Hippocampal and neocortical contributions to memory: advances in the complementary learning systems framework

Randall C. O'Reilly and Kenneth A. Norman

The complementary learning systems framework provides a simple set of principles, derived from converging biological, psychological and computational constraints, for understanding the differential contributions of the neocortex and hippocampus to learning and memory. The central principles are that the neocortex has a low learning rate and uses overlapping distributed representations to extract the general statistical structure of the environment, whereas the hippocampus learns rapidly using separated representations to encode the details of specific events while minimizing interference. In recent years, we have instantiated these principles in working computational models, and have used these models to address human and animal learning and memory findings, across a wide range of domains and paradigms. Here, we review a few representative applications of our models, focusing on two domains: recognition memory and animal learning in the fear-conditioning paradigm. In both domains, the models have generated novel predictions that have been tested and confirmed.

Randall C. O'Reilly
Kenneth A. Norman
Dept of Psychology,
University of Colorado
Boulder, 345 UCB,
Boulder, CO 80309, USA.
e-mail: oreilly@
psych.colorado.edu

How does the brain learn and memorize? Although we are very far from having a single theory that accounts for all learning and memory data, this article describes one parsimonious theory of how the neocortex and hippocampus contribute to learning and memory, implemented in biologically-based neural-network computer simulations, which has been used to address animal and human data in paradigms ranging from episodic recognition memory to discrimination learning to fear conditioning. Furthermore, this theory has generated several distinctive predictions that have been confirmed by subsequent experiments. Based on this initial success, we can optimistically and provocatively claim that this theory represents an important first step towards a comprehensive theory of the neural basis of memory.

Far from being a radical new theory of the biological basis of memory, our theory, which we call the Complementary Learning Systems (CLS) model [1–3], incorporates several widely-held ideas about hippocampal and neocortical contributions to memory, that have been developed over many years by many different researchers (e.g. [4–21]). Thus, we think of it as a 'consensual' model, with our primary contributions being: (1) to emphasize certain computational principles in describing the division of labor between hippocampus and neocortex; and (2) to implement these principles in working computational models, which can be used to explore whether the principles are sufficient to account for specific findings, and to generate novel predictions.
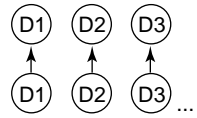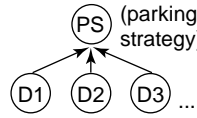
The complementary learning systems framework

The CLS framework is based on the logic of tradeoffs between mutually incompatible computational goals. The central tradeoff behind our framework involves two basic types of learning that an organism must engage in – learning about specifics versus generalities – which require conflicting neural architectures as summarized in Fig. 1.

Thus, each of these types of learning can be more optimally achieved by using two separate, but nonetheless highly interactive, memory systems with different specializations [3,24,25,2,1,10]): The hippocampus, which is specialized for rapidly and automatically encoding arbitrary conjunctions of existing cortical representations, and the neocortex, which is specialized for slowly developing representations of the general statistical structure of the environment. The hippocampus assigns distinct representations to input patterns to avoid interference across memories, whereas neocortex uses overlapping representations that encode shared structure across many different experiences. Furthermore, neocortex uses a small learning rate to gradually integrate new information with existing knowledge, whereas hippocampus uses a large learning rate to encode episodic memories of specific events as they unfold. This framework contrasts for example with the framework of [26] which attempts to accommodate stable category learning and rapid memorization within a single system.

These computational principles outlined here converge nicely with analyses of the physiological properties of the hippocampus and neocortex [2,3,27]. As just one example, the notion that the hippocampus should use distinct representations is supported the fact that hippocampal areas have much sparser levels of firing than those in the neocortex (e.g [28,29]); mathematical analyses (e.g. [5,27]) show that making activity more sparse results in reduced overlap between representations ('pattern separation').

