.

2.2.5 HSMM-MVPA applied to EEG

In our HSMM, we explicitly model the variability of endogenous ERP components that would otherwise be distorted or lost in the average waveforms. The HSMM-MVPA method identifies brief, distinctive profiles of scalp activity (i.e. bumps) with variable latencies in the single-trial EEG (Anderson et al., 2016). A bump is modeled as a half-sine multidimensional peak across the scalp that signifies a significant change in the information processing, followed by a flat period where the mean of the ongoing sinusoidal noise is 0. Our HSMM models the durations of the flats as gamma distributions.

Two steps of dimensionality reduction were carried out to simplify the analysis and make the computations more efficient and tractable. First, the data was down-sampled to 100 Hz (i.e., 10-ms samples). Second, to deal with the highly inter-correlated nature of the EEG sensors and to reduce the dimensionality of the signal, spatial PCA (i.e., across electrodes) was performed to generate orthogonal PCA dimensions. The first 10 PCA components were retained. These accounted for 69.2% of the variance in the signal. The PCA components were z-scored for each trial. As a result, the data for the analysis con-

sisted of 10 orthogonal PCA components sampled every 10 ms and with constant mean and variability across trials. Five samples (50 ms) beyond the response were also included in the analysis to ensure that the bump signifying the motor response was fully modeled, in the case that it occurred at the moment of trial completion. We only considered data from correct trials .

As described in more detail in our previous application of the HSMM-MVPA method (Anderson et al., 2016), several assumptions about the temporal structure of the signal are made to facilitate the analysis. First, the bumps were given a 50 ms width (i.e., 5 samples) with a half-sine shape. Such narrow bumps promote precision in the identification of stage boundaries, even if the bumps may actually be somewhat wider than 50 ms. Second, the analysis assumes that bumps do not overlap. Third, the flat durations are modeled as a gamma distribution with a fixed shape parameter 2 and a free scale parameter estimated to fit the data. See Anderson et al. (2016) for a detailed discussion of these assumptions, and tests of the robustness of the method against violations of each. Bumps in the HSMM are intended to account for the portion of the EEG signal corresponding to task-related processing; that is, variability arising from stimulus processing, memory retrieval and decision making, and response commission. Other sources of variability in the EEG signal that are not accounted for by bumps in the HSMM include noise from muscle movement, ambient electrical activity in the recording environment, stochasticity in neural responses to the same or related events, and additional neural processes unrelated to the task that take place in a non-stimulus locked fashion.

An n bump HSMM requires estimating n + 1 stage distributions to describe the durations of the flats plus the n 5-sample bumps for each PCA component. A different magnitude is estimated for each of the n bumps along each PCA dimension. A bump extends temporally across 5 samples (50 ms) and is multiplied by weights of 0.309, 0.809, 1.000, 0.809, and 0.309 (i.e., a 10 Hz half sin wave). The best model fit of such HSMMs is given by maximizing the summed log likelihood of the bumps and flats across all trials. For each trial, this log likelihood can be decomposed into two parts: the likelihood of the EEG data given that the bumps are centered at each time point, and the likelihoods that the bumps are centered at those time points given the gamma distributions that constrain their locations. In other words, the HSMM must select bump locations within a trial to

maximize the correspondence between the observed and the estimated EEG signal, while selecting relatively consistent flat durations across trials to maximize their fit to the gamma distributions. The estimation process has to consider all possible combinations of bump locations and this is what is efficiently calculated by the dynamic programming associated with hidden semi-Markov models (Yu, 2010).

The HSMM methods also return the probabilities of each bump occurring at each time point on a trial-by-trial basis. These probabilities can be used to calculate the most likely location of each bump in a trial, which is the sum of the time points in the trial multiplied by the corresponding probability that the bump occurred at that time. Mean stage durations for a particular subject can then be calculated as the average time between bumps across all trials within that subject.

2.3 Results

2.3.1 Behavioral Results

During the study phase, the number of times a triple was presented during each block before being correctly completed can be used to assess rate of learning (Table 2). With a repeated-measures ANOVA with fan and block as factors, mean frequency decreased across blocks, F(2, 36) = 30.382, $p < .001$. The frequency was higher for Fan 2 triples versus Fan 1 triples, F(1, 18) = 13.500, $p < .001$. The overall effect of fan decreased slightly across block, with an interaction between block and fan, F(2, 36) = 3.322, $p < .05$. By the third block, participants tended to respond correctly on their first attempt.

Table 2. Mean Frequency of Triples with *SEM*s in Parenthesis

	Block 1	*Block 2*	*Block 3*
Fan 1	2.34 (0.23)	1.58 (0.11)	1.25 (0.07)
Fan 2	2.66 (0.25)	1.62 (0.09)	1.39 (0.09)

During the test phase, data were trimmed by excluding trials with RTs shorter than 40 ms or longer than 3 sec (4% of all trials). The mean correct RTs and error rates appear in Table 3. RT and error rate were submitted to a repeated-measures ANOVA with fan and

probe type as factors. RT was longer and error rate was higher for Fan 2 triples versus Fan 1 triples, reflecting a main effect of fan on RT, F(1, 18) = 7.066, $p < .001$, and error rate, F(1, 18) = 29.561, $p < .001$. RT and error rate both varied with probe type (F(3, 54) = 18.571, $p < .001$ and F(3, 54) = 11.965, $p < .001$, respectively). Responses to dissimilar foils were fastest and most accurate. The effect of fan on performance was greatest for targets, reflecting an interaction between fan and probe type on RT, F(3, 54) = 5.683, $p < .01$, and error rate, F(3, 54) = 28.654, $p < .01$.

Table 3. Mean RT (in sec) and Error Rates (as a %) with *SEM*s in Parenthesis

	RT		*Error Rate*	
Probe Type	*Fan 1*	*Fan 2*	*Fan 1*	*Fan 2*
Target	1.26 (0.05)	1.42 (0.07)	13.3 (2.0)	27.9 (3.2)
Dissimilar	1.24 (0.06)	1.36 (0.07)	2.7 (0.1)	6.4 (1.2)
Similar 1	1.30 (0.06)	1.38 (0.08)	13.1 (2.7)	8.8 (1.4)
Similar 2	1.34 (0.06)	1.42 (0.07)	14.4 (2.2)	15.9 (1.8)

2.3.2 ERP Results

Stimulus Locked: 300 to 500 ms. Figure 2.2 shows ERPs from the four regions based on probe type and fan. We first analyzed the data from all four regions during 300 to 500 ms, the typical time window of the FN400 (Strozak, Abedzadeh, & Curran, 2016; Mollison & Curran, 2012; Speer & Curran, 2007). Consistent with the impression conveyed by the figures, there were no significant main effects or interactions involving the experimental factors during this window (Table 4).

Stimulus Locked: 600 to 1000 ms. We then analyzed the data from all four regions during 600 to 1000 ms, the typical time window of the parietal old/new effect. Waveforms appeared more positive for targets, reflecting a main effect of probe (Table 4). The main effect of fan was not significant, nor was the interaction between probe type and fan. The topographical distribution of the probe type effect is shown in Figure 2.3 .

Response Locked: -50 to 50 ms. Effects not otherwise observable in stimulus-locked waveforms can sometimes be seen in response-locked waveforms (Figure 2.4). A clear

effect of probe type along with an effect of fan emerged surrounding the response from -50 to 50 ms. This was reflected in main effects of fan and probe (Table 4). The effect of fan was smallest for targets, reflecting an interaction between fan and probe. Waveforms were most positive for fan 1 triples, and for targets. Figure 2.5 shows the topographical distribution.

Table 4. 4 (Probe) by 2 (Fan) Repeated Measures ANOVAs for Mean Voltages during Stimulus and Response Locked Intervals

	Stimulus Locked 300 to 500 msec	*Stimulus Locked 600 to 1000 msec*	*Response Locked −50 to 50 msec*
Probe: $F(3, 54)$	0.715, *ns*	8.626, $p < .001$	5.826, $p < .01$
Fan: $F(1, 18)$	0.375, *ns*	2.348, *ns*	7.531, $p < .05$
Probe × Fan: $F(3, 54)$	0.947, *ns*	0.757, *ns*	4.678, $p < .01$
Probe × Anterior/Posterior: $F(3, 54)$	0.456, *ns*	0.159, *ns*	4.430, $p < .05$
Probe × Laterality: $F(3, 54)$	1.361, *ns*	2.117, *ns*	3.862, $p < .05$
Fan × Anterior/Posterior: $F(1, 18)$	1.452, *ns*	0.757, *ns*	1.554, *ns*
Fan × Laterality: $F(1, 18)$	0.001, *ns*	0.179, *ns*	0.735, *ns*

In summary, we did not observe an FN400 in the stimulus-locked data. The time course, direction, and topographical distribution of the probe effect in the stimulus-locked data are consistent with the parietal old/new effect. The probe effect was also evident in the response-locked data, which is not surprising given that responses occurred shortly after the time window used to measure the parietal old/new effect. Finally, an effect of fan only appeared in the response-locked data. These results are consistent with our earlier study of paired associate recognition (Anderson et al., 2016). Now we turn from the traditional analysis of averaging over trials to an HSMM-MVPA analysis that parses each trial into its stages.

2.3.3 Identifying the Stage Durations and the Bump Profiles in HSMM-MVPA

HSMM-MVPA identifies bumps in the ongoing EEG signal related to significant changes in information processing. In this study, the number of stages in HSMM was decided jointly based on the EEG data and our theoretical understanding of the task model developed in ACT-R. Using LOOCV we confirmed that an HSMM with 4 bumps outperformed any model with fewer bumps. When we included more than four bumps, the HSMM

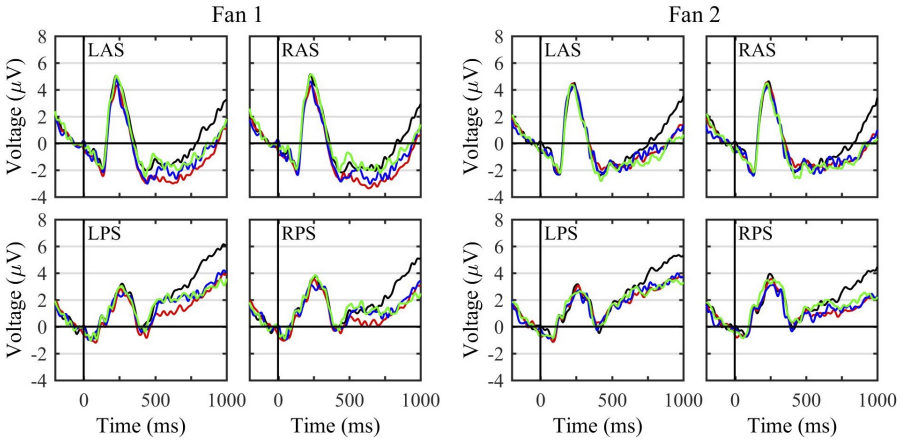

Figure 2.2: Stimulus locked waveforms from the four regions. Line colors correspond to target probes (black), dissimilar foils (red), similar 1 foils (blue), and similar 2 foils (red).

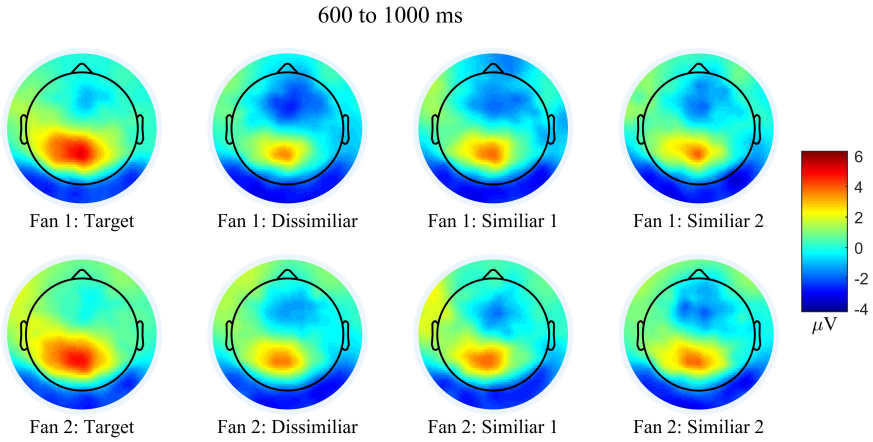

Figure 2.3: Mean voltage over scalp from 600 to 1000 ms after stimulus onset.

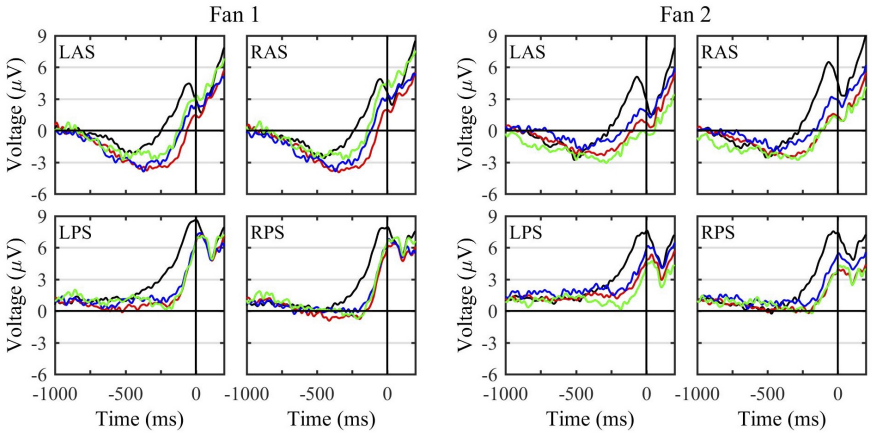

Figure 2.4: Response locked waveforms from the four regions. Line colors correspond to target probes (black), dissimiliar foils (red), similar 1 foils (blue), and similar 2 foils (red).

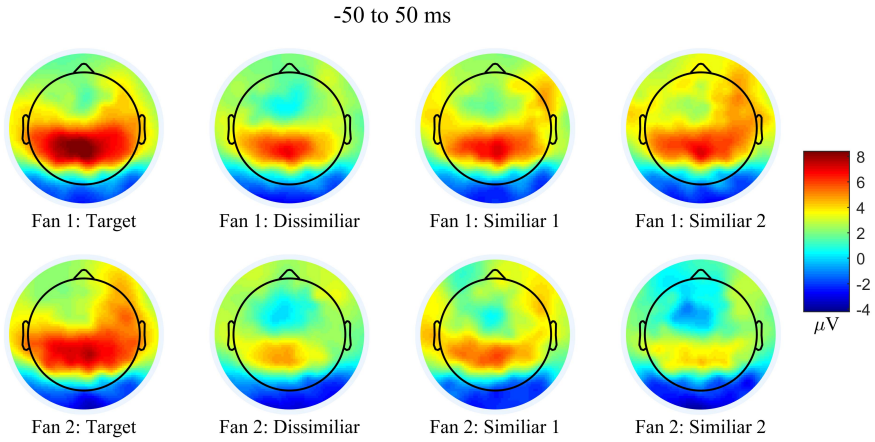

Figure 2.5: Mean voltage over scalp from -50 to 50 ms with respect to response.

placed the additional bumps adjacent to one another to capture a sustained positivity in the EEG preceding the response, as also observed by Anderson et al. (2016). Given the absence of a theoretical motivation for including more stages in the computational model, we stuck with the 4-bump model. Further, given that the current task involved associative recognition, as did the task reported in Anderson et al. (2016), their identification of stages motivated us to adopt a similar model here.

Based on the ACT-R model of the task, we obtained a total of four bumps and five stages in processing: 1. Pre-attention stage: The time from stimulus onset to the first bump reflecting the time for the visual signal to reach the brain and be attended to; 2. Encoding stage: The time from the first to the second bump reflecting the time to encode the stimulus; 3. Retrieval stage: The time from the second to the third bump reflecting the time to retrieve a memory for comparison; 4. Comparison stage: The time from the third to the fourth bump reflecting the comparison of the memory with the probe; 5. Response stage: The time from the fourth bump to the response reflecting response execution.

Figure 2.6 illustrates the durations of the five processing stages and the scalp topologies of the four bumps identified using the HSMM-MVPA method. Each of the bumps is modeled as a 50 ms half-sine multidimensional peak that can be projected back to the scalp given the known PCA projection weights. The Pre-attention, Encoding, and Response stages were relatively brief compared to the Retrieval and Comparison stages. The bump profiles and durations are similar to those in Anderson et al. (2016), with the exception of the comparison stage, which is considerably longer here. The longer comparison stage reflects the greater number of words that participants needed to compare and the demanding similarity manipulation used in this experiment.

Response times varied by condition (Table 3). These differences in response latency must show up in the durations of some of the stages. To determine which stages were affected by the experiment manipulations, we fit HSMMs with different stage durations to each condition. That is, we estimated scale parameters for the gamma-2 distributions of each stage separately for the different conditions while constraining the bump magnitudes to be the same. To focus on the effects of probe type, we restricted the effect of fan to the Retrieval stage (Stage 3) and investigated which stages were impacted by probe type. To do so we created an HSMM where the durations of all stages were allowed to vary with

probe, but only the duration of the third stage was allowed to vary with fan. Figure 2.7 shows the resulting estimated stage durations. We computed the mean stage durations for each subject as a function of probe type, and submitted them to a 4 (probe) repeated-measures ANOVA separately for each stage. Consistent with the impression conveyed by Figure 2.7 , the effect of probe type was significant during the third stage, F(3, 54) = 13.070, $p < .0001$, and the fourth stage, F(3, 54) = 21.720, $p < .0001$. We subsequently performed a 4 (probe) x 2 (stages three or four) repeated-measures ANOVA. The main effect of probe was significant, F(3, 54) = 7.489, $p < .001$, but the effect of stage was not, F(1, 18) = 3.949, $p < .1$. Most importantly, these two factors interacted, F(3, 54) = 48.417, $p < .0001$, owing to the opposing effects of probe on the durations of Stages 3 and 4.

The duration of Stage 3 was shortest for targets, and comparable for the three foil types. This is consistent with the prediction of the ACT-R model that retrieval time will be the shortest for targets, since the retrieved triple always receives two sources of spreading activation (Eq. 1 and Eq. 2). However, the ACT-R model also predicts the slowest retrieval time for dissimiliar foils (which receive one source of spreading activation), followed by similar 1 foils and similar 2 foils. Probe similarity also affected Stage 4. In contrast to the Stage 3, durations were longest for targets and similar 2 foils, followed by similar 1 foils, and finally by dissimilar foils. This is consistent with a model in which comparison terminates once a mismatch is detected.

Averaging across associative fan and probe type, response times decreases from the first to the second half of the experiment (1374 versus 1277 ms, t(18) = 3.19, $p < .01$). The ACT-R model predicts that practice should not affect the speed of encoding or motor planning, but rather the time to retrieve items from memory. Specifically, base-level activation (B_i in Eq. 2) increases with the number of repetitions. We compared the inferred stage durations from the first and second halves of the experiment in the HSMM. The results were partially in line with the model's predictions: the duration of the retrieval stage deceased from 527 to 498 ms, t(18) = 4.994, $p < .05$, but the duration of the comparison stage also decreased from 588 to 529 ms, t(18) = 13.219, $p < .01$.

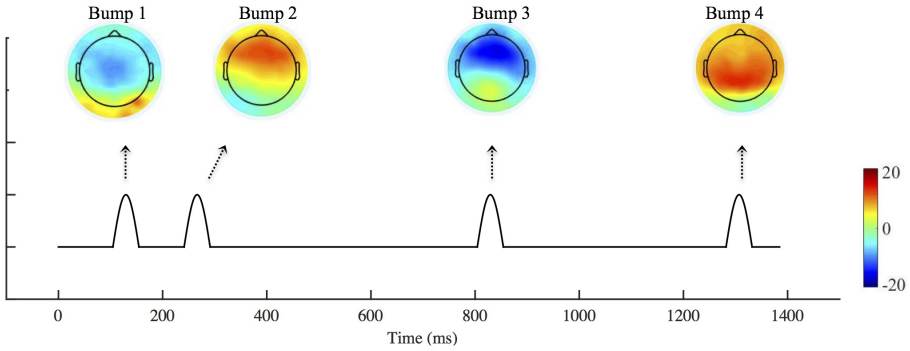

Figure 2.6: Mean electrode activity reconstructed for the four bumps by projecting the bump magnitudes back to the PCA weights (top). Mean durations of five stages interleaved by four 50-ms bumps (bottom).

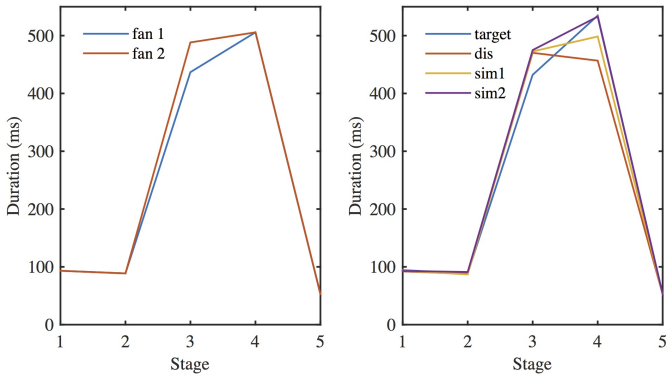

Figure 2.7: Condition-specific stage durations when localizing the effect of fan to Stage 3, and allowing probe type to vary across all stages.

2.3.4 Averaged Electrode Activity Anchored by Model Events

In standard ERP analyses, the EEG signal is anchored to observable events such as the presentation of a stimulus. The bumps obtained from the HSMM signal latent points of change. These events can be used, along with observable events, to align the EEG data. We anchored the EEG data from each trial according to stimulus onset, response, and the maximum likelihood locations of each of the four bumps during that trial. We then expanded or contracted the resulting five intervals in every trial of a condition to have durations equal to those specified by HSMM. In this way, the stimulus, the locations of the four bumps and the response are aligned across all trials in a condition. In the resulting displays (Figure 2.8 and 2.9), conditions with longer response times stretch further forward in the stimulus-locked waveforms.

In Figure 2.8 and 2.9, substantial differences among conditions are seen before and during the fourth bump over the parietal region. We performed a 4 (probe type) x 2 (fan) x 2 (flat/bump) repeated measures ANOVA using mean amplitude over the flat preceding the fourth bump and peak amplitude of the fourth bump. There is a significant interaction of the three factors (F(3,54)=6.072, $p = .002$). Then we performed separately a 4 (probe type) x 2 (fan) repeated measures ANOVA using mean amplitude over the flat or peak amplitude of the fourth bump. The main effect of probe was significant both during the preceding flat (F(3,54)=12.568, $p < .001$) and during the fourth bump (F(3,54)=6.593, $p = 0.002$), reflecting the greater positivity for targets versus foils. The main effect of fan on the peak amplitude of the fourth bump was also significant (F(1,18)=11.446, $p = 0.003$), owing to the greater positivity for Fan 1 triples versus Fan 2 triples.

2.3.5 Using the Neuroimaging Analysis to Inform the Task Model

Our model's retrieve-to-reject strategy consists of four stages: (1) Encoding–Randomly select two of the words to encode; (2) Retrieval–Retrieve a triple from memory that most closely matches the two encoded words; (3) Comparison–Compare the encoded and retrieved triples to determine whether any word differs; (4) Response–Press K ("not studied") if any word differs, and press J ("studied") otherwise. Response times equal the cumulative duration of the four stages. This can be represented as:

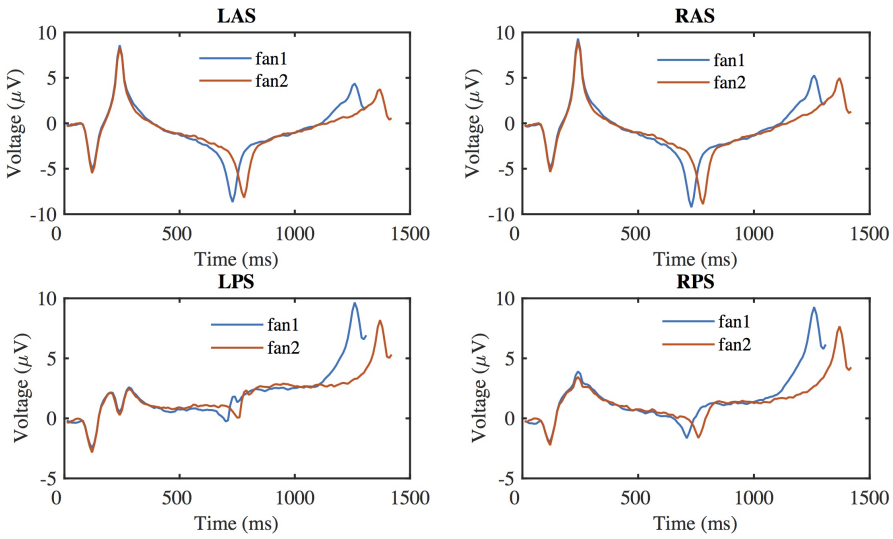

Figure 2.8: Average EEG data after warping every trial so that the maximum likelihood locations of the bumps correspond to the average locations for that condition. Data are shown in four regions (LAS, RPS, LPS, and RPS) for each fan condition.

