HIPPOCAMPUS 24:1672–1686 (2014)

Neurocomputational Account of Memory and Perception: Thresholded and Graded Signals in the Hippocampus

Kane W. Elfman, Mariam Aly, and Andrew P. Yonelinas*

ABSTRACT: Recent evidence suggests that the hippocampus, a region critical for long-term memory, also supports certain forms of high-level visual perception. A seemingly paradoxical finding is that, unlike the thresholded hippocampal signals associated with memory, the hippocampus produces graded, strength-based signals in perception. This article tests a neurocomputational model of the hippocampus, based on the complementary learning systems framework, to determine if the same model can account for both memory and perception, and whether it produces the appropriate thresholded and strength-based signals in these two types of tasks. The simulations showed that the hippocampus, and most prominently the CA1 subfield, produced graded signals when required to discriminate between highly similar stimuli in a perception task, but generated thresholded patterns of activity in recognition memory. A threshold was observed in recognition memory because pattern completion occurred for only some trials and completely failed to occur for others; conversely, in perception, pattern completion always occurred because of the high degree of item similarity. These results offer a neurocomputational account of the distinct hippocampal signals associated with perception and memory, and are broadly consistent with proposals that CA1 functions as a comparator of expected versus perceived events. We conclude that the hippocampal computations required for high-level perceptual discrimination are congruous with current neurocomputational models that account for recognition memory, and fit neatly into a broader description of the role of the hippocampus for the processing of complex relational information. © 2014 Wiley Periodicals, Inc.

KEY WORDS: computational model; perception; memory; recognition; recollection

INTRODUCTION

Neuropsychological studies have indicated that hippocampal damage leads to profound long-term memory impairments (Scoville and Milner, 1957; Baddeley and Warrington, 1970; Graf and Schacter, 1984; Zola-Morgan et al., 1986; Cohen and Eichenbaum, 1993; Vargha-Khadem et al., 1997; Reed and Squire, 1998; Yonelinas et al., 2002). Neuroimaging and electrophysiological studies have corroborated the role of the hippocampus in long-term memory by establishing a clear link between hippocampal activity and memory encoding and retrieval (Ranganath et al., 2004; Yonelinas et al., 2005; Montaldi et al., 2006; Sederberg et al., 2006; Fell et al., 2011; Axmacher et al., 2010).

Additionally, neuroimaging studies have indicated that memory-related hippocampal activity is characterized by relatively distinct states (Ranganath et al., 2004; Yonelinas et al., 2005; Montaldi et al., 2006; for a review, see Eichenbaum et al., 2007). For example, in recognition tests, previously studied items that are recognized on the basis of recollection are associated with high levels of hippocampal activity, whereas items that are not recollected are associated with very low levels of activity, even if they are successfully recognized on the basis of familiarity (Yonelinas et al., 2005; Montaldi et al., 2006). Similarly thresholded patterns of hippocampal activity have been found in relational memory tasks (e.g., source and associative memory; Eichenbaum et al., 2007), which are largely dependent on recollection (Yonelinas, 2002; Yonelinas et al., 2010).

Cognitive memory models describe recollection as a threshold retrieval process, in that qualitative information about a study event is either retrieved or not (Yonelinas, 2002). For example, in tests of relational memory, thresholded memory retrieval has been evidenced by linear receiver operating characteristics [ROCs[1]]. In addition, in a recent fMRI study, ROCs constructed from hippocampal activity values during a source memory task were found to be linear and therefore consistent with the threshold model (Slotnick and Thakral, 2013). The thresholded response of the hippocampus, and of recollection, can be contrasted with the perirhinal cortex, which responds in a more graded manner and is associated with familiarity, or strength-based, memory (Ranganath et al., 2004; Montaldi et al., 2006; Staresina et al., 2013), which in turn produces curved ROCs (Parks and Yonelinas, 2007).

Challenges to the Traditional View of Hippocampal Function

The neuropsychological and neuroimaging work on the hippocampus has, until recently, focused on its role

Department of Psychology, University of California, Davis
*Correspondence to: Andrew Yonelinas, Department of Psychology, University of California, Davis CA 95616, USA.
E-mail: apyonelinas@ucdavis.edu

Accepted for publication 5 August 2014.
DOI 10.1002/hipo.22345

Published online 12 August 2014 in Wiley Online Library (wileyonlinelibrary.com).

© 2014 WILEY PERIODICALS, INC.

---

[1] Linear ROCs reflect a thresholded signal, whereas curvilinear ROCs indicate a continuously graded signal; for review, see (Parks and Yonelinas, 2007; but for exceptions see Elfman et al., 2008; Parks et al., 2011).

in long-term memory. However, empirical findings in the past decade have challenged the traditional view that the hippocampus is critical for long-term memory but not other cognitive functions (Scoville and Milner, 1957; Baddeley and Warrington, 1970; Graf and Schacter, 1984; Zola-Morgan et al., 1986; Cohen and Eichenbaum, 1993; Reed and Squire, 1998). For example, it is now clear that hippocampal damage can impair working memory and high-level scene perception[2] (e.g., Bussey and Saksida, 2005; Olson et al., 2006; Graham et al., 2010; Lee et al., 2005, 2012; Warren et al., 2012; Aly et al., 2013; Yonelinas, 2013). Moreover, neuroimaging studies have linked perceptual processing with hippocampal activity in healthy adults (e.g, Barense et al., 2010; Lee and Rudebeck, 2010; Lee et al., 2008, 2010; Mundy et al., 2012; Aly et al., 2013).

As a result of such findings, it has been argued that the hippocampus serves a broader role than previously thought, contributing to perception and working memory as well as long-term memory (Cohen and Eichenbaum, 1993; Graham et al., 2010; Saksida and Bussey, 2010; Lee et al., 2012; Olsen et al., 2012; Nadel and Peterson, 2013; Shohamy and Turk-Browne, 2013; Yonelinas, 2013). The mechanisms that would account for the joint roles of the hippocampus across these cognitive domains are still not well established, although it has been suggested that the hippocampus is involved in relational (Cohen and Eichenbaum, 1993; Olsen et al., 2012) or complex conjunctive (Graham et al., 2010; Lee et al., 2012; Saksida and Bussey, 2010) processing in the service of perception, working memory, and long-term memory (also see Nadel and Peterson, 2013; Shohamy and Turk-Browne, 2013; Yonelinas, 2013).

Insights on Hippocampal Function from Neurocomputational Models

Neurocomputational models have furthered our understanding of how memories are encoded and retrieved in the hippocampus by exploring how distinct patterns of neural activity and task performance arise from the network's architecture (e.g., McNaughton and Morris, 1987; Treves and Rolls, 1994; McClelland et al., 1995; Rolls, 1996; Hasselmo and Wyble, 1997; Rudy and O'Reilly, 1999; Hasselmo et al., 2002; Norman and O'Reilly, 2003; Hasselmo and Eichenbaum, 2005). However, it is currently unclear whether these models can also account for the recent findings linking the hippocampus to perception. In addition, an important challenge for any hippocampal model that attempts to account for both the memory and perception findings is that the nature of the respective signals can be quite different: whereas in memory the contribution of the hippocampus appears to be thresholded, or state-based, there is evidence that in perception its contribution is graded, or strength-based (Aly et al., 2013).

In this article, we first provide an overview of how computational models of the hippocampus account for thresholded memory retrieval. Then, we outline the recent evidence implicating the hippocampus in high-level perception and highlight how the type of signal associated with perception differs from memory, and the challenge this poses for existing models. Finally, we report simulation results from a computational model of the hippocampus to determine whether it can simultaneously account for memory and perception, and specifically whether it can produce the thresholded memory signals and graded perception signals that have been observed in the empirical literature.

Thresholded Hippocampal Signals in Long-Term Memory

Computational models of long-term memory generally agree that the hippocampus is capable of forming distinct representations of events, even when the events are quite similar (McNaughton and Morris, 1987; O'Reilly and McClelland, 1994; McClelland et al., 1995; Rolls, 1996). This capability, sometimes referred to as pattern separation (or orthogonalization), is thought to result from high levels of lateral inhibition, effectively limiting the network to very sparse excitatory activity, especially within the dentate gyrus (DG). This constraint creates a competitive learning environment, so that if two stimuli differ even subtly, the hippocampal neurons that activate to encode those stimuli (via Hebbian learning) are likely to differ dramatically.

A common assumption is that the hippocampus is fundamental to the storage and reinstatement of episodic memories. Due to the singular nature of such events (e.g., remembering where you parked your car today), synaptic weights must be updated rapidly to capture unique associations (e.g., "Toyota," "Shields Avenue," and "this morning") as they occur (Tulving and Markowitsch, 1998; Aggleton and Brown, 1999; Eichenbaum and Cohen, 2001; Leutgeb et al., 2006, 2007). Pattern separation is therefore important as it allows new events to be quickly learned without damaging or conflating representations of other, similar events (e.g., remembering where you parked your car today versus yesterday).

At the time of retrieval, if the network is presented with part of the original stimulus event, activation will sometimes spread through the network and reinstate aspects of the original event that are not currently present; this is referred to as pattern completion. Some computational models and electrophysiological studies suggest that this process is facilitated by recurrent connections within CA3, which acts as an autoassociation network (Marr, 1971; O'Reilly and McClelland, 1994; Rudy and O'Reilly, 1999; but for evidence supporting an alternative hypothesis, see Colgin et al., 2010). However, due to the aforementioned effects of pattern separation, if a retrieval cue differs too greatly from the original stimulus, the network will be unable to reactivate the learned pattern. Thus, the hippocampus is expected to behave in a thresholded manner, pattern completing to produce a high level of activity, or else failing to pattern complete and thus producing much lower activity.

