Published in final edited form as:

Hippocampus. 2010 November ; 20(11): 1217–1227. doi:10.1002/hipo.20855.

How hippocampus and cortex contribute to recognition memory: Revisiting the Complementary Learning Systems model

Kenneth A. Norman

Department of Psychology and Princeton Neuroscience Institute, Princeton University, Princeton, NJ 08540 USA

Kenneth A. Norman: knorman@princeton.edu

Abstract

We describe how the Complementary Learning Systems neural network model of recognition memory (Norman & O'Reilly, 2003) can shed light on current debates regarding hippocampal and cortical contributions to recognition memory. We review simulation results illustrating three critical differences in how (according to the model) hippocampus and cortex contribute to recognition memory, all of which derive from the hippocampus' use of *pattern separated* representations. Pattern separation makes the hippocampus especially well-suited for discriminating between studied items and related lures; it makes the hippocampus especially poorly suited for computing global match; and it imbues the hippocampal ROC curve with a Y-intercept > 0 . We also describe a key boundary condition on these differences: When the average level of similarity between items in an experiment is very high, hippocampal pattern separation can fail, at which point the hippocampal model will start to behave like the cortical model. We describe the implications of these simulation results for extant debates over how to describe hippocampal vs. cortical contributions and how to measure these contributions.

1. Introduction

Over the past decade, there has been a heated debate over how to characterize the contributions of the hippocampus vs. medial temporal lobe cortex to recognition memory. Some researchers have argued that hippocampus and cortex implement recollection and familiarity, respectively (e.g., Yonelinas, 2002; Aggleton & Brown, 1999), whereas others have argued that this distinction does not adequately capture the differences in these structures' contributions (e.g., Squire, Wixted, & Clark, 2007). One reason that this debate has lasted so long is because of disagreement over how to measure hippocampal and cortical contributions. Yonelinas and colleagues have developed several techniques for estimating recollection and familiarity based on behavioral data (Yonelinas, 1994; Yonelinas, Kroll, Dobbins, Lazzara, & Knight, 1998); these procedures rely on a core set of assumptions about the properties of recollection and familiarity. If these assumptions adequately describe the (respective) contributions of hippocampus vs. cortex to recognition memory, then Yonelinas' measurement procedures – which were originally developed apart from neural considerations – can be used to estimate the contributions of these structures. However, the validity of the assumptions is still very much in question (e.g., Wixted, 2007; Wixted & Squire, 2008).

In this paper, we describe how a computational model that my colleagues and I have developed (the Complementary Learning Systems neural network model; Norman & O'Reilly, 2003) can help to address these controversies. Importantly, this model does not build in any specific claims about how hippocampus vs. cortex might contribute to recognition memory. Rather, our approach has been to start out with some widely-accepted ideas about key anatomical and physiological properties of hippocampus and cortex, and

then use our computational models to explore the implications of these ideas for how these structures might contribute to recognition judgments. The most important of these ideas is that hippocampus implements *pattern separation* (i.e., it is biased to assign distinct representations to input patterns, regardless of their similarity) but cortex does not. The model allows us to explore the implications of this difference in pattern separation for recognition performance.

In the Section 2 of this paper, we review the basic Complementary Learning Systems account of the division of labor between hippocampus and cortex (McClelland, McNaughton, & O'Reilly, 1995). In the Section 3, we describe how we have applied our models of hippocampus and cortex to recognition memory data. In the Section 4, we review simulation results illustrating three critical differences in how (according to the model) hippocampus and cortex contribute to recognition memory; all of these differences relate to the difference in pattern separation described above. We also describe a key boundary condition (i.e., limit) on these differences: When the average level of similarity between items in an experiment is very high (e.g., all of the stimuli in the experiment are computer-generated male faces; Yotsumoto, Kahana, Wilson, & Sekuler, 2007), hippocampal pattern separation can fail, at which point the hippocampal model will start to behave like the cortical model. Finally, in Section 5, we describe the implications of these simulation results for extant debates about the neural basis of recognition memory.

2. The Complementary Learning Systems framework

The Complementary Learning Systems (*CLS*) framework developed by McClelland et al. (1995) incorporates several widely-held ideas about the division of labor between hippocampus and neocortex that have been developed over many years by many different researchers (e.g., Scoville & Milner, 1957; Marr, 1971; Grossberg, 1976; O'Keefe & Nadel, 1978; Teyler & Discenna, 1986; McNaughton & Morris, 1987; Sherry & Schacter, 1987; Rolls, 1989; Sutherland & Rudy, 1989; Squire, 1992; Eichenbaum, Otto, & Cohen, 1994; Treves & Rolls, 1994; Burgess & O'Keefe, 1996; Wu, Baxter, & Levy, 1996; Moll & Miikkulainen, 1997; Hasselmo & Wyble, 1997; Aggleton & Brown, 1999; Becker, 2005). According to the CLS framework, neocortex forms the substrate of our internal model of the structure of the environment. In contrast, hippocampus is specialized for rapidly and automatically memorizing patterns of cortical activity, so they can be recalled later based on partial cues (*pattern completion*). The model posits that neocortex learns incrementally; each training trial results in relatively small adaptive changes in synaptic weights. These small changes allow cortex to gradually adjust its internal model of the environment in response to new information. The other key property of neocortex (according to the model) is that it assigns similar (overlapping) representations to similar stimuli. Use of overlapping representations allows cortex to represent the shared structure of events, and therefore makes it possible for cortex to generalize to novel stimuli based on their similarity to previously experienced stimuli. In contrast, the model posits that hippocampus assigns distinct, *pattern separated* representations to stimuli, regardless of their similarity. This property allows hippocampus to rapidly memorize arbitrary patterns of cortical activity without suffering unacceptably high (catastrophic) levels of interference.

3. Applying CLS to recognition memory data

Norman and O'Reilly (2003) built cortical and hippocampal neural networks that instantiate the CLS principles outlined above. Here, we describe how we applied these neural networks to recognition memory data.

3.1. Cortical contributions to recognition memory

The cortical model can not support pattern completion of novel associations after a single study exposure (due to its relatively low learning rate), but it can discriminate between studied and nonstudied items based on the sharpness of their representations in perirhinal cortex.