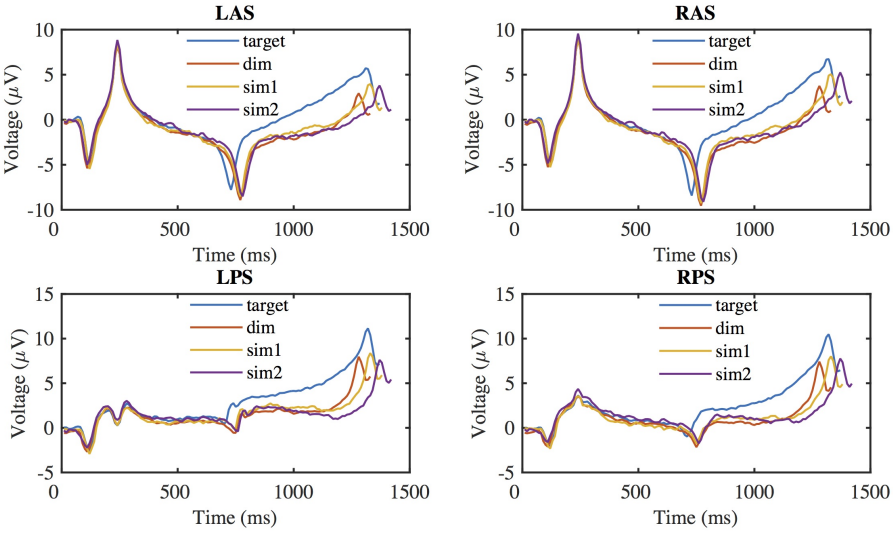

Figure 2.9: Average EEG data after warping every trial so that the maximum likelihood locations of the bumps correspond to the average locations for that condition. Data are shown in four regions (LAS, RPS, LPS, and RPS) for each probe condition.

$$RT_i = SharedTime + RetrievalTime_i + ComparisonTime_i \tag{2.4}$$

Shared Time includes the durations of the encoding and response stages. These encompass the time for the signal to reach the brain, the time to encode the words, and the time to program and perform a motor response.

Retrieval Time is computed from the set of equations that make up ACT-R's theory of declarative memory, and that are described in the introduction. In this task, retrieval time depends on the number of cues (i.e. encoded words) that spread activation to the retrieved triple (Eq 2.1). The number of sources of activation in targets always equaled two because all possible pairs of encoded words (P-V, P-L, or V-L) occurred together in a studied triple. The number of sources of activation for the retrieved triple in dissimiliar foils always equaled one because no two of the words occurred together in a studied triple. The number of sources for similar 1 foils equaled 1 with 66% (P-V or V-L encoded) and 2 with 33% (P-L encoded). Lastly, the number of sources for similar 2 foils equaled 1 with 33% (P-L encoded) and 2 with 66% (P-V or V-L encoded).

Retrieval Time also depended on the number of associates, or fan, of the encoded words in the probe (Eq 2.3). Person and location always had the same fan as one another (1 or 2), which varied by triple. Verb always had a fan of 3. Based on the spread of activation (Eq 2.2), and the degree of match (i.e. probe similarity) between current context and the retrieved triple (Eq 2.1), we computed activation for all combinations of two encoded words for the four probe types and the two fan conditions. Retrieval time for a condition was calculated as the average expected retrieval duration for all combinations of two encoded words (P-V, P-L or V-L encoded).

Comparison Time depended on the number of encoded words that matched words in the retrieved triple. We modeled the comparison stage as a serial process in which participant first compared the two encoded words to the corresponding words in the retrieved triple. If either word differed, the comparison process ended. If neither word differed, the participant compared the third word. The third word could conceivably be encoded before or during the final comparison.

The probability of comparing the third word varied by probe type. Targets always re-

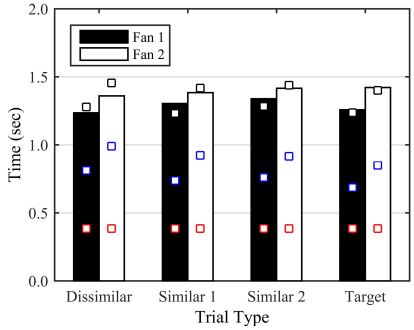

Figure 2.10: Observed response times (bars), and cumulative duration of model's shared time (red squares), retrieval time (blue squares), and comparison time (black squares).

quired comparing the third word because the first two words always matched the retrieved triple. Dissimilar foils never required the final comparison because the first two words never both matched the retrieved triple. Similar 1 foils required the final comparison 33% of the time (P-L initially encoded). Lastly, similar 2 foils required the final comparison 66% of the time (P-V or V-L initially encoded). Time to compare the third word (Final Comparison) was treated as a free parameter. Overall comparison time for a condition was calculated as the sum of a Comparison Intercept parameter, which was the same across conditions and accounted for the duration of comparing the first two words, and the Final Comparison weighted according to the probability that the final word was compared in each condition.

2.3.6 Model Fitting Procedure

The ACT-R model contained a total of four free parameters: Shared Time, F (latency scale for mapping activation onto retrieval time), Final Comparison, and Comparison Intercept. Maximum associative strength S (Eq 2.3) was set to a default value of 1.5 , and attentional weight W (Eq 2.2) was set to 1/2 because context was defined by the two encoded words. We estimated the free model parameters to maximize the correspondence between the durations of the model stages (i.e. Retrieval and Comparison) and the stages inferred from the HSMM using a simplex optimization algorithm. This is in contrast to the standard approach to parameter estimation which involves finding parameter values that maximize the correspondence between the model's output and the observed overall RTs. By isolating

the durations of the retrieval and comparison stages using the neuroimaging data, it is possible better estimate F (which only affects the retrieval stage) and Final Comparison (which only affects the comparison stage).

2.3.7 Model Results

The latency scalar parameter (F) was estimated to maximize the correspondence between the duration of the model's retrieval stage and the duration of the third stage in the HSMM-MVPA analysis. For the best fitting value of F (Table 5), the model accounted for the different durations of third stage across the 8 conditions formed by fan and probe, as shown in Figure 2.7 (r = 0.88, MSE = .005). The Final Comparison and Comparison Intercept parameters were estimated to maximize the correspondence between the duration of the model's comparison stage and the fourth stage in the HSMM-MVPA analysis. For the best fitting values (Table 5), the model accounted for the different durations of the fourth stage across the 8 conditions formed by fan and probe, as shown in Figure 2.7 (r = 0.94, MSE = .001).

Table 5. Model Parameter Estimates

Parameter Name	*Estimate*
Shared Time	430 msec
F	905 msec
Final Comparison	85 msec
Comparison Intercept	465 msec

Using the parameter estimates from the neuroimaging data (Table 5), we calculated the three durations in Eq. 4 (Shared Time, Retrieval Time, and Comparison Time) to generate model behavioral response times. The correspondence between the model's overall RTs and the observed RTs was fairly high (r = .81, MSE = .001). As illustrated in Figure 2.10, fan only affected the duration of the model's retrieval stage. Alternatively, probe type affected the duration of the model's retrieval and comparison stages in opposite ways (Retrieval Stage: Target = 382 ms, Similar 2 = 453 ms, Similar 1 = 445 ms, Dissimiliar = 515 ms; Comparison Stage: Target = 550 ms, Similar 2 = 522 ms, Similar 1 = 494 ms,

Dissimiliar = 465 ms). Owing to the strong negative correlation between the durations of the retrieval and comparison stages ($r^2 = -0.86$) the net effect of probe type on overall RT was a small (Figure 2.10).

2.4 Discussions

We conducted a study of associative recognition in which participants decided whether a probe made up of three words matched any previously studied triple. We varied the associative fan of words in the probes, as well as the degree of similarity between probes and studied triples. Our HSMM-MVPA method revealed that probe similarity substantially affected the durations of both the retrieval stage and the comparison stage. The retrieval stage was shortest when probes were more similar to a studied triple, whereas the comparison stage was longest when probes were more similar to a studied triple. The opposing ways in which probe similarity impacted retrieval and comparison stages explained why this factor had only a modest effect on overall RT; for instance, response times were nearly identical for similar 1 foils and targets, yet the durations of the retrieval and comparison stages clearly differed.

2.4.1 A Model of Associative Recognition

The results from the HSMM-MVPA analysis were largely consistent with the ACT-R model of associative recognition. The model predicts that when a studied triple shares more words with a probe, the triple will be retrieved more quickly because of the greater number of sources spreading activation to it (Eq 2.2). Alternatively, the model predicts that when a probe has more words in common with a studied triple, serial comparison of the words in the probe and in the retrieved triple will take longer. This is because more words on average must be compared before detecting a difference. By isolating the durations of the retrieval and comparison stages using HSMM-MVPA, we were able to measure the effect of probe similarity on retrieval latency. Targets were retrieved more quickly than foils as predicted from the ACT-R model.

The results from the retrieval stage were not perfectly consistent with the model, however. The model predicts that dissimilar foils, which share only one word with any studied

triple, will lead to slower retrievals than similar 1 or similar 2 foils, which share two words with some studied triple. In contrast, the HSMM-MVPA indicated that the duration of the retrieval stage was the same for all foils. This could reflect a simplifying assumption in our instantiation of the ACT-R model. Predictions based on (Eq 2.2) correspond to the case of noiseless activation during retrieval. In its complete instantiation, the ACT-R cognitive architecture adds continuously varying, logistically distributed noise to activation values. Because many studied triples overlap with a dissimilar foil (six triples averaged over fan for dissimilar foils versus one or two triples for the other foil types), the retrieved associate will be the most active of a larger set of candidates. As a consequence, the retrieved triple in the case of dissimilar foils will on average have a larger, positive noise term added to its activation value.

The HSMM-MVPA showed that the comparison stage was shortest for dissimilar foils, and was longest for similar 2 foils and targets. This result relates to a classic finding from the "same"–"different" perceptual judgment task (Sternberg, 1969). The more attributes two visual stimuli share, the longer it takes for people to determine that they differ (Farell, 1985). This is consistent with a serial comparison process, as instantiated in our ACT-R model of associative recognition. Interestingly, in the "same"–"different" task, matching stimuli produce faster responses than predicted by simple linear extrapolation of the number of comparisons required (Nickerson, 1967). As shown by the HSMM-MVPA, the duration of the comparison stage in our task increased with the number of comparisons for foils, but was equivalent for targets and similar 2 foils even though targets required more comparisons. To account for the differential effects of number of comparisons on "same" versus "different" judgments, some models assume that "same" judgments are based on a parallel holistic process rather than a serial analytical process (Farrell, 1985). The ACT-R model does not currently include a holistic comparison process, and so predicts a slightly longer comparison stage for targets.

2.4.2 Comparing ERP Components and HSMM-MVPA bumps of Associative Recognition

EEG studies of recognition memory reveal two ERP components, the FN400 and the parietal old/new effect. In many studies of recognition memory, participants view probes and

are asked to decide whether they previously studied the probe. Some of the probes perfectly match a studied item ("targets"), some are entirely novel ("dissimiliar lures"), and some merely resemble a studied item ("similar lures"). The FN400 is typically more negative for dissimilar lures than for targets or similar lures, suggesting that it is sensitive to item familiarity. The parietal old/new effect, on the other hand, is more positive for targets than for dissimilar or similar lures, indicating that it is sensitive to the retrieval of perfect matches (for a review, see Rugg & Curran, 2007).

Given that associative recognition involves remembering details about what an item appeared with, we expected that participants would display the standard parietal old/new effect. Indeed, as in previous studies of associative recognition (Diana, Van den Boom, Yonelinas, & Ranganath, 2011; Donaldson & Rugg, 1998), targets produced a late posterior positivity during the fourth flat that extended to the last bump, whereas foils did not. In the ACT-R model, the stage coinciding with the parietal old/new effect involves processing of the memory trace after its retrieval. We previously suggested that the amplitude of the sustained response reflected the different activations of the retrieved memories in the various conditions (Anderson et al., 2016). Activation is greater for targets than foils, and for Fan 1 triples than for Fan 2 triples. Consistent with this view, voltages over the parietal scalp were greater for Fan 1 triples than for Fan 2 triples in the moments preceding a response.

All of the words in the test phase of our experiment appeared earlier during the study phase. As such, it was unclear whether participants would display an FN400 for foils versus targets. Some studies of associative recognition show that the novelty of the association between probe items can produce an FN400 (Speer & Curran, 2007), while others do not (Anderson et al., 2016; Ecker et al., 2007). These discrepancies may relate to whether the individual elements in an associate are represented as a single unit. Conditions that favor unitization in paired associate learning may yield a distinct FN400 for re-arranged pairs relative to studied pairs (Ecker et al., 2007). These conditions include repeated study (Speer & Curran, 2007), semantically meaningful pairs (e.g. traffic-jam; Rhodes & Donaldson, 2007), and elaborative encoding (Rhodes & Donaldson, 2008). The conditions in our experiment did not promote unitization: participants learned triples (rather than pairs), they studied triples as few as four times, and the words in triples were unrelated.

All regions show activity related to each of the four bumps. The topographical distribution and time course of the bumps are largely consistent with those described in our previous application of the HSMM-MVPA method to a paired associate recognition task (Anderson et al., 2016). The first bump likely corresponds to the N1, given its early time course, its anterior distribution, and its insensitivity to the fan and probe type manipulations. This component is typically interpreted as an index of visual attention (Luck, Woodman, & Vogel, 2000). The intermediate time course and anterior distribution of the second bump are consistent with the P2 (Van Petten et al., 1991).

The third bump may relate to the N2 (c.f. Anderson et al, 2016), a frontocentral negativity caused by response conflict (Yeung, 2004). The N2 typically appears somewhat earlier in ERP waveforms. However, most studies of the N2 involve decisions far simpler than associative recognition. As the third bump in the ACT-R model initiates the comparison stage, it follows that this bump occurs at the moment of maximum response conflict within the trial. The late and variable latencies of the third bump may obscure the N2 in the conventional ERP waveforms of our study and in other studies of recognition memory. Finally, the time course, direction, and topographical distribution of the fourth bump are consistent with the parietal old/new effect. As in our previous experiment (Anderson et al., 2016), this bump was sensitive to perfect matches, and had higher amplitude for Fan 1 triples than for Fan 2 triples.

2.4.3 The Path Forward: HSMM-MVPA

HSMM-MVPA can be used to guide the development of new theories by providing a direct measure of the durations of information processing stages to make inferences about the effects of experimental factors. RT-based methods have also been used. However, if an experimental factor affects the durations of multiple stages, its impact on each cannot be directly observed from overall RTs. Rather, one must specify a model of how the factor(s) affect each stage, calculate the expected durations of all stages, and compare the summed stage durations to overall RTs. A discrepancy between expected and observed RTs does not indicate which stages and factors were modeled incorrectly. More problematically, the absence of a discrepancy does not exclude the possibility that the model over-predicted the duration of one stage, and under-predicted the duration of another. By isolating each

stage's duration, the HSMM method overcomes these limitations of RT-based methods.

HSMM-MVPA can be used for a second, related purpose; to obtain more accurate parameter estimates for a process model. An advantage of linking cognitive models to neural data would be the sheer wealth of additional information that neural data can provide in comparison with behavioral data (see a review of different linking approaches: Hollander, Forstmann, & Brown, 2016). The HSMM-MVPA method provides individual stage durations instead of the sum of them (i.e. RT) to better constrain model fitting. For example, the ACT-R retrieval latency scalar (F) only affects the duration of the retrieval stage, whereas Final Comparison only affects the duration of the comparison stage. Because both parameters modulate the effects of probe type, changes in F can be partially offset by changes in Final Comparison time to produce similar overall RTs. By estimating model parameters based on the stage durations in the HSMM-MVPA, no such parameter compensation occurs. Another example is the estimation of Comparison Intercept, which is not affected by fan or probe type. Without information about the duration of the comparison stage from the HSMM-MVPA, it would not be possible to estimate the Comparison Intercept separately from other processes (e.g. encoding and responding) that are also not affected by the experimental factors. Beyond just a common intercept, Figure 2.7 illustrates that the method can break this time out into periods of pre-stimulus attention (stage 1), encoding (stage 2), and responding (stage 5).

Chapter 3

Spatial dynamics of memory retrieval in working memory

In this study, we investigated the time course and neural correlates of the retrieval process underlying visual working memory. We made use of a rare dataset in which the same task was recorded using both scalp electroencephalography (EEG) and Electrocorticography (ECoG), respectively. This allowed us to examine with great spatial and temporal detail how the retrieval process works, and in particular how the medial temporal lobe (MTL) is involved. In each trial, participants judged whether a probe face had been among a set of recently studied faces. With a method that combines hidden semi-Markov models and multivariate pattern analysis, the neural signal was decomposed into a sequence of latent cognitive stages with information about their durations on a trial-by-trial basis. Analyzed separately, EEG and ECoG data yielded converging results on discovered stages and their interpretation, which reflected 1) a brief pre-attention stage, 2) encoding the stimulus, 3) retrieving the studied set, and 4) making a decision. Combining these stages with the high spatial resolution of ECoG suggested that activity in the temporal cortex reflected item familiarity in the retrieval stage; and that once retrieval is complete, there is active maintenance of the studied face set in the decision stage in the MTL. During this same period, the frontal cortex guides the decision by means of theta coupling with the MTL. These observations generalize previous findings on the role of MTL theta from long-term memory tasks to short-term memory tasks.

3.1 Background

Information can be preserved in working memory across a short delay without active maintenance (Ericsson and Kintsch, 1995; Lewis-Peacock, Drysdale, Oberauer, & Postle, 2012; Oberauer, 2002). Our interest is in the process by which this information is later retrieved. This can take place either through an attention-based process that refocuses and refreshes the memory traces (Lewis-Peacock at al., 2012; Souza, Rerko, & Oberauer, 2015), or a cue-driven process that is very similar to cue-based retrieval from long-term memory (Nairne, 2002). The objective of this study to map the time course and neural correlates of this retrieval process. We focus on a visual working memory task from a published dataset (van Vugt, Sekuler, Wilson, & Kahana, 2013). Two experiments using the same task were carried out using scalp electroencephalography (EEG) and electrocorticography (ECoG), which allows us to harness the complementary strengths of these two recording methods. In each trial, participants first studied a list of faces, then, after a short delay, they were cued with a probe face and asked to judge if it had been among the just-studied faces (i.e., a Sternberg task). To create a detailed mapping of the time course of the retrieval process, we modeled the trial-to-trial variability of this process with a novel method that combines hidden semi-Markov models with multivariate pattern analysis (HSMM-MVPA; e.g., Anderson et al., 2016), which we applied to both EEG and ECoG datasets. To create a detailed mapping of the neural correlates we relied on the ECoG dataset, which has superior spatial resolution. In the rest of this introduction, we will describe three main challenges to isolate the working memory retrieval process in space and in time, and our approaches to tackle these challenges. In the method section, we will provide the technical details of these approaches.

3.1.1 The first challenge: trial-to-trial variability in the timing of cognitive processes

We focus on the period of the working memory task where participates are asked to judge whether the presented probe is one of the recently studied faces. Each trial analyzed starts with encoding a probe face and ends with a motor response (the period marked as ”probe” in Figure 3.1), but the intermediate cognitive processes (including the retrieval process)

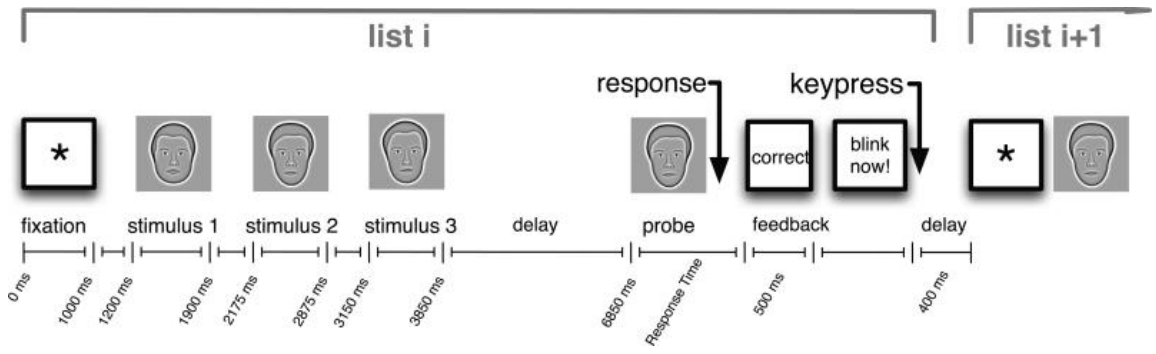

Figure 3.1: Trial structure of the Sternberg task. This figure illustrates the sequence and timing of events in a trial with a set size of 3. Adopted from (Van Vugt et al., 2013) with permission.

can vary greatly from trial to trial given the self-paced nature of the task. Event-related potentials (ERP) are commonly used in EEG literature to identify the occurrence of important cognitive events by averaging stimulus-locked or response-locked signal across trials. However, cognitive events that are further away from stimulus or response will have higher trial-to-trial temporal variability and are often distorted or lost in the averaged waveforms (Luck, 2014). Therefore, in order to isolate the full range of cognitive processes in the visual working memory task, we need to apply a method that can capture the trial-to-trial variability in when they occur.

In the past, we have combined hidden semi-Markov models and multivariate pattern analysis (i.e. HSMM-MVPA) to decompose the neural signal (i.e. EEG, MEG) into a sequence of latent stages (Anderson, Zhang, Borst, & Walsh, 2016; Borst & Anderson, 2015; Zhang, Walsh, & Anderson, 2017). This approach assumes that transitions from one processing stage to the next are accompanied by a significant change in the information processing and therefore in the underlying neural signal. HSMM-MVPA shares the same notion as the microstate analysis developed by Lehmann and colleagues, which is to extract a sequence of non-overlapping mental states from neural data during a specific task (Lehmann, 1987; Pascual-Marqui, Michel, & Lehmann, 1995). In addition, the HSMM-MVPA method also captures the trial-to-trial variability of the durations in each state. This is in contrast to the microstate analysis that was applied to average ERP data (Pascual-Marqui et al., 1995).

The HSMM-MVPA model identifies brief, distinctive profiles of scalp activity (i.e.,

bumps) with variable latencies in single trial EEG data. A bump is modeled as a half-sine multidimensional peak across the scalp that signifies a significant change in the information processing. This assumption is inspired by two theories of ERP generation (Yeung, Bogacz, Holroyd, & Cohen, 2004; Makeig et al., 2002; Shah et al., 2004; Basar, 1980). According to the classical theory, significant cognitive events generate bursts of activity in discrete brain regions (Shah et al., 2004). Therefore, the EEG signal can be described as a sum of sinusoidal peaks and ongoing neural signal of uncorrelated sinusoidal variation. According to the second theory of ERP generation, the synchronized oscillation theory, significant cognitive events reset the phase of the oscillation at a certain frequency (Basar, 1980). Under both theories, averaging neural signal across trials will reveal the peaks as averaged ERP waveforms that we see, and the ERP waveforms are indistinguishable between the two theories under simulated datasets (Yeung, Bogacz, Holroyd, Nieuwenhuis, & Cohen, 2007).

In the current work, we will apply this method to both the EEG and ECoG datasets to identify and isolate the cognitive processes in our visual working memory task. Details on parameter estimation and model fitting can be found in the Method section.