We have constructed neural-network models of the hippocampus and cortex (and their often complex interactions) that instantiate the complementary learning systems principles outlined above, and incorporate key aspects of the anatomy and physiology of these structures [24,2,32,1]

Two incompatible goals	Remember specifics	Extract generalities
Example: Need to:	*Where is car parked?* Avoid interference	*Best parking strategy?* Accumulate experience
Solution: (1)	Separate representations (keep days separate) D1 D2 D3 ↑ ↑ ↑ D1 D2 D3 ...	Overlapping representations (integrate over days) PS (parking strategy) D1 D2 D3 ...
(2)	Fast learning (encode immediately)	Slow learning (integrate over days)
(3)	Learn automatically (encode everything)	Task driven learning (extract relevant stuff)
System:	Hippocampus	Neocortex

TRENDS in Cognitive Sciences

Fig. 1. Computational motivation for two complementary learning and memory systems in the brain: there are two goals that such systems need to solve. One goal is to remember specific information, in this example where one's car is parked on a specific day (D1 = day 1, D2 = day 2, etc). The other goal is to extract generalities across many experiences, for example in developing the best parking strategy over several different days. The neural solutions to these goals are incompatible: Memorizing specifics requires separate representations that are learned quickly, and automatically, whereas extracting generalities requires overlapping representations and slow learning (to integrate over experiences) and is driven by task-specific constraints. Thus, it makes sense to have two separate neural systems separately optimized for each of these goals.

(different models implement different parts of cortex, but all use the same general parameters and learning mechanisms, which are held to be the same across most cortical areas, with possible exceptions for highly specialized areas such as primary sensory or motor cortex). All of these models are very similar and can be considered one unified model. In one line of research, we have demonstrated that the model can account for a wide range of findings in human episodic memory paradigms including cued recall and recognition [1]. In another line of research, we have applied our models to a similarly large number of animal learning paradigms including nonlinear discrimination, habituation, fear conditioning, and transitive inference [2,22,23]. Some representative examples of these applications are presented below.

A biologically-based dual-process recognition memory model
As applied to recognition memory, the CLS model belongs to a long tradition of 'dual-process' theories of recognition. The defining feature of dual-process theories is the idea that one can recognize a previously studied item as such based on (1) *recall* of specific details from the study episode, or (2) a non-specific sense of *familiarity* that tracks the global match between the test item and stored memory traces (for a review, see [21]).

We show that the hippocampal model can support recall of specific studied details, and that, although cortex cannot support recall of once-presented
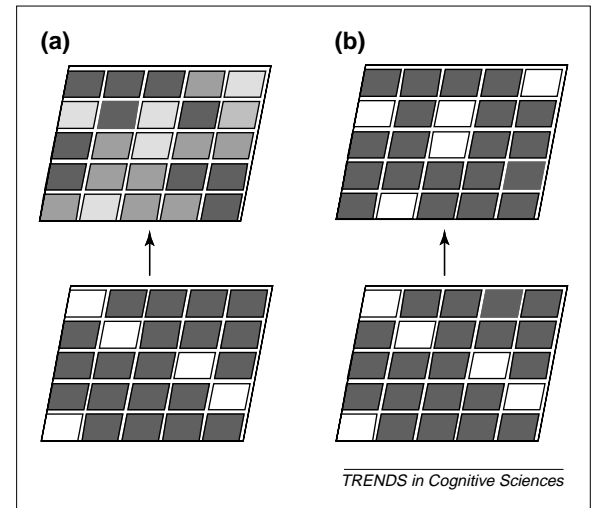

Fig. 2. Illustration of the sharpening of activation patterns over the hidden layer in our cortical model, which we use as measure of familiarity. (a) the hidden layer responding to a given input pattern before sharpening, where the activations (more active = lighter color) are relatively undifferentiated. (b) the hidden layer after Hebbian learning and inhibitory competition produce sharpening, where a subset of the units are strongly active and the remainder are inhibited.

stimuli, it is possible to extract a scalar familiarity signal from a model of medial temporal lobe cortex (MTLC) [1]. In this respect, our model resembles other dual-process theories that have linked recall to the hippocampus and familiarity to MTLC (e.g. [20,14,12]). However, this resemblance is only superficial. Whereas the other theories use simple, verbally stated dichotomies to characterize the difference between hippocampal and cortical contributions (e.g. associative versus item memory), the CLS model grounds its conception of hippocampal and cortical contributions in terms of graded architectural differences between these structures (e.g. representations are more sparse in the hippocampus). As a result, the CLS model's predictions frequently cut across the boundaries proposed by other researchers (e.g. Aggleton and Brown argue that cortex can support item recognition on its own, but – as discussed below – the CLS model predicts that item recognition will depend on the hippocampus when non-studied lure items are similar to studied items).