Figure 1 illustrates the structure of the CLS cortical model (note that the actual model had more simulated neurons). The model consists of an input layer (corresponding to lower regions of the cortical hierarchy) which projects in a feedforward fashion to a hidden layer (corresponding to perirhinal cortex). As mentioned earlier, the main function of cortex is to extract statistical regularities in the environment; the two-layer CLS cortical network (where perirhinal hidden units compete to encode regularities that are present in the input layer) is meant to capture this idea in the simplest possible fashion. Because the cortical model uses a small learning rate, it is not capable of pattern completion following limited exposure to a stimulus. However, it is possible to extract a scalar signal from the cortical model that discriminates between previously presented and novel stimuli. In the cortical model, as items are presented repeatedly, their representations in the upper (perirhinal) layer become *sharper*. Novel stimuli weakly activate a large number of perirhinal units, whereas previously presented stimuli strongly activate a relatively small number of units. Sharpening occurs in the model because Hebbian learning specifically tunes some perirhinal units to represent the stimulus. When a stimulus is first presented, some perirhinal units, by chance, will respond more strongly to the stimulus than other units. These “winning” units get tuned by Hebbian learning to respond even more strongly to the item then next time it is presented; this increased response triggers an increase in feedback inhibition to units in the layer, resulting in decreased activation of the “losing” units. This latter property (whereby some initially-responsive units drop out of the stimulus representation as it is repeated) is broadly consistent with the neurophysiological finding that some perirhinal neurons show decreased responding as a function of how often a stimulus has been presented (e.g., Xiang & Brown, 1998; Li, Miller, & Desimone, 1993; Viskontas, Knowlton, Steinmetz, & Fried, 2006). Figure 1 illustrates this sharpening dynamic.

In the Norman and O'Reilly (2003) paper, the cortical memory signal was operationalized by reading out the activation of the k winning units in the perirhinal layer (where k is a model parameter that defines the maximum number of units that are allowed to be strongly active at once), although other methods of operationalizing familiarity are possible (e.g., we can measure how quickly activation spreads to the perirhinal layer of the network from the input layer; Norman & O'Reilly, 2003). We should also note that sharpening is not a novel property of the CLS model. Rather, it is a general property of competitive-learning networks with graded unit-activation values in which there is some kind of contrast enhancement within a layer (see, e.g., Grossberg, 1986, Section 23, and Grossberg & Stone, 1986, Section 16).

While the Norman & O'Reilly (2003) cortical model provides a straightforward implementation of sharpening, the model falls short in several ways. For example, Bogacz and Brown (2003) demonstrated that the simple Hebbian learning rule used in the model yields grossly inadequate storage capacity (i.e., the number of novel vs. familiar patterns that can be successfully discriminated does not scale well as a function of network size). The model also fails to account for cortex's ability to complete well-learned patterns. Extant evidence from amnesic patients suggests that, while the hippocampus may be needed to complete novel patterns that have only been studied once, patients with relatively spared medial temporal lobe cortex can learn to recall novel associations (e.g., new semantic facts) after extensive exposure to these novel associations (e.g., Gardiner et al., 2008). Because the Norman & O'Reilly (2003) cortical model only has feedforward connections (i.e., it lacks

feedback connections from the upper layer to the lower layer, and it lacks recurrent connections within layers), the model lacks the ability to fill in missing pieces of well-learned input patterns. To remedy these deficits, we have explored a variant of the model that includes feedback and recurrent connections and uses a new learning algorithm (the *Oscillating Learning Algorithm*; Norman, Newman, & Perotte, 2005; Norman, Newman, Detre, & Polyn, 2006). The new version of the model has much better storage capacity than the original version, and it also has the ability to fill in missing pieces of stored patterns after extensive exposure to those patterns. Importantly, all of the cortical model properties that are described in Section 4 (below) also apply to this updated version of the model.

Another shortcoming of the Norman & O'Reilly (2003) cortical model is that it does not directly address questions of functional specialization within cortex: In our description of the model (above), we asserted that that judgments of the familiarity of object stimuli are based on the sharpness of representations in perirhinal cortex, but we did not explain *why* perirhinal cortex (out of all cortical regions) plays this role. With regard to this question, the CLS framework (as originally formulated by McClelland, McNaughton, & O'Reilly) is broadly compatible with the representational-hierarchical view of cortical organization expressed by Cowell, Bussey, & Saksida (this issue). According to this view, cortex is arranged into a hierarchy of layers, where each layer learns to represent conjunctions of features on the layer below it. Perirhinal cortex represents highly complex conjunctions of object features by virtue of its position at the top of the ventral visual stream hierarchy. As described by Cowell et al. (this issue; see also Bussey, Saksida, & Murray, 2002; Bussey & Saksida, 2002; Barense et al., 2005; Cowell, Bussey, & Saksida, 2006, 2009), this ability to represent complex feature conjunctions makes perirhinal cortex especially useful for high-level object perception (i.e., discriminating between objects that contain different combinations of complex features); it also makes perirhinal cortex especially useful for object recognition memory, compared to regions lower in the hierarchy. The simple visual features (e.g., straight lines) that are represented in lower cortical regions are likely to be present in both studied items and lures, rendering them useless for discriminating between studied items and lures on a recognition test; by contrast, the more complex visual features that are represented in perirhinal cortex are more diagnostic (i.e., they are less likely to be shared across studied items and lures), thereby making them more useful for discriminating between studied items and lures (Cowell et al., 2006, this issue).

The main difference between the CLS framework and the representational-hierarchical framework relates to the degree of conjunctivity attributed to representations in perirhinal cortex. Cowell et al. (2006, this issue) posit that each layer in the cortical hierarchy learns conjunctions of features in the preceding layer, regardless of how useful (or useless) these conjunctive representations are for predicting meaningful outcomes. This process culminates in perirhinal neurons that represent conjunctions at the “whole object” level, such that different objects (i.e., objects that differ in at least one feature) are assigned non-overlapping representations in perirhinal cortex. The CLS view of conjunctive learning in cortex is somewhat different; this view is most fully described in a paper by O'Reilly and Rudy (2001) that incorporated error-driven learning into the cortical model (in addition to Hebbian learning). The key idea from O'Reilly and Rudy (2001) is that, while cortex has the ability to develop conjunctive representations, it does not learn these representations automatically – rather, it acquires conjunctive representations in a task-driven fashion (i.e., whenever the conjunctive representation is needed to reduce prediction error; see also O'Reilly and Norman, 2002). Because conjunctive representations are learned on an “as needed” basis, representations in higher-level cortical regions in the O'Reilly & Rudy (2001) model tend to be relatively less conjunctive and more overlapping than the representations posited by Cowell et al. (2006, this issue). We should emphasize that this difference is largely a matter of degree (both frameworks posit the existence of conjunctive representations in perirhinal

cortex). Nonetheless, we think that the difference is consequential: The CLS framework predicts that representations in perirhinal cortex should overlap to the extent that objects share features, whereas the framework posited by Cowell et al. does not. This “graded representational similarity” property is very important for the global match predictions described in Section 4 below.

One last issue relates to the difference between perirhinal and parahippocampal contributions to recognition memory. While this difference has not been explicitly modeled in our work thus far, we should note that the CLS framework is broadly compatible with the BIC framework described by Diana, Yonelinas, & Ranganath (2007) and Ranganath (this issue). We believe that the same basic computational operation (hierarchical feature extraction) is taking place in ventral stream structures (perirhinal cortex) and dorsal stream structures (parahippocampal cortex); the only difference is in the types of information that are represented. Ventral-stream structures represent object information, making this stream useful for judging the familiarity of individual objects; dorsal-stream structures represent other types of information (e.g., relating to spatial configuration) that can be used to contextualize object representations.