The above computational account of the hippocampus is consistent with theories that propose that the hippocampus is

---

[2] By "high-level perception", we refer to perceptual discriminations or judgments about complex scenes, faces, or objects (i.e., at a level of stimulus complexity that is higher than that of simple features such as brightness, orientation, color, etc.).

critical for recollection (Schacter et al., 1996; Aggleton and Brown, 1999; Eldridge et al., 2000; Yonelinas, 2002; Eichenbaum et al., 2007; Yonelinas et al., 2010). Recollection is proposed to be thresholded in the sense that individuals retrieve qualitative information about some events, but for other events recollection fails entirely (Yonelinas, 1994, 2001; Yonelinas et al., 2010). Thus, the thresholded output from computational models of the hippocampus is consistent with the thresholded nature of recollection, and converges with neuroimaging findings that relate thresholded hippocampal signals to encoding and retrieval of recollection-based memories (Ranganath et al., 2004; Yonelinas et al., 2005; Montaldi et al., 2006; see Eichenbaum et al., 2007 for review), as well as patient findings that show that damage to the hippocampus selectively impairs recollection (Yonelinas et al., 2002; see Yonelinas et al., 2010 for review).

It is worthwhile emphasizing that the term “threshold” does not imply that memory retrieval, or pattern completion, is “all-or-none.” That is, successful retrieval may involve reinstating any number of features from a past episode; consequently, hippocampal activity levels related to memory retrieval may vary across different recollected episodes. The term “threshold” is used to indicate that only some studied items are recollected (e.g., items exceeding a recollective threshold, or leading to pattern completion), whereas others fail to be recollected and thus remain indistinguishable from nonstudied items. This definition of “threshold” is consistent with both the hippocampal model instantiated by Norman and O’Reilly (2003), which was adopted for the current simulations, and behavioral models of memory such as the dual-process signal detection model (e.g., Yonelinas, 2001). Such thresholded memory signals give rise to ROCs that approximate a linear function when all but the most stringent response criteria are used (Elfman et al., 2008). Importantly, a thresholded signal does not necessitate that all items that exceed the threshold be associated with the same memory strength.

Strength-Based Hippocampal Signals in Perception

As mentioned above, a growing body of literature indicates that the hippocampus is involved in tasks assessing high-level visual perception (e.g., Lee et al., 2005; Saksida and Bussey, 2005; Barense et al., 2010; Lee and Rudebeck, 2010; Mundy et al., 2012; Warren et al., 2012; see Graham et al., 2010 and Lee et al., 2012 for review). In addition, a recent behavioral study with hippocampal lesion patients and a neuroimaging study with healthy adults have suggested that the hippocampal signal in perception tasks is strength-based, or continuously graded (Aly et al., 2013), rather than reflecting a thresholded signal as in recollection-based memory.

In the behavioral task in Aly et al. (2013), the stimuli were pairs of scenes that were either identical or differed, in that the scenes were slightly contracted or expanded relative to one another. The specific manipulation was a “pinching” or “spherizing,” which kept the sizes of the images the same but contracted (pinched) or expanded (spherized) the scenes, with the largest changes at the center and gradually decreasing changes toward the periphery (also see Aly and Yonelinas, 2012). These changes altered the configural, or relational, information within the scenes (i.e., the relative distance between component parts) without adding or removing any objects.

Patients with selective hippocampal or more extensive MTL damage were briefly presented with these scene pairs, simultaneously, and rated their confidence that the two scenes were the same or different. The patients were significantly impaired relative to healthy controls on this task. Importantly, however, this impairment was specific to one kind of perceptual judgment. Patients were not impaired in their ability to make high-confidence judgments associated with identifying specific details that had changed (i.e., state-based perception; see Aly and Yonelinas, 2012). Rather, they were selectively impaired in graded, lower-confidence judgments associated with a sense of match/mismatch (i.e., strength-based perception, Aly and Yonelinas, 2012). Importantly, this pattern of results was observed even when the analysis was restricted to patients with selective hippocampal damage.

The role of the hippocampus in graded, strength-based perception was corroborated in a follow-up neuroimaging study with healthy adults. Aly et al. found that BOLD activity in the hippocampus linearly tracked graded, lower-confidence responses, but was not differentially sensitive to high-confidence responses associated with a state of identifying specific, detailed differences.

Thus, in contrast to studies of long-term memory, where the hippocampus supports high-confidence responses and shows thresholded activity, in perception tasks, the current evidence suggests that the hippocampus is critical for strength-based responses at lower confidence levels, and shows graded levels of activity. How can these conflicting findings be reconciled?

To our knowledge, no previous computational modeling work has attempted to apply a hippocampal memory model to perception. However, if the existing memory models accurately capture the functions of the hippocampus, then they may also produce the type of strength-based signals that have been observed in perception. In this study, we examined the hippocampal component of Norman and O’Reilly’s (2003) complementary learning systems (CLS) model to test whether (1) this model could be used to make accurate perception judgments, and (2) the same model would produce a graded signal in perception while simultaneously producing a thresholded signal in memory.

METHOD

The Hippocampal Model

The network design is based on the hippocampal component of the CLS model, as implemented by Norman and

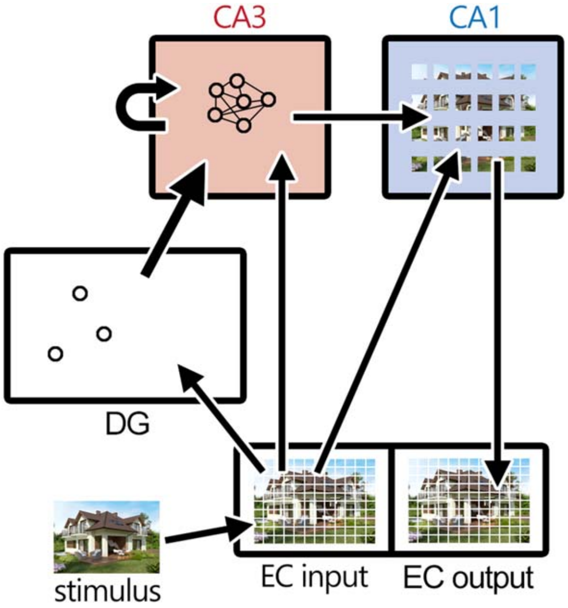

FIGURE 1. A schematic illustration of the hippocampal network presented with a new stimulus event. An input pattern (i.e., the stimulus) is represented in the EC input layer as a matrix of features (e.g., color, shape, etc.). Activity spreads to the CA subfields, directly and indirectly via the sparsely active DG. Patterns are learned via Hebbian weight changes between coactive units, including recurrent collaterals within CA3, which help bind the stimulus features together. During a cued recall task, the network is presented with a subset of features from a studied item, and CA3 is able to "complete" the pattern that corresponds to that item. This in turn activates the associated CA1 pattern, which converts the retrieved representation back to its original form at the EC output layer. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

O'Reilly (2003). We used the software package Emergent (Aisa et al., 2008) version 6 that incorporates the Leabra neural network algorithm (O'Reilly and Munakata, 2000). The algorithm includes a continuous rate-code output function, conditional principal component analysis Hebbian learning, and a competitive inhibition function that can be tuned to produce distributed representations with varying levels of sparseness.

The model comprises entorhinal cortex (EC) input and output layers, the DG, CA3, and CA1 layers (see Fig. 1), and has distinct encoding and retrieval modes of operation. Stimuli are presented to the network via the EC input layer, which has projections to DG and CA3 (the perforant pathway), and to CA1. DG, which has very sparse activation, has strong projections to CA3 (the Mossy fibers). During encoding (i.e., in *encoding mode*), recurrent collaterals in CA3 help bind together the different stimulus features. At the same time, CA1 learns to associate direct input from the EC with input from diffuse CA3 projections (the Schaffer collaterals). During retrieval (i.e., in *retrieval mode*), the pathway from EC input to CA1, and the mossy fiber pathway (DG-CA3), are inactivated so that any pattern retrieved by CA3 becomes the dominant output signal (Rolls, 1996; Hasselmo et al., 2002).

A subtle departure from past simulation work with the CLS model is that, in the current work, the pathway from EC input to CA1 was diminished in strength during retrieval mode, rather than inactivated completely. This allowed CA1 to act as a "comparator" of the current stimulus from the EC input and the retrieved pattern from CA3 (see Results and Discussion for more). For consistency, this modification was used in both the memory task and the perception task. Importantly, this modification did not affect our ability to replicate previous results with memory stimulations using the CLS model (see Results for more details). Specific model parameters are included in Appendix A.

Perception Task

For the perception task, a single trial comprised two stimulus patterns (items) that were presented to the network sequentially. A single item was composed of 36 feature dimensions, each made up of four units (groups of neurons) with one active unit per dimension. Consequently, any two randomly generated stimulus patterns overlapped by an average of 25%.

The task consisted of 20 trials, with 10 matching and 10 mismatching pairs presented in random order. The pairs were created from 20 randomly generated items. Matching pairs were created by simply duplicating half of the items, and mismatching pairs were created by duplicating the remaining items and then pseudo-randomly changing one or more (depending on the condition) features of the duplicate, ensuring that the same units were never reselected. Simplified illustrations (i.e., fewer features than were actually used) of matching and mismatching item pairs are presented in Figure 2A. The task was run 50 times at each of 36 possible levels of feature mismatch.