3.1.2 The second challenge: isolate the retrieval process from the discovered stages

The HSMM-MVPA method gives a sequence of latent stages that correspond to different cognitive processes. We are particularly interested in when the retrieval process takes place. In order to isolate this retrieval process among the set of stages uncovered by HSMM-MVPA, we need to find the mapping between the obtained model stages and the series of cognitive processes thought to be involved in our task. Extending an existing stage model based on an HSMM-MVPA analysis of a Sternberg task involving digits (Anderson et al., 2016), we expect to find the following five stages of processing:

1) Pre-attention: the time for the visual signal to reach the brain and to be attended to; 2) Encoding: encode the probe presented; 3) Retrieval: reactivate face(s) from the studied memory set; 4) Decision: compare and decide if the probe face is part of the reactivated set; 5) Motor response: press a key that reflects the decision.

To confirm this mapping for the current task, we will also examine how stage dura-

tions vary given the manipulation of different experimental factors. If a particular stage corresponds to a given cognitive process, the way that different experimental factors alter the duration of this stage should be consistent with our knowledge of the corresponding cognitive process. We use regular scalp EEG with normal populations to identify such condition effects, because there is more power associated with larger number of subjects.

3.1.3 The third challenge: identify the neural correlates associated with the retrieval process

Processing stages obtained in the HSMM-MVPA method provide us with the fine temporal resolution of when the retrieval process takes place on single trials. We can further examine which brain regions are activated during these discovered periods in a combined spatio-temporal analysis. With EEG data, we can identify the mapping from HSMM-MVPA stages to individual cognitive processes with high temporal precision, but poor spatial resolution limits our ability to determine how different brain regions are engaged during these periods. To achieve a better spatial resolution, we look into the ECoG activity during the periods of interest. ECoG recordings from epileptic patients not only give finer spatial resolution in the cortical regions, but also make it possible to examine subcortical activity such as that of MTL, which plays an important role in visual (working) memory tasks (van Vugt, Schulze-Bonhage, Litt, Brandt, & Kahana, 2010; Ranganath, 2006). Because recording sites in ECoG data vary across subjects in both numbers and locations, we need a method of subject alignment before pooling data from different subjects. Multi-set canonical correlation analysis (M-CCA) is used for this purpose to transform electrode activity in each subject to a common neural representational space, where the inter-subject correlations of the transformed data are maximized across subjects. We have previously demonstrated the reliability of M-CCA in aligning subjects in MEG data (Zhang, Borst, Kass, & Anderson, 2017). In the current work, we will apply M-CCA to align ECoG data from individual subjects; we will also apply M-CCA to the EEG data, which allows us to compare the obtained common dimensions and the HSMM-MVPA results across two experiments in two different measurement modalities.

3.2 Methods

3.2.1 Experimental paradigm

Participants completed a Sternberg task: Following the appearance of a fixation stimulus, participants viewed a short series of faces in sequence (Figure 3.1). After a retention interval, a probe item appeared and participants indicated with a key press whether the probe was a member of the just presented set (target) or not (foil). After each trial, participants were given accuracy feedback. Figure 3.1 illustrates the sequence and timing of trial events in the original report of the data used in this study (van Vugt et al., 2013). 16 synthetic faces, which varied along four perceptual dimensions, were used as stimuli (Wilson, Loffler, & Wilkinson, 2002). A multidimensional scaling (MDS) study was carried out on 23 participants to characterize the psychological perceived similarities among these synthetic faces (see van Vugt et al., 2013; Kahana & Bennett, 1994, for details). Therefore, we can characterize the overall similarity between a probe and the just studied set as the average similarity between the probe and each face on the set (i.e., set similarity). We define two levels of set similarity among the foils (i.e., low-sim and high-sim), dividing foils into two groups with equal numbers of trials. In the datasets we will use, there are two versions of the visual working memory task, adapted to the cognitive capacities of the respective subject populations. One was administered to healthy undergraduates while their EEG was recorded (Experiment 1), while the other was adapted to epileptic patients from who ECoG was recorded (Experiment 2). The two experiments differed in the number of faces participants viewed in the studied set before a probe in each series (set size). In Experiment 1, set size could be two, three or four; in Experiment 2, set size could be one, two or three (adapted to the capacity of the patients). Sets were constructed so that items could not be repeated on successive sets, and targets were equally likely to match a study item from each serial position. Incorrect trials and trials with RTs shorter than 400 ms or longer than 4000 ms were removed from the analysis.

3.2.2 Participants

In Experiment 1, 29 adults (ages 20–32) were recruited from the University of Pennsylvania student community. Informed consent was obtained from all participants. Each

participant completed two sessions, with each session involving ten blocks of 30 trials. Participants with mean accuracy lower than 60 % or with mean RT longer than 4s were excluded. All 29 participants were retained for further analysis. Participants in Experiment 2 were 16 neurosurgical patients being treated for medically refractory epilepsy and were monitored with arrays of subdural and/or depth electrodes. Patients were recruited from Brigham and Women's Hospital in Boston, the Hospital of the University of Pennsylvania in Philadelphia, and Universitts Klinikum Freiburg in Germany. Informed consent was obtained from all participants. Under the same accuracy and RT criteria as Experiment 1, 12 out of 16 patients were retained for further analysis. Each patient completed different number of trials for the experiment (μ = 202, SEM = 41), depending on their availability and willingness to participate in the experiments.

3.2.3 Scalp EEG recordings in Experiment 1

Scalp EEG signals were recorded using a 129-channel EGI Inc. system, with an AC-coupled, high-input-impedance amplifier (200 MΩ, Net Amps, Electrical Geodesics, Inc., Eugene, OR). The sampling rate was 500 Hz, and data were recorded with a 0.1–250 Hz bandpass filter. Individual channels were adjusted until impedances were below 50 kΩ. EEG signals were filtered with a bandpass of 0.1 to 70.0 Hz, and then decomposed into independent components using the EEGLAB FastICA algorithm (Delorme & Makeig, 2004). Components associated with eye blinks were automatically identified and projected out of the EEG recording. Epochs (from probe presentation to motor response in each trial) were then extracted from the continuous recording. Epochs containing voltages above +100 μV or below -100 μV were excluded. Data were down-sampled to 100Hz before further analysis. More details of the experiment can be found in the original report of the EEG data (van Vugt et al., 2013).

Table 1
Number of patients and electrodes in each of the five ROIs.

Numbers	Electrodes	Patients
Temporal Cortex	269	12
Medial Temporal Lobe	69	10
Frontal Cortex	161	9
Parietal Cortex	93	9
Occipital Cortex	22	7

3.2.4 ECoG and depth electrode recordings in Experiment 2

The local field potential was amplified and digitally recorded at sampling rates between 250 and 1024 Hz , and bandpass-filtered between 0.1 and 100 Hz. Data were subsequently notch-filtered with a Butterworth filter with zero phase distortion to eliminate line noise. For all participants, the locations of implanted electrodes were determined by means of co-registered postoperative computed tomographies and preoperative magnetic resonance imaging (MRI) or from postoperative MRIs by an indirect stereotactic technique and converted into MNI (Montreal Neurological Institute) coordinates. Localization of depth electrode contacts in the medial temporal lobe (MTL) was done manually through clinician's inspection of the postoperative MRIs, which includes areas in hippocampus and parahippocampal gyrus. Analysis was done in five pre-defined regions of interest (ROIs; see Table 1), and data were down-sampled to 100Hz. More details of the experiment can be found in the original reports of the ECoG data (van Vugt et al., 2013; van Vugt et al., 2010).

3.2.5 Alignment in the Common Representational Space

While it is common practice to assume that the sensors for different subjects in EEG recordings correspond, recording sites in ECoG data vary from subject to subject in both number and locations since localization is driven solely by clinical considerations. Therefore, in Experiment 2, a method of subject alignment is required before pooling data from all subjects together. Previously, we have found success in aligning subjects with multi-set canonical correlations analysis (M-CCA) using only functional information of the neural data in MEG datasets (Zhang et al., 2017). This method is used to find the optimal trans-

formation for each subject from electrode activity to a common neural representational space, where the inter-subject correlations of the transformed data are maximized across subjects.

Application of M-CCA

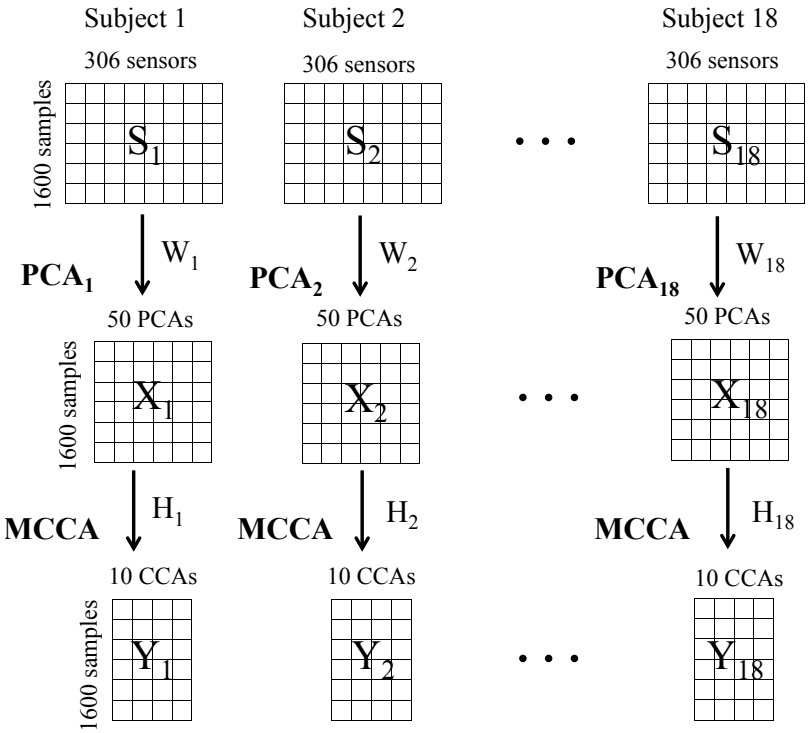

Figure 3.2: This figure illustrates application of M-CCA to 12 subjects. S_k is the averaged data (across all trials for each condition) from sensor data for subject k, each with 120 time points; X_k has 10 PCA components for subject k, each with 120 time points; Y_k has 5 CCA components, each with 120 time points; $W_1, W_2, \ldots, W_{12}$ are PCA weights obtained for each subject independently; and $H_1, H_2, \ldots, H_{12}$ are CCA weights obtained jointly from all subjects by maximizing all of the inter-subject correlations.

This section describes the pipeline to apply M-CCA to EEG and ECoG datasets. We start by pre-processing the sensor data for each subject. To overcome noise in the sensor data of individual trials, multiple trials are averaged to obtain a highly reliable representation of the change in sensor activity, which is similar to obtaining event-related potential waveforms in the EEG literature (Picton et al., 2000). When trials have a fixed trial length,

this averaging procedure is straightforward. However, trials are quite variable in their duration, and temporal alignment is lost when a time sample is further away from stimulus presentation or response emission. Samples locked to response emission in Experiment 2 also have poor temporal alignment, as there is potential delay in response timing given the condition of neurosurgical patients. Therefore, we only use samples from the first 600 ms (60 samples given the sampling rate of 100Hz) of a trial when applying M-CCA in Experiment 2. This averaging process is repeated for the target and foil conditions separately, as we potentially have different latent components for different conditions after averaging. This gives rise to 120 samples (60 samples x 2 conditions) per subject as the input S of M-CCA shown in Figure 3.2.

To reduce dimensionality and remove subject-specific noise, the next step after obtaining S_k is to perform spatial PCA. M-CCA is then applied to the top 10 PCA components from each subject instead of directly to the sensor data S_k. This results in 12 matrices of dimension 120×60, which are the inputs X_k to the M-CCA analysis for subjects $k = 1, 2, \ldots, 12$. As is illustrated in Figure 3.2, W_k are the PCA weights for subject k which are obtained independently for each subject. H_k are the CCA weights for subject k which are obtained jointly from all subjects resulting in common CCA dimensions $Y_k = X_k H_k$. Subject data do not align in either the sensor space or the PCA space, with S_k and X_k processed for each subject independently. Rather, subject data align in the common representational space after M-CCA, with Y_k maximally correlated across subjects. Once Wk and Hs are obtained, we can go back to full-time-course data of individual trials and transform them from representation in electrodes to representation in CCA dimensions. The top 5 CCA dimensions are retained for each subject. M-CCA was also applied to the EEG dataset in Experiment 1 using the exact same procedure, allowing us to better compare the two experiments.

Mathematical details of M-CCA

This section discusses the mathematical details of how H_k for subjects $k = 1, 2, \ldots, 18$ are obtained, so that after the transformation $Y_k = X_k H_k$, the new representation of data Y_k is more correlated across subjects than X_k is.

We first illustrate the simplest case where we look for correspondence over datasets

from two subjects instead of many subjects. Let $X_1 \in \mathbb{R}^{T\times m_1}$ and $X_2 \in \mathbb{R}^{T\times m_2}$ be PCA components from two subjects, with the same number of time points T, and PCA dimensions m_1 and m_2, respectively ($T = 1600$, $m_1 = m_2 = 50$ in our case). Each PCA component stored in X_1 and X_2 has mean 0. The objective in canonical correlation analysis (CCA) is to find two vectors $h_1 \in \mathbb{R}^{m_1\times 1}$ and $h_2 \in \mathbb{R}^{m_2\times 1}$ such that after the projection $y_1 = X_1h_1$ and $y_2 = X_2h_2$, y_1 and y_2 are maximally correlated. This is equivalent to:

$$\arg\max_{h_1,h_2} \rho = \frac{y_1^T y_2}{\|y_1\|\|y_2\|} = \frac{h_1^T R_{12} h_2}{\sqrt{h_1^T R_{11} h_1 h_2^T R_{22} h_2}} = h_1^T R_{12} h_2, \text{ where } R_{ij} = X_i^T X_j.$$

There are N solutions, $h^{(i)} = (h_1, h_2) \in \mathbb{R}^{m_1\times 1} \times \mathbb{R}^{m_2\times 1}$, obtained collectively in a generalized eigenvalue problem with $i = 1, \ldots, N$, subject to the constraints $h_1^T R_{11} h_1 = h_2^T R_{22} h_2 = 1$ (Borga, 1998). This results in N dimensions (each referred as a CCA component) in the common representational space with the transformed data $Y_1 = [y_1^{(1)}, y_1^{(2)}, \ldots, y_1^{(N)}]$ and $Y_2 = [y_2^{(1)}, y_2^{(2)}, \ldots, y_2^{(N)}]$. The value of N does not exceed the smaller of m_1 and m_2. The resulting CCA components in the common representational space are ranked in a decreasing order of between-subject correlations. The earlier CCA components are the more important ones and the later components can be removed. In other words, canonical correlation analysis finds the shared low-dimensional representation of data from different subjects.

M-CCA is an extension of CCA which considers more than 2 subjects. The objective is similar to before, but now it needs to maximize the correlations between every pair of subjects (i.e. inter-subject correlations) simultaneously. Let $X_k \in \mathbb{R}^{T\times m_k}$ with $k = 1, \ldots, M$ be datasets from M subjects ($M > 2$), each with mean 0 for all columns. The objective in M-CCA is to find M vectors $h_k \in \mathbb{R}^{m_k\times 1}$, where $k = 1, \ldots, M$, such that after the projection $y_k = X_k h_k$, the canonical variates y_k are maximally pairwise-correlated. The objective function to maximize is formulated as:

$$\begin{aligned}\arg\max_{h_1,\ldots,h_M} \rho &= \frac{1}{M(M-1)} \sum_{k,l=1,k\neq l}^{M} y_k^T y_l \\ &= \frac{1}{M(M-1)} \sum_{k,l=1,k\neq l}^{M} h_k^T R_{kl} h_l,\end{aligned}$$

where $R_{kl} = X_k^T X_l$, and $\frac{1}{M}\sum_{k=1}^{M} h_k^T R_{kk} h_k = 1$. The solution is given by solving a generalized eigenvalue problem (**?**). This formulation is not an exact maximization but an approximation of the pairwise correlations, given the complexity of the problem when $M > 2$. It is equivalent to the Maximum Variance (MAXVAR) generalization of CCA proposed by Kettenring (1971). See the proof of this equivalence in (Vía, Santamaría, & Pérez, 2005). Other ways of formulating the objective function in M-CCA yield similar results (Li, Eichele, Calhoun, & Adali, 2011).

3.2.6 HSMM-MVPA

The HSMM-MVPA method explicitly models the variability of endogenous ERP components that would otherwise be distorted or lost in the average waveforms. Previous applications of the HSMM-MVPA method to EEG data were effective in recovering the durations of the underlying processing stages (e.g., recollection, decision) and showed predictable changes with experimental factors in an associative recognition task for word pairs and a Sternberg task (Anderson et al., 2016; Zhang, Walsh, & Anderson, 2017). The HSMM-MVPA model identifies brief, distinctive profiles of scalp activity (i.e., bumps) with variable latencies in the single trial EEG data. A bump is modeled as a half-sine multidimensional peak across the scalp that signifies a significant change in the information processing, followed by a flat period where the signal appears as ongoing sinusoidal noise around a mean of 0. HSMM-MVPA models the durations of the flats as gamma distributions. The HSMM-MVPA method was applied to the first 5 CCA components, separately over two experiments. The CCA components were z-scored for each trial. As a result, the data for the analysis consisted of 5 orthogonal CCA components in each experiment, sampled every 10 ms and with constant mean and variability across trials. Only correct trials are considered in the analyses presented below. As described in more detail in our previous application of the HSMM-MVPA method (Anderson et al., 2016), a n-bump HSMM requires estimating n + 1 stage distributions to describe the durations of the flats plus the n 5-sample bumps for each CCA component. A different magnitude is estimated for each of the n bumps along each CCA dimension. A bump extends temporally across 5 samples (50 ms) and is multiplied by weights of 0.309, 0.809, 1.000, 0.809, and 0.309 (i.e., a 10-Hz half sine wave). The best model fit of such HSMMs is given

by maximizing the summed log likelihood of the bumps and flats across all trials. For each trial, this log likelihood reflects the combination of two factors: the likelihood of the EEG data given that the bumps are centered at each time point, and the likelihoods that the bumps are centered at those time points given the gamma distributions that constrain their locations. In other words, the HSMM must select bump locations within a trial to maximize the correspondence between the observed and the estimated EEG/ECoG signal, while selecting relatively consistent flat durations across trials to maximize their fit to the gamma distributions.

3.2.7 Brain synchrony analysis

Quantification of phase synchrony between two neural signals was done by means of a phase locking value (PLV). PLV measures the consistency in the phase difference at a frequency of interest between two recording sites across multiple trials at the corresponding time points (Lachaux, Rodriguez, Martinerie, & Varela, 1999). Phases for single-trial neural signal at different brain regions were measured using Hilbert transform (Tass et al., 1998; Lachaux et al., 1999). It is equivalent to an alternative method by using convolution with a complex wavelet, as demonstrated in a direct comparison study (Le Van Quyen et al., 2001). If the phase difference between signals from two recording sites is very similar from trial to trial at the corresponding time point, then it is considered that at that particular time point, the two brain sites are well synchronized in phase with PLV close to 1; if the phase difference is very variable across trials, PLV is close to 0. Calculation of phase locking values requires establishing corresponding time points across multiple trials. Typically one assumes that samples correspond when they are at the same delay from the stimulus. This assumption is only approximately correct given the trial-to-trial variability of where the same cognitive event occurs. Alternatively, one can assume that samples correspond when they at the same offset from the same bump in the HSMM-MVPA analysis (Portoles, Borst, & Van Vugt, 2018). In this way, phase locking value calculates how synchronized two brain sites are around the cognitive event signified by a particular bump. In this experiment, we are interested in the synchronization between the MTL and cortical regions around when retrieval completes. Following the procedure outlined by Lachaux (1999), we bandpass-filtered ECoG signal with finite impulse response (i.e., theta: 4-9 Hz, alpha:

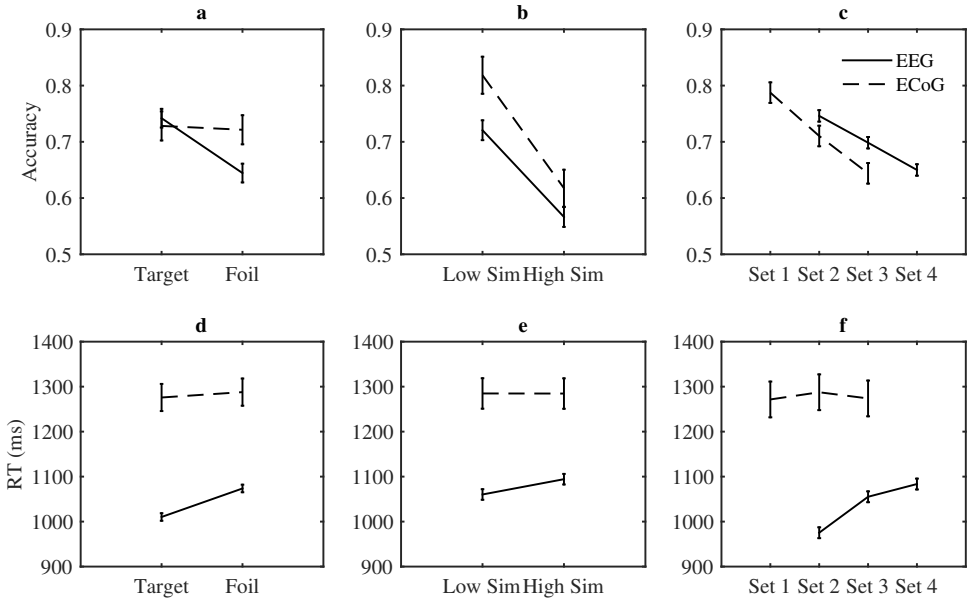

Figure 3.3: Accuracy (a-c) and RT (d-f) of EEG and ECoG data as a function of probe type, set similarity and set size. SEMs are shown in the error bars with between-subjects variance removed (Loftus and Masson, 1994).

9-15 Hz, and beta: 15-30 Hz) and extracted instantaneous phases using Hilbert transform. Then we focused on the 100-ms time period right before and the 100-ms time period right after the bump that signifies completion of retrieval, and compared the phase locking values averaged within each period to examine changes upon transitioning of cognitive processing stages. This procedure was done for each subject separately, and repeated for every combination of electrode pairs, with one electrode from the MTL and one from the cortical region.

3.3 Results

3.3.1 Behavioral analyses (Figure 3.3)

Probe type. Participants are faster and more accurate when the presented probe is a target than a foil. The effect is significant in Experiment 1, as revealed by a repeated measures ANOVA (accuracy: F(1,28) = 17.5, $p < .001$; RT: F(1,28) = 28.6, $p < .001$) but not in

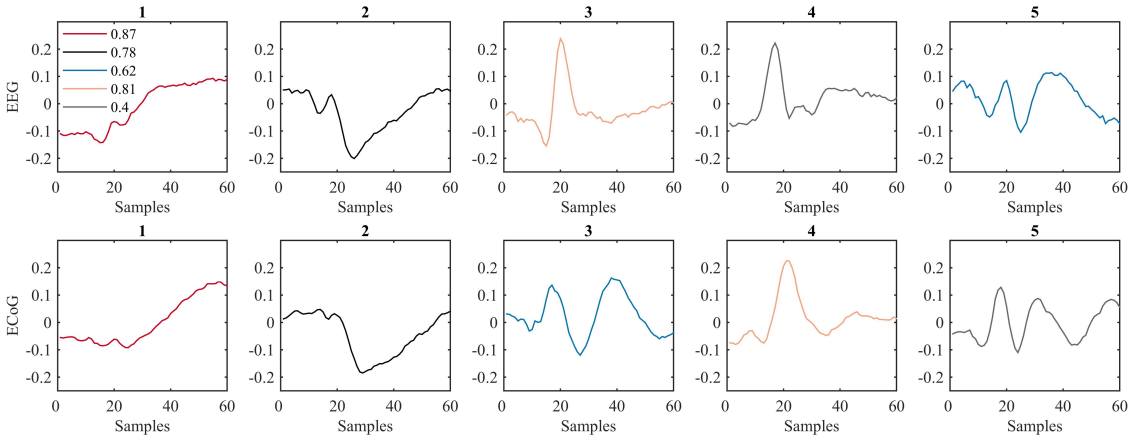

Figure 3.4: First 5 CCA components obtained over EEG data in Experiment 1 (top row) and over ECoG data in Experiment 2 (bottom row). 60 samples correspond to 600 ms of stimulus-locked data in target condition. Corresponding CCAs are highlighted in the same color, with their absolute correlations noted in the legend. CCAs in EEG with large negative correlations with CCAs in ECoG have been flipped, so that all signs are positive.