3.2. Hippocampal contributions to recognition memory

The hippocampal model, unlike the cortical model, can support recall (pattern completion) of stored patterns after a single study exposure. This allows the model to make recognition decisions based on the content of retrieved information, not just the amount of retrieved information.

Figure 2 shows the hippocampal model. To apply the hippocampal model to recognition, Norman and O'Reilly (2003) compared the test cue (presented on the EC in layer, which corresponds to superficial layers of entorhinal cortex) to the pattern of retrieved information (activated over the EC out layer, which corresponds to deep layers of entorhinal cortex). When recalled information matches the test cue, this constitutes evidence that the item was studied; conversely, mismatch between recalled information and the test cue constitutes evidence that the test cue was not studied (e.g., study “rats”, test “rat”; if the hippocampal model recalls that rats-plural was studied, not rat-singular, this can serve as grounds for rejection of rat). For every item, we computed a hippocampal “memory evidence” score by taking the number of recalled features that matched the test cue and subtracting out the number of recalled features that mismatched the test cue. Note that none of the simulation results described below depend on our use of this *match - mismatch* measure. For example, the properties described below are still present when we use the average activity of the hippocampal network (in region CA1 of the model) as our measure of memory evidence.

Importantly, while the CLS model provides a detailed account of the properties of the hippocampal and cortical memory signals, it does not directly address the question of how subjects *use* these signals to make recognition memory decisions. In the Norman & O'Reilly (2003) paper, recognition decision-making was treated as a “black box” that was external to the network model itself (i.e., the decision process was not itself simulated by a neural network). The simulations from that paper used a simple decision rule whereby test items were called “old” if the hippocampal signal exceeded a certain value, otherwise, the decision was made based on the cortical signal (Jacoby, Yonelinas, & Jennings, 1997). This decision rule is almost certainly oversimplified (e.g., we agree with Wixted, 2007 that subjects probably combine the signals in some way before making their decision, instead of treating the two signals separately). Developing a richer account of recognition memory decision rules and the neural mechanisms that implement these rules is an important direction for future research (see Norman, Detre, & Polyn, 2008, for further discussion of these issues).

4. Key results from our CLS recognition memory simulations

In this section, we review three critical differences in how (according to our model) hippocampus and cortex contribute to recognition memory. We then describe an important boundary condition on these differences.

4.1. Critical difference 1: Differences in how the cortical and hippocampal signals are distributed

The cortical memory signal is well described by a Gaussian signal detection model. The ratio of studied variance to lure variance depends on which sources of variability are present in the paradigm being simulated.

Figure 3 shows the distributions of memory scores associated with studied items and lures from a representative cortical-model simulation. Both of these distributions have an approximately Gaussian shape (note that this is an emergent property of the model rather than something that was built into the model in the first place). In Figure 3, the distributions are shown as having equal variance, but this is not always the case. As discussed in Norman and O'Reilly (2003), the ratio of the variances will depend on the particular sources of variability that are present in a particular simulation. Some kinds of variability, e.g., encoding variability (variability in how well items are stored during the study phase) affect studied items more than lures, and should therefore lead to greater variability in the studied-item memory strength distribution than in the lure memory strength distribution (Wixted, 2007; Hilford, Glanzer, Kim, & Decarlo, 2002). Other forms of variability, e.g., pre-experimental variability (variability in how much the features of a test item overlap with features of items studied prior to the experiment) should affect studied items and lures equally. If the ratio of pre-experimental variability to encoding variability is high (i.e., encoding variability is making a relatively small contribution to performance) then studied-item variance and lure variance will be approximately equal. Conversely, if encoding variability is more prominent than pre-experimental variability in a particular simulation, then variance in cortical memory strength (across items) will be larger for studied items than for lures.

The memory signal generated by the hippocampus is both continuous and non-Gaussian. The maximal level of confirmatory evidence generated by studied items typically exceeds the maximal level of evidence generated by lures. This property results in a positive Y intercept on ROC curves.

Figure 4 shows (for a representative simulation taken from Norman & O'Reilly, 2003) the distributions of hippocampal memory scores associated with studied items and lures. This particular simulation had a moderate level of feature overlap between input patterns (20 percent on average). As is clear from Figure 4, the hippocampal memory signal distributions do not adhere to a simple Gaussian model. The bulk of the lure distribution is located at the zero point, although some lures trigger above-zero memory signal. The studied distribution is bimodal, and crucially, it extends further to the right than the lure distribution, so there are some (high) memory scores that are sometimes triggered by studied items but never by lures.

The low level of lure memory strength shown in Figure 4 is a consequence of hippocampal pattern separation. In the hippocampal model, recall (pattern completion) occurs when the test item activates the hippocampal neurons that were involved in representing an item at study. Because of pattern separation, the hippocampal representations of lures and studied items typically do not overlap strongly; this low level of overlap prevents lures from accessing stored hippocampal memory traces and triggering recall. Because the hippocampal

memory strength distribution for studied items extends past the hippocampal memory strength distribution for lures, the model predicts that the hippocampal memory process will generate an ROC curve with a positive Y-intercept (that is, it should be possible to place a decision boundary so there are some hits but no false alarms).

These properties of the CLS hippocampal model are broadly consistent with the idea (promoted by Yonelinas and colleagues) that the hippocampus generates a *thresholded* memory signal that is sometimes triggered by studied items but never by lures (Yonelinas, 1994; Yonelinas et al., 1998). If we treat the upper limit of the lure distribution in Figure 4 as a “threshold”, it is true that studied items sometimes exceed the threshold but lures never do. Above-threshold memory strength values (in Figure 4, memory strength values $> .2$) are perfectly diagnostic as to whether the test item was studied (since these high values are only ever generated by studied items), whereas below-threshold memory strength values carry very little information about whether the test item was studied. To be clear, the idea that the CLS hippocampal memory signal is a threshold process (in the technical, signal-detection sense; Macmillan & Creelman, 2005) does *not* imply that the hippocampus generates a discrete, all-or-none signal, and it does *not* imply that qualitatively different processes are engaged on trials where memory exceeds (vs. does not exceed) the threshold. In the CLS model, studied items and lures both trigger continuously varying levels of evidence; it just happens to be the case that (because of pattern separation) the maximum level of evidence triggered by studied items exceeds the maximum level of evidence triggered by lures.

It is also useful to clarify how the CLS model relates to the memory model recently proposed by Shimamura and Wickens (2009). Like our model, the Shimamura and Wickens (2009) model posits that the hippocampus can generate a non-Gaussian, positively skewed memory signal distribution whereas regions lower in the MTL processing hierarchy generate less-skewed (more Gaussian) memory signal distributions. A crucial difference between the Shimamura and Wickens (2009) model and the CLS model is that the CLS model is a *mechanistic* model (i.e., it specifies a procedure for computing memory strength scores for individual simulated items), whereas the Shimamura model is a *measurement* model (i.e., it specifies a procedure for estimating model parameters from behavioral summary statistics, but it does not specify a procedure for computing memory strength scores for individual items). In the CLS model, the non-Gaussian shape of the hippocampal distribution is an emergent property of assumptions that we have made about how the hippocampal memory signal is computed; in the Shimamura model, the non-Gaussian shape of the hippocampal distribution is captured by their use of ex-Gaussian distributions, which contain a parameter for positive skew.