The current approach to modeling perceptual differences was to change a small number of features in the mismatch condition. In this way, we were able to vary perceptual similarity in a simple and tractable way, and no changes to the original model's architecture were required. An alternative approach would have been to make smaller changes to a large number of features, which might appear to be more analogous to the "pinching" or "spherizing" manipulation used in previous empirical studies (Aly and Yonelinas, 2012; Aly et al., 2013, 2014). However, because it is not clear how features in the model are mapped to visual features in real-world stimuli, the current feature change approach can be used as an approximation to the empirical manipulation.

Two simulation approaches were examined; an *encode/retrieve* model and an *always-encode* model. For the encode/retrieve model, an item was presented to the network in encoding mode, and then, for the second (matching or mismatching) item, the network was switched to retrieval mode. In taking this approach, we used the network in a way that is comparable to typical long-term memory simulations (Norman and O'Reilly, 2003; Elfman et al., 2008). The always-encode model differed in that both items were presented to the network in encoding mode. We included the latter approach to ensure that

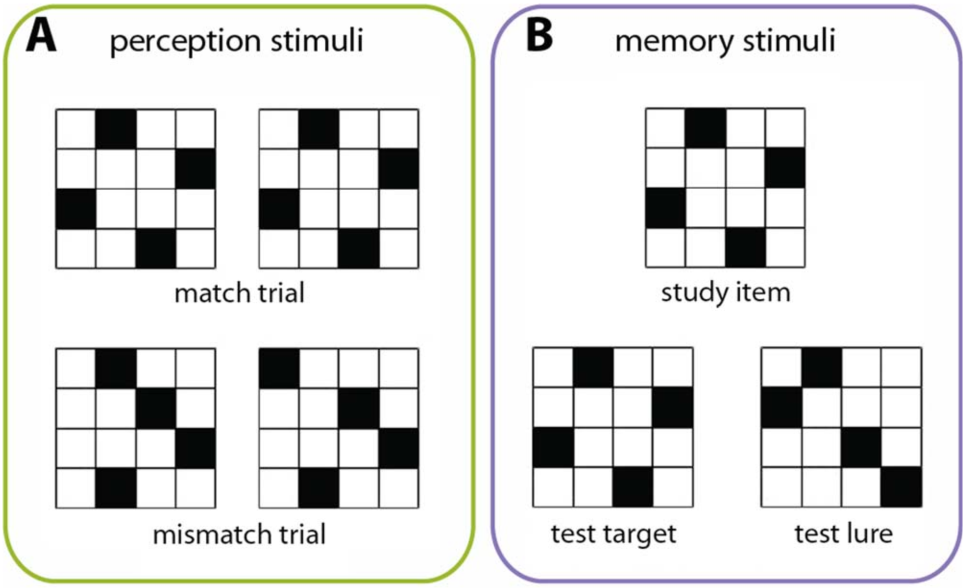

FIGURE 2. Simplified representations of the perception and recognition memory stimuli. Each pattern represents a stimulus and each row is a feature dimension; a black square represents a maximally active unit (or group of neurons). A: Perception (match/mismatch) task: Examples of matching (top) and mismatching (bottom) item pairs. The mismatching pair has one feature difference (top feature). The degree of mismatch was manipulated by changing the number of feature differences. B: Recognition memory task (complete retrieval cue): One study item (top) and two test items (bottom). Half the test items were learned patterns (targets) and half were new (lures). There was also a partial retrieval cue memory test (not shown here), which differed in that 50% of features were missing at test. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

the results would replicate if one assumes that perception does not involve any kind of (even very short-term) "retrieval."

To measure performance, the overall level of activity was recorded from CA1 and CA3 when the second item in each pair was presented. We also measured activity at the EC output layer, but since this layer receives input from CA1 alone and is in this sense redundant with CA1, we did not analyze the EC data. An additional reason for this approach was that CA1, as it is instantiated in the model, supports comparatively more fine-grained activity levels than EC.

Memory Tasks

Recognition memory was simulated in two ways: using complete retrieval cues and partial retrieval cues. For both simulations, 20 randomly generated items—the same items used in the perception task—comprised the study list. The items were sequentially presented to the network in encoding mode. The network was then switched to retrieval mode for the test phase. The test list comprised the 20 study items and 20 randomly generated lures (see Fig. 2B) presented in random order. For the complete retrieval cue simulations, the entire studied item (or lure) was presented at test. In this respect, the memory task was kept as similar as possible to the perception task (in which complete stimulus patterns, i.e., all feature dimensions, were always presented). For the partial retrieval cue simulations, half the feature dimensions in each test item were left blank. The partial retrieval cue simulation approach is typical of the memory simulations used to probe hippocampal pattern completion (e.g., Norman and O'Reilly, 2003). In addition, it may more accurately reflect episodic item recognition tests, in which participants are required to indicate if the test item was presented in an earlier study context (i.e., the earlier study list, which is represented by the missing features). In either case, as indicated below, the two simulations led to similar conclusions.

RESULTS

Figure 3A shows the average activity of CA1 units for matching item pairs and mismatching pairs in the perception task. The activity distributions were overlapping and Gaussian in shape, and activity strength was on average predictive of match or mismatch condition. Additional simulations showed that, as the number of mismatching features was systematically varied in equal intervals from 0 (complete match) to 36 (complete mismatch), CA1 activity distributions gradually shifted from high to low activity (Fig. B1, Appendix B). That is, on average, the match trials produced higher activity levels than mismatch trials. The initial simulations were conducted using separate encoding and retrieval modes, but the same pattern of results was found using the always-encode model (see below for additional analyses). Together, the results show that the

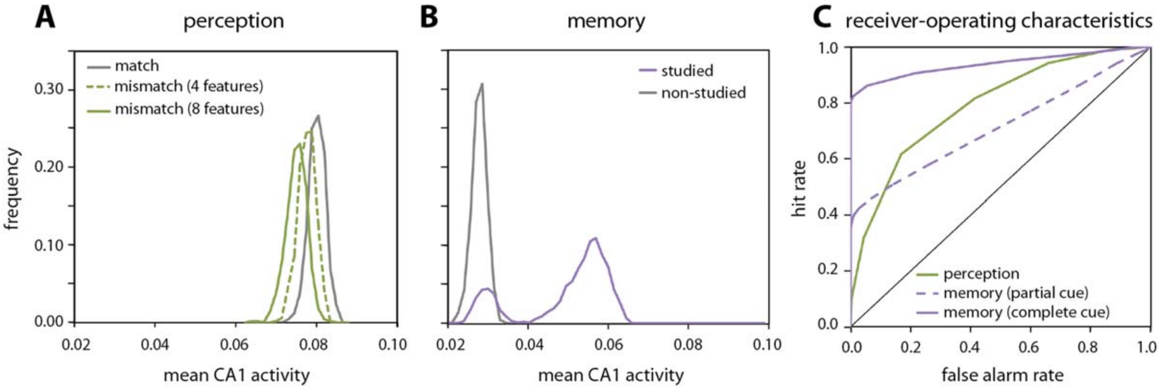

FIGURE 3. Comparison of hippocampal network performance during the perception and recognition memory tasks. A: Distributions of average activity levels for CA1 in the perception task (shown for matching items and items with 4 and 8 mismatching features out of a possible 36). B: Distributions of CA1 activity for studied (target) versus nonstudied (lure) items in the complete retrieval cue recognition memory task. In perception (A), the match and mismatch items produced overlapping strength distributions. The match items produced slightly higher activity, on average, than mismatch items, but critically, both types of items led to pattern completion. In contrast, in recognition memory (B), some studied items led to pattern completion (the hump on the right) whereas others failed to lead to pattern completion (the smaller hump on the left). Nonstudied items never led to pattern completion. Thus, the result was a thresholded distribution. C. ROCs, plotted using the average CA1 activity in the perception task and the memory tasks. The ROC in the perception task is curvilinear, whereas the memory ROCs are relatively linear. Note that in the memory tasks, overall performance is lower when a partial retrieval cue is used (vs. a complete cue), but in both cases, the resulting ROCs are relatively linear. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

hippocampus produces continuous strength distributions when used to make perceptual discriminations.

In contrast, the recognition memory simulations produced distributions that were consistent with a threshold process. Studied versus nonstudied item recognition distributions for the complete retrieval cue simulations are presented in Figure 3B, and the distributions for the partial retrieval cue simulations are shown in Figure C1 (Appendix C). In both cases, the studied items had a bimodal distribution, producing either strong activity (indicating pattern completion) or weak activity (indicating retrieval failure). In contrast, nonstudied items produced only weak activity (retrieval failure).

To verify that pattern completion corresponded to accurate item retrieval, we compared each EC output pattern with its respective target pattern using a match-mismatch rule (see Norman and O'Reilly, 2003; correct retrieval was assigned to scores greater than 0.67). Using a binary logistic regression, we found that CA1 activity predicted accurate item recognition (i.e., target versus lure) for all trials of the complete retrieval cue test; for the partial retrieval cue test, 99.8% of target trials and 100% of lures were accurately predicted.