Given our earlier statement that cortex learns in small increments, how can it support familiarity judgments after a single exposure to a stimulus? We show that familiarity signals arise in the cortical model via a 'sharpening' dynamic (Fig. 2), where the contrast between the more active and the less active neurons is increased over presentations of a given input pattern (see also [33]). On the first presentation of an item, the hidden neurons are not very well tuned to the input pattern, producing a broad distribution of relatively weak activations across the hidden layer. However, the most active units will tune their weights the most (via Hebbian learning), and will therefore become more activated upon subsequent
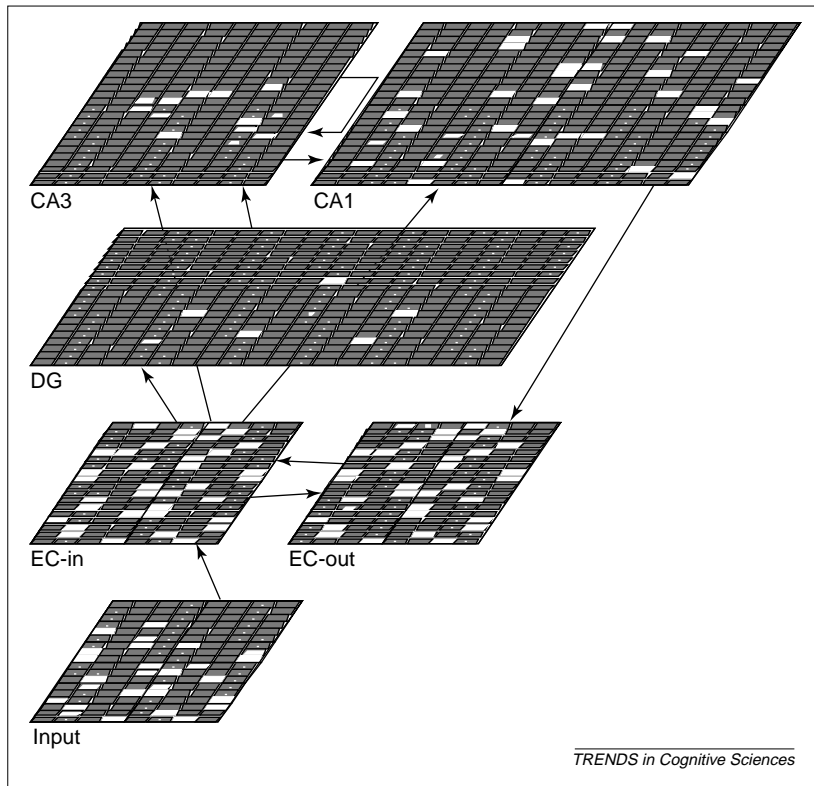
Fig. 3. The hippocampus model with a example pattern of activity (active units shown in white) showing the relative levels of sparseness across the different layers. Input represents converging cortical input into entorhinal cortex (superficial layers, EC_in), which project to the dentate gyrus (DG) and areas CA3 and CA1. CA1 projects back to the deep EC layers (EC_out), which then in principle project back out to the cortex to drive behavior (although in the model we just monitor EC_out directly). The cortical model shown in Fig. 2 represents cortical processing taking place in surrounding cortical areas providing input to the hippocampus (e.g. the perirhinal cortex). In some simulations, the cortical and hippocampal components are combined, in which case the cortical familiarity signal is read off of EC_in, and in other cases we simulate the models separately to more carefully evaluate their distinctive properties.

presentations of the item. Critically, these more active neurons will produce greater inhibitory competition, which depresses the activations of the rest of the more weakly activated neurons. This dynamic is consistent with neural recording data showing decreased responding of some perirhinal neurons with repeated presentation of stimuli (e.g. [34,35]). The evidence for increased firing of other neurons is not as clear, but this is not strictly required by the model – increased inhibition could cancel out increased activations of the 'winners'.