4.2. Critical difference 2: The hippocampus performs better than cortex at discriminating between studied items and related lures (except if the test gives a forced choice between studied items and corresponding related lures)

The CLS model predicts that cortex and hippocampus can both support good recognition performance when lures are not closely related to studied items. However, when lures are closely related to specific studied items (e.g., study “rats” and test with “rat”; Hintzman, Curran, & Oppy, 1992), hippocampally-based recognition performance should be higher than cortically-based recognition performance, because of the hippocampus’ ability to assign distinct representations to similar stimuli. As mentioned above, this property implies that most lures will not trigger a strong hippocampal signal. Furthermore, if a lure is similar enough to a studied item to trigger pattern completion of that item, the lure can be rejected based on mismatch between the test probe and the recalled information.

The CLS model also predicts that effects of target-lure similarity should interact with test format. Most recognition tests use a yes-no (YN) format where test items are presented one

at a time, and subjects are asked to label them as old or new. The model predicts that cortex should perform very poorly on YN tests with related lures (because the distributions of cortical memory scores associated with studied items and related lures overlap strongly). However, the model predicts that cortex should perform much better when given a forced choice between studied items and corresponding related lures (e.g., “rat” and “rats” are presented simultaneously, and subjects have to choose which item was studied). In this situation, the model predicts that the mean difference in cortical memory strength between the studied item and the related lure will be small, but the studied item should reliably trigger a slightly larger cortical memory signal than the corresponding related lure (thereby allowing for correct responding). This prediction comes about because there is extensive covariance in the cortical memory scores triggered by studied items and corresponding related lures; this high level of covariance makes the variance of the difference in these cortical memory scores very small (see Hintzman, 1988 for further discussion of this point about covariance).

Taken together, these predictions imply that patients with hippocampal damage should perform very poorly on YN tests with related lures. However, the same patients should show relatively spared performance on tests with unrelated lures, or when they are given a forced choice between targets and corresponding related lures (since cortex can pick up the slack in both cases). Holdstock et al. (2002) and Mayes et al. (2001) tested these predictions in a patient with focal hippocampal damage and obtained the predicted pattern of results; for additional evidence in support of these predictions, see also Westerberg et al. (2006). Contrary to this prediction, Bayley, Wixted, Hopkins, and Squire (2008) observed similar levels of impairment in hippocampally-lesioned patients on YN tests and forced-choice tests with corresponding related lures – however, it is possible (as is the case in all patient studies) that these patients had dysfunction in cortex that was not detected with standard MRI techniques.

To be clear, the model does not predict that all types of forced-choice responding will be relatively spared in patients with focal hippocampal damage. The prediction of relatively spared performance is related to a particular type of forced-choice test (i.e., tests that pair studied items with corresponding lures: study A, B; test A vs. A', B vs. B'). If we set up a forced-choice test that pairs studied items with lures that are related to other studied items (e.g., study A, B; test A vs. B') the model predicts that performance will be impaired after focal hippocampal damage. This prediction has yet to be tested in patients (see Migo, Montaldi, Norman, Quamme, & Mayes, 2009 for evidence from normal subjects that speaks to this prediction).

4.3. Critical difference 3: The cortical signal tracks global match. The hippocampal signal does not

Numerous behavioral studies have demonstrated that subjects' recognition judgments are sensitive to global match, where global match is defined as the summed similarity of the test probe to all of the studied items (Clark & Gronlund, 1996). Studies have found that the probability of calling a nonstudied test item “old” increases as a function of 1) the number of similar items that have been studied (e.g., Koutstaal, Schacter, & Jackson, 1999; Shiffrin, Huber, & Marinelli, 1995; Zaki & Nosofsky, 2001; Criss & Shiffrin, 2004) and 2) the degree of similarity between studied items and the test item, holding the number of studied items constant (e.g., Kahana & Sekuler, 2002; Yotsumoto et al., 2007).

The memory signal generated by the cortical network tracks, in a graded fashion, the global match between the test cue and all stored memory traces. This property arises because cortex assigns similar representations to similar stimuli — if the representation of one item is sharpened, this benefit accrues to other stimuli in proportion to their similarity to the initial

stimulus. Importantly, the hippocampal model does not exhibit this global match property. Because of hippocampal pattern separation, the test probe has to be very similar to a studied item before it triggers any recall. If the test probe is a studied item, adding other items to the study list that overlap moderately with that item will reliably boost the cortical memory signal, but the hippocampal signal will not increase, and (depending on how the model is parameterized) the hippocampal signal can even decrease due to interference between all of the similar memory traces. This is a manifestation of the well-known principle that increasing the number of associates to a given cue can hurt recall of these associates; this principle has been described elsewhere in terms of *cue overload effects* (Watkins & Watkins, 1975) and *fan effects* (Anderson & Reder, 1999; Reder et al., 2000).

4.4. Boundary conditions on the critical differences listed above

The distinctive properties of the hippocampal signal go away when the overlap between memory traces is extremely high; in this situation, the operating characteristics of the hippocampal signal come to resemble the operating characteristics of the cortical signal.

The three critical properties of the hippocampal model listed above (non-Gaussian memory strength distributions leading to a positive ROC Y-intercept; good discrimination between studied items and related lures; lack of sensitivity to global match) are straightforward consequences of the hippocampal model's ability to assign relatively distinct representations to overlapping inputs. Importantly, there are limits on this ability – if the inputs coming into the hippocampus are extremely similar to one another, the hippocampus will fail to assign distinct representations to these patterns. In this pattern separation failure regime, the properties of the hippocampal memory signal come to resemble the properties of the cortical memory signal – the memory strength distributions for studied items and lures start to become more Gaussian and overlapping, and (correspondingly) the ROC becomes more symmetrical and the Y-intercept approaches zero. This difference in operating characteristics can be seen by comparing Figure 5 (which shows the distribution of the hippocampal memory signal given 40.5% feature overlap between input patterns) and Figure 4 (which shows the distribution of the hippocampal memory signal given 20% feature overlap between input patterns). Elfman, Parks, and Yonelinas (2008) tested the model's predictions about effects of feature overlap using a source memory experiment where each item was presented either on the left or the right side of the screen at study. They manipulated the level of item feature overlap by using photos of houses as stimuli (in the “high feature overlap” condition) or else using randomly selected photos as stimuli (in the “low feature overlap” condition). At test, subjects were shown photos and asked to say whether they were presented on the left or right side of the screen at study. Elfman et al. (2008) plotted normalized source memory ROC curves (under conditions of high vs. low item feature overlap) and measured the shape of the z-ROC – if the underlying distributions are Gaussian, then the z-ROC will be linear; conversely, curvature in the z-ROC implies that the underlying distributions are not Gaussian (but see Malmberg & Xu, 2006 for some cautions regarding how to interpret z-ROC curvature). In keeping with the CLS model's predictions, Elfman et al. (2008) found that the z-ROC was significantly U-shaped (indicating non-Gaussian distributions) in the low feature overlap condition, and that the z-ROC was significantly less U-shaped in the high feature overlap condition, suggesting that the underlying distributions in this condition were closer to being Gaussian.