Experiment 2 (accuracy: F(1,11) = 0.04, $p = .84$; RT: F(1,11) = .08, $p = .78$.

Set similarity. There is an effect of set similarity on both RT and accuracy in Experiment 1, as revealed by a repeated measures ANOVA (accuracy: F(1,28) = 34.8, p¡.001; RT: F(1,28) = 4.83, $p < .05$). In Experiment 2, set size has an significant effect on accuracy but not on RT (accuracy: F(1,11) = 43.7, $p < .001$; RT: F(1,11) = .078, $p = .77$).

Set size. There is a significant effect of set size on both RT and accuracy in Experiment 1, as revealed by a repeated measures ANOVA (accuracy: F(2,56) = 28.0, p¡.001; RT: F(2,56) = 27.8, $p < .001$). In Experiment 2, set size had a significant effect on accuracy but not on RT (accuracy: F(2,22) = 25.4, $p < .001$; RT: F(2,22) = .08, $p = .92$.

3.3.2 Consistency in CCA dimensions across two experiments

Before pooling data across different ECoG subjects, we first identified dimensions that correspond across subjects by applying the M-CCA method. M-CCA serves to transform the data for each subject from the number of recording sites into a reduced number (5 in our case) of common dimensions (i.e., CCA components) shared across subjects. Two M-CCAs were applied to the two experiments independently. We then compare the temporal dynamics of the top 5 CCAs from the two experiments. Being recorded from the same

task, the two experiments are expected to share similar CCAs.

Although the recording sites from individual subjects in ECoG data were quite varied, the M-CCA obtained common dimensions of variation. Moreover, Figure 3.4 shows that the top 5 CCA components are comparable to the top 5 CCAs obtained from the EEG dataset – only the order is different. We matched each CCA component in ECoG with the one in EEG that had the highest absolute correlation (value noted in the legend; matching indicated with colors). There appears to be a one-to-one mapping from the 5 CCAs in ECoG to the 5 CCAs in EEG, with their corresponding correlations of time courses ranging from 0.4 to 0.87.

3.3.3 Identification of the stage durations and the bump profiles in HSMM-MVPA

First, the number of stages in the HSMM-MVPA was determined. Two HSMM-MVPAs were applied to the top 5 CCAs of the two experiments independently. HSMM-MVPA identifies bumps in the ongoing EEG signal related to significant changes in information processing. In this study, the number of stages in HSMM was decided on the basis of between-experiment predictions. Model parameters of the stage distributions were obtained from the 29 subjects in Experiment 1, and used to calculate the likelihood for each of the 12 subjects in Experiment 2 while re-estimating the bump magnitudes; and vice versa. We prefer a more parsimonious model with fewer bumps: we only select a model with m+1 bumps over that with m bumps if there is improvement in model likelihood over a significant proportion of the total number of subjects. A 3-bump model is significantly ($p < .0001$; two-tailed sign test) better than a 2-bump model in 34 out of 41 subjects (specifically, 24 of 29 in Exp 1 and 10 of 12 in Exp 2). A 4-bump model is significantly ($p = .01$; two-tailed sign test) better than a 3-bump model in 29 out of 41 subjects (specifically, 19 of 29 in Exp 1 and 10 of 12 in Exp 2), but a 5-bump model is only better ($p = 1.0$; two-tailed sign test) than a 4-bump model in 21 (17 in Exp 1 and 4 in Exp 2) out of 41 subjects. Therefore, the 4-bump model is the preferred solution.

Second, we compared consistency in HSMM-MVPA results across two experiments. Two 4-bump HSMMs were applied to the top 5 CCAs of the two experiments independently. We then compared bump magnitudes and stage durations of the resulting two

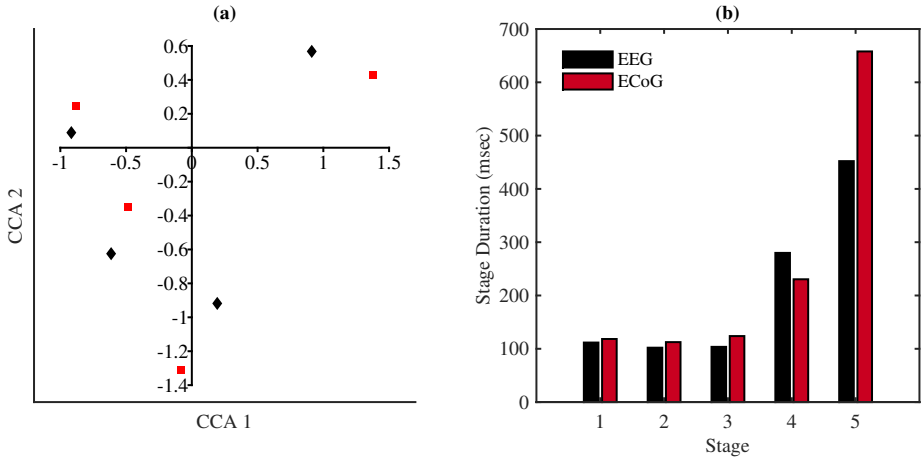

Figure 3.5: Representation of the 4 bumps in the top two CCA dimensions in the obtained HSMMs for both EEG and ECoG data (a). Durations of the five processing stages identified using the HSMM-MVPA method for both EEG and ECoG data (b).

models. Despite that two HSMMs were applied to EEG and ECoG separately, there is considerable consistency in bump magnitudes with the correlation of the 20 values (5 CCAs x 4 bumps/CCA) being 0.62. Figure 3.5a further demonstrates this consistency by comparing representations of the 4 bumps in the top two CCA dimensions across two experiments, which are highly similar.

Next, we interpreted the recovered stages and mapped them to corresponding cognitive processes. Figure 3.5b shows the stage durations of the two HSMMs. Consistent across the two experiments are three briefer periods at the beginning of the trial followed by two longer periods towards the end. ECoG subjects, who have overall longer RTs, are markedly slower in the last stage. Figure 3.6a shows the reconstructed scalp profiles of the four bumps in Experiment 1 (EEG), which were created by averaging the observed voltages at the time of the maximum-likelihood samples for each bump and during each trial. The scalp profiles are plotted against the five stage durations.

Guided by the process model of Anderson et al. (2016), which decomposed the recognition memory task into a encoding process, a retrieval process, a decision process, and a motor response process, we can interpret the HSMM-MVPA stages as follows: In both latency and topographical distribution, the first bump resembles the N1, which is typically interpreted as an index of visual attention (Luck, Woodman, & Vogel, 2000). Therefore,

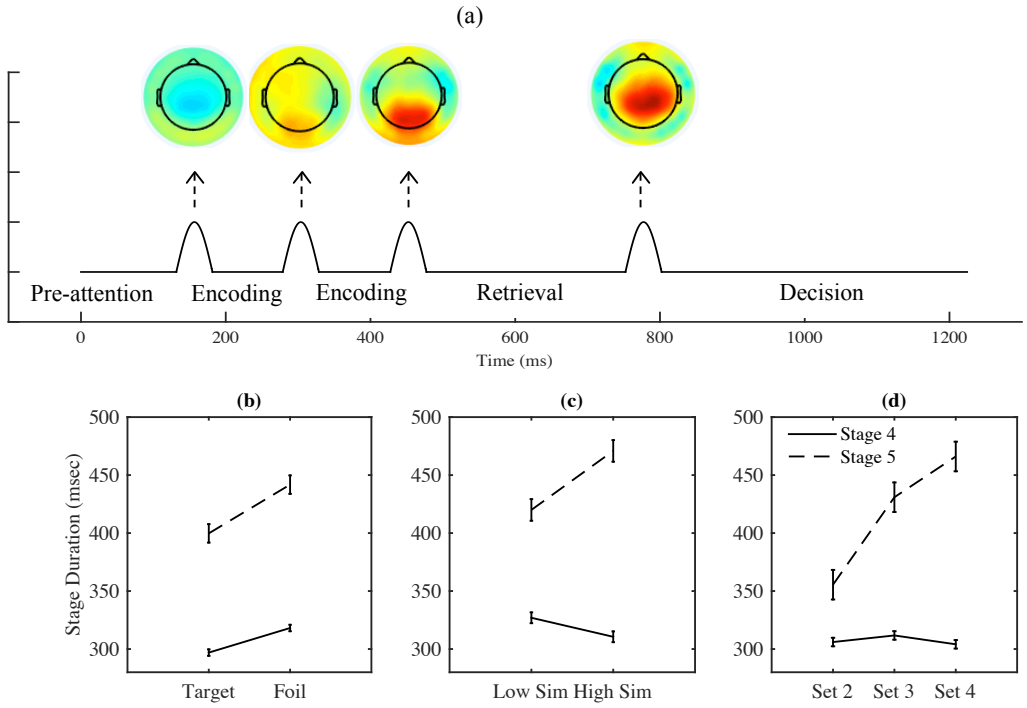

Figure 3.6: (a). HSMM stages labeled with corresponding cognitive processes, interleaved by four reconstructed bumps for EEG data. Mean EEG electrode activity of the reconstructed bumps were obtained by averaging the observed voltages at the time of the maximum-likelihood samples for each bump and during each trial. Electrode voltages have been normalized for each trial. (b-d). Duration of Stage 4 (Retrieval stage) and Stage 5 (Decision stage) in EEG as a function of probe type, set similarity and set size. SEMs are shown in the error bars with between-subjects variance removed (Loftus et al., 1994).

we interpret the first stage is a pre-attention stage before actual encoding takes place. It is then followed by two brief encoding stages associated with the second and the third bumps that show a posterior positivity. We believe these two bumps mark the encoding of the probe face similar to a pair of encoding bumps identified in a word recognition task (Anderson et al., 2016). Given its stage duration (i.e. around 200 ms) and the process model, Stage 4 is identified as the Retrieval stage. This is consistent with the frontal-central distribution of the fourth bump. Of the 5 stages, the last stage is not likely to be a motor response stage given its long duration of around 600 ms. It is more likely that this stage represents the combination of the decision and motor response stage and that we failed to detect a bump separating the decision and motor stages. The interpretations of the five stages as labeled in Figure 3.6a will be further verified in the following sections by examining the effect of different experimental factors (probe type/set similarity/set size) on the durations and brain activity associated with each stage.

3.3.4 Stage durations by condition with EEG

RTs varied by condition for EEG subjects in Experiment 1. These differences must show up in the durations of some of the stages. To determine which stages were affected by the experimental manipulations and examine if these effects are consistent with our interpretation of the stages, we fit HSMMs with different stage durations to each condition. That is, we estimated parameters for the gamma distributions of each stage separately for the different conditions while constraining the bump magnitudes to be the same. The HSMM methods return the probabilities of each bump occurring at each time point on a trial-by-trial basis. These probabilities can be used to calculate the most likely location of each bump in a trial. Mean stage durations for a particular subject can then be calculated as the average time between bumps across all trials within that subject. Figure 3.6b-d shows the resulting mean time durations across all EEG subjects for Stage 4 and Stage 5. We submitted them to a repeated-measures ANOVA for each condition (probe type/set similarity/set size) and for each stage. The stage durations do not differ between conditions for the first three stages. Consistent to the impression conveyed in the figure, there is an effect of probe type in Stage 4 - the Retrieval stage ($F(1,28) = 29.210$; $p < .001$) and in Stage 5 - the Decision stage ($F(1,28) = 14.013$; $p = 0.001$). There is also effect of set similarity in

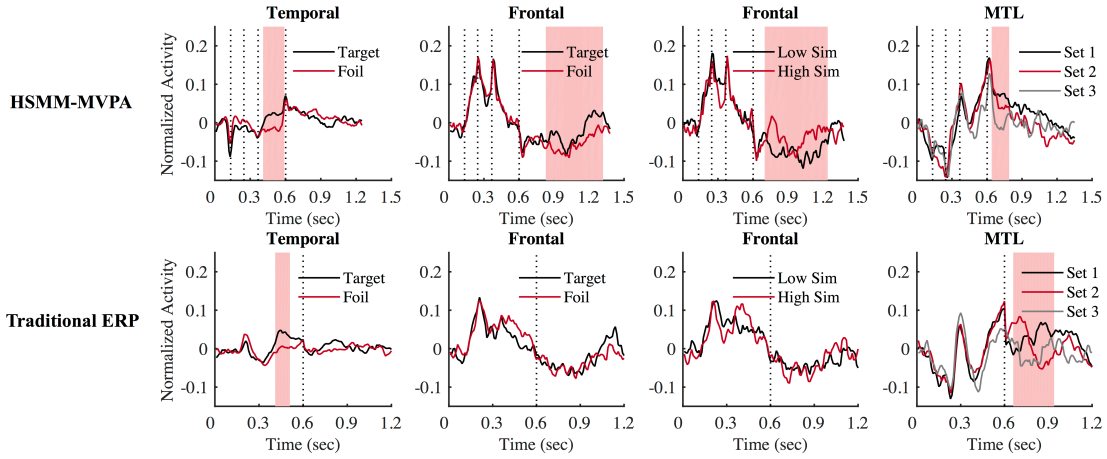

Figure 3.7: Time periods (over sliding window of 100 ms) over stage-locked average activity (top row) and stimulus/response-locked average activity (bottom row) where there are significant differences for different probe types, set similarity and set sizes in ECoG data, adjusted for multiple comparisons with Bonferroni correction. The dotted lines mark average positions of four bumps in the stage-locked activity (top row), and separation between stimulus-locked activity and response-locked activity (bottom row).

the Retrieval stage (F(1,28) = 6.227; $p = .019$) and the Decision stage (F(1,28) = 14.943; $p = 0.001$). There is effect of set size in the Decision stage (F(2,56) = 29.923; $p < 0.001$) but not in the Retrieval stage (F(2,56) = 1.260; $p = .29$). To summarize, the Retrieval stage is shorter for targets than for foils, consistent with the idea that it is easier to retrieve the memory set given a member of that set. In addition, the Retrieval stage is faster when the similarity between foil and memorized items is high. The duration of the Decision stage is affected by all three experimental factors: probe type, set similarity and set size.

3.3.5 Stage-locked brain activity by condition with ECoG

We exploit the spatial resolution of the ECoG data to examine where the differences in brain activity across different conditions occur. In event-related potential (ERP) analyses, the neural signal is anchored to observable events such as the presentation of a stimulus. The bumps obtained from the HSMM signal latent points of change in information processing. These events can be used, along with observable events, to align the neural

data. We anchored ECoG data from each trial according to stimulus onset, response, and the maximum likelihood locations of each of the four bumps during that trial. We then expanded or contracted the resulting five intervals in every trial to have durations equal to mean durations specified by HSMM-MVPA. In this way, the stimulus, the locations of the four bumps, and the response are aligned across all trials. Over such stage-locked data, we can examine where and when different conditions lead to different patterns of brain activity. Figure 3.7 (top row) highlights time periods and brain regions where there are significant differences of stage-locked brain activity between any two conditions using paired t-tests over a sliding window of 100 ms (dashed line represents average bump positions). Multiple comparisons across regions and time windows were corrected with the Bonferroni correction. The bottom row in Figure 3.7 is a comparison using traditional ERP analysis with 600-ms stimulus-locked and 600-ms response-locked data (separated by the dashed line).

In the stage-locked brain activity, we observe more positivity for targets than foils in the temporal cortex before the Retrieval stage completes. The temporal cortex has previously been associated with face familiarity (Gainotti, 2007) and is known to be instrumental in determining stimulus familiarity (e.g., Borst, Ghuman, & Anderson, 2016; Diana, Yonelinas, & Ranganath, 2007; Henson, Shallice, & Dolan, 2005; Gonsalves, Kahn, Curran, Norman, & Wagner, 2005; Rugg & Yonelinas, 2003). During the Decision stage, when the items are in active maintenance, targets are associated with more positive deflections than foils in the frontal cortex, which is known to be involved in post-retrieval monitoring and maintenance (e.g., Achim & Lepage, 2005; Borst et al., 2016; Mitchell, Johnson, Raye, & Greene, 2004; Rugg et al., 2003). Similar patterns in the frontal cortex also show up in different levels of set similarity among foils, where a higher set similarity corresponds to more positive ECoG amplitudes. During the same period in the Decision stage, there is an ordered effect of set size in MTL, with set size 1 being the most positive. This observation extends what we know about the MTL in maintaining items in working memory. Increasing working memory load (set size) was associated with elevated negativity of evoked response potentials of hippocampus during the delay period (Axmacher et al., 2007). In our study, this pattern also extends to the period after memory retrieval (Bump 4), when the retrieved face(s) need to be actively maintained during the Decision

stage.

Comparison with the traditional ERP analysis demonstrates the power of the HSMM-MVPA method in modeling cognitive events on a trial-by-trial basis. In the ERP analysis, we observe more positivity for targets than foils in the temporal cortex before the Retrieval stage completes, similar to the stage-locked activity with the HSMM-MVPA method. Such consistency is attributed to low trial-to-trial variability when data are closely time-locked to the stimulus. However, when moving further away from the stimulus, in contrast to the stage-locked brain activity, there are no significant effects across conditions in the frontal cortex. There is a significant effect of set size across conditions in the MTL in the Decision stage, but this difference is not ordered by the set size.

3.3.6 Phase synchrony between medial temporal lobe and cortical areas

ECoG data provides the unique opportunity to study the properties of oscillations in subcortical regions such as the MTL. In this section, we examine the synchrony between MTL and different cortical regions during the retrieval process.

Brain synchrony measures the relation between the temporal structures of the brain signals, and is considered as an important mechanism for integrating activity from across distributed brain areas into coherent perception and behavior (Varela, Lachaux, Rodriguez, & Martinerie, 2001; Fries, 2009; Fell & Axmacher, 2011). Simultaneous recordings in the hippocampus and prefrontal regions in animal studies have revealed synchronized theta oscillations during working memory tasks (Siapas, Lubenov, & Wilson, 2005). To examine whether those also occur in humans around the memory retrieval stage, we focus on phase synchrony between MTL and different cortical areas in the current experiment. Specifically, we examine if there is any transient change in phase synchrony upon completion of memory retrieval (i.e. when transitioning from the Retrieval stage to the Decision stage; cf. Portoles et al., 2018).

We measured synchrony with phase locking values (Lachaux et al, 1999). Figure 3.8a plots the phase locking values between MTL and 4 cortical regions in different frequency bands. We compared synchronization during two distinct periods (each of a 100 ms duration) across all trials: 1) samples right before the completion of the Retrieval stage; 2)

samples right after completion of the Retrieval stage (i.e. the onset of the Decision stage). Synchrony for a particular cortical electrode was averaged across phase locking values calculated with each of the MTL electrodes of the same subject. Electrodes in a cortical region from different subjects were then pooled together to obtain the standard error of the means shown in the figure. We were able to include 7 out of 12 subjects with more than 2 MTL electrodes in this analysis. Comparing the period right before and right after when the retrieval completes, there is increased theta phase synchrony between MTL and frontal recording sites (34 out of 47 electrodes; $p = .003$, binomial test). This increase is specific to theta band, and is not significant to alpha band (19 out of 47 electrodes; $p = .24 => .05$, binomial test) or beta band (22 out of 47 electrodes; $p = .77 > .05$, binomial test). We also observed increased phase synchrony between MTL and recording sites in the temporal cortex in the theta band (125 out of 190 electrodes; $p < .0001$, binomial test), but not in alpha band (105 out of 190 electrodes; $p = .17 > .05$, binomial test) or beta band (101 out of 190 electrodes; $p = .42 > .05$, binomial test). There is no significant theta coupling between MTL with either the occipital cortex (22 out of 55 electrodes; $p = .18 > .05$, binomial test) or the parietal cortex (10 out of 18 electrodes; $p = .81 > .05$, binomial test).

To further test if the theta coupling with MTL is significant on the level of individual electrodes, we built surrogate data by randomly shuffling the two time periods examined (i.e. before and after retrieval completes), and calculated the resulting phase locking values. This procedure is repeated 1000 iterations for between each pair of electrodes examined, and its corresponding increase in theta coupling with MTL is only considered significant if the amount of increase is larger than 95% of the times in the surrogate data. In total, there are 40 out of 190 electrodes in the temporal cortex that demonstrate significant increase in theta coupling with MTL, and 16 out of 47 electrodes in the frontal cortex. Figure 3.8b shows that, despite that the 16 frontal electrodes are from 5 out of the 7 subjects examined, there is considerable across-subject consistency in their locations (after being mapped to a common brain).

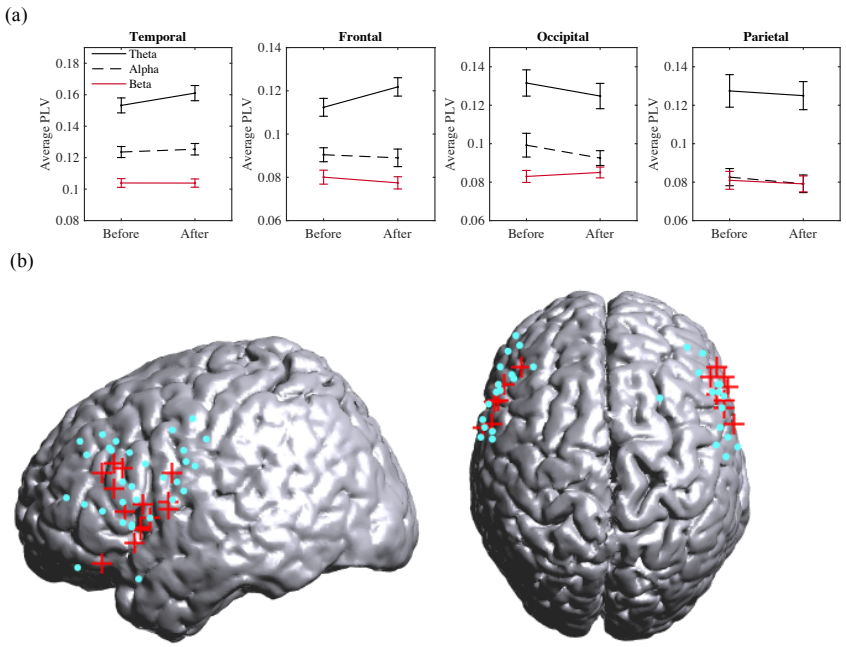

Figure 3.8: (a). Phase synchrony in theta, alpha and beta bands between MTL and cortical regions in 100-ms periods before and after retrieval completes. Synchrony for a particular cortical electrode is averaged across phase locking values calculated with each of the MTL electrodes of the same subject. Cortical electrodes across 7 subjects are then pooled together to obtain the standard error of the means. (b). Frontal electrodes with significant increase in phase-locking values with MLT in the theta band during the 100-ms after retrieval compared with before (red marker). The remaining non-significant frontal electrodes are in blue.

Table 2
A summary of identified effects on the duration and the brain activity across different processing stages.

	Retrieval Stage	Decision Stage
Duration	i Probe type	i Probe type
	ii Set similarity ↓	ii Set size ↑
		iii Set similarity ↑
Temporal Cortex	i Probe type	i Synchrony with MTL
MTL		i Set size ↓
		ii Synchrony with frontal cortex
		iii Synchrony with temporal cortex
Frontal Cortex		i Probe type
		ii Set similarity ↑
		iii Synchrony with MTL

↑ indicates more positivity or longer duration; ↓ indicates the opposite.

3.4 Discussions

In this study, we provided a detailed mapping of the time course and neural correlates of the retrieval process underlying visual working memory. Detailed mapping of the time course was achieved by capturing the trial-to-trial variability of different cognitive processes using the HSMM-MVPA method, instead of examining averaged neural activity across trials locked to observable events in traditional ERP analysis. HSMM-MVPA decomposed each trial into a sequence of latent stages. By examining how the duration of the HSMM-MVPA stages differed between task conditions, combined with our knowledge of a process model, we found evidence for the existence of the following cognitive stages in the visual working memory task: After a brief period for the visual signal to reach the brain (i.e. pre-attention), participants first encode the probe face (i.e. encoding), then retrieve faces from the memorized set (i.e. retrieval), and lastly, compare the retrieved set of faces with the probe to make a decision (i.e. decision). Once a detailed temporal mapping of the task was achieved with HSMM-MVPA, fine spatial resolution in ECoG was used to examine the neural correlates associated with each identified cognitive stage. Main effects centered around the transition from the Retrieval stage to the Decision stage (after completion of

retrieval). These effects are summarized in Table 2 and will be discussed in the following sections.