The hippocampal model (Fig. 3) supports recall in the following manner: when stimuli are presented at study, the hippocampus develops relatively non-overlapping (pattern-separated) representations in region CA3 (pattern separation is strongly facilitated by the very sparse dentate gyrus (DG) inputs). Active units in CA3 are linked to one another (again via Hebbian learning), and to a re-representation of the input pattern in region CA1. At test, presentation of a partial version of a studied pattern leads to reconstruction of the complete original CA3 representation (i.e. 'pattern completion') and, through this, reconstruction of the entire studied pattern on the output layer (and then

to cortex) via area CA1. As reviewed in [1,2], our hippocampal model closely resembles other neural-network models of the hippocampus set forth by [15,36,16–19] – there are differences but the family resemblance between these models far outweighs the differences. We apply the hippocampal model to recognition by computing how well retrieved information matches the recall cue.

Equipped with these basic mechanisms, we summarize two representative applications of these models to recognition memory that are described in detail, along with many others, in [1].

Effects of lure similarity
We can use manipulations of lure similarity to contrast the properties of the cortex and hippocampus (Fig. 4a). For dissimilar stimuli (e.g. study the word *cat* and test with *cat* and *potato*) both the cortical and hippocampal systems should provide a memory signal that clearly discriminates between studied words (*cat*) and non-studied lures (*potato*). However, when stimuli are highly similar (e.g. study *cat* and test with *cat* and *cats*), the cortical familiarity signal will not support good discrimination because the lures (*cats*) will tend to activate the sharpened (familiar) representations of the studied items (*cat*). By contrast, hippocampal pattern separation (coupled with its ability to recall the actual studied item, if pattern separation fails) should enable it to still distinguish between these similar items, providing good recognition memory performance.

Interestingly, the model also predicts that cortex will be able to discriminate well between studied items and similar lures, if memory is tested using a forced-choice (FC) paradigm instead of the standard yes-no (Y/N) paradigm (Fig. 4b). In the FC test, participants must choose between a studied item and corresponding related lures (e.g. *rat* versus *rats*). The model predicts that even though the cortical representations overlap considerably, studied items should reliably be slightly more familiar than corresponding lures. In collaboration with Andrew Mayes and Juliet Holdstock, we were able to test these predictions on a patient (YR) who has selective hippocampal damage [37]. YR exhibited deficits relative to matched controls only in the related lure Y/N condition, and not in unrelated lure conditions [38] or related lure conditions with FC testing. This is just as predicted by our model.

Interference and the list strength effect
As noted earlier, an important motivation for the unique properties of the hippocampus is to avoid interference from different memories. Nevertheless, interference cannot be completely avoided, and, paradoxically, our models suggest that in certain recognition memory paradigms hippocampally-based discrimination actually suffers *more* interference than the cortically-based discrimination. Specifically, we simulated a 'list-strength' paradigm, which
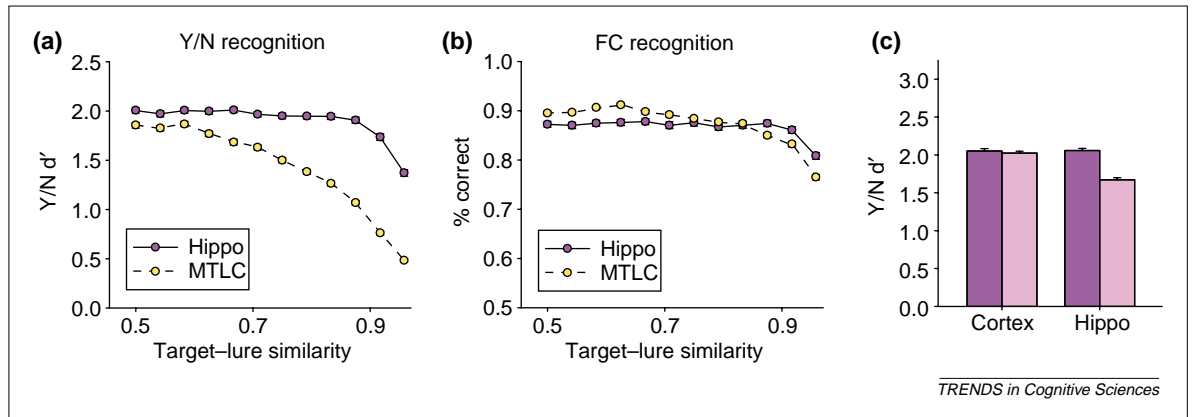
Fig. 4. (a) Yes-No (Y/N) recognition performance in the hippocampal and cortical models (Hippo, hippocampus, MTLC, medial temporal lobe cortex) as a function of similarity between target and lure (proportion of shared input features). As lures become more similar to studied items, lures become increasingly familiar, and recognition performance in MTLC decreases in a graded fashion. The hippocampal model is more robust to the lure similarity manipulation because of its pattern-separation abilities; lures have to be almost identical to studied items to trigger appreciable (false) recall. (b) Forced-choice (FC) recognition performance (percentage correct) in the two models, as a function of target–lure similarity. The cortical model shows good FC performance even with similar lures. (c) The effect of list strength on recognition accuracy in the two models. The cortical model exhibits no list-strength effect: the strong interference condition is the same as the weak one. The hippocampal model does exhibit a list-strength effect: the strong interference condition is worse than the weak one.