5. Implications for the current debate

5.1. Characterizing hippocampal vs. cortical contributions

Recollection vs. familiarity—The question of how well recollection/familiarity map onto the hippocampus/cortex distinction depends critically on how we define these terms. If

we define familiarity very broadly as “a scalar signal that can be used to support recognition memory judgments”, then our simulations suggest that hippocampus can support familiarity – the average level of activity in region CA1 of the hippocampal model discriminates between old and new items, and this signal could conceivably be used to support recognition memory judgments. However, within the memory modeling community, practically all researchers have adopted a more specific definition of familiarity centered around the idea of summed similarity or global match (e.g., Gillund & Shiffrin, 1984; Hintzman, 1988; Nosofsky, 1988; Humphreys, Bain, & Pike, 1989; Murdock, 1993; Shiffrin & Steyvers, 1997; for reviews see Clark & Gronlund, 1996; Malmberg, 2008; Norman, Detre, & Polyn, 2008).

Our simulations show that the cortical model can support global match judgments but the hippocampal model can not support these judgments, except under conditions where stimuli are extremely similar to one another. Some of the behavioral studies that have demonstrated sensitivity to global match have used extremely similar stimuli (e.g., Kahana & Sekuler, 2002 used grayscale textures generating by superimposing horizontal and vertical gratings, and Yotsumoto et al., 2007 used highly similar synthetically-generated faces), but global match effects can be observed even when stimuli are not highly similar to one another. For example, Shiffrin et al. (1995) purposefully disguised their main manipulation (of the number of items in a semantic category) by using loosely-defined categories and intermixing items from different categories in the study list; subjects did not notice the category structure of the list but they still showed a robust effect of category length on false alarms to items from that category. If we follow the lead of the memory modeling community and treat “sensitivity to global match” as an essential property of familiarity, our CLS model simulations suggest that cortex can support familiarity but hippocampus can not support familiarity (except in the aforementioned case where the stimuli being used in the experiment are all extremely similar to one another).

Moving beyond recollection vs. familiarity—It is important to not let definitional questions about recollection and familiarity distract us from the fundamental question at hand: coming up with a precise, testable description of the essential differences between the two systems’ contributions to recognition memory. Generating these precise computational descriptions was the central goal of the Norman and O’Reilly (2003) *Psychological Review* paper. However, in retrospect, it appears that we muddied the waters by continuing to use the terms “familiarity” and “recollection” (with all of their baggage) to describe cortical and hippocampal contributions. For many readers, the take-home message of our paper was that it was basically OK to keep using the existing recollection/familiarity framework. Now that we are re-visiting these issues, we would like to try to refocus discussion on the most fundamental computational idea to come out of our simulations, namely: *Hippocampal pattern separation imbues the hippocampal signal with different operating characteristics than the cortical signal*. Pattern separation makes the hippocampus especially well-suited for discriminating between studied items and related lures; it makes the hippocampus especially poorly suited for computing global match; and it imbues the hippocampal ROC curve with a Y-intercept > 0.

The basic idea that hippocampus does pattern separation is widely accepted in the field, and the past several years have seen the publication of some extremely clever and powerful experimental demonstrations of hippocampal pattern separation (e.g., Bakker, Kirwan, Miller, & Stark, 2008; Leutgeb, Leutgeb, Moser, & Moser, 2007; Clelland et al., 2009). However, we think that the implications of pattern separation differences (in hippocampus vs. cortex) for recognition performance have not been fully explored. For example, while there have been fMRI studies demonstrating that hippocampus does an especially good job of separating out studied items from similar lures (Bakker et al., 2008), there have not (to

our knowledge) been any imaging studies that test the model's predictions about the respective contributions of cortex vs. hippocampus to global matching effects. These predictions could easily be tested using fMRI or intracranial EEG. Specifically, one could run experiments that parametrically manipulate the number of studied items related to each test probe (as per Shiffrin et al., 1995) and also the similarity of these items to the test probe (as per Kahana & Sekuler, 2002 and Yotsumoto et al., 2007). The key prediction is that, in every situation where subjects show a behavioral global match effect, perirhinal cortex activity should monotonically track global match. By contrast, hippocampal activity should only track global match in situations where all of the items are very similar to one another. If we can establish that cortex is more useful for global match computations than the hippocampus, this would be a huge step forward in fractionating the contributions of MTL substructures to recognition memory. Whereas the claim that "hippocampus does not implement familiarity" is subject to interpretive ambiguity relating to the meaning of "familiarity", the claim that "hippocampus can not compute summed similarity (unless the inputs show extensive overlap)" is much more clear.

5.2. Implications for measurement models

The CLS model simulations described above can help us to assess the validity of techniques that are commonly used to measure the contributions of hippocampus and cortex to recognition memory. Here, we will focus on the dual-process signal-detection (DPSD) model set forth by Yonelinas et al. (Yonelinas, 1994; Yonelinas et al., 1998; Yonelinas, 2001; Yonelinas et al., 2002), since this is the most widely-used measurement model in the literature. The most important assumptions underlying the DPSD measurement model are 1) that familiarity is an equal-variance Gaussian signal detection process (yielding a symmetric ROC with a zero y-intercept) and 2) recollection is a thresholded process (in the technical sense that the studied-item evidence distribution extends to the right of the lure evidence distribution, yielding an asymmetric ROC with a positive Y-intercept). Given these assumptions, any asymmetry in the ROC curve can be attributed to the recollection process. If we make the further assumption that recollection maps onto the hippocampus (Yonelinas, 2002), we can use asymmetry in the ROC to index the degree to which the hippocampus is contributing to recognition performance.

According to the CLS model, the assumptions about familiarity and recollection outlined above will hold in some circumstances but not others. In our simulations, the cortical signal is well-described by an equal-variance Gaussian signal-detection model unless contribution of encoding variability (which affects targets more than lures) is large relative to the contribution of other forms of variability (e.g., variability in pre-experimental familiarity with stimuli) that affect targets and lures equally. Also, in our simulations, the hippocampal signal yields a positive ROC Y-intercept unless there is an extremely high level of feature overlap between items. If encoding variability is controlled, e.g., using an encoding task, and between-item similarity is not too high, then estimates of recollection and familiarity obtained using the DPSD model procedure should approximately reflect the contributions of the hippocampal and cortical processes (respectively). More specifically: The DPSD estimate of hippocampal recollection corresponds to the height of the (estimated) ROC Y-intercept, which indicates the proportion of the studied item distribution that lies to the right of the maximum of the lure distribution. The hippocampal distributions shown in Figure 4 are too complex to be fully described with this single number – knowing the proportion of the studied-item distribution that lies above a particular memory strength value does not tell us how the distribution is shaped below that point. Nonetheless, the Y-intercept is a useful proxy for the overall strength of the hippocampal signal in the model.