3.4.1 Isolation of a retrieval process prior to the decision-making

It has been widely assumed that retention of information in the working memory relies on maintenance of an active memory trace (Fuster and Alexander, 1971). However, the maintained information can be fragile if attention is temporally directed away or if the information is not amenable to rehearsal (Jeneson & Squire, 2011). In that case, memory can be preserved across a short delay without active maintenance (Lewis-Peacock et al., 2012), and retrieved later.

In this study, we were able to identify such a retrieval process, with EEG and ECoG data yielding converging results on a Retrieval stage isolated from the Decision stage. In the Retrieval stage, we observed more positivity for targets than foils in the temporal cortex, reflecting a fast and automatic process that does not require item details (Clark & Gronlund, 1996; Raaijmakers & Shiffrin, 1992). Once retrieval is complete, in the decision stage, we observed evidence of active maintenance of the just-retrieved items, with an effect of set size in MTL similar to that in Axmacher et al. (2007), and an effect of probe type in the frontal cortex in supporting post-retrieval monitoring (e.g., Achim & Lepage, 2005; Borst et al., 2016; Mitchell, Johnson, Raye, & Greene, 2004; Rugg et al., 2003). This provides support for cue-based retrieval theories of working memory.

3.4.2 Duration of the Retrieval stage and the Decision stage

How the duration of a particular cognitive stage changes across different experimental conditions provides important information on the nature of the underlying process. Typically, the effect of an experiment factor on a particular stage is reflected in the overall reaction times (RTs). However, given that RT only provides a cumulative measure of all the cognitive stages involved in a particular trial, there is not enough information in RT alone to isolate the effect of a particular experimental factor when there is more than one cognitive stage affected. In this study, we applied the HSMM-MVPA method to obtain durations of individual cognitive stages, and identified effects in stage durations when both

the Retrieval stage and the Decision stage vary across conditions.

In particular, we observed from the EEG data that the duration of the Retrieval stage depends on both probe type and set similarity. When the probe is one of the previously memorized faces (i.e., targets), retrieval is faster than when the probe is a foil because it is easier to reactivate the memory set given a member of that set. In addition, retrieval speed also depends on the set similarity. According to the ACT-R theory, greater similarity between foil and memorized item will result in greater activation for the retrieved item, and consequently shorter retrieval time (Anderson, 2007). The Retrieval stage does not depend on set size, which is comparable to the Sternberg task analyzed in Anderson et al. (2016).

The duration of the Decision stage is affected by all three factors: probe type, set similarity and set size. There are two possible mechanisms underlying the Decision stage: One possibility is that there is a serial comparison procedure where the probe is compared with each of the items in the retrieved set to find a match. The other possibility is that there is an evidence-accumulation procedure where a decision is driven by the strength of the overall similarity between the probe and the retrieved set of items – basically the summed similarity discussed in van Vugt et al. (2013), van Vugt et al. (2009), Kahana and Sekuler (2002) and Examplar-Based Random Walk models in Nosofsky et al. (2011).

A serial self-terminating decision process would be consistent with the data: Targets will yield a faster Decision stage due to earlier termination once a match is found (Figure 3.6b). In foils, higher set similarity slows down each comparison thus giving rise to a longer total Decision time (Figure 3.6c). Larger set size corresponds to a larger number of comparisons, and therefore a longer Decision stage (Figure 3.6d). A 2 (probe type) x 3 (set size) repeated-measures ANOVA shows significant effect of interaction between probe type and set size ($F(2,56) = 4.68$, $p = 0.01$). This is also consistent with the assumption that the serial comparison is self-terminating because the difference between targets and foils increases with set size.

If the Decision stage involves evidence accumulation, that would also be consistent with the data: Low similarity will lead to faster responses in correct foils (Nosofsky et al., 2011; Ratcliff, 1978; Bogacz, Brown, Moehlis, Holmes, &Cohen, 2006). The effect of set size on the duration of the Decision stage can also be explained by set similarity, as a smaller set size corresponds to higher set similarity which leads to faster response

in correct targets. In addition, smaller set size also increases average memory strength with briefer time lags, which gives rise to faster decision time under the evidence accumulation account (Nosofsky et al., 2011). Under either mechanism, the observation that increased set similarity speeds up the Retrieval stage while slowing down the Decision stage is consistent with an earlier study with an associative recognition task (Zhang et al., 2017).

3.4.3 Role of frontal cortex and MTL during the Decision stage

fMRI studies have identified that the MTL and prefrontal cortex support the general retrieval process of working memory tasks (Bledowski et al., 2006; ztekin, McElree, Staresina, & Davachi, 2009). However, given the low temporal resolution of fMRI, isolation of detailed processing stages such as the retrieval stage and the decision stage is difficult (but see Borst & Anderson, 2013, 2017, for an application of model-based fMRI analysis that attempts to tease apart working memory updating and retrieval). In contrast, the high temporal resolution of applying the HSMM-MVPA method to EEG and ECoG data enables us to divide the working memory task into a sequence of stages. In particular, retrieval of the studied face set marks the transition to a period of active maintenance of items, where information in the retrieved set is compared with the probe before a decision is made (Figure 3.6a). The active maintenance is supported by MTL, where we observed more negativity when the set size is large. Negativity in hippocampal activity has been associated with increasing workload (i.e., set size) in visual working memory task during the delay period, when there is active maintenance of items right after they are encoded (Axmacher, 2007). With the finer temporal resolution of HSMM-MVPA, the current experiment extended this result to the decision period of a working memory task after the items are just retrieved.

During the Decision stage, we also observed more positive activity in frontal cortex for targets than foils (comparable to the process described in Borst et al., 2016). In working memory tasks, there is debate concerning whether the involvement of prefrontal cortex contributes to the maintenance of items or the selection of an item from memory to guide a response (Curtis & D'Esposito, 2003; Rowe, 2000). Our experiment supports the latter, with an effect for target/foil but not for different set sizes in frontal cortex. This interpretation is consistent with the process during the Decision stage, where an item needs to

be selected to match with the probe. When there is a match (i.e., targets), a more positive response is triggered in the frontal cortex compared with that of non-matches or foils. An alternative mechanism for the Decision stage could be that instead of selecting and comparing each item in the retrieved face set against the probe, the frontal cortex guides the decision in an evidence-accumulation procedure driven by the similarity between the probe and the entire retrieved face set (i.e., set similarity; Nosofsky et al., 2011). This latter account is supported by more positivity in frontal cortex for high-similarity trials than for low-similarity trials, in addition to more positivity for targets than for foils.

We also found evidence that the frontal cortex guides the decision by means of theta coupling with the MTL once retrieval has been completed. Previously, it has been suggested that 4–9Hz theta power in the hippocampus is associated with encoding and retrieval of episodic memories (Lega, Jacobs, & Kahana, 2011). Functional coupling between prefrontal cortex and medial temporal cortex is considered one of the key connections in the neural circuitry underlying working memory tasks, with theta oscillations proposed to mediate this interaction (Mitchell, McNaughton, Flanagan, & Kirk, 2008; Anderson, Rajagovindan, Ghacibeh, Meador, & Ding, 2009). This is supported by multiple animal studies that demonstrate phase-locking in the theta band between prefrontal cortex and hippocampus (Siapas et al., 2005; Hyman, 2010; for a review see Colgin, 2011). In human studies, PFC-hippocampal coupling, both structural (Cohen, 2011) and functional (Campo at al., 2011), has been shown to correlate with individual differences in task performance. However, PFC-hippocampal communication through theta oscillations in human working memory tasks has not been directly observed, though it is shown to be important in free recall and associative recognition memory tasks (Anderson et al., 2009). The observation of frontal-MTL theta coupling generalizes the role of MTL theta from retrieval in episodic memory to retrieval of previous items in working memory. This is consistent with recent studies that MTL is not uniquely involved in long-term memory, but also critical to short-term memory even when the retention period is as short as 2-10s (Holdstock, Shaw, & Aggleton, 1995; Owen, Sahakian, Semple, Polkey, & Robbins, 1995; Holdstock, Gutnikov, Gaffan, &Mayes, 2000; Aggleton, Shaw, & Gaffan, 1992; Hannula, Tranel, & Cohen, 2006, van Vugt et al., 2010).

There are several limitations to our research. First, the analysis of cortical-subcortical

interaction was limited by the number and locations of recording sites in each ECoG subject. Therefore, only pairs of electrodes were examined at one time and pooled across all subjects in the end. Second, the HSMM-MVPA method successfully isolated different cognitive processing stages, but it did not provide enough evidence to distinguish whether the decision stage undergoes a serial comparison process or an evidence accumulation process. It would take more experimental studies with targeted manipulations to make this distinction in the future.

3.5 Conclusion

In this study, we extended a previous account of the visual working memory task using a summed-similarity model (van Vugt et al., 2013) to one that comprises multiple sequential cognitive stages. Combining the temporal resolution of the EEG data, the spatial resolution of the ECoG data, and the application of the HSMM-MVPA method, we were able to identify the time duration and brain activity associated with these stages. In contrast to traditional ERP analyses which only models the effects that are closely locked to the beginning and the end of a trial, the HSMM-MVPA method applied to both experiments isolated a Retrieval stage where memorized items were re-activated, followed by a Decision stage. In contrast to examining only the overall RT, the HSMM-MVPA method applied to EEG data revealed how durations of the Retrieval stage and the Decision stage vary across different experimental conditions. Combined with fine spatial resolution of ECoG data, it was identified that frontal cortex and MTL play a key role in response selection and item maintenance respectively. The effect of set size observed in MTL generalizes its role in actively maintaining items from the delay period in working memory tasks to the decision period once items are re-activated. The theta coupling between frontal cortex and the MTL generalizes previous findings that were considered unique to long-term memory tasks to working memory tasks. In addition, they provide support for a cue-based retrieval account of visual working memory.

Chapter 4

A rational account of human semantic search

A central paradigm to study human memory search is the semantic fluency task, where participants are asked to retrieve as many items as possible from a category in a fixed amount of time. Observed responses tend to be clustered semantically. To understand when the mind decides to stop exploiting the current patch and to start exploring the next, a strategic memory search process using the marginal value theorem has been proposed, similar to optimal foraging in a spatial environment. Later, however, it has been shown that behavior patterns consistent with the marginal value theorem can also emerge from a random walk over a semantic network, without any strategic switch. The fact that a strategic switching mechanism and a non-strategic switching mechanism give rise to similar behavioral patterns poses challenges in understanding the exact cognitive mechanism used by humans. In the current work, in addition to comparing existing mechanisms over observed human data, we propose a rational analysis of the problem by examining what would be an optimal patch-switching policy under the framework of reinforcement learning. Built upon the random walk model and features of the local semantic patch, the resulted switching mechanism under reinforcement learning performs better than the marginal value theorem and gives a better fit to single-trial human behavioral data. Our results provide theoretical justification of switching mechanisms used in human memory research, and shed light on how a rational account of the task can generate alternative hypotheses about human

cognitive mechanism in the same task.

4.1 Introduction

How do we recall items on a shopping list? Performing a daily task like this requires searching through our memories. There is rich dynamics in during the memory search when we move from recalling one item to the next. Uncovering the underlying cognitive mechanism is important for better understanding of how knowledge and information are represented in our memory, and the processes that we use to efficiently navigate through them.

A central paradigm to study human memory search is the semantic fluency task, where participants are asked to retrieve as many items as possible from a category (e.g. animals) in a fixed time period. It is observed that responses tend to be clustered semantically (e.g. "cat"' follows "dog") (Hills et al., 2012). Evidence over human behavioral data shows that the marginal value theorem accounts for how our minds decide to switch from one cluster/patch to the next (Hills et al., 2012). Under this mechanism, people make strategic decisions to search the semantic space, similar to how animals optimally forage in a patchy spatial environment: one forages locally in one food patch, then switches to a new patch when the resources in the current patch is depleted. It was observed that participants leave a patch in memory search when current rate of finding items is near the average rate for the entire task (Hills et al., 2012), consistent with what the marginal value theorem predicts in optimal foraging (Charnov, 1976). Recent work, however, demonstrated that similar behavioral patterns, that are consistent with the marginal value theorem, can also emerge from a random walk simulation on a semantic network generated by human word-association experiments (Abbott et al., 2015). The random walk model moves from one patch to another stochastically, without basing its decision on the information about the current patch. The fact that a strategic switching mechanism and a non-strategic switching mechanism predict similar temporal profiles around the switches poses challenges in understanding the exact cognition mechanism used by humans.

The complication in the literature motivates us to further examine if humans use the marginal value theorem during the memory search. To provide more evidence in the com-

parison across alternative mechanisms, in addition to examining human data, we will also consider the abstract computational problem posed by searching the memory and explore what would be an optimal mechanism in this task (i.e. which mechanism can generate the most items in a semantic fluency task). This approach is based on the principle of rationality, which explain human behavior as an optimal solution to the computational problems posed by our environment (J. Anderson, 1990; J. R. Anderson & Milson, 1989); see also bounded rationality in (Simon, 1978), resource-rational analysis in (Griffiths, Lieder, & Goodman, 2015) and ecological rationality in (Todd & Gigerenzer, 2007)). Examining which cognitive mechanism best solves the computational problem gives additional justification on why it should be used by humans.

The principle of rationality can be applied not only to compare existing cognitive mechanisms, i.e. the marginal value theorem versus the non-strategic random walk, but also to propose new hypothesis of alternative mechanisms. This is especially the case when existing mechanisms are not optimal for the given task. Marginal value theorem decides whether to switch by comparing an instantaneous reward to an overall average rewards; it is optimal under a set of conditions, including the assumption that local but not global resources are depleted during the course of the search (Charnov, 1976). This assumption does not hold in the semantic fluency task, as humans can re-enter the same patches, and they are depleted over time as more items are recalled. To perform well in this task requires a decision policy that takes into account such dynamics in the environment. We propose an alternative mechanism based on reinforcement learning that directly optimizes performance under this specific task environment. We will apply temporal-differencing algorithm during reinforcement learning, as there is considerable neural and behavioral support for TD learning in humans and other animals (Houk & Adams, 1995; Montague, Dayan, & Sejnowski, 1996; Niv, 2009; O'Doherty, Dayan, Friston, Critchley, & Dolan, 2003). In contrast to the marginal value theorem that relies on estimating the immediate reward, TD learning is capable of capturing non-immediate values of switching/non-switching in sequentially encountered time steps. Instead of considering the value of switching as unknown and approximating it as the overall average reward rate in the marginal value theorem, TD learns the value of switching directly, similarly to learning the value of non-switching (Sutton, 1988). To assume minimally the amount of

information humans can access, the policy obtained from reinforcement learning utilizes only information from immediate neighbors of current location in the semantic network.

The principle of rationality would predict that the cognitive mechanism that has an advantage in achieving better task performance is also more likely to be used by human cognition. Therefore, once we compare performance of different mechanisms in simulations, we will test if the switching mechanism with better performance in the semantic fluency task will also account for patch switches in the human behavioral data better.

The remaining paper is organized as below. We will first describe the random walk model that is used to simulate the semantic fluency task. We will then describe alternative switching mechanisms, and how to optimize task performance under each switching mechanism. Then we will evaluate them by comparing both their relative performances in simulations and their fit to human behavioral data. We will show that the proposed mechanism with reinforcement learning based on local patch quality outperforms the marginal value theorem in how well it achieves the computational goal of the task, and in its fit to single-trial human behavioral data.

4.2 Methods

In this section, we will first describe how to simulate the semantic fluency task under different switching mechanisms, closely following the procedure used in Abbott et al. (2015). Then we will discuss how to optimize task performance under each switching mechanism. Lastly, we will describe methods to compare different switching mechanisms over human behavioral data.

4.2.1 Simulate the semantic fluency task

The random walk considered by (Abbott et al., 2015) operates over a semantic network with 5018 words (165 of them are animal names) (Nelson, McEvoy, & Schreiber, 2004). This semantic network is obtained from human behavior database in a word-association task, where more than 6,000 participants responded with the first word that came into their mind when cued with another word. The weight of the edge from word A to word B is the proportion of participants that responded B when cued with A. Following closely the pro-

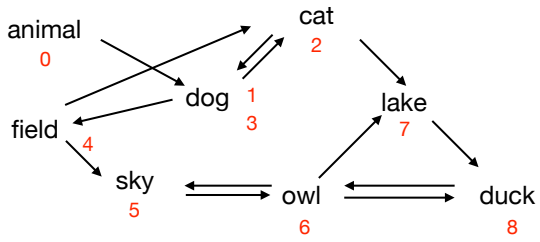

Figure 4.1: An example of random walk without jumping back to "animal" node. It always starts from the node "animal" as the initial cue. The numbers mark the sequence of nodes visited. A response is generated only when the visited node is both an animal and has not been visited before. In this example, the sequence of responses generated are: "dog", "cat", "owl', "duck".

cedure described in (Abbott et al., 2015), a random walk starts from the node "animal" as the initial cue, and generates a behavioral response every time it visits an animal name for the first time. At every time step, it goes from current node to the next node under two options. Under the first option, it transits from the current node X_i to the next node X_j with the probability being the edge weight: $P(X_j|X_i) = w(X_i, X_j)$. Walking to the neighboring nodes is equivalent to locally exploiting the resources. Figure 4.1 demonstrates an example of random walk under the first option. Under the second option, it jumps back to the node "animal" and continues the random walk from there. Opportunity to switch away from the current node allows for exploring animals at other parts of the semantic network. The three different switching mechanisms we will compare specify three different policies in deciding whether to switch at each time step.

4.2.2 The first switching mechanism: marginal value theorem

The first switching mechanism is based on the marginal value theorem, which is an important model that characterizes optimal foraging behavior in the literature of animal foraging (Charnov, 1976). It is based on the assumption that local resources are monotonically depleted during foraging. Under this environment, animals seek to maximize the gain per unit time. Marginal value theorem describes that the optimal foraging policy is to leave a patch when the instantaneous rate (i.e. marginal value) of gain is equal to the long-term average intake under the current environment. When it transfers to memory search,

the marginal value theorem predicts that individuals should leave the current patch and switch to another when the gain falls below the expected benefits of searching elsewhere in memory (see more details in Hills et al. (2012)). Given fixed reward per retrieved item, the current switching mechanism describes that one should leave a patch once the instantaneous time cost (i.e. time cost of the just-retrieved item) exceeds the average time cost under the environment. Equivalently, given fixed cost per time step, a different formulation of the marginal value theorem describes that one should leave a patch once the instantaneous gain estimated from the local semantic patch falls below the global average gain under the environment.

Assume that we do not know the global average of the task in advance. To optimize the task performance under this switching mechanism, we will explore a range of time thresholds, and select the one that demonstrates the best task performance.

4.2.3 The second switching mechanism: non-strategic switching

The second switching mechanism describes a scenario where switching patches takes place in a random and non-strategic way (Abbott et al., 2015). At each time step, there is a fixed probability p that the random walk jumps back to the start node "animal"; the other times under probability $(1 - p)$, it transits from the current node to the next node, with a probability equal to the edge weight. To optimize the task performance under this switching mechanism, we will explore a range of values p, and select the one that gives the best task performance.

4.2.4 The third switching mechanism: reinforcement learning based on local patch quality

In the current work, we propose a third switching mechanism based on the local patch quality, described by local statistics around the currently visited node in the random walk.

It is unclear how the information of various local statistics of the currently visited node converts to a switch/not-switch decision. To optimize the task performance under this switching mechanism, we obtain this mapping under reinforcement learning. We consider that a random walk agent interacts with the environment in a sequence of actions,

observations and rewards. At each time step, the agent selects an action a_t from a binary set $A = \{0, 1\}$, which represents whether to switch or not. The entire semantic graph that the random walk operates on is not visible to the agent. Instead, we make minimal assumptions of the information accessed by the agent by exposing it only to the information of its nearest k neighbors. There can be three types of information associated with each neighbor node, including 1) the edge weight $w \in \mathbb{R}$, 2) whether it is a non-visited animal node $I = \{0, 1\}$, and 3) how many times it has already been visited $n \in \mathbb{Z}_{\geq 0}$. These information are concatenated as the state information for the model $S = \{w_1, w_2, \ldots, w_k, I_1, I_2, \ldots, I_k, n_1, n_2, \ldots, n_k\}$.

One simulated trial is considered as one episode, which terminates when the time is up. As in Abbott et al. (2015), the time limit is set to be 2000 steps, which corresponds to performance level of the agent at the same scale as the human experiments (Hills et al., 2012). At each time step t, the agent receives a reward $r_t = \{0, 1\}$, with 1 representing an animal name is successfully generated and 0 otherwise.

The goal of the agent is to find the switching mechanism that selects switch/not-switch in a fashion that maximizes cumulative future reward (discounted by a factor of γ per time step). This rule can be obtained from the optimal action-value function:

$$Q^*(s, a) = \max_{\pi} E[r_t + \gamma r_{t+1} + \gamma^2 r_{t+2} + ...|s_t = s, a_t = a, \pi]$$

which is the maximum sum of rewards r_t achievable by a behavior policy $\pi = P(a|s)$, after making an observation s and taking an action a. $Q^*(s, a)$ obeys the Bellman equation:

$$Q^*(s, a) = E[r + \gamma \max_{a'} Q^*(s', a')|s, a]$$

where s' at the next time step is known for all possible actions a'. Using the Bellman equation as an iterative update, i.e. $Q_{i+1}(s, a) = E[r + \gamma \max_{a'} Q_i(s', a')|s, a]$, it will converge to the optimal action-value function $Q_i \rightarrow Q^*$ as $i \rightarrow \infty$ (Sutton & Barto, 1998). To allow the representation to emerge flexibly from experience, we use a deep neural network to approximate the action-value function $Q(s, a; \theta) \approx Q^*(s, a)$. This results in a deep Q-Network (DQN) with weights represented by θ (Mnih et al., 2015). It is trained by adjusting the parameters θ_i at iteration i to reduce the mean-squared error in the bell-

man equation. We then use the Q-learning algorithm to update weights by optimizing the loss function with stochastic gradient descent Watkins and Dayan (1992). The behavior distribution during training uses a ϵ-greedy strategy, with follows greedy strategy with probability (1-ϵ) and selects a random action with probability ϵ. The value of ϵ represents the amount of exploration, and is set to decrease gradually and reaches 0 at the end of the training. To model the switching mechanism with a simple heuristics, we adopt a small neural network architecture composed of two hidden layers with ten and five hidden units respectively.

To summarize, the obtained DQN describes the third switching mechanism. At each time step, the action to take is the one with larger Q-network value given the local patch quality. The resulting policy is not guaranteed to be optimal for the semantic fluency task, given that the full state of the environment is not available to the agent. However, we are interested in a policy that has the potential to outperform the marginal value theorem.

4.2.5 Experimental data

We focus on a semantic fluency task from a public dataset Hills et al. (2012). There were 141 participants in this experiment at Indiana University, Bloomington. With a time limit of three minutes, participants were asked to type in as many items as possible for the given category "animal". There were 373 unique animals produced by these 141 participants in the experiment. More details about the experiment can be found in the original report in Hills et al. (2012).

4.2.6 Predict switches over human behavioral data

Next, we will compare which switching mechanism accounts for human behavioral data better. Human participants give responses whenever they think of an animal name; however, they do not give the full path of the random walk. In other words, human memory search is an initial-visiting emitting (INVITE) random walk (Jun, Zhu, Rogers, Yang, & Yuan, 2015), with animaml labels generated only when the node is visited for the first time. Given a behavioral response of a just-retrieved item, the dependent variable Y is whether there is a switch until the next response. The independent variable is a score X obtained

differently for each switching mechanism, characterizing information only concerning the previous response. The goal is to examine which switching mechanism can better predict the switch. We add a softmax decision noise to the switching mechanism (J. D. Cohen, McClure, & Yu, 2007; Mehlhorn et al., 2015):

$$p(Y = 1) = f(\beta_0 + \beta_1 X)$$

where $f(x) = \frac{1}{1+e^{-x}}$ is the logistic function. β_0 allows for a bias term in each switching mechanism. Larger β_1 values give more deterministic behavioral under the decision rule. Y is approximated with pre-defined animal categories, similar to the approach in Hills et al. (2012). Below outlines how X is obtained differently for each switching mechanism.