measures the effect of strengthening some list items on memory for other, non-strengthened list items. For example, compare studying: *apple bicycle* versus *apple bicycle bicycle bicycle*, with testing on: *apple robot*. Does strengthening some list items (*bicycle*) affect participants' ability to discriminate between other (non-strengthened) studied items (*apple*) and non-studied lure items (*robot*)? We found that increasing list strength impairs subjects' ability to discriminate between studied items and lures based on hippocampal recall, but it does not impair discrimination based on cortical familiarity (see Fig. 4c).

How can we explain these results? The key is to appreciate that recognition discrimination depends on the *difference* in responding to studied and lure items. We found that interference does in fact degrade the cortical network's responding to studied items, but it degrades responding to lures as well. Thus, the difference between studied items and lures does not decrease, and overall recognition is unharmed.

In contrast with the cortex, the response of the hippocampus to lures is essentially zero because of its strong pattern separation abilities. Nevertheless, strengthening of interference items still degrades the memory traces of other studied items. Thus, the studied item recognition signal goes down in the hippocampus, and is not compensated for by a similar decrease in the signal for lure items. This results in a smaller difference between studied and lure items and a list strength interference effect.

Up to this point, nearly all experiments that have looked for a list strength effect for recognition have failed to find one (e.g. [39]). According to the model, this finding can be explained in terms of subjects relying on familiarity; conversely, it should be possible to find a list strength effect using recognition tests and measures that load more heavily on the recall process. This prediction was confirmed in [40].

Pattern completion and contextual fear conditioning in rats

The same models and principles described above have been applied to a wide variety of animal learning phenomena [2,22,23]. Here, we summarize one application of the model that captures its general principles particularly well, and has been tested with both intact and hippocampally lesioned rats [22,41]. The key idea behind this model is that the hippocampus can rapidly and automatically encode a conjunctive representation of an environmental context as a result of simply being exposed to the environment, which provides the basis for the hippocampal involvement in encoding episodic memories in humans. By contrast, the cortex requires more learning trials and specific task demands to encode novel conjunctive representations. In rats, we can use the fear conditioning paradigm to demonstrate the existence of this conjunctive representation by taking advantage of hippocampal pattern completion triggered by a cue that was reliably associated with the environmental context.