Our CLS simulations suggest that it should not be difficult to engineer situations where the assumptions of DPSD are invalid. For example, strong encoding variability will lead to

violations of the equal-variance assumption for the cortical process (strong encoding variability will also lead to violations of the assumption that the two signals are stochastically independent; see and Norman & O'Reilly, 2003 for further discussion of this point). It seems clear to us that blindly using models like DPSD without regard to boundary conditions can lead to trouble. However, while the assumptions of DPSD will not always be met, it seems likely to us that many – if not most – of the studies that have used DPSD fit within the boundary conditions outlined above (low encoding variability, not too much overlap between items in the learning set). For example, in the studies that have applied ROC analysis to data from hippocampally lesioned animals (Fortin, Wright, & Eichenbaum, 2004; Sauvage, Fortin, Owens, Yonelinas, & Eichenbaum, 2008), the rats had to sniff each odor at training (thereby reducing encoding variability) and the odors were very distinctive (thereby ensuring that there would not be too much overlap between the cortical representations of these odors).

Speaking generally, a measurement model does not have to be exactly correct to be a useful source of converging evidence, and it does not have to be applicable in every situation. For these reasons, it seems like we would be depriving ourselves of a useful tool if we banished measurement models like DPSD. We just need to be aware of the boundary conditions, and mechanistic models like CLS give us a way of developing intuitions about what these boundary conditions will be.

6. Summary

The high-level goal of this paper has been to reassert the importance of mechanistic models in the debate over how hippocampus vs. cortex contribute to recognition memory. It is worth noting that practically all of the recent debates over the neural basis of recognition memory have focused on the relative validity of different measurement models – the dual-process signal-detection model vs. the unequal-variance signal-detection model described by Wixted (2007) and others vs. the Shimamura and Wickens (2009) model vs. countless other models (for a review, see Yonelinas & Parks, 2007). Mechanistic models like CLS have played a relatively small role in the debate thus far. In this paper, we have described two related areas where mechanistic models can contribute: First, mechanistic models can be used to proactively generate predictions about how particular manipulations will affect hippocampal and cortical memory processes. Second, mechanistic models can be used to evaluate the assumptions of measurement models.

The specific goal of the paper has been to reassert the claim from Norman and O'Reilly (2003) that pattern separation differences are the key to understanding differences in hippocampal vs. cortical differences. We listed several implications of these pattern separation differences for recognition performance, including the idea that cortex is better-suited for computing global match. We believe that fully exploring these implications, using models like CLS as a guide, will help us to arrive at a richer understanding of how the contributions of hippocampus vs. cortex to recognition memory differ.

Acknowledgments

This work was supported by NIH R01 MH069456 awarded to KAN.

References

- Aggleton JP, Brown MW. Episodic memory, amnesia, and the hippocampal-anterior thalamic axis. *Behavioral and Brain Sciences*. 1999; 22:425–490. [PubMed: 11301518]
- Anderson JR, Reder LM. The fan effect: New results and new theories. *Journal of Experimental Psychology: General*. 1999; 128:186.

- Bakker A, Kirwan CB, Miller M, Stark CEL. Pattern separation in the human hippocampal CA3 and dentate gyrus. *Science (New York, NY)*. 2008; 319(5870):1640–2.
- Barense MD, Bussey TJ, Lee AC, Rogers TT, Davies RR, Saksida LM, et al. Functional specialization in the human medial temporal lobe. *Journal of Neuroscience*. 2005; 25(44):10239–46. [PubMed: 16267231]
- Bayley PJ, Wixted JT, Hopkins RO, Squire LR. Yes/no recognition, forced-choice recognition, and the human hippocampus. *Journal of Cognitive Neuroscience*. 2008; 20:505–512. [PubMed: 18004951]
- Becker S. A computational principle for hippocampal learning and neurogenesis. *Hippocampus*. 2005; 15(6):722–38. [PubMed: 15986407]
- Bogacz R, Brown MW. Comparison of computational models of familiarity discrimination in the perirhinal cortex. *Hippocampus*. 2003; 13:494–524. [PubMed: 12836918]
- Burgess N, O'Keefe J. Neuronal computations underlying the firing of place cells and their role in navigation. *Hippocampus*. 1996; 6:749–762. [PubMed: 9034860]
- Bussey TJ, Saksida LM. The organisation of visual object representations: A connectionist model of effects of lesions in perirhinal cortex. *European Journal of Neuroscience*. 2002; 15:355–364. [PubMed: 11849301]
- Bussey, TJ.; Saksida, LM.; Murray, EA. The role of perirhinal cortex in memory and perception: Conjunctive representations for object identification. In: Witter, MP.; Waterlood, FG., editors. *The Parahippocampal Region: Organisation and Role in Cognitive Functions*. New York: Oxford; 2002.
- Clark SE, Gronlund SD. Global matching models of recognition memory: How the models match the data. *Psychonomic Bulletin and Review*. 1996; 3:37–60.
- Clelland CD, Choi M, Romberg C, Clemenson GD, Fragniere A, Tyers P, et al. A functional role for adult hippocampal neurogenesis in spatial pattern separation. *Science (New York, NY)*. 2009; 325(5937):210–3.
- Cowell RA, Bussey TJ, Saksida LM. Why does brain damage impair memory? A connectionist model of object recognition memory in perirhinal cortex. *Journal of Neuroscience*. 2006; 26:12186–12197. [PubMed: 17122043]
- Cowell RA, Bussey TJ, Saksida LM. Functional dissociations within the ventral object processing pathway: Cognitive modules or a hierarchical continuum? *Journal of Cognitive Neuroscience*. 2009
- Cowell RA, Bussey TJ, Saksida LM. Components of recognition memory. dissociable cognitive processes or just differences in representational complexity? *Hippocampus*. this issue.
- Criss AH, Shiffrin RM. Context noise and item noise jointly determine recognition memory: A comment on dennis and humphreys (2001). *Psychological Review*. 2004; 111:800–807. [PubMed: 15250785]
- Diana RA, Yonelinas AP, Ranganath C. Imaging recollection and familiarity in the medial temporal lobe: a three-component model. *Trends in Cognitive Sciences*. 2007; 11:379–86. [PubMed: 17707683]
- Eichenbaum H, Otto T, Cohen NJ. Two functional components of the hippocampal memory system. *Behavioral and Brain Sciences*. 1994; 17(3):449–518.
- Elfman KW, Parks CM, Yonelinas AP. Testing a neurocomputational model of recollection, familiarity, and source recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2008; 34(4):752–68.
- Fortin NJ, Wright SP, Eichenbaum HB. Recollection-like memory retrieval in rats is dependent on the hippocampus. *Nature*. 2004; 431:188–191. [PubMed: 15356631]
- Gardiner JM, Brandt KR, Baddeley AD, Vargha-Khadem F, Mishkin M. Charting the acquisition of semantic knowledge in a case of developmental amnesia. *Neuropsychologia*. 2008; 46:2865–8. [PubMed: 18589461]
- Gillund G, Shiffrin RM. A retrieval model for both recognition and recall. *Psychological Review*. 1984; 91:1–67. [PubMed: 6571421]
- Grossberg S. Adaptive pattern classification and universal recoding I: Parallel development and coding of neural feature detectors. *Biological Cybernetics*. 1976; 23:121–134. [PubMed: 974165]