ROC curves were then plotted to compare the performance characteristics of the perception and memory tasks, using activity from CA1 (Fig. 3C). The perception ROC plots the proportion of correctly identified matching items ("hit rate") against incorrectly identified mismatching items ("false alarm rate"; using four mismatching features) as the criterion for attributing activity to a "matching" pair was relaxed from the maximum observed activity to zero. Similarly, the recognition ROCs plot correctly identified target items ("hit rate") against incorrectly identified lure items ("false alarm rate") as a function of the criteria based on activity.

The perception ROC function was curved and had a *y*-intercept of 0; this shape is consistent with a continuous, signal detection process (Swets, 1988). By comparison, the memory ROC functions were relatively linear and had a non-zero *y*-intercept, consistent with a threshold process (Swets, 1988; Slotnick and Thakral, 2013). This latter pattern held for both the partial cue and full cue memory simulations.

To examine the roles of different hippocampal subfields in the perception task, we examined measures of perceptual sensitivity in CA1 and CA3. Sensitivity (d') of the model to matching versus mismatching pairs was calculated as the difference in the mean activity for match (μ_{act_match}) and mismatch ($\mu_{act\_mismatch}$) trials divided by the average standard deviations (σ), as follows:

$$d' = \frac{\mu_{act\_match} - \mu_{act\_mismatch}}{\sqrt{(\sigma_{act\_match}{}^2 - \sigma_{act\_mismatch}{}^2)/2}} \tag{1}$$

Using this metric, we compared CA1 and CA3 as the number of mismatching features was increased from 1 to 36 (Fig. 4). The DG layer was not included in this analysis because it was inactive during retrieval mode, and the EC output was excluded because it is redundant as a strength signal in the model (i.e., its activity is determined only by CA1).

For the encode/retrieve model (Fig. 4A), CA1 sensitivity increased gradually as the number of feature differences was increased, and eventually tapered off. CA3 showed a similar trend, but was much less sensitive overall. Using the always-

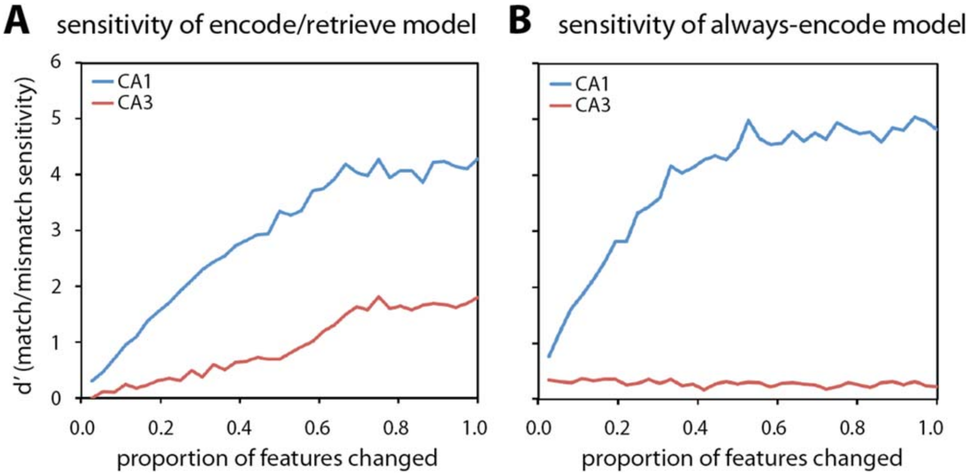

FIGURE 4. Sensitivity (d') of CA1 and CA3 to matching versus mismatching item pairs as the number of mismatching features was varied from 1 to 36 (i.e., all 36 features mismatch). Sensitivity is based on overall levels of activity. A: Sensitivity using an encode/retrieve model. The network alternated between an encoding mode for the first item and a retrieval mode for the second item. Activity was measured from the response to the second item. B: Sensitivity using an always-encode model. Activity was again measured from the response to the second item. In both models, CA1 was more sensitive to perceptual mismatch than CA3. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

encode model, CA1 demonstrated a similarly graded increase in sensitivity (Fig. 4B), while CA3 again showed weak sensitivity compared to CA1, and did not improve as the number of feature differences was increased.

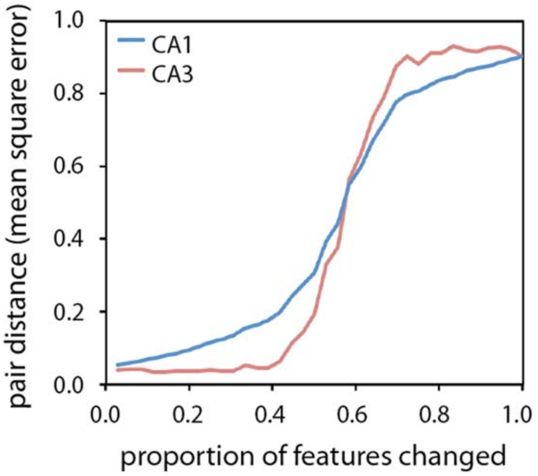

FIGURE 5. The average pattern distance (mean square error) between items in mismatching pairs in the perception task, represented in CA1 and CA3 as a function of the number of feature differences. The steep rise near the center of the figure indicates an increasing proportion of items that failed to exceed the threshold for pattern completion, leading to the formation of a new pattern. This threshold is particularly pronounced for CA3, in the sense that the function remains flat and does not respond to stimulus changes until the threshold is reached (about 0.40 in the current simulations). In contrast, CA1 responds more continuously and is able to differentiate between small stimulus changes. Thus, CA3 is relatively more thresholded compared with CA1, which shows a more graded pattern and is able to represent subtler stimulus changes. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

The results of the simulations suggest that, whether one adopts the encode/retrieve or the always-encode model, CA1 is particularly sensitive to perceptual changes, and its sensitivity gradually increases as stimulus pairs become more different.

We then explored the network dynamics that produced continuous strength-based perception by examining how items were represented relative to one another as the number of mismatching features increased. To compare perception to previous findings in the memory literature, we examined the encode/retrieve model and plotted the average pattern distance—measured as the mean square error (MSE)—between mismatching representations for CA1 and CA3 (Fig. 5).

The distance plots reveal two interesting transitions that speak to both the continuous perception signal and the thresholded memory signal. First, CA1 shows a steady rise in pattern distance at the low end of mismatch (proportions of 0 to 0.4 features changed). This indicates that the two items initially overlapped completely in CA1, but gradually diverged as the number of feature differences increased. Thus, pattern similarity in CA1 tracks the objective similarity of the two items. Importantly, the shift in overlap corresponded with increasing strength-based sensitivity over the same period (see Fig. 4A, proportions of 0 to 0.4 features changed). In contrast, CA3 pattern distances remained minimal over the same period, indicating that CA3 was more likely to simply reactivate the pattern associated with the first item. Thus, pattern similarity in CA3 does not track objective similarity very well, and is less sensitive than CA1 to very small changes between stimuli (see Fig. 4A, proportions of 0 to 0.4 features changed).

Figure 5 also indicates that CA3 had two distinct retrieval states, delimited by a steep transition near the center of the similarity scale. On the left side of the transition, distances are small, indicating pattern completion. That is, the second item was drawn into the attractor created by the first item. In contrast, on the right side, much larger pattern distances indicate that CA3 was no longer reactivating the first item pattern (c.f. Yassa and Stark, 2011).

Because CA1 receives projections directly from CA3, it showed a similar transition as CA3 shifted from one state to the other. Critically, however, CA1 showed gradually increasing pattern distance with small stimulus differences, and it did so before CA3 started to discriminate similar items (i.e., < 0.4 proportion of features changed). This indicates that CA1 can produce a continuously graded signal that is diagnostic at low levels of mismatch, while CA3 is considerably less effective at discriminating between similar stimuli until a threshold is reached.

Why does CA1 show a more graded response than CA3, particularly given that CA3 is one of the major inputs to CA1? There are several reasons for this. First, because CA3 is subject to greater lateral inhibition than CA1, CA3 produces proportionately sparser activation. Greater inhibition reduces the probability of pattern overlap except at very high levels of input similarity, the result being a rapid transition between high and low overlapping states. In addition, recurrent excitatory projections within CA3 act to bind the elements of a stimulus representation together as a cohesive unit, together forming an attractor. When an input is sufficiently close to the attractor, activation spreads along the strengthened recurrent connections to reproduce the original pattern. In contrast, the lack of recurrent connectivity within CA1, combined with more diffuse activation patterns, supports more graded levels of pattern overlap and activation strength.

An interesting finding of the current simulations is that there was greater activity for matching than mismatching trials in the perception task. Why might such a pattern arise? When presented with the second item in a mismatch trial, CA3 pattern completes to the previous stimulus, and this signal is projected to CA1. Concurrently, CA1 receives information about the current, slightly mismatching stimulus from the EC input layer. Excitation of CA1 is therefore spread out over more units compared with a trial consisting of matching pairs. Although this might lead to the prediction of greater overall activity in CA1 for mismatching trials, competitive inhibition within CA1 counteracts such a result by dampening the activity of units that are less strongly excited (i.e., that are stimulated by EC alone or CA3 alone). In contrast, on matching trials, the same CA1 units are jointly and strongly reactivated by both EC and CA3, and inhibitory dynamics do not dampen the overall level of activity. (A simplified model that demonstrates the effect is included with the hippocampal model project file referenced in Appendix A.)