Specifically, the model was applied to the following experimental paradigm (Fig. 5a). First, rats were repeatedly pre-exposed by transporting them in a distinctive black ice bucket to a distinct environmental context (A). Thus, we expected this bucket to become a reliable cue associated with the pre-exposure environment. Then, we transported the rats in the bucket into a different environment (C), where they were immediately shocked and then removed from the environment. The question is, did they associate this shock with the actual environment where they were shocked (C), or to the pre-exposure environment (A), the memory of which could have been cued by the bucket? In the model, we found that the simulated bucket cue triggered hippocampal

pattern completion to the pre-exposure environment, which was then associated with the shock (Fig. 5b). In the rats, we tested for this effect by later putting the rats (using a novel transport container) in either the conditioning environment (C) or in the pre-exposure environment A, and measuring freezing. We found that rats pre-exposed to A exhibited freezing in A, not C. We also ran a control with rats pre-exposed to another environment B that showed no freezing in either condition. Furthermore, as predicted by the model, these results depend on an intact hippocampus: hippocampal lesions selectively eliminated the freezing of the A pre-exposure group to environment A (Fig. 5c).

Consistent with the complementary learning systems framework, the cortex can nevertheless learn to associate shock with a conditioning context, if it is given repeated exposures to the environment in the presence of the shock. We demonstrated this in an additional control experiment where sham and hippocampally lesioned rats were shocked three times over 6 min; there were no differences between the lesioned and control group in freezing under this condition [41]. Thus, the unique contribution of the hippocampus is specifically with rapid and automatic learning of conjunctive information, not simply with any form of conjunctive learning (see [2] for many other demonstrations of this point).

Conclusions

How well do we understand hippocampal and neocortical contributions to memory? The representative applications of the CLS model described above, coupled with the much larger range of applications described in [1,2], show that networks based on CLS principles can be used to explain and predict results in highly disparate learning and memory domains. Because our models incorporate many widely held views, this should be good news for many researchers.

However, it is important to note that the CLS model differs in subtle but important ways from other prominent frameworks that specifically contrast hippocampus and cortex. For example, Aggleton and Brown [20] stipulate that medial temporal cortex can only support item recognition, whereas hippocampus is required for making associative links between items (which is similar in many respects to [12]). By contrast, we find such a distinction too rigid to account for both animal nonlinear discrimination learning data, and human recognition memory data, as discussed in detail in [1,2]. Gluck and Myers postulate that the hippocampus can perform a form of pattern separation called 'predictive differentiation', whereas the cortex cannot [42]. The focus on pattern separation is consonant with our own model, but their model specifies that separation only occurs over repeated trials in response to specific task demands, as implemented via error-driven backpropagation
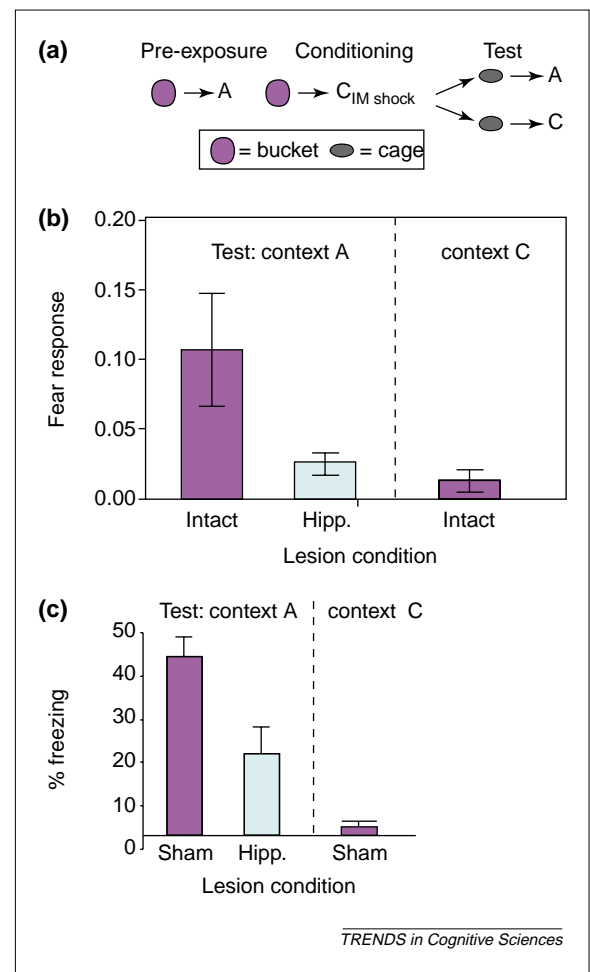

