

Computational Principles of Learning in the Neocortex and Hippocampus

Randall C. O'Reilly* and Jerry W. Rudy

Department of Psychology, University of Colorado
at Boulder, Boulder, Colorado

ABSTRACT: We present an overview of our computational approach towards understanding the different contributions of the neocortex and hippocampus in learning and memory. The approach is based on a set of principles derived from converging biological, psychological, and computational constraints. The most central principles are that the neocortex employs a slow learning rate and overlapping distributed representations to extract the general statistical structure of the environment, while the hippocampus learns rapidly, using separated representations to encode the details of specific events while suffering minimal interference. Additional principles concern the nature of learning (error-driven and Hebbian), and recall of information via pattern completion. We summarize the results of applying these principles to a wide range of phenomena in conditioning, habituation, contextual learning, recognition memory, recall, and retrograde amnesia, and we point to directions of current development. *Hippocampus* 2000;10:389–397. © 2000 Wiley-Liss, Inc.

KEY WORDS: neocortex; computational models; neural networks; learning; memory

INTRODUCTION

This paper presents a computational approach towards understanding the different contributions of the neocortex and hippocampus in learning and memory. This approach uses basic principles of computational neural network learning mechanisms to understand both *what* is different about the way these two neural systems learn, and *why* they should have these differences. Thus, the computational approach can go beyond mere description towards understanding the deeper principles underlying the organization of the cognitive system. These principles are based on a convergence of biological, psychological, and computational constraints, and serve to bridge the gaps between these different levels of analysis.

The set of principles discussed in this paper were first developed in McClelland et al. (1995), and have been refined several times since then (O'Reilly et al., 1998; O'Reilly and Rudy, 2000). The computational principles have been applied to a wide range of learning and memory phenomena across several species (rats, monkeys, and humans). For example, they can account for impaired and preserved learning capacities with hippocampal lesions in conditioning, habituation, contextual learning, recognition memory, recall, and retrograde amnesia. This paper provides a concise summary of previous work, and a discussion of current and future directions.

*Correspondence to: Randall C. O'Reilly, Department of Psychology, University of Colorado at Boulder, Campus Box 345, Boulder, CO 80309.
E-mail: oreilly@psych.colorado.edu
Accepted for publication 1 May 2000

THE PRINCIPLES

There are several levels of principles that can be distinguished by their degree of specificity in characterizing the nature of the underlying mechanisms. We begin with the most basic principles and proceed towards greater specificity.

Learning Rate, Overlap, and Interference

The most basic set of principles can be motivated by considering how subsequent learning can interfere with prior learning. A classic example of this kind of interference can be found in the AB-AC associative learning task (e.g., Barnes and Underwood, 1959). The A represents one set of words that are associated with two different sets of other words, B and C. For example, the word *window* will be associated with the word *reason* in the AB list, and associated with *locomotive* on the AC list. After studying the AB list of associates, subjects are tested by asking them to give the appropriate B associate for each of the A words. Then, subjects study the AC list (often over multiple iterations), and are subsequently tested on both lists for recall of the associates after each iteration of learning the AC list. Subjects exhibit some level of interference on the initially learned AB associations as a result of learning the AC list, but they still remember a reasonable percentage (see Fig. 1a for representative data).

The first set of principles concerns the effects of overlapping representations (i.e., shared units between two different distributed representations) and rate of learning on the ability to rapidly learn new information with a level of interference characteristic of human subjects:

- Overlapping representations lead to interference (conversely, separated representations prevent interference).
- A faster learning rate causes more interference (conversely, a slower learning rate causes less interference).

The mechanistic basis for these principles within a neural network perspective is straightforward. Interference is caused when weights used to encode one association are disturbed by the encoding of another (Fig. 2a). Overlapping patterns share more weights, and therefore lead to greater amounts of interference. Clearly, if en-

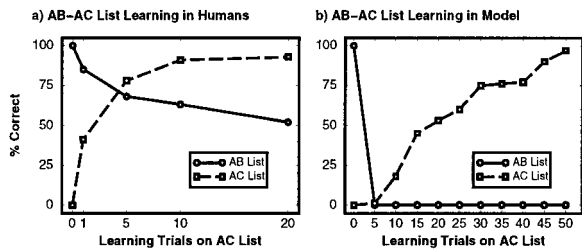


FIGURE 1. Human and model data for AB-AC list learning. a: Humans show some interference for the AB list items as a function of new learning on the AC list items. b: Model shows a catastrophic level of interference (data reproduced from McCloskey and Cohen, 1989).

tirely separate representations are used to encode two different associations, then there will be no interference whatsoever (Fig. 2b). The story with learning rate is similarly straightforward. Faster learning rates lead to more weight change, and thus greater interference (Fig. 3). However, a fast learning rate is necessary for rapid learning.

Integration and Extracting Statistical Structure

Figure 3 shows the flip side of the interference story, i.e., *integration*. If the learning rate is low, then the weights will integrate over many experiences, reflecting the *underlying statistics* of the environment (White, 1989; McClelland et al., 1995). Furthermore, overlapping representations facilitate this integration process, because the same weights need to be reused across many different experiences to enable the integration produced by a slow learning rate. This leads to the next principle:

- Integration across experiences to extract underlying statistical structure requires a slow learning rate and overlapping representations.