- Grossberg, S. The adaptive self-organization of serial order in behavior: Speech, language, and motor control. In: Swab, EC.; Nusbaum, HC., editors. Pattern recognition in humans and machines, volume I: Speech perception. New York: Academic Press; 1986.
- Grossberg S, Stone G. Neural dynamics of word recognition and recall: Attentional priming, learning, and resonance. *Psychological Review*. 1986; 93:46–74. [PubMed: 3961051]
- Hasselmo ME, Wyble B. Free recall and recognition in a network model of the hippocampus: Simulating effects of scopolamine on human memory function. *Behavioural Brain Research*. 1997; 89:1–34. [PubMed: 9475612]
- Hilford A, Glanzer M, Kim K, Decarlo LT. Regularities of Source Recognition : ROC Analysis. *Journal of Experimental Psychology: General*. 2002; 131(4):494–510. [PubMed: 12500860]
- Hintzman DL. Judgments of frequency and recognition memory in a multiple-trace memory model. *Psychological Review*. 1988; 95:528–551.
- Hintzman DL, Curran T, Oppy B. Effects of similarity and repetition on memory: Registration without learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1992; 18:667–680.
- Holdstock JS, Mayes AR, Roberts N, Cezayirli E, Isaac CL, O'Reilly RC, et al. Under what conditions is recognition spared relative to recall after selective hippocampal damage in humans? *Hippocampus*. 2002; 12:341–351. [PubMed: 12099485]
- Humphreys MS, Bain JD, Pike R. Different ways to cue a coherent memory system: A theory for episodic, semantic, and procedural tasks. *Psychological Review*. 1989; 96:208–233.
- Jacoby, LL.; Yonelinas, AP.; Jennings, JM. The relation between conscious and unconscious (automatic) influences: A declaration of independence. In: Cohen, JD.; Schooler, JW., editors. *Scientific Approaches to Consciousness*. Mahwah, NJ: Erlbaum; 1997. p. 13-47.
- Kahana MJ, Sekuler R. Recognizing spatial patterns: A noisy exemplar approach. *Vision Research*. 2002; 42:2177–92. [PubMed: 12207978]
- Koutstaal W, Schacter DL, Jackson EM. Perceptually based false recognition of novel objects in amnesia: Effects of category size and similarity to category prototypes. *Cognitive Neuropsychology*. 1999; 16:317.
- Leutgeb JK, Leutgeb S, Moser M-B, Moser EI. Pattern separation in the dentate gyrus and CA3 of the hippocampus. *Science (New York, NY)*. 2007; 315(5814):961–6.
- Li L, Miller EK, Desimone R. The representation of stimulus familiarity in anterior inferior temporal cortex. *Journal of Neurophysiology*. 1993; 69:1918–1929. [PubMed: 8350131]
- Macmillan, NA.; Creelman, CD. *Detection Theory: A User's Guide*. 2. Mahwah, NJ: Erlbaum; 2005.
- Malmberg KJ. Recognition memory: a review of the critical findings and an integrated theory for relating them. *Cognitive Psychology*. 2008; 57(4):335–84. [PubMed: 18485339]
- Malmberg KJ, Xu J. The influence of averaging and noisy decision strategies on the recognition memory ROC. *Psychonomic Bulletin & Review*. 2006 Mar; 13(1):99–105. [PubMed: 16724775]
- Marr D. Simple memory: A theory for archicortex. *Philosophical Transactions of the Royal Society (London) B*. 1971; 262:23–81.
- Mayes AR, Isaac CL, Downes JJ, Holdstock JS, Hunkin NM, Montaldi D, et al. Memory for single items, word pairs, and temporal order in a patient with selective hippocampal lesions. *Cognitive Neuropsychology*. 2001; 18:97–123. [PubMed: 20945208]
- McClelland JL, McNaughton BL, O'Reilly RC. Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*. 1995; 102:419–457. [PubMed: 7624455]
- McNaughton BL, Morris RGM. Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends in Neurosciences*. 1987; 10(10):408–415.
- Migo E, Montaldi D, Norman KA, Quamme J, Mayes A. The contribution of familiarity to recognition memory is a function of test format when using similar foils. *Quarterly Journal of Experimental Psychology (2006)*. 2009; 62(6):1198–215. [PubMed: 19096990]
- Moll M, Mikkilainen R. Convergence-zone episodic memory: Analysis and simulations. *Neural Networks*. 1997; 10:1017–1036. [PubMed: 12662496]