Fig. 5. (a) Design of the conditioning experiment in rats. The black ice bucket was used to transport rats into the pre-exposure environment (A), and then served as a retrieval cue when rats were immediately shocked in the conditioning environment (C). Testing in context A or C (with a different, cage or transport cue) determines which environment was associated with shock. (b) Predictions from the model, showing that rats pre-exposed to A should associate this environment with the shock actually delivered in environment C. (c) Results from intact (Sham) and hippocampally lesioned rats (Hipp.) confirm the predictions of the model.

learning. By contrast, we think that this kind of task-driven differentiation is more characteristic of cortex, whereas the hippocampus can automatically and rapidly form pattern separated representations. Finally, Eichenbaum and Cohen postulate that the hippocampus is specialized for flexible, relational processing, in contrast with a cortex that relies on rote, habit learning [14]. We find this characterization appealing in some respects (e.g. hippocampal pattern completion might produce some kinds of 'flexibility' that the cortex lacks), but it tends to suffer from imprecision. Furthermore, we have recently found that detailed simulations of the transitive inference paradigm in our model are inconsistent with Eichenbaum's characterization of hippocampal involvement in this task [23]. We believe that the CLS model overcomes some of the limitations of other models, and plan to continue to expand the scope of data it can address.

References

1 Norman, K.A. and O'Reilly, R.C. Modeling hippocampal and neocortical contributions to recognition memory: a complementary learning systems approach. *Psychol. Rev.* (in press)

2 O'Reilly, R.C. and Rudy, J.W. (2001) Conjunctive representations in learning and memory: principles of cortical and hippocampal function. *Psychol. Rev.* 108, 311–345

3 McClelland, J.L. *et al.* (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

4 Scoville, W.B. and Milner, B. (1957) Loss of recent memory after bilateral hippocampal lesions. *J. Neurol. Neurosurg. Psychiatry* 20, 11–21

5 Marr, D. (1971) Simple memory: a theory for archicortex. *Philos. Trans. R. Soc. London Ser. B* 262, 23–81

6 Grossberg, S. (1976) Adaptive pattern classification and universal recoding: I. Parallel development and coding of neural feature detectors. *Biol. Cybern.* 23, 121–134

7 O'Keefe, J. and Nadel, L. 1978. *The Hippocampus as a Cognitive Map*, Oxford University Press

8 Teyler, T.J. and Discenna, P. (1986) The hippocampal memory indexing theory. *Behav. Neurosci.* 100, 147–154

9 McNaughton, B.L. and Morris, R.G.M. (1987) Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends Neurosci.* 10, 408–415

10 Sherry, D.F. and Schacter, D.L. (1987) The Evolution of multiple memory systems. *Psychol. Rev.* 94, 439–454

11 Rolls, E.T. (1989) Functions of neuronal networks in the hippocampus and neocortex in memory. In *Neural Models of Plasticity: Experimental and Theoretical Approaches* (Byrne, J. *et al.* eds), pp. 240–265, Academic Press

12 Sutherland, R.J. and Rudy, J.W. (1989) Configural association theory: the role of the hippocampal formation in learning, memory, and amnesia. *Psychobiology* 17, 129–144

13 Squire, L.R. (1992) Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. *Psychol. Rev.* 99, 195–231

14 Eichenbaum, H. *et al.* (1994) Two functional components of the hippocampal memory system. *Behav. Brain Sci.* 17, 449–518

15 Treves, A. and Rolls, E.T. (1994) A computational analysis of the role of the hippocampus in memory. *Hippocampus* 4, 374–392

16 Burgess, N. and O'Keefe, J. (1996) Neuronal computations underlying the firing of place cells and their role in navigation. *Hippocampus* 6, 749–762

17 Wu, X. *et al.* (1996) Context codes and the effect of noisy learning on a simplified hippocampal CA3 model. *Biol. Cybern.* 74, 159–165

18 Moll, M. and Miikkulainen, R. (1997) Convergence-zone episodic memory: analysis and simulations. *Neural Netw.* 10, 1017–1036

19 Hasselmo, M.E. and Wyble, B. (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

20 Aggleton, J.P. and Brown, M.W. (1999) Episodic memory, amnesia, and the hippocampal–anterior thalamic axis. *Behav. Brain Sci.* 22, 425–490