DISCUSSION

The involvement of the hippocampus in memory is well established, and neurocomputational models of the hippocampus have been shown to explain a wide range of extant memory data (e.g., Norman and O'Reilly, 2003; Elfman et al., 2008). Recent evidence suggests the hippocampus also plays an important role in high-level scene perception (e.g., Graham et al., 2010; Lee et al., 2005, 2012; Olsen, et al., 2012; Warren et al., 2012; Aly et al., 2013). Whether existing computational models are able to account for these perception findings has not been established. To address this issue, we conducted simulations using a hippocampal model that was adapted from the CLS framework (Norman and O'Reilly, 2003) and found that the same network was able to reproduce findings from both memory and perception paradigms.

Consistent with the patient and neuroimaging work implicating the hippocampus in strength-based perception (Aly et al., 2013), the model produced graded, overlapping levels of activity that discriminated between matching and mismatching pairs in the perception task. The same model showed a thresholded activity function for recognition memory, in agreement with prior work indicating that the hippocampus is involved in recollection in long-term memory (Yonelinas et al., 2002, 2005; Ranganath et al., 2004; Montaldi et al., 2006; see Eichenbaum et al., 2007 and Yonelinas et al., 2010 for review). In addition, the continuous perception signal was most strongly related to mismatch magnitude in the CA1 subfield.

The finding that CA1 was the primary locus of the perceptual signal is consistent with a number of models in which CA1 acts as a "comparator" of perceived versus remembered events (Eichenbaum and Buckingham 1990; Hasselmo and Wyble, 1997; Meeter et al., 2004; Lisman and Grace, 2005; Kumaran and Maguire, 2007; Yassa and Stark, 2011; Duncan et al., 2012). The current results extend this work by indicating that CA1 can also serve as a comparator in perceptual tasks. In the case of perception, the comparison is between a stimulus that is currently in the environment (or in the focus of attention), and one that was recently perceived (or recently attended).

Importantly, the current simulations with separate encoding/retrieval modes showed that the hippocampus can produce both continuous signals and thresholded signals without changing any properties of its architecture or basic mode of operation. The continuous output, however, was not dependent on having separate encoding and retrieval modes: a continuous perceptual signal was also observed in simulations in which the network was kept in an encoding mode. Thus, we need not assume that the hippocampus rapidly alternates between encoding and retrieval modes (e.g., Hasselmo et al., 2002) in order to support perceptual discriminations and output a graded signal.

Why are There Thresholded Signals in Memory and Graded Signals in Perception?

A critical point of success for the model is that we did not need to modify the way the network functioned—such as adjusting connectivity parameters—between the perception and recognition memory tasks in order for the network to perform both tasks successfully, or for it to produce the respective strength-based (continuous) and state-based (thresholded) signals. Rather, each signal emerged naturally as a result of the differing task demands; that is, the demand to differentiate between items with a high proportion of feature differences, as in the memory task, and the demand to differentiate between items with relatively few feature differences, as in the perception task.

The threshold nature of memory retrieval is not entirely attributable to any single subfield or parameter of the hippocampal model. However, one aspect of the model that is critical is the high level of inhibition within the hippocampal layers, particularly the DG, which leads to sparse activity and minimal representational overlap between encoded events. Consequently, only test cues that are very similar to a specific studied item will support successful retrieval. In addition, recurrent connectivity in CA3 produces nonlinear neural attractors that support a high level of pattern completion when a test cue is sufficiently similar to a stored pattern, but cues that do not match any specific studied item very well will effectively fail to lead to pattern completion. It is this "retrieval failure" that produces the thresholded hippocampal output.

In contrast, in our tests of perception, because the cue stimulus (i.e., the second item) is highly similar—or identical in the case of matching item trials—to the initial stimulus in the pair, the cue stimulus invariably leads to pattern completion of the initial item and produces an output that is highly correlated with the cue stimulus. There is variability in how closely the completed pattern matches the cue and this allows for the continuous signal that is useful in discriminating between match and mismatch trials. Importantly, because the two items in a perception trial are so similar to one another, the hippocampus is effectively always operating at a level well above its pattern completion threshold; thus, the perceptual signal reflects the continuous variability beyond the threshold.

The Role of the Hippocampus in Perception

A growing body of work suggests that the hippocampus is critical for perceptual decisions that require representations of detailed relational or spatial information (e.g., Graham et al., 2010; Lee et al., 2005, 2012; Olsen, et al., 2012; Warren et al., 2012; Aly et al., 2013; see Lee et al., 2012 and Yonelinas, 2013 for review). In the scene perception task used by Aly et al., (2013), the perceptual change was a subtle distortion designed to globally change the relational or configural information within the scene. The discriminations therefore required a detailed representation of how components of the scene were bound together within the overall spatial layout. We suggest that it is the need for this kind of representation that makes this task hippocampally dependent.

In contrast, if the task involved a manipulation of items rather than relational information (e.g., a tree that is present in one scene but absent in another), performance could depend on item-level representations and then presumably would be less dependent on the hippocampus (Baxter, 2009; Lee et al., 2012). Rather, regions in ventral temporal cortex, or perhaps perirhinal cortex, could be capable of supporting this discrimination (Lee et al., 2005; Saksida and Bussey, 2005; Graham et al., 2010). Thus, it is not our view that the hippocampus is needed for all perceptual discriminations; only that it is critically important for perceptual decisions based on assessments of precise, or high-resolution, relational information (Yonelinas, 2013). This is in keeping with a large number of studies that have found intact perceptual discrimination of colors, faces, and simple objects in patients with hippocampal damage (see Graham et al., 2010 and Lee et al., 2012 for review).

The Role of Other Brain Regions in Strength-Based Perception

Is the hippocampus the only brain region that is able to perform computations that support strength-based perception, or can other regions also support these judgments? Aly et al. (2013) reported that selective damage to the hippocampus is associated with significant impairments in strength-based perception, indicating that the hippocampus itself makes an important contribution. However, strength-based perception was not completely eliminated in hippocampal patients, suggesting that other brain regions might also be capable of supporting this kind of perceptual decision. Moreover, an fMRI study provided evidence that regions outside of the hippocampus are related to strength-based perception, including the parahippocampal cortex, lateral occipital complex, and fusiform gyrus (Aly et al., 2013, 2014).

The current simulations showed how the hippocampus can support strength-based perception via Hebbian weight changes, and additional simulations indicated that an adaptation-based mechanism can produce comparable results in the absence of weight-based learning (Appendix D, Fig. D1). Can a model of cortex also support perceptual discriminations via Hebbian weight changes and/or adaptation-based mechanisms? To investigate these questions, we conducted additional simulations with a simple, two-layer model of cortex, based on the CLS model (Norman and O'Reilly, 2003). We found that, using either Hebbian weight-based learning or neural adaptation, the cortex could reliably discriminate between match and mismatch trials (see Fig. E1, Appendix E). These results are consistent with the above-mentioned hypothesis that, although the hippocampus makes an important contribution to strength-based perception, it may not be the only region that can support these judgments.

The findings with the cortical model raise an important question, namely, does the hippocampus play a role in strength-based perception that is distinct from other regions? We have suggested previously that the hippocampus is critical in supporting complex, high-resolution bindings (Yonelinas, 2013). That is, the hippocampus is critical in linking together the complex, qualitative aspects that make up events, and in a way that maintains high-resolution, detailed information (e.g., the specific color of the shirt you wore in a specific location at a specific time). Detecting the global changes used in the Aly et al. (2013) study would presumably benefit from such high-resolution, relational representations—representations that contain detailed information about the components of the scenes as well as their precise relations to one another.

The hippocampus may contribute to strength-based perception by virtue of its representation of high-resolution, relational information; but other regions may contribute to strength-based perception in other ways. Although the manipulation used in Aly et al., (2013) is "global" in the sense of extending over the entire image, relatively local components are also affected. For example, "pinching" an image may alter the perception of individual objects within the image. Alternatively, global changes over the image may have effects on the perceived depth of the scene. The detection of these kinds of changes (i.e., to individual objects, or to image depth) need not require the hippocampus. Rather, visual or spatial representations in medial temporal, occipito-temporal, or parietal cortical regions may be sufficient. Thus, hippocampal damage reduces performance, but does not entirely eliminate accurate strength-based perception (Aly et al., 2013). If these hypotheses are correct, it should be possible to take the current scene perception task and reduce the extent to which the hippocampus will be critical by decreasing the complexity of relational bindings or by varying the extent to which the required discrimination relies on high-resolution (detailed) versus low-resolution (less precise) representations (see Yonelinas, 2013). Such studies are ongoing.

Limitations and Future Directions

A matter for future consideration concerns the direction of the relationship between hippocampal activity and match/mismatch detection. In the study by Aly et al. (2013), the level of hippocampal activity—inferred from the blood-oxygen-level dependent (BOLD) signal—was positively correlated with mismatch confidence. That is, the more that individuals were confident that a pair of items was different, the greater the hippocampal activity. By contrast, the current model showed greater overall activity for matching items than for mismatching items. The relationship between the BOLD signal, neural firing rates, and local field potentials is not yet fully understood (Ekstrom et al., 2009; Ekstrom, 2010; Logothetis, 2003), so the reason BOLD activity increased whereas activity in the model decreased cannot be confidently addressed. One possibility is that increases in the BOLD signal do not reflect increases in activation strength per se, but rather the time it takes the network to settle into a stable pattern. In fact, we found that the model took longer to settle for mismatching pairs, which is at least consistent with this explanation.