Under marginal value theorem, X is the experimental time cost of the just-retrieved item. Under the reinforcement learning framework, X is the probability that there is not a switch before retrieving the next animal, calculated as below:

Consider a random walk on a set of states S with an initial distribution π and transition matrix T, where T_{ij} is the probability of going from state i to state j. A surfer starts from an initial state drawn from π and outputs an animal state whenever visiting it for the first time. The naive method to compute the likelihood of starting from i and ending at j is intractable given that there can be an infinite number of trajectories in between. Jun et al. (2015) introduces a method to compute INVITE likelihood by considering the process of the initial-visit emitting random walk into a series of absorbing random walks (Jun et al., 2015) In absorbing random walks, transition matrix **T** can be written in a canonical form below:

$$\mathbf{T} = \begin{pmatrix} \mathbf{Q} & \mathbf{R} \\ \mathbf{0} & \mathbf{I} \end{pmatrix}$$

where **Q** is the transition between transient states, **R** the transition from non-absorbing states to absorbing states. The transition matrix from absorbing states to transient states is always **0**, and the transition matrix from absorbing states to absorbing states is always **I**. Under absorbing random walks, an essential tool of *fundamental matrix* can be applied:

Theorem. *(Doyle & Snell, 2000) The fundamental matrix of the Markov chain is* $N = (I - Q)^{-1}$. $N_{i,k}$ *is the expected number of times that a chain visits state* k *before absorption when starting from* i. *Furthermore, define* $B = (I - Q)^{-1}R$. *Then,* B_{ij} *is the*

probability of a chain starting from the transient state i being absorbed by the absorbing state j. In other words, $B_{i\cdot}$ is the absorption distribution of a chain starting from i.

Jun et al. (2015) applies the fundamental matrix by considering any two overt behavioral responses from the human behavioral data as a starting state and an ending state of an absorbing Markov chain; as a result, B_{ij} is the probability of generating item j given the previous response i. In the current work, we need to consider the additional possibility of the random walk making a switch before generating the next item. We therefore set up the absorbing Markov chain differently: the absorbing states include both the remaining animal items, as that in (Jun et al., 2015), and additionally, items that are identified as switching points by the RL model. The absorbing Markov chain terminates either when it walks to a non-visited animal item without encountering a switch (with probability P_1, which X is set equal to), or encountering a switch before reaching a non-visited animal item (with probability P_2). We have $P_1 + P_2 = 1$, as the absorbing Markov chain will eventually terminate in one of the absorbing states. The rest of the nodes in the network are transient states, over which the random walk can keep visiting. Given what animal names have already been recalled, each pair of adjacent responses is considered a different absorbing Markov chain, with their unique set of absorbing states and transient states. Under this formulation, we can obtain: $X = P_1 = \sum_{j \in S} B_{ij}$, where transient state i is the just-retrieved item, and S is the set of absorbing states that are non-visited animal items among the neighbors of i.

4.3 Results

The first section is the simulation results that detail how performance under each mechanism varies across parameter settings, and how the best model under different switching mechanisms compare to each other. According to rational analysis, the switching mechanism that gives the best performance is more likely to be used by human participants. The second section further tests this hypothesis and examines how well different switching mechanisms predict switches in human behavioral data.

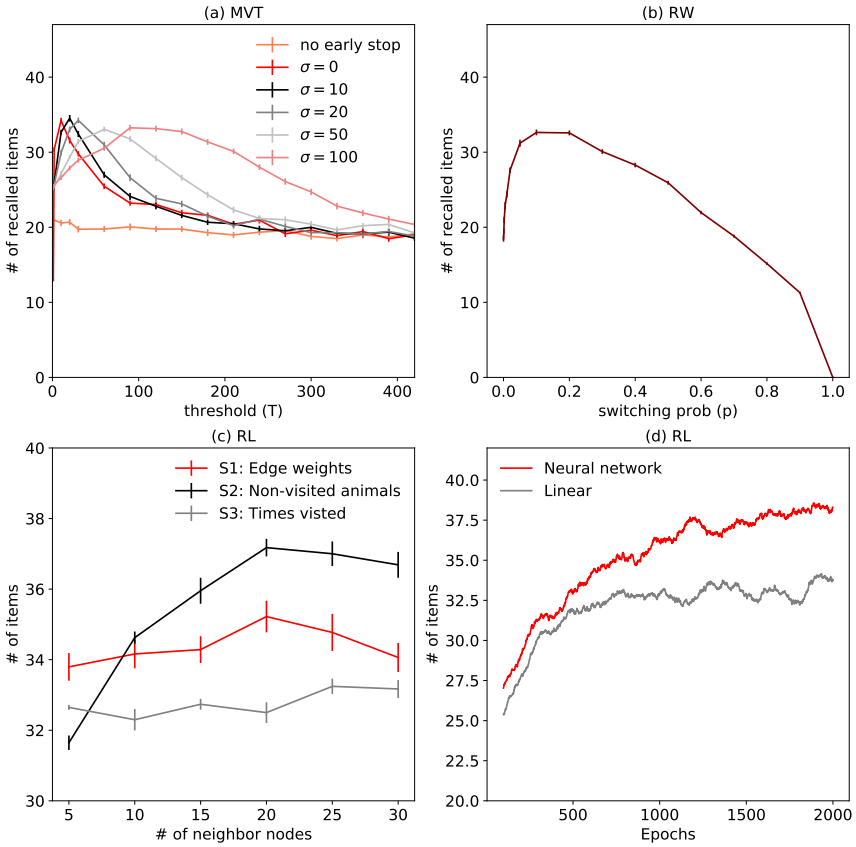

Figure 4.2: Simulation results of different switching mechanisms, measured as the number of animal names recalled in one trial. (a) Marginal value theorem with instantaneous time cost estimated as time steps passed since retrieving the last item, over 100 simulations, allowing for early stopping or not. Thresholds used in each decision is stochastic drawn from a normal distribution, with time threshold as the mean μ, examined under different values of standard deviations σ. (b) Random walk model under different switching probability p, over 100 runs. (c) Deep reinforcement learning model that takes different sources of input from the local patch quality (edge weights connecting current node to the neighbor node , the neighbor node being a non-visited animal or not, how many times the neighbor node has been visited), under increasing number of neighbor nodes considered, each over 10 runs of Q-learning algorithm. Learning curves of the best model over 10 runs of Q-learning algorithm, smoothed by averaging a sliding window of 100 episodes, for the reinforcement learning model that uses neural network (based on combining S1 and S2 for 10 neighbor nodes) to approximate the Q-function and one that uses a linear function (d). Parameters for the best models are $\mu = 20$, $\sigma = 10$ for MVT, $p = 0.1$ for RW, and the neural network model that combines S1 and S2 of 10 neighbor nodes for RL. Hyper-parameters used in the DQN under the reinforcement learning model: learning rate 0.01, number of episodes 2000, $\gamma = 0.95$, $\epsilon_{\text{max}} = 1$, $\epsilon_{\text{min}} = 0.0001$, $\epsilon_{\text{decay}} = 0.998$. All error bars are standard error of the mean.

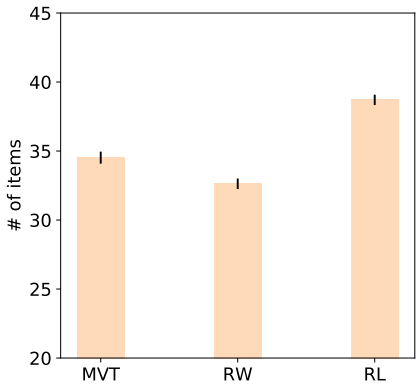

Figure 4.3: A comparison of the best models under each switching mechanism, over 100 simulations. Parameters for the best models are $\mu = 20$, $\sigma = 10$ for MVT, $p = 0.1$ for RW, and the neural network model that combines S1 and S2 of 10 neighbor nodes for RL. Hyper-parameters used in the DQN under the reinforcement learning model: learning rate 0.01, number of episodes 2000, $\gamma = 0.95$, $\epsilon_{\text{max}} = 1$, $\epsilon_{\text{min}} = 0.0001$, $\epsilon_{\text{decay}} = 0.998$. All error bars are standard error of the mean.

4.3.1 Performance in achieving the task goal

Figure 4.2 describes the simulation results of different switching mechanisms. Under a rational analysis, we compare which mechanism is more optimal in the memory search task, measured as the number of animal items recalled given a fixed number of time steps.

Figure 4.2a shows how performance under the marginal value theorem varies as a function of the time threshold used. MVT makes a switch based on how many time steps it has been since the last retrieval. This evaluation only takes place when a new item is just retrieved (i.e. no early stop), or at every single time step. A stochastic component is added by considering the time threshold drawn from a normal distribution at each time step of decision. Adding a stochastic component can be potentially beneficial when the environment itself is noisy. MVT performs better when evaluation takes place at every time step. The best MVT model corresponds to a time threshold of 20 with standard deviation of 10.

Figure 4.2b shows how performance under the non-strategic switching mechanism with a random walk model varies as a function of probability of switching (p) at a given time step. Performance when p is 0 corresponds to a random walk model without any

switches, whereas performance when p is 1 never leaves the starting node "animal", leading to a number of zero animal names generated. The model reaches its best performance when p is 0.1.

Figure 4.2c and Figure 4.2d examines the factors that affect the performance of the reinforcement learning model. In Figure 4.2c, three sources of information are considered to characterize the neighborhood of the current node. The first source of information (S_1) is the edge weight connecting the current node to a neighbor node. The second source of information (S_2) is whether a neighbor node is a yet-to-be-visited animal name. The third source of information (S_3) is a count of how many times a neighbor node has been visited. Figure 4.2c also examines the performance as a function of k nearest neighbors considered. One can observe that the first two sources of information but not the third gives performance that is better than the best model under non-strategic switching (i.e. 32.63). We therefore constrain further analysis to utilize only the first two sources of information, S_1 and S_2. Performance increases by the number of neighbor nodes for S_2 but less for S_1. A number of 10 is considered sufficient for further analysis, as it appears to be the smallest set of neighbors necessary to produce performance better than the other switching mechanisms. Figure 4.2d is a model that contains both S_1 and S_2 of 10 neighbor nodes. We run the Q-learning algorithm 10 times, and plot the training curve of the best model out of the 10 runs. This same procedure is repeated twice, approximating the Q-function with either a two-hidden-layer (10 nodes and 5 nodes) neural network or one that uses a linear function. The model performs better when a neural network is used, which will be referred to as the best RL model during subsequent analyses.

Figure 4.3 is a comparison across the best model under different switching mechanisms, each simulated in 100 episodes, representing 100 trials. The RL model, based on local patch quality, outperforms other switching mechanisms in the semantic fluency task. Note that the difference between the marginal value theorem and the RL model does not merely lie in the information utilized. An equivalent formulation of the marginal value theorem, based on instantaneous reward estimated from local patch quality instead of instantaneous time cost, gives the best performance of 22.83 items over a range of thresholds used. The instantaneous reward is the expected reward at the current time step, derived from the same set of information utilized in the RL model, S_1 and S_2.

4.3.2 Evidence over human behavioral data

In this section, the goal is to examine which switching mechanism can better predict if there is a switch before the next response, given information associated with the previous response. We add a softmax decision noise to the switching mechanism in a logistic regression. Figure 4.4a plots the instantaneous time cost, as the regressor in the logistics regression, for switches and non-switches in human data. The switches in the human behavioral data are identified from pre-defined categories (Troyer, Moscovitch, & Winocur, 1997). Figure 4.4b-c plots the non-switching probability used to predict switches with RL, for switches and non-switches in human data. As illustrated in Figure 4.4e, given the just-retrieved response i, one can calculate the probability of reaching the next animal name j before reaching a switching point defined by the RL model (P_1). When the non-switching probability is not zero, Figure 4.4b plots their corresponding log likelihood. Distribution of the non-switching probability for non-switches in Figure 4.4b is more skewed to the right compared with that of the switches. Similarly, when non-switching probability is zero (this happens when the previous response is one of the switching points defined in RL policy), Figure 4.4c shows that there is higher proportion of switches than non-switches. We formally test these differences in examining how well it can predict switches in a logistic regression. Using the non-switching probability in RL as a regressor, it returns a significant coefficient of -1.06210 ($p < .0001$). This fit is better than that of a logistic regression model using instantaneous time cost in MVT as a regressor (ΔBIC = 50.69). There is strong evidence to favor the RL model than MVT with log of the Bayes factor being 25.34 (Kass & Raftery, 1995), approximated from ΔBIC (Wagenmakers, 2007). Figure 4.4d further plots the difference in the log likelihood of the RL model and the MVT model at the level of individual subjects. There are 95 out of 141 subjects that favor RL model, compared with 46 out of 141 subjects that favor MVT.

4.4 Discussion

In a semantic fluency task, we compared existing mechanisms and proposed new mechanism that give rise to patterns of semantically clustered responses. The contribution of this work is twofold. First, we directly compare how optimal each switching mechanism

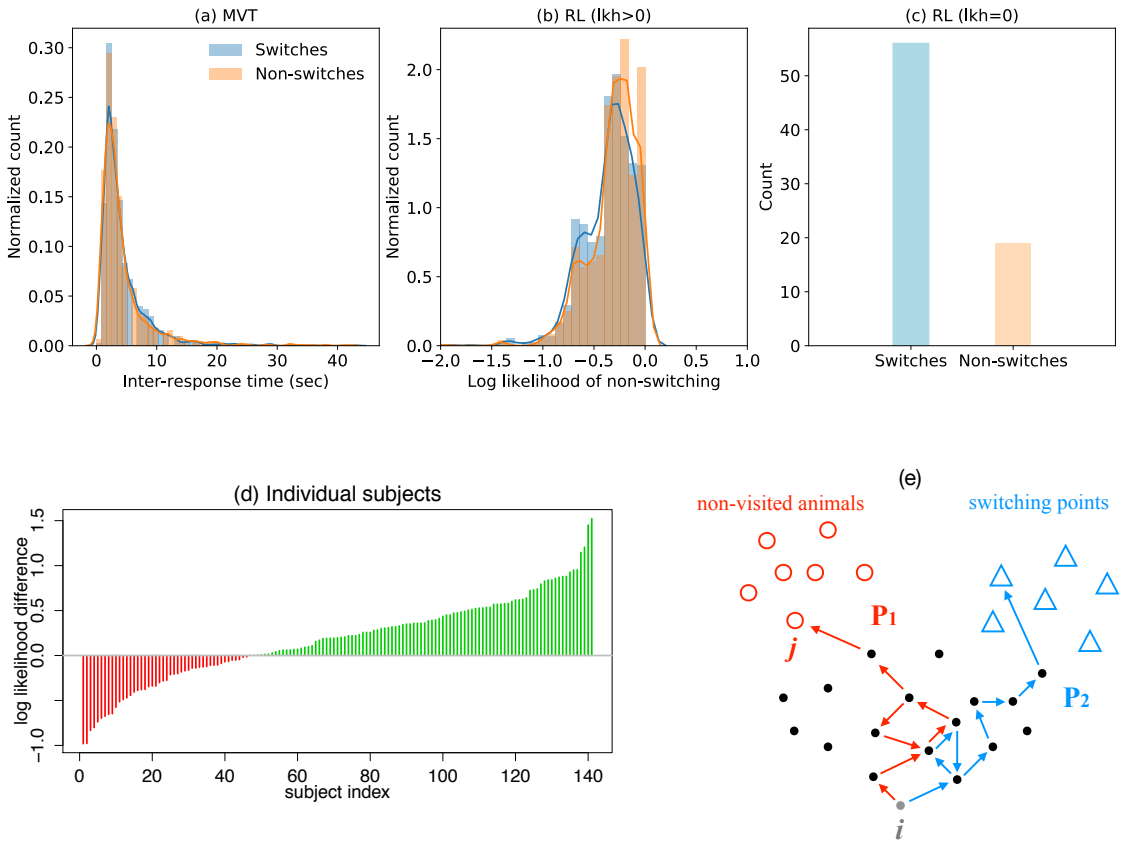

Figure 4.4: Predict switches over human behavioral data. (a) Instantaneous time cost used to predict switches with MVT, for switches and non-switches in human data. (b) Log of non-switching probability used to predict switches with RL, for switches and non-switches in human data, when the non-switching probability is non-zero. (c) Count of switches and non-switches in human data, when the non-switching probability is zero. (d) Likelihood difference in a logistic regression model between RL and MVT for each subject, sorted by the amount of difference. (e) Non-switching probability (P_1) in RL is obtained from an absorbing Markov chain that starts at the transient state i, the just-recalled item, and ends at an absorbing state either being a non-visited animal item or a switching point defined by the RL model. P_1 is the probability for the absorbing chain to end at any of the non-visited animal names. P_2 is the probability for the absorbing chain to end at any of the switching points.

is in simulations in a rational analysis. Second, instead of comparing different switching mechanisms over averaged patterns of human behavioral data, we examine how well they account for switches in single-trial human behavioral data. According to rational analysis (J. Anderson, 1990; J. R. Anderson & Milson, 1989), human cognition optimally solves the problems it faces; therefore, the cognitive mechanism that has an advantage in achieving better task performance is more likely to be used by human. Our results over simulations and human behavioral data provided such evidence: the proposed switching mechanism derived under a framework of reinforcement learning outperformed the marginal value theorem, both in how well it achieves the computational goal of the task, and in its fit over the human behavioral data. Below we will discuss assumptions underlying the marginal value theorem, and the plausibility of representation and learning mechanism underlying the proposed switching mechanism. Lastly, we will discuss further implication of the work and its connection to related literature of memory search.

4.4.1 Assumptions underlying the marginal value theorem

Decision in the marginal value theorem is based on comparing the instantaneous reward of staying with a threshold that measures overall environmental richness. First, it assumes that local resources are monotonically depleted during foraging. However, in the semantic fluency task, the instantaneous reward can vary from moment to moment. In contrast, the RL model does not base its decision on the instantaneous reward, and is capable of capturing accumulated rewards including sequentially encountered time steps. The marginal value theorem also assumes that the value of switching is unchanged throughout the task. This holds in cases where one enters a new patch of resources after each switch. However, in the semantic fluency task, humans can re-enter the same patches that are depleted over time as more items are recalled. Therefore, to perform well in this task, it requires a learning mechanism that takes into consideration of such characteristics in the environment: the RL model learns the value of switching directly, instead of considering the value of switching as unknown and approximating it as a constant value. To assume minimally the amount of information humans can access, the policy obtained from reinforcement learning utilizes only information from immediate neighbors of current location in the semantic network. This is similar to human foraging tasks in non-memory domains, where partic-

ipants are sensitive to the absence and presence of the resources locally (Hills, Kalff, & Wiener, 2013; Kalff, Hills, & Wiener, 2010).

Using a reinforcement learning framework that maximizes the cumulative rewards, we explored the set of information that can most effectively characterize local patch quality, and the policy that best describes how local patch quality maps to a switch/non-switch decision. The reinforcement learning model based on local patch quality (i.e. edge weight, whether being a non-visited animal) outperformed the marginal value theorem in terms of the total number of responses generated in simulations, and in accounting for switches in the human behavioral data.

4.4.2 Plausibility of the proposed switching mechanism base on local patch quality

One may wonder if the human mind is capable of obtaining the representation like that in the deep Q-network. We will discuss in turn three components associated with the deep Q-network obtained under reinforcement learning.

The first component is the temporal differencing that learning Q-network is based on. In the literature of human reinforcement learning, there has been neural evidence of temporal differencing identified, suggesting that human are capable of using temporal differencing to learn from sparse future rewards (Houk & Adams, 1995; Montague et al., 1996; Niv, 2009; O'Doherty et al., 2003).

The second component is the neural network that represents Q-function. We use a neural network instead of other specified structures, as it allows for a flexible and non-linear representation of the policy to emerge from experiences and learning. Similar approaches are used in connectionist models of semantic learning and language acquisition (Elman, 1990; McClelland et al., 2010; Rohde & Plaut, 1999). It is possible that there exists other presentation of the Q-function that is simpler than the deep Q-network obtained currently. Our main goal here is to demonstrate the existence of such mapping from local patch quality to switch/non-switch decisions that outperforms the marginal value theorem. Better fits over human behavioral data under RL further support the idea that the obtained deep Q-network captures realistic features of the local patch quality that are important during the memory search.

The third component is the input data, including both the amount of training experiences and the information utilized in the model. First, we assume that human have been exposed to a large amount of experiences of similar memory search tasks. The Q-network that is eventually obtained through this exposure can be a simple heuristics used in memory search in general. Second, the proposed mechanism does not assume access to the entire semantic graph at once, but only the neighbors of the currently visited node. This is equivalent to a memory search right after recalling item"dog", one is searching the next item based on information associated with "dog". One should be able to access the edge weights connecting to the neighbor nodes: in fact, the edge weights of the semantic network were generated from large-scale human data as probabilities of recalling one item given the previous item (Nelson et al., 2004). One should also be able to access whether a node has been visited before, as people are good at censoring already generated responses (Hills et al., 2012).

4.4.3 Further implications

In the current work, we investigated the cognitive mechanisms that give rise to the switching behavior during human memory search. Particularly, the decision to continue engaging with a current option versus switching to a new one falls into a larger class of problems that has a rich theoretical and experimental history in ecology and human decision-making (Charnov, 1976; Cuthill, Kacelnik, Krebs, Haccou, & Iwasa, 1990; McNickle & Cahill, 2009). While the marginal value theorem holds over a range of animal and human foraging tasks, we show that in situations when optimality of the marginal value theorem are not guaranteed due to the conditions in the task environment, it is outperformed by a model based on local patch quality under reinforcement learning.

Related literature in free recall examines various heuristics to approximate the probability of search termination, including the total time spent (Davelaar, Yu, Harbison, Hussey, & Dougherty, 2013), proportion of items left to be recalled (Kragel, Morton, & Polyn, 2015), time since last recall (Rundus, 1973), and number of retrieval failures (Harbison, Dougherty, Davelaar, & Fayyad, 2009; J. G. Raaijmakers & Shiffrin, 1980). Evaluating search termination for episodic memories is closely related to evaluating patch switch for semantic memories. Approaches in deriving switching mechanism in the cur-

rent semantic fluency task can potentially be extended to evaluating search termination over episodic memories.

References

Abbott, J. T., Austerweil, J. L., & Griffiths, T. L. (2015). Random walks on semantic networks can resemble optimal foraging. 1.3, 4.1, 4.2, 4.2.1, 4.2.3, 4.2.4

Achim, A. M., & Lepage, M. (2005, Feb). Dorsolateral prefrontal cortex involvement in memory post-retrieval monitoring revealed in both item and associative recognition tests. *NeuroImage*, *24*(4), 1113–1121. Retrieved from http://dx.doi.org/10.1016/j.neuroimage.2004.10.036 doi: 10.1016/j.neuroimage.2004.10.036

Aggleton, J., Shaw, C., & Gaffan, E. (1992, Sep). The performance of postencephalitic amnesic subjects on two behavioural tests of memory: Concurrent discrimination learning and delayed matching-to-sample. *Cortex*, *28*(3), 359–372. Retrieved from http://dx.doi.org/10.1016/S0010-9452(13)80146-3 doi: 10.1016/s0010-9452(13)80146-3

Anderson, J. (1990). The adaptive character of thought. *Hillsdale, NJ: Earlbaum*. 1.3, 4.1, 4.4

Anderson, J. R. (1974, Oct). Retrieval of propositional information from long-term memory. *Cognitive Psychology*, *6*(4), 451–474. Retrieved from http://dx.doi.org/10.1016/0010-0285(74)90021-8 doi: 10.1016/0010-0285(74)90021-8

Anderson, J. R. (2007). How can the human mind occur in the physical universe?

Anderson, J. R. (2009). *How can the human mind occur in the physical universe?* Oxford University Press.

Anderson, J. R., Borst, J. P., Fincham, J. M., Ghuman, A. S., Tenison, C., & Zhang, Q. (2018, Jul). The common time course of memory processes revealed. *Psychological Science*, *29*(9), 1463–1474. Retrieved from http://dx.doi.org/10.1177/0956797618774526 doi: 10.1177/0956797618774526

Anderson, J. R., Bothell, D., Lebiere, C., & Matessa, M. (1998, May). An integrated theory of list memory. *Journal of Memory and Language*, *38*(4), 341–380. Retrieved from http://dx.doi.org/10.1006/jmla.1997.2553 doi: 10.1006/jmla.1997.2553

Anderson, J. R., & Fincham, J. M. (2013, Aug). Discovering the sequential structure of thought. *Cognitive Science*, *38*(2), 322–352. Retrieved from http://dx.doi.org/10.1111/cogs.12068 doi: 10.1111/cogs.12068

Anderson, J. R., & Fincham, J. M. (2014, Nov). Extending problem-solving procedures through reflection. *Cognitive Psychology*, *74*, 1–34. Retrieved from http://dx.doi.org/10.1016/j.cogpsych.2014.06.002 doi: 10.1016/j.cogpsych.2014.06.002
1.1

Anderson, J. R., & Milson, R. (1989). Human memory: An adaptive perspective. *Psychological Review*, *96*(4), 703.
1.3, 4.1, 4.4

Anderson, J. R., & Reder, L. M. (1999). The fan effect: New results and new theories. *Journal of Experimental Psychology: General*, *128*(2), 186.