- Murdock BB. TODAM2: A model for the storage and retrieval of item, associative, and serial-order information. *Psychological Review*. 1993; 100:183–203. [PubMed: 8483981]
- Norman, KA.; Detre, GJ.; Polyn, SM. Computational models of episodic memory. In: Sun, R., editor. *The Cambridge Handbook of Computational Psychology*. New York: Cambridge University Press; 2008.
- Norman KA, Newman EL, Detre GJ, Polyn SM. How inhibitory oscillations can train neural networks and punish competitors. *Neural Computation*. 2006; 18:1577–610. [PubMed: 16764515]
- Norman KA, Newman EL, Perotte AJ. Methods for reducing interference in the complementary learning systems model: Oscillating inhibition and autonomous memory rehearsal. *Neural Networks*. 2005; 18:1212–1228. [PubMed: 16260116]
- Norman KA, O'Reilly RC. Modeling hippocampal and neocortical contributions to recognition memory: A complementary-learning-systems approach. *Psychological Review*. 2003; 104:611–646. [PubMed: 14599236]
- Nosofsky R. Exemplar-based accounts of relations between classification, recognition, and typicality. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1988; 14(4):700708.
- O'Keefe, J.; Nadel, L. *The Hippocampus as a Cognitive Map*. Oxford, England: Oxford University Press; 1978.
- O'Reilly RC, Rudy JW. Conjunctive representations in learning and memory: Principles of cortical and hippocampal function. *Psychological Review*. 2001; 108:311–345. [PubMed: 11381832]
- O'Reilly RC, Norman KA. Hippocampal and neocortical contributions to memory: Advances in the Complementary Learning Systems framework. *Trends in Cognitive Sciences*. 2002; 6:505–510. [PubMed: 12475710]
- Ranganath C. Hippocampus. this issue.
- Reder LM, Nhouyvanisvong A, Schunn CD, Ayers MS, Angstadt P, Hiraki KA. A mechanistic account of the mirror effect for word frequency: A computational model of remember-know judgments in a continuous recognition paradigm. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2000; 26:294–320.
- Rolls, ET. Functions of neuronal networks in the hippocampus and neocortex in memory. In: Byrne, JH.; Berry, WO., editors. *Neural Models of Plasticity: Experimental and Theoretical Approaches*. San Diego, CA: Academic Press; 1989. p. 240-265.
- Sauvage MM, Fortin NJ, Owens CB, Yonelinas AP, Eichenbaum H. Recognition memory: opposite effects of hippocampal damage on recollection and familiarity. *Nature Neuroscience*. 2008; 11(1): 16–8.
- Scoville WB, Milner B. Loss of recent memory after bilateral hippocampal lesions. *Journal of Neurology, Neurosurgery, and Psychiatry*. 1957; 20:11–21.
- Sherry DF, Schacter DL. The evolution of multiple memory systems. *Psychological Review*. 1987; 94(4):439–454.
- Shiffrin RM, Huber DE, Marinelli K. Effects of category length and strength on familiarity in recognition. *Journal of Experimental Psychology: Learning, Memory and Cognition*. 1995; 21:267–287.
- Shiffrin RM, Steyvers M. A model for recognition memory: REM – retrieving effectively from memory. *Psychonomic Bulletin and Review*. 1997; 4:145–166. [PubMed: 21331823]
- Shimamura AP, Wickens TD. Superadditive memory strength for item and source recognition: the role of hierarchical relational binding in the medial temporal lobe. *Psychological Review*. 2009; 116(1):1–19. [PubMed: 19159146]
- Squire LR. Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psychological Review*. 1992; 99:195–231. [PubMed: 1594723]
- Squire LR, Wixted JT, Clark RE. Recognition memory and the medial temporal lobe: a new perspective. *Nature reviews Neuroscience*. 2007; 8(11):872–83.
- Sutherland RJ, Rudy JW. Configural association theory: The role of the hippocampal formation in learning, memory, and amnesia. *Psychobiology*. 1989; 17(2):129–144.
- Teyler TJ, Discenna P. The hippocampal memory indexing theory. *Behavioral Neuroscience*. 1986; 100:147–154. [PubMed: 3008780]

- Treves A, Rolls ET. A computational analysis of the role of the hippocampus in memory. *Hippocampus*. 1994; 4:374–392. [PubMed: 7842058]
- Viskontas IV, Knowlton BJ, Steinmetz PN, Fried I. Differences in mnemonic processing by neurons in the human hippocampus and parahippocampal regions. *Journal of Cognitive Neuroscience*. 2006; 18(10):1654–62. [PubMed: 17014370]
- Watkins OC, Watkins MJ. Buildup of proactive inhibition as a cue-overload effect. *Human Learning*. 1975; 104(4):442–452.
- Westerberg CE, Paller KA, Weintraub S, Mesulam MM, Holdstock J, Mayes A, et al. When memory does not fail: Familiarity-based recognition in mild cognitive impairment and Alzheimer's disease. *Neuropsychology*. 2006; 20:193–205. [PubMed: 16594780]
- Wixted JT. Dual-process theory and signal-detection theory of recognition memory. *Psychological Review*. 2007; 114(1):152–76. [PubMed: 17227185]
- Wixted JT, Squire LR. Constructing receiver operating characteristics (ROCs) with experimental animals: cautionary notes. *Learning & Memory*. 2008; 15(9):687–90. [PubMed: 18772256]
- Wu X, Baxter RA, Levy WB. Context codes and the effect of noisy learning on a simplified hippocampal CA3 model. *Biological Cybernetics*. 1996; 74:159–165. [PubMed: 8634367]
- Xiang JZ, Brown MW. Differential encoding of novelty, familiarity, and recency in regions of the anterior temporal lobe. *Neuropharmacology*. 1998; 37:657–676. [PubMed: 9705004]
- Yonelinas AP. Receiver-operating characteristics in recognition memory: Evidence for a dual-process model. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1994; 20:1341–1354.
- Yonelinas AP. Consciousness, control, and confidence: The 3 Cs of recognition memory. *Journal of Experimental Psychology: General*. 2001; 130:361–379. [PubMed: 11561915]
- Yonelinas AP. The nature of recollection and familiarity: A review of 30 years of research. *Journal of Memory and Language*. 2002; 46:441–517.
- Yonelinas AP, Kroll NE, Quamme JR, Lazzara MM, Sauve MJ, Widaman KF, et al. Effects of extensive temporal lobe damage or mild hypoxia on recollection and familiarity. *Nature Neuroscience*. 2002; 5(11):1236–1241.
- Yonelinas AP, Kroll NEA, Dobbins I, Lazzara M, Knight RT. Recollection and familiarity deficits in amnesia: Convergence of remember-know, process dissociation, and receiver operating characteristic data. *Neuropsychology*. 1998; 12:323–339. [PubMed: 9673991]
- Yonelinas AP, Parks CM. Receiver operating characteristics (ROCs) in recognition memory: a review. *Psychological Bulletin*. 2007; 133(5):800–32. [PubMed: 17723031]
- Yotsumoto Y, Kahana MJ, Wilson HR, Sekuler R. Recognition memory for realistic synthetic faces. *Memory & Cognition*. 2007 Sep; 35(6):1233–44.
- Zaki S, Nosofsky R. Exemplar accounts of blending and distinctiveness effects in perceptual old-new recognition. *Journal of Experimental Psychology: Learning, Memory and Cognition*. 2001; 27(4):1022–1041.

Figure 1.

Illustration of the sharpening of hidden (perirhinal) layer activity patterns in a miniature version of the CLS cortical model. (a) shows the network prior to sharpening; perirhinal activity (more active = lighter color) is relatively undifferentiated. (b) shows the network after Hebbian learning and inhibitory competition produce sharpening; a subset of the units are strongly active, while the remainder are inhibited. Figure adapted from Norman and O'Reilly (2003).

Figure 2.

Diagram of the CLS hippocampal network. The hippocampal network links input patterns in entorhinal cortex (EC) to relatively non-overlapping (pattern-separated) sets of units in region CA3. The dentate gyrus (DG) serves to facilitate pattern separation in region CA3. Recurrent connections in CA3 bind together all of the units involved in representing a particular EC pattern; the CA3 representation is linked back to EC via region CA1. Learning in the CA3 recurrent connections, and in projections linking EC to CA3 and CA3 to CA1, makes it possible to recall entire stored EC patterns based on partial cues. Figure reprinted from Norman and O'Reilly (2003).

Figure 3.

Histogram of the memory signal generated by the CLS cortical network, given 20% average overlap between input patterns. Figure adapted from Norman and O'Reilly (2003).

Figure 4. Histogram of the memory signal generated by the CLS hippocampal network, given 20% average overlap between input patterns. Figure adapted from Norman and O'Reilly (2003).

Figure 5. Histogram of the memory signal generated by the CLS hippocampal network, given 40.5% average overlap between input patterns. Figure adapted from Norman and O'Reilly (2003).