Another possible reason for the difference in signals is the inherent limitations in the algorithm used in the current simulations. Specifically, competitive inhibition was simulated as a fairly firm limit on the number of neurons that can be activated above a certain level (O'Reilly and Munakata, 2000). In our simulations, mismatching stimulus cues caused excitation to be spread out over more CA1 units relative to matching cues. With a more flexible inhibition rule it is possible that this would have led to an increase, rather than decrease, in overall CA1 activation.

Finally, another factor that may affect whether perceptual matches lead to an increase or decrease in neural activity is adaptation. Adaptation (or repetition suppression) is frequently reported in fMRI studies that involve repeating stimuli, and this adaptation has been observed in the MTL (Goh, et al., 2004; Howard et al., 2011; Diana et al., 2012; see Krekelberg et al., 2006). In these studies, presenting the same stimulus twice in a row is found to produce a diminished neural response (note, however, that repetition enhancement is also sometimes observed; see Turk-Browne et al., 2007).

In post-hoc simulations, we found that incorporating neural adaptation into the hippocampal model reversed the sign of match/mismatch difference, such that mismatching items were associated with greater activity than matching items (Fig. D1, Appendix D). That is, as the level of mismatch increased, activity also increased in CA1 (but not CA3, consistent with our simulations without adaptation). We also found that the network produced similar results with or without Hebbian learning, that is, using adaptation alone (data not shown). This suggests that activation-based mechanisms (adaptation), in contrast to weight-based mechanisms (Hebbian learning), might be an important component of how the hippocampus performs tasks in which information need not persist over long intervals. Nevertheless, the inclusion of adaptation did not improve the ability of the model to account for graded perceptual signals; CA3 alone was vastly less effective at performing adaptation-based mismatch detection compared with CA1, and CA1 was effective with or without the inclusion of adaptation. Importantly, CA3 did not produce graded signals in the simulations with adaptation. Nonetheless, future research examining the effects of neural adaptation on the model will be useful.

The current simulations lead to some potentially interesting predictions of how different hippocampal subfields might respond to changes in perceptual similarity that could be assessed using high-resolution fMRI. For example, by incrementally manipulating the number of mismatching features between pairs of items in a perceptual matching task, different subregions of the hippocampus should respond in comparatively unique ways. That is, based on the simulations, one would expect CA3 to show a consistent response at low levels of mismatch, indicative of pattern completion, and a fairly steep transition beyond some critical level of mismatch. Conversely, the CA1 subfield should exhibit a more gradual and continuous change in signal. Results that are broadly consistent with this were reported in a recent high-resolution fMRI memory study (Lacy et al., 2011), in which participants viewed a series of images that were either first presentations of an image, low-similarity lures, high-similarity lures, or repetitions of previous items. In CA1, activity was found to vary in a graded fashion with the level of change, whereas in CA3/DG, activity transitioned more steeply for small changes.

CONCLUSIONS

Several lines of evidence suggest the hippocampus contributes to a range of cognitive functions, but how it performs seemingly disparate tasks under a single neural architecture is a question of great interest. The modeling work presented here provides an important step towards a unified understanding of hippocampal functioning across the domains of memory and visual perception, by showing how the hippocampal network is able to produce contrasting signals—that is, strength-based perception and state-based memory—using a common set of core assumptions. Future research will undoubtedly add to this emerging picture.

APPENDIX A

Basic Network Parameters

The following notes are a selective description of the rules and parameters used in the model simulations. Table A1 shows each layer size (i.e., number of units) and percentage of activity determined by the *k*-winners-take-all (Norman and O'Reilly, 2003). Table A2 shows the properties of the main projections, including the mean initial weight strengths (Mean), variances of the weight distribution (Var), relative strengths of the projections during encoding (Scale enc) and retrieval (Scale retr), and the proportions of receiving units that each sending unit is connected to (% Con). The learning rate was 0.1. The project file is available for download at http://goo.gl/4XkKya.

TABLE A1.

Layer Sizes and Activity Levels

Layer/area	Units	Activity (%)
Entorhinal Cortex (in/out)	144	25.0
DG	1600	1.0
Area CA3	480	3.8
Area CA1	384	9.4

Note. In/out = input and output layers, respectively.

TABLE A2.

Properties of Network Projections

Projection	Mean	Var	Scale (Enc)	Scale (Retr)	% Con
EC to DG, CA3	0.5	0.1	1	1	25
CA3 recurrent	0.5	0.1	2	2	100
CA3 to CA1	0.5	0.1	0.3	0.3	90
DG to CA3	0.9	0.01	15	0	4
EC to CA1	0.5	0.25	1	0.15	8.3
CA1 to EC	0.5	0.25	0.25	0.25	8.3

Note: Mean: mean initial weight strength; Var: variance of initial weight distribution; Scale: scaling of this projection relative to other projections; Enc: encoding mode, Retr: retrieval mode; % Con: percentage connectivity; EC: entorhinal cortex; DG: dentate gyrus.

APPENDIX B

CA1 Activity Distributions Ranging 0–100% Feature Mismatch

We repeated the perception simulations with a broader range of match/mismatch levels, systematically varying the number of mismatching features from 0 (complete match) to 36 (complete mismatch). Activity distributions for CA1 across the varying levels of mismatch are presented in Figure B1. These simulations show that, as the number of mismatching features is systematically increased, CA1 activity distributions gradually shift from high to low activity. Moreover, with very high levels of mismatch, the "match" and "mismatch" distributions are essentially nonoverlapping, suggesting consistent pattern completion for "match" items and a failure of pattern completion for "mismatch" items. The latter finding is analogous to recognition memory simulations, in which pattern completion never occurs for nonstudied (lure) items (compare 24, 30, and 36 feature mismatch items in Fig. B1 to the nonstudied item distribution in Fig. 3).

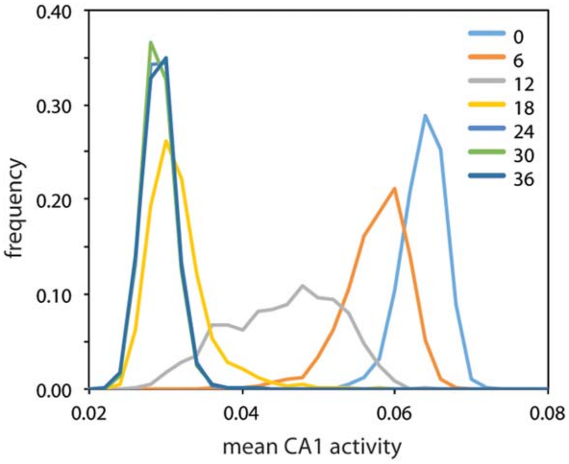

FIGURE B1. **Perception strength distributions of CA1 activity as the number of feature differences is varied from 0 to 36 (of 36), at intervals of 6. CA1 produces continuous strength distributions that gradually move farther apart as the number of feature differences is increased. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]**

APPENDIX C

Partial Retrieval Cue Memory Simulation

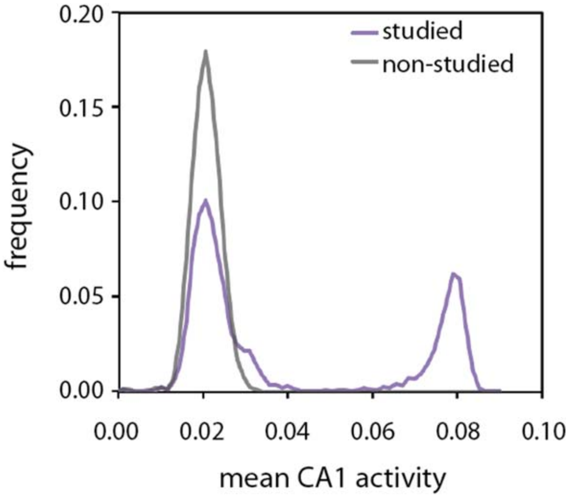

FIGURE C1. Recognition memory distributions of CA1 activity for studied versus nonstudied items using partial retrieval cues (50% of features were blank at test). CA1 produces a bimodal distribution for studied items, indicating a pattern completion threshold in the hippocampus. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

APPENDIX D

Match/Mismatch Sensitivity With Adaptation

The simulations reported in this paper generally showed greater activity for matching versus mismatching items pairs in the perception task. However, Aly, et al. (2013) found evidence for the opposite effect (i.e., greater hippocampal activity for mismatching versus matching item pairs). Additional simulations were conducted to test the hypothesis that neural adaptation can account for this discrepancy. Figure D1A shows distributions of average activity in CA1. As expected, mismatching pairs are now associated with an increase in the overall level of activity. Figure D1B compares match/mismatch sensitivity for CA1 and CA3 as mismatch was varied from 0 to 100%. CA3 was comparatively very poor at discriminating between matching and mismatching pairs.

Adaptation was included using the "accommodation" parameters of cell activation from the leabra algorithm. There was a firing-rate threshold of onset of .1 and an update rate of 0.7. Computations were made at the end of each trial.

APPENDIX E

Match/Mismatch Strength in the Cortical Model

Additional simulations were run to test whether a simple two-layer model of the cortex (comprising an input layer and a neocortical, or hidden, layer of 800 units) could perform the perception task. Simulations were run using Hebbian-based learning (Fig. E1A) with a learning rate of 0.004, and separately using the adaptation parameters listed in Appendix D (Fig. E1B). Using four feature differences for the mismatching pairs, both approaches were successful in discriminating between matching and mismatch item pairs, and produced overlapping, Gaussian-like distributions that were comparable to those observed in area CA1 of the hippocampal model.