Anderson, J. R., Zhang, Q., Borst, J. P., & Walsh, M. M. (2016a). The discovery of processing stages: Extension of sternberg's method. *Psychological Review*, *123*(5), 481–509. Retrieved from http://dx.doi.org/10.1037/rev0000030 doi: 10.1037/rev0000030
1, 1.1

Anderson, J. R., Zhang, Q., Borst, J. P., & Walsh, M. M. (2016b). The discovery of processing stages: Extension of sternberg's method. *Psychological review*, *123*(5), 481.
1.1

Anderson, K. L., Rajagovindan, R., Ghacibeh, G. A., Meador, K. J., & Ding, M. (2009, Oct). Theta oscillations mediate interaction between prefrontal cortex and medial temporal lobe in human memory. *Cerebral Cortex*, *20*(7), 1604–1612. Retrieved from http://dx.doi.org/10.1093/cercor/bhp223 doi: 10.1093/cercor/bhp223

Aron, A. R., Robbins, T. W., & Poldrack, R. A. (2004, Apr). Inhibition and the right inferior frontal cortex. *Trends in Cognitive Sciences*, *8*(4), 170–177. Retrieved from http://dx.doi.org/10.1016/j.tics.2004.02.010 doi: 10.1016/j.tics.2004.02.010

Axmacher, N., Mormann, F., Fernandez, G., Cohen, M. X., Elger, C. E., & Fell, J. (2007, Jul). Sustained neural activity patterns during working memory in the human medial temporal lobe. *Journal of Neuroscience*, *27*(29), 7807–7816. Retrieved from http://dx.doi.org/10.1523/JNEUROSCI.0962-07.2007 doi: 10.1523/jneurosci.0962-07.2007

Axmacher, N., Schmitz, D. P., Wagner, T., Elger, C. E., & Fell, J. (2008, Jul). Interactions between medial temporal lobe, prefrontal cortex, and inferior temporal regions during visual working memory: A combined intracranial eeg and functional magnetic resonance imaging study. *Journal of Neuroscience*, *28*(29), 7304–7312. Retrieved from http://dx.doi.org/10.1523/JNEUROSCI.1778-08.2008 doi: 10.1523/jneurosci.1778-08.2008

Bader, R., Mecklinger, A., Hoppstädter, M., & Meyer, P. (2010). Recognition memory for one-trial-unitized word pairs: Evidence from event-related potentials. *NeuroImage*, *50*(2), 772–781.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, *113*(4), 700–765. Retrieved from http://dx.doi.org/10.1037/0033-295X.113.4.700 doi: 10.1037/0033-295x.113.4.700

Borga, M. (1998). *Learning multidimensional signal processing* (Unpublished doctoral dissertation). Linköping University Electronic Press.
3.2.5

Borst, J. P., Ghuman, A. S., & Anderson, J. R. (2016, Nov). Tracking cognitive processing stages with meg: A spatio-temporal model of associative recognition in the brain. *NeuroImage*, *141*, 416–430. Retrieved from http://dx.doi.org/10.1016/j.neuroimage.2016.08.002 doi: 10.1016/j.neuroimage.2016.08.002

Borst, J. P., Schneider, D. W., Walsh, M. M., & Anderson, J. R. (2013, Dec). Stages of processing in associative recognition: Evidence from behavior, eeg, and classification. *Journal of Cognitive Neuroscience*, *25*(12), 2151–2166. Retrieved from http://dx.doi.org/10.1162/jocn_a_00457 doi: 10.1162/jocn_a_00457

Brincat, S. L., & Miller, E. K. (2015, Feb). Frequency-specific hippocampal-prefrontal interactions during associative learning. *Nature Neuroscience*, *18*(4), 576–581. Retrieved from http://dx.doi.org/10.1038/nn.3954 doi: 10.1038/nn.3954

Buzsáki, G. (2002, Jan). Theta oscillations in the hippocampus. *Neuron*, *33*(3), 325–340. Retrieved from http://dx.doi.org/10.1016/S0896-6273(02)00586-X doi: 10.1016/s0896-6273(02)00586-x

Campo, P., Garrido, M. I., Moran, R. J., Maestu, F., Garcia-Morales, I., Gil-Nagel, A., ... Friston, K. J. (2011, Aug). Remote effects of hippocampal sclerosis on effective connectivity during working memory encoding: A case of connectional diaschisis? *Cerebral Cortex*, *22*(6), 1225–1236. Retrieved from http://dx.doi.org/10.1093/cercor/bhr201 doi: 10.1093/cercor/bhr201

Charnov, E. L. (1976). Optimal foraging, the marginal value theorem. *Theoretical population biology*, *9*(2), 129–136.
1.3, 4.1, 4.2.2, 4.4.3

Clark, S. E., & Gronlund, S. D. (1996, Mar). Global matching models of recognition memory: How the models match the data. *Psychonomic Bulletin & Review*, *3*(1), 37–60. Retrieved from http://dx.doi.org/10.3758/BF03210740 doi: 10.3758/bf03210740

Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should i stay or should i go? how the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *362*(1481), 933–942.
4.2.6

Cohen, M. X. (2011, Nov). Hippocampal-prefrontal connectivity predicts midfrontal oscillations and long-term memory performance. *Current Biology*, *21*(22), 1900–1905. Retrieved from http://dx.doi.org/10.1016/j.cub.2011.09.036 doi: 10.1016/j.cub.2011.09.036

Colgin, L. L. (2011, Jun). Oscillations and hippocampal–prefrontal synchrony. *Current Opinion in Neurobiology*, *21*(3), 467–474. Retrieved from http://dx.doi.org/10.1016/j.conb.2011.04.006 doi: 10.1016/j.conb.2011.04.006

CONWAY, M., & FTHENAKI, A. (2003). Disruption of inhibitory control of memory following lesions to the frontal and temporal lobes. *Cortex*, *39*(4-5), 667–686.

Retrieved from http://dx.doi.org/10.1016/S0010-9452(08)70859-1 doi: 10.1016/s0010-9452(08)70859-1

Correa, N. M., Adali, T., Li, Y.-O., & Calhoun, V. D. (2010). Canonical correlation analysis for data fusion and group inferences. *IEEE signal processing magazine*, *27*(4), 39–50.

Curran, T. (2000, Nov). Brain potentials of recollection and familiarity. *Memory & Cognition*, *28*(6), 923–938. Retrieved from http://dx.doi.org/10.3758/BF03209340 doi: 10.3758/bf03209340

Curran, T., & Cleary, A. M. (2003, Jan). Using erps to dissociate recollection from familiarity in picture recognition. *Cognitive Brain Research*, *15*(2), 191–205. Retrieved from http://dx.doi.org/10.1016/S0926-6410(02)00192-1 doi: 10.1016/s0926-6410(02)00192-1

Curtis, C. E., & D'Esposito, M. (2003, Sep). Persistent activity in the prefrontal cortex during working memory. *Trends in Cognitive Sciences*, *7*(9), 415–423. Retrieved from http://dx.doi.org/10.1016/S1364-6613(03)00197-9 doi: 10.1016/s1364-6613(03)00197-9

Cuthill, I. C., Kacelnik, A., Krebs, J. R., Haccou, P., & Iwasa, Y. (1990, Oct). Starlings exploiting patches: the effect of recent experience on foraging decisions. *Animal Behaviour*, *40*(4), 625–640. Retrieved from http://dx.doi.org/10.1016/S0003-3472(05)80692-X doi: 10.1016/s0003-3472(05)80692-x
4.4.3

Danker, J. F., Gunn, P., & Anderson, J. R. (2008). A rational account of memory predicts left prefrontal activation during controlled retrieval. *Cerebral Cortex*, *18*(11), 2674–2685.

Davelaar, E. J., Yu, E. C., Harbison, J. I., Hussey, E. K., & Dougherty, M. R. (2013, Sep). A rational approach to memory search termination. *Cognitive Systems Research*, *24*, 96–103. Retrieved from http://dx.doi.org/10.1016/j.cogsys.2012.12.012 doi: 10.1016/j.cogsys.2012.12.012
4.4.3

de Hollander, G., Forstmann, B. U., & Brown, S. D. (2016, Mar). Different ways of linking behavioral and neural data via computational cognitive models. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, *1*(2), 101–109. Retrieved from http://dx.doi.org/10.1016/j.bpsc.2015.11.004 doi: 10.1016/j.bpsc.2015.11.004

Delorme, A., & Makeig, S. (2004, Mar). Eeglab: an open source toolbox for analysis of single-trial eeg dynamics including independent component analysis. *Journal of Neuroscience Methods*, *134*(1), 9–21. Retrieved from http://dx.doi.org/10.1016/j.jneumeth.2003.10.009 doi: 10.1016/j.jneumeth.2003.10.009

Depue, B. E. (2012, May). A neuroanatomical model of prefrontal inhibitory modulation of memory retrieval. *Neuroscience & Biobehavioral Reviews*, *36*(5), 1382–1399. Retrieved from http://dx.doi.org/10.1016/j.neubiorev.2012.02.012 doi: 10.1016/j.neubiorev.2012.02.012

D'Esposito, M., Postle, B. R., Jonides, J., & Smith, E. E. (1999, Jun). The neural substrate and temporal dynamics of interference effects in working memory as re-

vealed by event-related functional mri. *Proceedings of the National Academy of Sciences*, *96*(13), 7514–7519. Retrieved from http://dx.doi.org/10.1073/pnas.96.13.7514 doi: 10.1073/pnas.96.13.7514

Diana, R. A., Reder, L. M., Arndt, J., & Park, H. (2006, Feb). Models of recognition: A review of arguments in favor of a dual-process account. *Psychonomic Bulletin & Review*, *13*(1), 1–21. Retrieved from http://dx.doi.org/10.3758/BF03193807 doi: 10.3758/bf03193807

Diana, R. A., Van den Boom, W., Yonelinas, A. P., & Ranganath, C. (2011, Jan). Erp correlates of source memory: Unitized source information increases familiarity-based retrieval. *Brain Research*, *1367*, 278–286. Retrieved from http://dx.doi.org/10.1016/j.brainres.2010.10.030 doi: 10.1016/j.brainres.2010.10.030

Diana, R. A., Yonelinas, A. P., & Ranganath, C. (2007, Sep). Imaging recollection and familiarity in the medial temporal lobe: a three-component model. *Trends in Cognitive Sciences*, *11*(9), 379–386. Retrieved from http://dx.doi.org/10.1016/j.tics.2007.08.001 doi: 10.1016/j.tics.2007.08.001

Donaldson, D. I., & Rugg, M. D. (1998, May). Recognition memory for new associations: electrophysiological evidence for the role of recollection. *Neuropsychologia*, *36*(5), 377–395. Retrieved from http://dx.doi.org/10.1016/S0028-3932(97)00143-7 doi: 10.1016/s0028-3932(97)00143-7

Donders, F. (1969). On the speed of mental processes. *Acta Psychologica*, *30*, 412–431. Retrieved from http://dx.doi.org/10.1016/0001-6918(69)90065-1 doi: 10.1016/0001-6918(69)90065-1
1.1

Doyle, P. G., & Snell, J. L. (2000). Random walks and electric networks. *arXiv preprint math/0001057*.
4.2.6

Duzel, E., Yonelinas, A. P., Mangun, G. R., Heinze, H.-J., & Tulving, E. (1997, May). Event-related brain potential correlates of two states of conscious awareness in memory. *Proceedings of the National Academy of Sciences*, *94*(11), 5973–5978. Retrieved from http://dx.doi.org/10.1073/pnas.94.11.5973 doi: 10.1073/pnas.94.11.5973

Ecker, U. K., Zimmer, H. D., Groh-Bordin, C., & Mecklinger, A. (2007, May). Context effects on familiarity are familiarity effects of context — an electrophysiological study. *International Journal of Psychophysiology*, *64*(2), 146–156. Retrieved from http://dx.doi.org/10.1016/j.ijpsycho.2007.01.005 doi: 10.1016/j.ijpsycho.2007.01.005

Elman, J. L. (1990). Finding structure in time. *Cognitive science*, *14*(2), 179–211.
4.4.2

Farell, B. (1985). " same"–" different" judgments: A review of current controversies in perceptual comparisons. *Psychological Bulletin*, *98*(3), 419.

Farovik, A., Dupont, L. M., Arce, M., & Eichenbaum, H. (2008, Dec). Medial prefrontal cortex supports recollection, but not familiarity, in the rat. *Journal of Neuroscience*, *28*(50), 13428–13434. Retrieved from http://dx.doi.org/10.1523/

JNEUROSCI.3662-08.2008 doi: 10.1523/jneurosci.3662-08.2008

Fell, J., & Axmacher, N. (2011, Feb). The role of phase synchronization in memory processes. *Nature Reviews Neuroscience*, *12*(2), 105–118. Retrieved from http://dx.doi.org/10.1038/nrn2979 doi: 10.1038/nrn2979

Fletcher, P. (1998, Jul). The functional roles of prefrontal cortex in episodic memory. ii. retrieval. *Brain*, *121*(7), 1249–1256. Retrieved from http://dx.doi.org/10.1093/brain/121.7.1249 doi: 10.1093/brain/121.7.1249

Fries, P. (2009, Jun). Neuronal gamma-band synchronization as a fundamental process in cortical computation. *Annual Review of Neuroscience*, *32*(1), 209–224. Retrieved from http://dx.doi.org/10.1146/annurev.neuro.051508.135603 doi: 10.1146/annurev.neuro.051508.135603

Gainotti, G. (2007, Nov). Face familiarity feelings, the right temporal lobe and the possible underlying neural mechanisms. *Brain Research Reviews*, *56*(1), 214–235. Retrieved from http://dx.doi.org/10.1016/j.brainresrev.2007.07.009 doi: 10.1016/j.brainresrev.2007.07.009

Gonsalves, B. D., Kahn, I., Curran, T., Norman, K. A., & Wagner, A. D. (2005, Sep). Memory strength and repetition suppression: Multimodal imaging of medial temporal cortical contributions to recognition. *Neuron*, *47*(5), 751–761. Retrieved from http://dx.doi.org/10.1016/j.neuron.2005.07.013 doi: 10.1016/j.neuron.2005.07.013

Griffiths, T. L., Lieder, F., & Goodman, N. D. (2015). Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in cognitive science*, *7*(2), 217–229.
4.1

Hannula, D. E., Tranel, D., & Cohen, N. J. (2006, Aug). The long and the short of it: Relational memory impairments in amnesia, even at short lags. *Journal of Neuroscience*, *26*(32), 8352–8359. Retrieved from http://dx.doi.org/10.1523/JNEUROSCI.5222-05.2006 doi: 10.1523/jneurosci.5222-05.2006
1.2

Harbison, J. I., Dougherty, M. R., Davelaar, E. J., & Fayyad, B. (2009, Jun). On the lawfulness of the decision to terminate memory search. *Cognition*, *111*(3), 397–402. Retrieved from http://dx.doi.org/10.1016/j.cognition.2009.03.002 doi: 10.1016/j.cognition.2009.03.002
4.4.3

Henson, R. (2005, Oct). A mini-review of fmri studies of human medial temporal lobe activity associated with recognition memory. *The Quarterly Journal of Experimental Psychology: Section B*, *58*(3-4), 340–360. Retrieved from http://dx.doi.org/10.1080/02724990444000113 doi: 10.1080/02724990444000113

Henson, R. N. A., Shallice, T., & Dolan, R. J. (1999, Jul). Right prefrontal cortex and episodic memory retrieval: a functional mri test of the monitoring hypothesis. *Brain*, *122*(7), 1367–1381. Retrieved from http://dx.doi.org/10.1093/brain/122.7.1367 doi: 10.1093/brain/122.7.1367

Hills, T. T., Jones, M. N., & Todd, P. M. (2012). Optimal foraging in semantic memory.

Psychological review, *119*(2), 431.
1.3, 4.1, 4.2.2, 4.2.4, 4.2.5, 4.2.6, 4.4.2

Hills, T. T., Kalff, C., & Wiener, J. M. (2013). Adaptive lévy processes and area-restricted search in human foraging. *PLoS One*, *8*(4), e60488.
4.4.1

Hintzman, D., & Curran, T. (1994, Feb). Retrieval dynamics of recognition and frequency judgments: Evidence for separate processes of familiarity and recall. *Journal of Memory and Language*, *33*(1), 1–18. Retrieved from http://dx.doi.org/10.1006/jmla.1994.1001 doi: 10.1006/jmla.1994.1001

Holdstock, J., Gutnikov, S., Gaffan, D., & Mayes, A. (2000, Jan). Perceptual and mnemonic matching-to-sample in humans: Contributions of the hippocampus, perirhinal and other medial temporal lobe cortices. *Cortex*, *36*(3), 301–322. Retrieved from http://dx.doi.org/10.1016/S0010-9452(08)70843-8 doi: 10.1016/s0010-9452(08)70843-8
1.2

Holdstock, J., Shaw, C., & Aggleton, J. (1995, Dec). The performance of amnesic subjects on tests of delayed matching-to-sample and delayed matching-to-position. *Neuropsychologia*, *33*(12), 1583–1596. Retrieved from http://dx.doi.org/10.1016/0028-3932(95)00145-X doi: 10.1016/0028-3932(95)00145-x
1.2

Houk, J. C., & Adams, J. L. (1995). 13 a model of how the basal ganglia generate and use neural signals that. *Models of information processing in the basal ganglia*, 249.
4.1, 4.4.2

Hyman. (2010). Working memory performance correlates with prefrontal-hippocampal theta interactions but not with prefrontal neuron firing rates. *Frontiers in Integrative Neuroscience*. Retrieved from http://dx.doi.org/10.3389/neuro.07.002.2010 doi: 10.3389/neuro.07.002.2010

Hyman, J. M., Zilli, E. A., Paley, A. M., & Hasselmo, M. E. (2005). Medial prefrontal cortex cells show dynamic modulation with the hippocampal theta rhythm dependent on behavior. *Hippocampus*, *15*(6), 739–749. Retrieved from http://dx.doi.org/10.1002/hipo.20106 doi: 10.1002/hipo.20106

Jacobs, J., & Kahana, M. J. (2010, Apr). Direct brain recordings fuel advances in cognitive electrophysiology. *Trends in Cognitive Sciences*, *14*(4), 162–171. Retrieved from http://dx.doi.org/10.1016/j.tics.2010.01.005 doi: 10.1016/j.tics.2010.01.005

Jacoby, L. L. (1991, Oct). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language*, *30*(5), 513–541. Retrieved from http://dx.doi.org/10.1016/0749-596X(91)90025-F doi: 10.1016/0749-596x(91)90025-f

Jeneson, A., & Squire, L. R. (2011, Dec). Working memory, long-term memory, and medial temporal lobe function. *Learning & Memory*, *19*(1), 15–25. Retrieved from http://dx.doi.org/10.1101/lm.024018.111 doi: 10.1101/lm.024018.111

Jones, M. N., & Mewhort, D. J. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological review*, *114*(1), 1.

Jonides, J., Smith, E. E., Marshuetz, C., Koeppe, R. A., & Reuter-Lorenz, P. A. (1998). Inhibition in verbal working memory revealed by brain activation. *Proceedings of the National Academy of Sciences*, *95*(14), 8410-8413. Retrieved from http://www.pnas.org/content/95/14/8410.abstract

Jun, K.-S., Zhu, J., Rogers, T. T., Yang, Z., & Yuan, M. (2015). Human memory search as initial-visit emitting random walk. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, & R. Garnett (Eds.), *Advances in neural information processing systems 28* (pp. 1072–1080). Curran Associates, Inc. Retrieved from http://papers.nips.cc/paper/5878-human-memory-search-as-initial-visit-emitting-random-walk.pdf 4.2.6

Kahana, M. J., & Bennett, P. J. (1994). Classification and perceived similarity of compound gratings that differ in relative spatial phase. *Attention, Perception, & Psychophysics*, *55*(6), 642–656.

Kahana, M. J., & Sekuler, R. (2002, Aug). Recognizing spatial patterns: a noisy exemplar approach. *Vision Research*, *42*(18), 2177–2192. Retrieved from http://dx.doi.org/10.1016/S0042-6989(02)00118-9 doi: 10.1016/s0042-6989(02)00118-9

Kalff, C., Hills, T., & Wiener, J. M. (2010). Human foraging behavior: A virtual reality investigation on area restricted search in humans. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 32). 4.4.1

Kaplan, R., Bush, D., Bonnefond, M., Bandettini, P. A., Barnes, G. R., Doeller, C. F., & Burgess, N. (2014, Feb). Medial prefrontal theta phase coupling during spatial memory retrieval. *Hippocampus*, *24*(6), 656–665. Retrieved from http://dx.doi.org/10.1002/hipo.22255 doi: 10.1002/hipo.22255

Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the american statistical association*, *90*(430), 773–795. 4.3.2

Kettenring, J. R. (1971). Canonical analysis of several sets of variables. *Biometrika*, *58*(3), 433–451. 3.2.5

King, D. R., & Anderson, J. R. (1976, Dec). Long-term memory search: An intersecting activation process. *Journal of Verbal Learning and Verbal Behavior*, *15*(6), 587–605. Retrieved from http://dx.doi.org/10.1016/0022-5371(76)90053-0 doi: 10.1016/0022-5371(76)90053-0

Koechlin, E. (2003, Nov). The architecture of cognitive control in the human prefrontal cortex. *Science*, *302*(5648), 1181–1185. Retrieved from http://dx.doi.org/10.1126/science.1088545 doi: 10.1126/science.1088545

Komorowski, R. W., Manns, J. R., & Eichenbaum, H. (2009, Aug). Robust conjunctive item-place coding by hippocampal neurons parallels learning what happens where. *Journal of Neuroscience*, *29*(31), 9918–9929. Retrieved from http://dx.doi.org/10.1523/JNEUROSCI.1378-09.2009 doi: 10.1523/jneurosci.1378-09.2009

Kragel, J. E., Morton, N. W., & Polyn, S. M. (2015, Feb). Neural activity in the me-

dial temporal lobe reveals the fidelity of mental time travel. *Journal of Neuroscience*, *35*(7), 2914–2926. Retrieved from http://dx.doi.org/10.1523/jneurosci.3378-14.2015 doi: 10.1523/jneurosci.3378-14.2015
4.4.3

Lachaux, J.-P., Rodriguez, E., Martinerie, J., & Varela, F. J. (1999). Measuring phase synchrony in brain signals. *Human Brain Mapping*, *8*(4), 194–208. Retrieved from http://dx.doi.org/10.1002/(SICI)1097-0193(1999)8:4<194::AID-HBM4>3.0.CO;2-C doi: 10.1002/(sici)1097-0193(1999)8:4⟨194::aid-hbm4⟩3.0.co;2-c

Lega, B. C., Jacobs, J., & Kahana, M. (2011, Apr). Human hippocampal theta oscillations and the formation of episodic memories. *Hippocampus*, *22*(4), 748–761. Retrieved from http://dx.doi.org/10.1002/hipo.20937 doi: 10.1002/hipo.20937

Le Van Quyen, M., Foucher, J., Lachaux, J.-P., Rodriguez, E., Lutz, A., Martinerie, J., & Varela, F. J. (2001, Sep). Comparison of hilbert transform and wavelet methods for the analysis of neuronal synchrony. *Journal of Neuroscience Methods*, *111*(2), 83–98. Retrieved from http://dx.doi.org/10.1016/S0165-0270(01)00372-7 doi: 10.1016/s0165-0270(01)00372-7

Lewis-Peacock, J. A., Drysdale, A. T., Oberauer, K., & Postle, B. R. (2012, Jan). Neural evidence for a distinction between short-term memory and the focus of attention. *Journal of Cognitive Neuroscience*, *24*(1), 61–79. Retrieved from http://dx.doi.org/10.1162/jocn_a_00140 doi: 10.1162/jocn_a_00140
1.2

Li, Y.-O., Eichele, T., Calhoun, V. D., & Adali, T. (2011, Jan). Group study of simulated driving fmri data by multiset canonical correlation analysis. *Journal of Signal Processing Systems*, *68*(1), 31–48. Retrieved from http://dx.doi.org/10.1007/s11265-010-0572-8 doi: 10.1007/s11265-010-0572-8
3.2.5

Luck, S. J. (2014). *An introduction to the event-related potential technique*.
1

Luck, S. J., Woodman, G. F., & Vogel, E. K. (2000a). Event-related potential studies of attention. *Trends in cognitive sciences*, *4*(11), 432–440.