21 Yonelinas, A.P. (2002) The nature of recollection and familiarity: a review of 30 years of research. *J. Mem. Lang.* 46, 441–517

22 Rudy, J.W. and O'Reilly, R.C. (2001) Conjunctive representations, the hippocampus, and contextual fear conditioning. *Cognit. Affect. Behav. Neurosci.* 1, 66–82

23 Frank, M.J. *et al.* Transitivity, flexibility, conjunctive representations and the hippocampus: II. A computational analysis. *Hippocampus* (in press)

24 O'Reilly, R.C. *et al.* (1998) A hippocampal model of recognition memory. In *Advances in Neural Information Processing Systems* (Vol. 10) (Jordan, M.I. *et al.*, eds), pp. 73–79, MIT Press

25 O'Reilly, R.C. and Rudy, J.W. (2000) Computational principles of learning in the neocortex and hippocampus. *Hippocampus* 10, 389–397

26 Carpenter, G.A. and Grossberg, S. (1993) Normal and amnesic learning, recognition and memory by a neural model of cortico-hippocampal interactions. *Trends Neurosci.* 16, 131–137

27 O'Reilly, R.C. and McClelland, J.L. (1994) Hippocampal conjunctive encoding, storage, and recall: avoiding a tradeoff. *Hippocampus* 4, 661–682

28 Boss, B.D. *et al.* (1987) On the numbers of neurons in fields CA1 and CA3 of the hippocampus of Sprague–Dawley and Wistar rats. *Brain Res.* 406, 280–287

29 Barnes, C.A. *et al.* (1990) Comparison of spatial and temporal characteristics of neuronal activity in sequential stages of hippocampal processing. *Prog. Brain Res.* 83, 287–300

30 Rudy, J.W. and Sutherland, R.W. (1995) Configural association theory and the hippocampal formation: an appraisal and reconfiguration. *Hippocampus* 5, 375–389

31 Wickelgren, W.A. (1979) Chunking and consolidation: a theoretical synthesis of semantic networks, configuring in conditioning, S–R versus cognitive learning, normal forgetting, the amnesic syndrome, and the hippocampal arousal system. *Psychol. Rev.* 86, 44–60

32 O'Reilly, R.C. and Munakata, Y. (2000) *Computational Explorations in Cognitive Neuroscience: Understanding the Mind by Simulating the Brain*, MIT Press

33 Grossberg, S. and Stone, G. (1986) Neural dynamics of word recognition and recall: Attentional priming, learning, and resonance. *Psychol. Rev.* 93, 46–74

34 Brown, M.W. and Xiang, J.Z. (1998) Recognition memory: neuronal substrates of the judgement of prior occurrence. *Prog. Neurobiol.* 55, 149–189

35 Li, L. *et al.* (1993) The representation of stimulus familiarity in anterior inferior temporal cortex. *J. Neurophysiol.* 69, 1918–1929

36 Touretzky, D.S. and Redish, A.D. (1996) A theory of rodent navigation based on interacting representations of space. *Hippocampus* 6, 247–270

37 Holdstock, J.S. *et al.* (2002) Under what conditions is recognition spared relative to recall after selective hippocampal damage in humans? *Hippocampus* 12, 341–351

38 Mayes, A.R. *et al.* (2002) Relative sparing of item recognition memory in a patient with adult-onset damage limited to the hippocampus. *Hippocampus* 12, 325–340

39 Ratcliff, R. *et al.* (1990) The list strength effect: I. Data and discussion. *J. Exp. Psychol. Learn. Mem. Cogn.* 16, 163–178

40 Norman, K.A. Differential effects of list strength on recollection and familiarity. *J. Exp. Psychol. Learn. Mem. Cogn.* (in press)

41 Rudy, J.W. *et al.* (2002) Hippocampal formation supports conditioning to memory of a context. *Behav. Neurosci.* 116, 530–538

42 Gluck, M.A. and Myers, C.E. (1993) Hippocampal mediation of stimulus representation: a computational theory. *Hippocampus* 3, 491–516