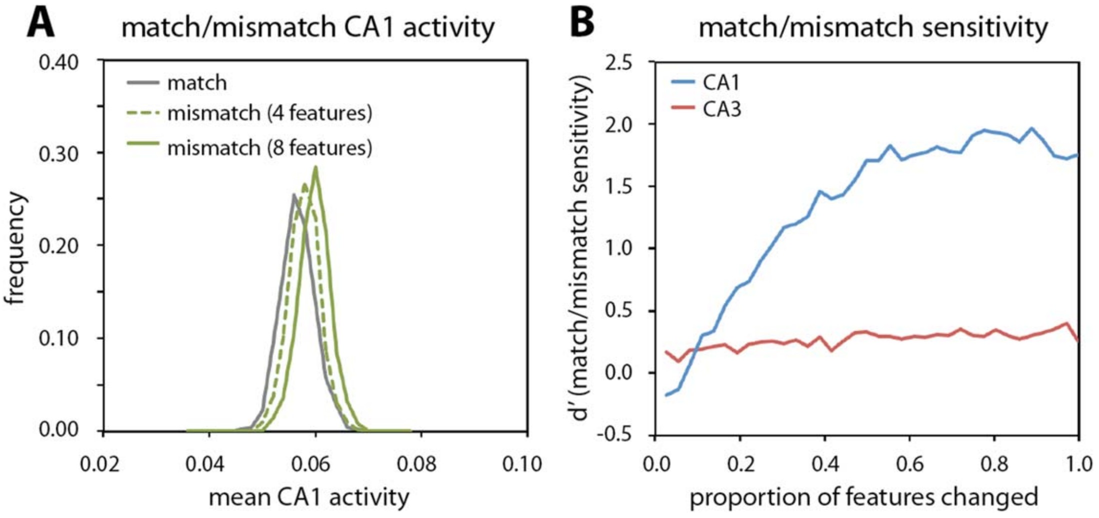

FIGURE D1. Hippocampal network performance for the match/mismatch perception task with the inclusion of neural adaptation (using the encode/retrieve model). A: Distributions of average activity levels for CA1 for matching items and mismatching items with 4 and 8 features changed. B: Sensitivity (d') of CA1 and CA3 as the number of mismatching features was varied from 1 to 36. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

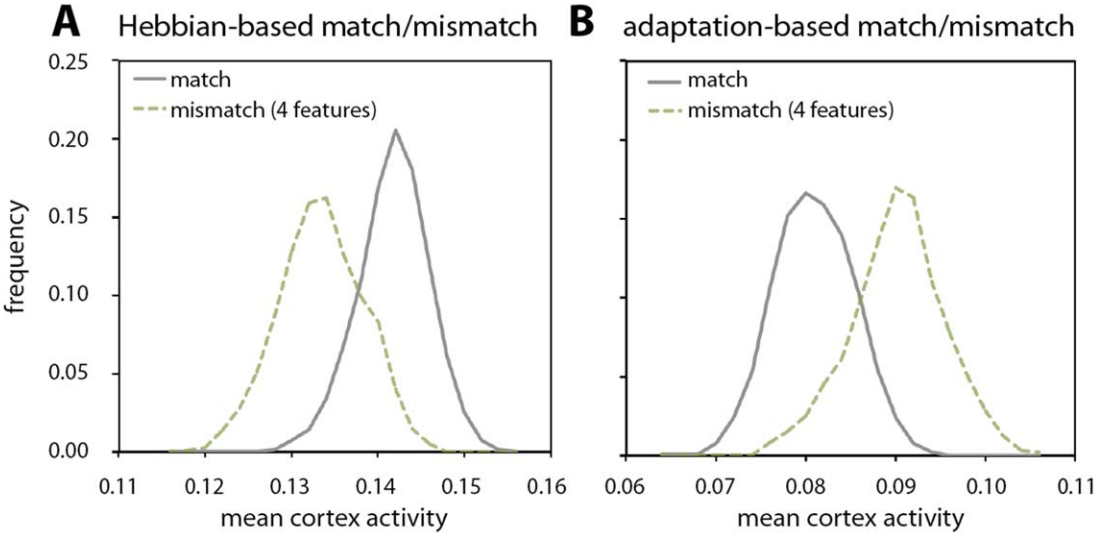

FIGURE E1. Match/mismatch task performance using a two-layer cortical model. The input layer is identical to the input layer for the hippocampal model; the neocortical layer is composed of 800 units with 20% active at a given time. A: Distributions of average activity levels with four mismatching features, using Hebbian weight-based learning. B: Distributions using adaptation and no weight-based learning. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

REFERENCES

Aggleton JP, Brown MW. 1999. Episodic memory, amnesia, and the hippocampal-anterior thalamic axis. Behav Brain Sci 22:425–444.

Aisa B, Mingus B, O'Reilly R. 2008. The Emergent neural modeling system. Neural Netw 21:1146–1152.

Aly M, Yonelinas AP. 2012. Bridging consciousness and cognition in memory and perception: Evidence for both state and strength processes. PLoS one 7:e30231.

Aly M, Ranganath C, Yonelinas AP. 2013. Detecting changes in scenes: The hippocampus is critical for strength-based perception. Neuron 78:1127–1137.

Aly M, Ranganath C, Yonelinas AP. 2014. Neural correlates of state- and strength-based perception. J Cogn Neurosci. 26:792–809.

Axmacher N, Cohen MX, Fell J, Haupt S, Dümpelmann M, Elger CE, Ranganath C. 2010. Intracranial EEG correlates of expectancy and memory formation in the human hippocampus and nucleus accumbens. Neuron 65:541–549.

Baxter MG. 2009. Involvement of medial temporal lobe structures in memory and perception. Neuron 61:667–677.

Baddeley AD, Warrington EK. 1970. Amnesia and the distinction between long-and short-term memory. J Verbal Learn Verbal Behav 9:176–189.

Barense MD, Henson RN, Lee ACH, Graham KS. 2010. Medial temporal lobe activity during complex discrimination of faces, objects, and scenes: Effects of viewpoint. Hippocampus 20:389–401.

Bussey TJ, Saksida LM. 2005. Object memory and perception in the medial temporal lobe: An alternative approach. Curr Opin Neurobiol 15:730–737.

Cohen NJ, Eichenbaum H. 1993. Memory, Amnesia, and the Hippocampal System. Cambridge, MA: MIT Press.

Colgin LL, Leutgeb S, Jezek K, Leutgeb JK, Moser EI, McNaughton BL, Moser MB. 2010. Attractor-map versus autoassociation based attractor dynamics in the hippocampal network. J Neurophysiol 104:35–50.

Diana RA, Yonelinas AP, Ranganath C. 2010. Medial temporal lobe activity during source retrieval reflects information type, not memory strength. J Cogn Neurosci. 22:1808–1818.

Duncan K, Ketz N, Inati SJ, Davachi L. 2012. Evidence for area CA1 as a match/mismatch detector: A high-resolution fMRI study of the human hippocampus. Hippocampus 22:389–398.

Eichenbaum H, Buckingham J. 1990. Studies on hippocampal processing: Experiment, theory, and model. In Learning and Computational Neuroscience: Foundations of Adaptive Networks (Eds: Michael Gabriel and John Moore). pp. 171–231. Cambridge, MA: MIT Press.

Eichenbaum H, Cohen NJ. 2001. From conditioning to conscious recollection: Memory systems of the brain. New York, NY: Oxford University Press.

Eichenbaum H, Yonelinas AR, Ranganath C. 2007. The medial temporal lobe and recognition memory. Annu Rev Neurosci 30:123.

Ekstrom A, Suthana N, Millett D, Fried I, Bookheimer S. 2009. Correlation between BOLD fMRI and theta-band local field potentials in the human hippocampal area. J. Neurophysiol., 101, 2668–2678.

Ekstrom AD. 2010. How and when the fMRI BOLD signal relates to underlying neural activity: The danger in dissociation. Brain Res Rev 62:233–244.

Eldridge LL, Knowlton BJ, Furmanski CS, Bookheimer SY, Engel SA. 2000. Remembering episodes: A selective role for the hippocampus during retrieval. Nature Neurosci 3:1149–1152.

Elfman KW, Parks CM, Yonelinas AP. 2008. Testing a neurocomputational model of recollection, familiarity, and source recognition. J Exp Psychol Learn Mem Cogn 34:752.

Fell J, Ludowig E, Staresina BP, Wagner T, Kranz T, Elger CE, Axmacher N. 2011. Medial temporal theta/alpha power enhancement precedes successful memory encoding: Evidence based on intracranial EEG. J Neurosci 31:5392–5397.

Goh JOS, Siong SC, Park D, Gutchess A, Hebrank A, Chee MWL. 2004. Cortical areas involved in object, background, and object-background processing revealed with functional magnetic resonance adaptation. J Neurosci 24:10223–10228.

Graf P, Schacter DL. 1984. Implicit and explicit memory for new associations in normal and amnesic subjects. J Exp Psychol Learn Mem Cogn 11:501–518.

Graham KS, Barense MD, Lee ACH. 2010. Going beyond LTM in the MTL: A synthesis of neurospsychological and neuroimaging finds on the role of the medial temporal lobe in memory and perception. Neuropsychologia 48:831–853.