Luck, S. J., Woodman, G. F., & Vogel, E. K. (2000b, Nov). Event-related potential studies of attention. *Trends in Cognitive Sciences*, *4*(11), 432–440. Retrieved from http://dx.doi.org/10.1016/S1364-6613(00)01545-X doi: 10.1016/s1364-6613(00)01545-x
1

MacKenzie, G., & Donaldson, D. I. (2007, Jun). Dissociating recollection from familiarity: Electrophysiological evidence that familiarity for faces is associated with a posterior old/new effect. *NeuroImage*, *36*(2), 454–463. Retrieved from http://dx.doi.org/10.1016/j.neuroimage.2006.12.005 doi: 10.1016/j.neuroimage.2006.12.005

Makeig, S., Westerfield, M., Jung, T.-P., Enghoff, S., Townsend, J., Courchesne, E., & Sejnowski, T. J. (2002). Dynamic brain sources of visual evoked responses. *Science*, *295*(5555), 690–694.

Malmberg, K. J. (2008, Dec). Recognition memory: A review of the critical findings and an

integrated theory for relating them. *Cognitive Psychology*, *57*(4), 335–384. Retrieved from http://dx.doi.org/10.1016/j.cogpsych.2008.02.004 doi: 10.1016/j.cogpsych.2008.02.004

McClelland, J. L. (1979). On the time relations of mental processes: an examination of systems of processes in cascade. *Psychological review*, *86*(4), 287.
1.1

McClelland, J. L., Botvinick, M. M., Noelle, D. C., Plaut, D. C., Rogers, T. T., Seidenberg, M. S., & Smith, L. B. (2010). Letting structure emerge: connectionist and dynamical systems approaches to cognition. *Trends in cognitive sciences*, *14*(8), 348–356.
4.4.2

McNickle, G. G., & Cahill, J. F. (2009, Mar). Plant root growth and the marginal value theorem. *Proceedings of the National Academy of Sciences*, *106*(12), 4747–4751. Retrieved from http://dx.doi.org/10.1073/pnas.0807971106 doi: 10.1073/pnas.0807971106
4.4.3

Mehlhorn, K., Newell, B. R., Todd, P. M., Lee, M. D., Morgan, K., Braithwaite, V. A., ... Gonzalez, C. (2015). Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision*, *2*(3), 191.
4.2.6

Mitchell, D. J., McNaughton, N., Flanagan, D., & Kirk, I. J. (2008, Nov). Frontal-midline theta from the perspective of hippocampal "theta". *Progress in Neurobiology*, *86*(3), 156–185. Retrieved from http://dx.doi.org/10.1016/j.pneurobio.2008.09.005 doi: 10.1016/j.pneurobio.2008.09.005

Mitchell, K. J., Johnson, M. K., Raye, C. L., & Greene, E. J. (2004, Jul). Prefrontal cortex activity associated with source monitoring in a working memory task. *Journal of Cognitive Neuroscience*, *16*(6), 921–934. Retrieved from http://dx.doi.org/10.1162/0898929041502724 doi: 10.1162/0898929041502724

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... others (2015). Human-level control through deep reinforcement learning. *Nature*, *518*(7540), 529.
4.2.4

Mollison, M. V., & Curran, T. (2012, Sep). Familiarity in source memory. *Neuropsychologia*, *50*(11), 2546–2565. Retrieved from http://dx.doi.org/10.1016/j.neuropsychologia.2012.06.027 doi: 10.1016/j.neuropsychologia.2012.06.027

Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive hebbian learning. *Journal of neuroscience*, *16*(5), 1936–1947.
4.1, 4.4.2

Nairne, J. S. (2002, Feb). Remembering over the short-term: The case against the standard model. *Annual Review of Psychology*, *53*(1), 53–81. Retrieved from http://dx.doi.org/10.1146/annurev.psych.53.100901.135131 doi: 10.1146/annurev.psych.53.100901.135131

Navawongse, R., & Eichenbaum, H. (2013, Jan). Distinct pathways for rule-based retrieval and spatial mapping of memory representations in hippocampal neurons. *Journal of Neuroscience*, *33*(3), 1002–1013. Retrieved from http://dx.doi.org/

10.1523/JNEUROSCI.3891-12.2013 doi: 10.1523/jneurosci.3891-12.2013

Nee, D. E., & Jonides, J. (2008, Aug). Neural correlates of access to short-term memory. *Proceedings of the National Academy of Sciences*, *105*(37), 14228–14233. Retrieved from http://dx.doi.org/10.1073/pnas.0802081105 doi: 10.1073/pnas.0802081105

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, & Computers*, *36*(3), 402–407.
4.2.1, 4.4.2

Nickerson, R. S. (1969). 'same'-'different'response times: A model and a preliminary test. *Acta Psychologica*, *30*, 257–275.

Niv, Y. (2009). Reinforcement learning in the brain. *Journal of Mathematical Psychology*, *53*(3), 139–154.
4.1, 4.4.2

Nosofsky, R. M. (1991). Tests of an exemplar model for relating perceptual classification and recognition memory. *Journal of Experimental Psychology: Human Perception and Performance*, *17*(1), 3–27. Retrieved from http://dx.doi.org/10.1037/0096-1523.17.1.3 doi: 10.1037/0096-1523.17.1.3

Nosofsky, R. M., Little, D. R., Donkin, C., & Fific, M. (2011). Short-term memory scanning viewed as exemplar-based categorization. *Psychological Review*, *118*(2), 280–315. Retrieved from http://dx.doi.org/10.1037/a0022494 doi: 10.1037/a0022494

Nyhus, E., & Curran, T. (2010, Jun). Functional role of gamma and theta oscillations in episodic memory. *Neuroscience & Biobehavioral Reviews*, *34*(7), 1023–1035. Retrieved from http://dx.doi.org/10.1016/j.neubiorev.2009.12.014 doi: 10.1016/j.neubiorev.2009.12.014

O'Doherty, J. P., Dayan, P., Friston, K., Critchley, H., & Dolan, R. J. (2003). Temporal difference models and reward-related learning in the human brain. *Neuron*, *38*(2), 329–337.
4.1, 4.4.2

Owen, A. M., Sahakian, B. J., Semple, J., Polkey, C. E., & Robbins, T. W. (1995, Jan). Visuo-spatial short-term recognition memory and learning after temporal lobe excisions, frontal lobe excisions or amygdalo-hippocampectomy in man. *Neuropsychologia*, *33*(1), 1–24. Retrieved from http://dx.doi.org/10.1016/0028-3932(94)00098-A doi: 10.1016/0028-3932(94)00098-a
1.2

Petten, C. V., Kutas, M., Kluender, R., Mitchiner, M., & McIsaac, H. (1991, Apr). Fractionating the word repetition effect with event-related potentials. *Journal of Cognitive Neuroscience*, *3*(2), 131–150. Retrieved from http://dx.doi.org/10.1162/jocn.1991.3.2.131 doi: 10.1162/jocn.1991.3.2.131

Picton, T., Bentin, S., Berg, P., Donchin, E., Hillyard, S., Johnson, R., ... et al. (2000, Mar). Guidelines for using human event-related potentials to study cognition: Recording standards and publication criteria. *Psychophysiology*, *37*(2), 127–152. Retrieved from http://dx.doi.org/10.1111/1469-8986.3720127 doi: 10.1111/

1469-8986.3720127

Pirolli, P. L., & Anderson, J. R. (1985). The role of practice in fact retrieval. *Journal of experimental psychology: Learning, memory, and cognition*, *11*(1), 136.

Poch, C., & Campo, P. (2012). Neocortical-hippocampal dynamics of working memory in healthy and diseased brain states based on functional connectivity. *Frontiers in Human Neuroscience*, *6*. Retrieved from http://dx.doi.org/10.3389/fnhum.2012.00036 doi: 10.3389/fnhum.2012.00036

Preston, A. R., & Eichenbaum, H. (2013, Sep). Interplay of hippocampus and prefrontal cortex in memory. *Current Biology*, *23*(17), R764–R773. Retrieved from http://dx.doi.org/10.1016/j.cub.2013.05.041 doi: 10.1016/j.cub.2013.05.041

Raaijmakers, J. G., & Shiffrin, R. M. (1980). Sam: A theory of probabilistic search of associative memory. *Psychology of Learning and Motivation*, 207–262. Retrieved from http://dx.doi.org/10.1016/s0079-7421(08)60162-0 doi: 10.1016/s0079-7421(08)60162-0
4.4.3

Raaijmakers, J. G. W., & Shiffrin, R. M. (1992, Jan). Models for recall and recognition. *Annual Review of Psychology*, *43*(1), 205–234. Retrieved from http://dx.doi.org/10.1146/annurev.ps.43.020192.001225 doi: 10.1146/annurev.ps.43.020192.001225

Raghavachari, S., Kahana, M. J., Rizzuto, D. S., Caplan, J. B., Kirschen, M. P., Bourgeois, B., ... Lisman, J. E. (2001). Gating of human theta oscillations by a working memory task. *Journal of Neuroscience*, *21*(9), 3175–3183. Retrieved from http://www.jneurosci.org/content/21/9/3175

Ranganath, C. (2006, Apr). Working memory for visual objects: Complementary roles of inferior temporal, medial temporal, and prefrontal cortex. *Neuroscience*, *139*(1), 277–289. Retrieved from http://dx.doi.org/10.1016/j.neuroscience.2005.06.092 doi: 10.1016/j.neuroscience.2005.06.092
1, 1.2

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*(2), 59–108. Retrieved from http://dx.doi.org/10.1037/0033-295X.85.2.59 doi: 10.1037/0033-295x.85.2.59

Rhodes, S. M., & Donaldson, D. I. (2007). Electrophysiological evidence for the influence of unitization on the processes engaged during episodic retrieval: Enhancing familiarity based remembering. *Neuropsychologia*, *45*(2), 412–424. Retrieved from http://dx.doi.org/10.1016/j.neuropsychologia.2006.06.022 doi: 10.1016/j.neuropsychologia.2006.06.022

Rhodes, S. M., & Donaldson, D. I. (2008, Jan). Electrophysiological evidence for the effect of interactive imagery on episodic memory: Encouraging familiarity for non-unitized stimuli during associative recognition. *NeuroImage*, *39*(2), 873–884. Retrieved from http://dx.doi.org/10.1016/j.neuroimage.2007.08.041 doi: 10.1016/j.neuroimage.2007.08.041

Rissman, J., Gazzaley, A., & D'Esposito, M. (2007, Nov). Dynamic adjustments in pre-

frontal, hippocampal, and inferior temporal interactions with increasing visual working memory load. *Cerebral Cortex*, *18*(7), 1618–1629. Retrieved from http://dx.doi.org/10.1093/cercor/bhm195 doi: 10.1093/cercor/bhm195

Roach, B. J., & Mathalon, D. H. (2008, Jul). Event-related eeg time-frequency analysis: An overview of measures and an analysis of early gamma band phase locking in schizophrenia. *Schizophrenia Bulletin*, *34*(5), 907–926. Retrieved from http://dx.doi.org/10.1093/schbul/sbn093 doi: 10.1093/schbul/sbn093

Rohde, D. L., & Plaut, D. C. (1999). Language acquisition in the absence of explicit negative evidence: How important is starting small? *Cognition*, *72*(1), 67–109.
4.4.2

Rotello, C. M., & Heit, E. (2000). Associative recognition: A case of recall-to-reject processing. *Memory & Cognition*, *28*(6), 907–922.

Rotello, C. M., Macmillan, N. A., & Van Tassel, G. (2000). Recall-to-reject in recognition: Evidence from roc curves. *Journal of Memory and Language*, *43*(1), 67–88.

Rowe, J. B. (2000, Jun). The prefrontal cortex: Response selection or maintenance within working memory? *Science*, *288*(5471), 1656–1660. Retrieved from http://dx.doi.org/10.1126/science.288.5471.1656 doi: 10.1126/science.288.5471.1656

Rugg, M. D., & Curran, T. (2007, Jun). Event-related potentials and recognition memory. *Trends in Cognitive Sciences*, *11*(6), 251–257. Retrieved from http://dx.doi.org/10.1016/j.tics.2007.04.004 doi: 10.1016/j.tics.2007.04.004

Rugg, M. D., Henson, R. N., & Robb, W. G. (2003, Jan). Neural correlates of retrieval processing in the prefrontal cortex during recognition and exclusion tasks. *Neuropsychologia*, *41*(1), 40–52. Retrieved from http://dx.doi.org/10.1016/S0028-3932(02)00129-X doi: 10.1016/s0028-3932(02)00129-x

Rugg, M. D., & Yonelinas, A. P. (2003, Jul). Human recognition memory: a cognitive neuroscience perspective. *Trends in Cognitive Sciences*, *7*(7), 313–319. Retrieved from http://dx.doi.org/10.1016/S1364-6613(03)00131-1 doi: 10.1016/s1364-6613(03)00131-1

Rundus, D. (1973, Feb). Negative effects of using list items as recall cues. *Journal of Verbal Learning and Verbal Behavior*, *12*(1), 43–50. Retrieved from http://dx.doi.org/10.1016/s0022-5371(73)80059-3 doi: 10.1016/s0022-5371(73)80059-3
4.4.3

Schneider, D. W., & Anderson, J. R. (2012, May). Modeling fan effects on the time course of associative recognition. *Cognitive Psychology*, *64*(3), 127–160. Retrieved from http://dx.doi.org/10.1016/j.cogpsych.2011.11.001 doi: 10.1016/j.cogpsych.2011.11.001

Siapas, A. G., Lubenov, E. V., & Wilson, M. A. (2005, Apr). Prefrontal phase locking to hippocampal theta oscillations. *Neuron*, *46*(1), 141–151. Retrieved from http://dx.doi.org/10.1016/j.neuron.2005.02.028 doi: 10.1016/j.neuron.2005.02.028

Simon, H. A. (1978). Rationality as process and as product of thought. *The American economic review*, *68*(2), 1–16.
1.3, 4.1

Simons, J. S., & Spiers, H. J. (2003, Aug). Prefrontal and medial temporal lobe interactions in long-term memory. *Nature Reviews Neuroscience*, *4*(8), 637–648. Retrieved from http://dx.doi.org/10.1038/nrn1178 doi: 10.1038/nrn1178

Sohn, M.-H., Goode, A., Stenger, V. A., Jung, K.-J., Carter, C. S., & Anderson, J. R. (2005, Mar). An information-processing model of three cortical regions: evidence in episodic memory retrieval. *NeuroImage*, *25*(1), 21–33. Retrieved from http://dx.doi.org/10.1016/j.neuroimage.2004.11.001 doi: 10.1016/j.neuroimage.2004.11.001

Souza, A. S., Rerko, L., & Oberauer, K. (2014, Dec). Refreshing memory traces: thinking of an item improves retrieval from visual working memory. *Annals of the New York Academy of Sciences*, *1339*(1), 20–31. Retrieved from http://dx.doi.org/10.1111/nyas.12603 doi: 10.1111/nyas.12603

Speer, N. K., & Curran, T. (2007, Oct). Erp correlates of familiarity and recollection processes in visual associative recognition. *Brain Research*, *1174*, 97–109. Retrieved from http://dx.doi.org/10.1016/j.brainres.2007.08.024 doi: 10.1016/j.brainres.2007.08.024

Sternberg, S. (1969). The discovery of processing stages: Extensions of donders' method. *Acta Psychologica*, *30*, 276–315. Retrieved from http://dx.doi.org/10.1016/0001-6918(69)90055-9 doi: 10.1016/0001-6918(69)90055-9
1.1

Sternberg, S. (2016, Aug). In defence of high-speed memory scanning. *The Quarterly Journal of Experimental Psychology*, *69*(10), 2020–2075. Retrieved from http://dx.doi.org/10.1080/17470218.2016.1198820 doi: 10.1080/17470218.2016.1198820

Stróżak, P., Abedzadeh, D., & Curran, T. (2016, Mar). Separating the fn400 and n400 potentials across recognition memory experiments. *Brain Research*, *1635*, 41–60. Retrieved from http://dx.doi.org/10.1016/j.brainres.2016.01.015 doi: 10.1016/j.brainres.2016.01.015

Sutton, R. S. (1988). Learning to predict by the methods of temporal differences. *Machine learning*, *3*(1), 9–44.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction* (Vol. 1) (No. 1). MIT press Cambridge.
4.2.4

Swanson, L. W., Wyss, J. M., & Cowan, W. M. (1978, Oct). An autoradiographic study of the organization of intrahippocampal association pathways in the rat. *The Journal of Comparative Neurology*, *181*(4), 681–715. Retrieved from http://dx.doi.org/10.1002/cne.901810402 doi: 10.1002/cne.901810402

Tass, P., Rosenblum, M. G., Weule, J., Kurths, J., Pikovsky, A., Volkmann, J., ... Freund, H.-J. (1998, Oct). Detection ofn:mphase locking from noisy data: Applica-

tion to magnetoencephalography. *Physical Review Letters*, *81*(15), 3291–3294. Retrieved from http://dx.doi.org/10.1103/PhysRevLett.81.3291 doi: 10.1103/physrevlett.81.3291

Thompson-Schill, S. L., Bedny, M., & Goldberg, R. F. (2005, Apr). The frontal lobes and the regulation of mental activity. *Current Opinion in Neurobiology*, *15*(2), 219–224. Retrieved from http://dx.doi.org/10.1016/j.conb.2005.03.006 doi: 10.1016/j.conb.2005.03.006

THOMPSON-SCHILL, S. L., JONIDES, J., MARSHUETZ, C., SMITH, E. E., D'ESPOSITO, M., KAN, I. P., ... SWICK, D. (2002, Jun). Effects of frontal lobe damage on interference effects in working memory. *Cognitive, Affective, & Behavioral Neuroscience*, *2*(2), 109–120. Retrieved from http://dx.doi.org/10.3758/CABN.2.2.109 doi: 10.3758/cabn.2.2.109

Todd, P. M., & Gigerenzer, G. (2007). Environments that make us smart: Ecological rationality. *Current directions in psychological science*, *16*(3), 167–171.
1.3, 4.1

Troyer, A. K., Moscovitch, M., & Winocur, G. (1997). Clustering and switching as two components of verbal fluency: evidence from younger and older healthy adults. *neuropsychology*, *11*(1), 138.
4.3.2

van Vugt, M. K., Schulze-Bonhage, A., Litt, B., Brandt, A., & Kahana, M. J. (2010, Feb). Hippocampal gamma oscillations increase with memory load. *Journal of Neuroscience*, *30*(7), 2694–2699. Retrieved from http://dx.doi.org/10.1523/JNEUROSCI.0567-09.2010 doi: 10.1523/jneurosci.0567-09.2010
1, 1.2

van Vugt, M. K., Schulze-Bonhage, A., Sekuler, R., Litt, B., Brandt, A., Baltuch, G., & Kahana, M. J. (2009, Nov). Intracranial electroencephalography reveals two distinct similarity effects during item recognition. *Brain Research*, *1299*, 33–44. Retrieved from http://dx.doi.org/10.1016/j.brainres.2009.07.016 doi: 10.1016/j.brainres.2009.07.016

van Vugt, M. K., Sekuler, R., Wilson, H. R., & Kahana, M. J. (2013). An electrophysiological signature of summed similarity in visual working memory. *Journal of Experimental Psychology: General*, *142*(2), 412–425. Retrieved from http://dx.doi.org/10.1037/a0029759 doi: 10.1037/a0029759
1.2

Varela, F., Lachaux, J.-P., Rodriguez, E., & Martinerie, J. (2001, Apr). The brainweb: Phase synchronization and large-scale integration. *Nature Reviews Neuroscience*, *2*(4), 229–239. Retrieved from http://dx.doi.org/10.1038/35067550 doi: 10.1038/35067550

Vía, J., Santamaría, I., & Pérez, J. (2005). Canonical correlation analysis (cca) algorithms for multiple data sets: Application to blind simo equalization. In *2005 13th european signal processing conference* (pp. 1–4).
3.2.5

Wagenmakers, E.-J. (2007). A practical solution to the pervasive problems ofp values. *Psychonomic bulletin & review*, *14*(5), 779–804.
4.3.2

Walsh, M. M., Paynter, C. A., Zhang, Y., & Reder, L. M. (2016, Jul). Hitting the reset button:

An erp investigation of memory for temporal context. *Brain Research*, *1642*, 524–531. Retrieved from http://dx.doi.org/10.1016/j.brainres.2016.04.047 doi: 10.1016/j.brainres.2016.04.047

Watkins, C. J., & Dayan, P. (1992). Q-learning. *Machine learning*, *8*(3-4), 279–292. 4.2.4

Wilson, H. R., Loffler, G., & Wilkinson, F. (2002, Dec). Synthetic faces, face cubes, and the geometry of face space. *Vision Research*, *42*(27), 2909–2923. Retrieved from http://dx.doi.org/10.1016/S0042-6989(02)00362-0 doi: 10.1016/s0042-6989(02)00362-0

Yeung, N., Bogacz, R., Holroyd, C. B., & Cohen, J. D. (2004). Detection of synchronized oscillations in the electroencephalogram: an evaluation of methods. *Psychophysiology*, *41*(6), 822–832.

Yeung, N., Botvinick, M. M., & Cohen, J. D. (2004, Oct). The neural basis of error detection: Conflict monitoring and the error-related negativity. *Psychological Review*, *111*(4), 931–959. Retrieved from http://dx.doi.org/10.1037/0033-295X.111.4.931 doi: 10.1037/0033-295x.111.4.931

Yonelinas, A. P. (2002, Apr). The nature of recollection and familiarity: A review of 30 years of research. *Journal of Memory and Language*, *46*(3), 441–517. Retrieved from http://dx.doi.org/10.1006/jmla.2002.2864 doi: 10.1006/jmla.2002.2864

Yu, S.-Z. (2010). Hidden semi-markov models. *Artificial intelligence*, *174*(2), 215–243.

Zhang, Q., & Anderson, J. R. (2018, May). A rational account of human memory search. Retrieved from http://dx.doi.org/10.1101/326397 doi: 10.1101/326397 1

Zhang, Q., Borst, J. P., Kass, R. E., & Anderson, J. R. (2017, Jun). Inter-subject alignment of meg datasets in a common representational space. *Human Brain Mapping*. Retrieved from http://dx.doi.org/10.1002/hbm.23689 doi: 10.1002/hbm.23689

Zhang, Q., van Vugt, M., Borst, J. P., & Anderson, J. R. (2018, Jul). Mapping working memory retrieval in space and in time: A combined electroencephalography and electrocorticography approach. *NeuroImage*, *174*, 472–484. Retrieved from http://dx.doi.org/10.1016/j.neuroimage.2018.03.039 doi: 10.1016/j.neuroimage.2018.03.039 1

Zhang, Q., Walsh, M. M., & Anderson, J. R. (2017, Feb). The effects of probe similarity on retrieval and comparison processes in associative recognition. *Journal of Cognitive Neuroscience*, *29*(2), 352–367. Retrieved from http://dx.doi.org/10.1162/jocn_a_01059 doi: 10.1162/jocn_a_01059 1, 1.2

Zhang, Q., Walsh, M. M., & Anderson, J. R. (2018, Mar). The impact of inserting an additional mental process. *Computational Brain Behavior*, *1*(1), 22–35. Retrieved from http://dx.doi.org/10.1007/s42113-018-0002-8 doi: 10.1007/s42113-018-0002-8 1

Appendix I.

Person (18)		*Verb (8)*	*Location (18)*	
actor	coach	paint	airport	attic
cowboy	dancer	work	bank	barn
chef	engineer	walk	factory	garage
farmer	musician	talk	hotel	kitchen
maid	judge	sing	castle	library
pilot	queen	laugh	museum	office
lawyer	sheriff	sleep	church	prison
soldier	teacher	drink	stadium	studio
tourist	doctor	-	temple	theatre