Hasselmo ME, Eichenbaum H. 2005. Hippocampal mechanisms for the context-dependent retrieval of episodes. Neural Netw 18:1172–1190.

Hasselmo ME, Wyble BP. 1997. Free recall and recognition in a network model of the hippocampus: Simulating effects of scopolamine on human memory function. Behav Brain Res 89:1–34.

Hasselmo ME, Bodelon C, Wyble BP. 2002. A proposed function for hippocampal theta rhythm: Separate phases of encoding and retrieval enhance reversal of prior learning. Neural Comput 14:793–818.

Howard LR, Kumaran D, Ólafsdóttir HF, Spiers HJ. 2011. Double dissociation between hippocampal and parahippocampal responses to object–background context and scene novelty. J Neurosci 31: 5253–5261.

Krekelberg B, Boynton GM, van Wezel RJA. 2006. Adaptation: From single cells to BOLD signals. Trends Neurosci 29:250–256.

Kumaran D, Maguire EA. 2007. Which computational mechanisms operate in the hippocampus during novelty detection? Hippocampus 17:735–748.

Lacy JW, Yassa MA, Stark SM, Muftuler LT, Stark CE. 2011. Distinct pattern separation related transfer functions in human CA3/dentate and CA1 revealed using high-resolution fMRI and variable mnemonic similarity. Learn Mem 18:15–18.

Lee ACH, Rudebeck SR. 2010. Investigating the interaction between spatial perception and working memory in the human medial temporal lobe. J Cogn Neurosci 22:2823–2835.

Lee ACH, Buckley MJ, Pegman SJ, Spiers H, Scahill VL, Gaffan D, Bussey TJ, Davies RR, Kapur N, Hodges JR, Graham KS. 2005. Specialization in the medial temporal lobe for processing of objects and scenes. Hippocampus 15:782–797.

Lee ACH, Scahill VL, Graham KS. 2008. Activating the medial temporal lobe during oddity judgment for faces and scenes. Cereb Cortex 18:683–696.

Lee ACH, Yeung L-K, Barense MD. 2012. The hippocampus and visual perception. Front Hum Neurosci 6, Article 91:1–17.

Leutgeb JK, Leutgeb S, Moser MB, Moser EI. 2007. Pattern separation in the dentate gyrus and CA3 of the hippocampus. Science 315:961–966.

Leutgeb S, Leutgeb JK, Moser EI, Moser MB. 2006. Fast rate coding in hippocampal CA3 cell ensembles. Hippocampus 16:765–774.

Lisman JE, Grace AA. 2005. The hippocampal-VTA loop: Controlling the entry of information into long-term memory. Neuron 46:703–713.

Logothetis NK. 2003. The underpinnings of the BOLD functional magnetic resonance imaging signal. J Neurosci 23:3963–3971.

Marr D. 1971. Simple memory: A theory for archicortex. Philos Trans R Soc London Ser B 262:23–81.

McClelland JL, McNaughton BL, O'Reilly RC. 1995. Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. Psychol Rev 102:419–457.

McNaughton BL, Morris RG. 1987. Hippocampal synaptic enhancement and information storage within a distributed memory system. Trends Neurosci 10:408–415.

Meeter M, Murre JMJ, Talamini LM. 2004. Mode shifting between storage and recall based on novelty detection in oscillating hippocampal circuits. Hippocampus 14:722–741.

Montaldi D, Spencer TJ, Roberts N, Mayes AR. 2006. The neural system that mediates familiarity memory. Hippocampus 16:504–520.

Mundy ME, Downing PE, Graham KS. 2012. Extrastriate cortex and medial temporal lobe regions respond differentially to visual feature overlap within preferred stimulus category. Neuropsychologia 50: 3053–3061.

Nadel L, Peterson MA. 2013. The hippocampus: Part of an interactive posterior representational system spanning perceptual and memorial systems. J Exp Psychol Gen 142:1242–1254.

Norman KA, O'Reilly RC. 2003. Modeling hippocampal and neocortical contributions to recognition memory: A complementary-learning-systems approach. Psychol Rev 110:611.

O'Reilly RC, McClelland JL. 1994. Hippocampal conjunctive encoding, storage, and recall: Avoiding a trade-off. Hippocampus 4: 661–682.

O'Reilly RC, Munakata Y. 2000. Computational explorations in cognitive neuroscience. Cambridge: MIT Press.

Olson IR, Page K, Moore KS, Chatterjee A. Verfaellie M. 2006. Working memory for conjunction relies on the medial temporal lobe. J Neurosci 26:4596–4601.

Olsen RK, Moses SN, Riggs L, Ryan JD. 2012. The hippocampus supports multiple cognitive processes through relational binding and comparison. Front Hum Neurosci 6, Article 146: 1–13.

Parks CM, Yonelinas AP. 2007. Moving beyond pure signal-detection models: Comment on Wixted. Psych Rev 114:188–202.

Parks CM, Murray LJ, Elfman KW, Yonelinas AP. 2011. Variations in recollection: The effects of complexity on source recognition. J Exp Psychol Learn Mem Cogn 37:861.

Ranganath C, Yonelinas AP, Cohen MX, Dy CJ, Tom SM, D'Esposito M. 2004. Dissociable correlates of recollection and familiarity within the medial temporal lobes. Neuropsychologia 42:2–13.

Reed JM, Squire LR. 1998. Retrograde amnesia for facts and events: Findings from four new cases. J Neurosci 18:3943–3954.

Rolls ET. 1996. A theory of hippocampal function in memory. Hippocampus 6:601–620.

Treves A, Rolls ET. 1994. Computational analysis of the role of the hippocampus in memory. Hippocampus 4:374–391.

Rudy JW, O'Reilly RC. 1999. Contextual fear conditioning, conjunctive representations, pattern completion, and the hippocampus. Behav Neurosci 113:867–880.

Saksida LM, Bussey TJ. 2010. The representational-hierarchical view of amnesia: Translation from animal to human. Neuropsychologia 48:2370–2384.

Schacter DL, Verfaellie M, Pradere D. 1996. The neuropsychology of memory illusions: False recall and recognition in amnesic patients. J Mem Lang 35:319–334.

Scoville WB, Milner B. 1957. Loss of recent memory after bilateral hippocampal lesions. J Neurol Neurosurg Psychiatry 20:11.

Sederberg PB, Gauthier LV, Terushkin V, Miller JF, Barnathan JA, Kahana MJ. 2006. Oscillatory correlates of the primacy effect in episodic memory. NeuroImage 32:1422–1431.

Shohamy D, Turk-Browne NB. 2013. Mechanisms for widespread hippocampal involvement in cognition. J Exp Psychol Gen 142: 1159–1170.

Slotnick SD, Thakral PP. 2013. The hippocampus operates in a threshold manner during spatial source memory. NeuroReport 24: 265–269.

Staresina BP, Fell J, Dunn JC, Axmacher N, Henson RN. 2013. Using state-trace analysis to dissociate the functions of the human hippocampus and perirhinal cortex in recognition memory. Proc Natl Acad Sci 110:3119–3124.

Swets JA. 1988. Measuring the accuracy of diagnostic systems. Science 240:1285–1293.

Tulving E, Markowitsch HJ. 1998. Episodic and declarative memory: Role of the hippocampus. Hippocampus 8:198–204.

Turk-Browne NB, Yi DJ, Leber AB, Chun MM. 2007. Visual quality determines the direction of neural repetitions effects. Cereb Cortex 17:425–433.

Vargha-Khadem F, Gadian DG, Watkins KE, Connelly A, Van Paesschen W, Mishkin M. 1997. Differential effects of early hippocampal pathology on episodic and semantic memory. Science 277: 376–380.

Warren DE, Duff MC, Jensen U, Tranel D, Cohen NJ. 2012. Hiding in plain view: Lesions of the medial temporal lobe impair online representation. Hippocampus 22:1577–1588.

Yassa MA, Stark CE. 2011. Pattern separation in the hippocampus. Trends Neurosci 34:515–525.

Yonelinas AP. 1994. Receiver-operating characteristics in recognition memory: Evidence for a dual-process model. J Exp Psychol Learn Mem Cogn 20:1341.

Yonelinas AP. 2001. Components of episodic memory: The contribution of recollection and familiarity. Philos Trans R Soc London Ser B 356:1363–1374.

Yonelinas AP. 2002. The nature of recollection and familiarity: A review of 30 years of research. J Mem Lang 46:441–517.

Yonelinas AP. 2013. The hippocampus supports high-resolution binding in the service of perception, working memory and long-term memory. Behav Brain Res 254:34–44.

Yonelinas AP, Kroll NE, Quamme JR, Lazzara MM, Sauvé MJ, Widaman KF, Knight RT. 2002. Effects of extensive temporal lobe damage or mild hypoxia on recollection and familiarity. Nat Neurosci 5:1236–1241.

Yonelinas AP, Otten LJ, Shaw KN, Rugg MD. 2005. Separating the brain regions involved in recollection and familiarity in recognition memory. J Neurosci 25:3002–3008.

Yonelinas AP, Aly M, Wang WC, Koen JD. 2010. Recollection and familiarity: Examining controversial assumptions and new directions. Hippocampus 20:1178–1194.

Zola-Morgan S, Squire LR, Amaral DG. 1986. Human amnesia and the medial temporal region: Enduring memory impairment following a bilateral lesion limited to field CA1 of the hippocampus. J Neurosci 6:2950–2967.