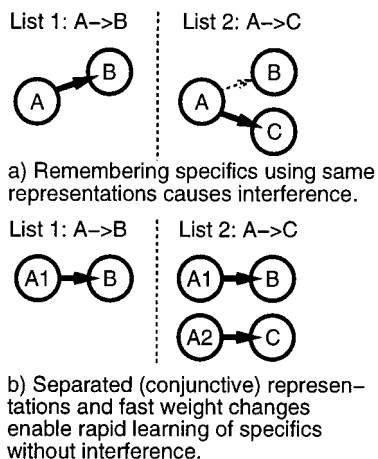


FIGURE 2. Interference as a function of overlapping (same) representations vs. separated representations. a: Using the same representation to encode two different associations ($A \rightarrow B$ and $A \rightarrow C$) causes interference: the subsequent learning of $A \rightarrow C$ interferes with the prior learning of $A \rightarrow B$ because the A stimulus must have stronger weights to C than to B for the second association, as is reflected in the weights. b: A separated representation, where A is encoded separately for the first list (A1) vs. the second list (A2) prevents interference.

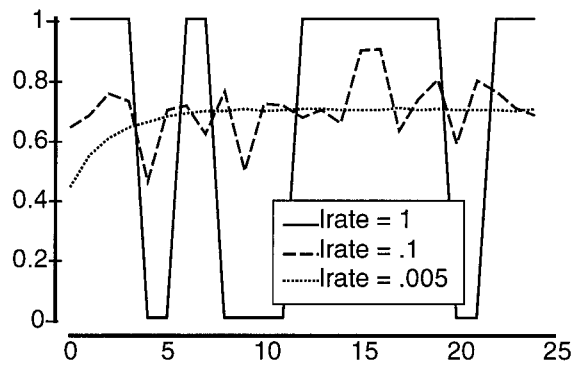


FIGURE 3. Weight value learning about a single input unit that is either active or not. The weight increases when the input is on, and decreases when it is off, in proportion to the size of the learning rate. The input has an overall probability of being active of 0.7. Larger learning rates (0.1 or 1) lead to more interference on prior learning, resulting in a weight value that bounces around substantially with each training example. In the extreme case of a learning rate of 1, the weight only reflects what happened on the previous trial, retaining no memory for prior events at all. As the learning rate gets smaller (0.005), the weight smoothly averages over individual events and reflects the overall statistical probability of the input being active.

Episodic Memory and Generalization: Incompatible Functions

Thus, focusing only on pattern overlap for the moment, we can see that networks can be optimized for two different, and incompatible, functions: avoiding interference, or integrating across experiences to extract generalities. Avoiding interference requires separated representations, while integration requires overlapping representations. These two functions each have clear functional advantages, leading to a further set of principles:

- Interference avoidance is essential for *episodic* memory, which requires learning about the specifics of individual events and keeping them separate from other events.
- Integration is essential for encoding the general statistical structure of the environment, abstracted away from the specifics of individual events, which enables *generalization* to novel situations.

The incompatibility between these functions is further evident in these descriptions (e.g., encoding specifics vs. abstracting away from them). Also, episodic memory requires relatively rapid learning: an event must be encoded as it happens, and does not typically repeat itself for further learning opportunities. This completes a pattern of opposition between these functions: episodic learning requires rapid learning, while integration and generalization require slow learning. This is summarized in the following principle:

- Episodic memory and extracting generalities are in opposition. Episodic memory requires rapid learning and separated patterns, while extracting generalities requires slow learning and overlapping patterns.

Applying the Principles to the Hippocampus and Neocortex

Armed with these principles, the finding that neural network models that have highly overlapping representations exhibit catastrophic levels of interference (McCloskey and Cohen, 1989, Fig. 1b) should not be surprising. A number of researchers showed that this interference can be reduced by introducing various factors that result in less pattern overlap (e.g., Kortge, 1993; French, 1992; Sloman and Rumelhart, 1992; McRae and Hetherington, 1993). Thus, instead of concluding that neural networks are fundamentally flawed, as McCloskey and Cohen (1989) argued (and a number of others have uncritically accepted), McClelland et al. (1995) argued that this catastrophic failure serves as an important clue into the structure of the human brain.

Specifically, we argued that because of the fundamental incompatibility between episodic memory and extracting generalities, the brain should employ two separate systems that each optimize these two objectives individually, instead of having a single system that tries to strike an inferior compromise. This line of reasoning provides a strikingly good fit to the known properties of the hippocampus and neocortex, respectively. The details of this fit in various contexts is the substance of the remainder of the paper, but the general idea is that:

- The hippocampus rapidly binds together information, using pattern-separated representations to minimize interference.
- The neocortex slowly learns about the general statistical structure of the environment, using overlapping distributed representations.

See also Sherry and Schacter (1987) for a similar conclusion. Before discussing the details, a few more principles need to be developed first.

Conjunctive Representations and Nonlinear Discrimination Learning

The *conjunctive* or *configural* representations theory provides a converging line of thinking about the nature of hippocampal function (Sutherland and Rudy, 1989; Rudy and Sutherland, 1995; Wickelgren, 1979; O'Reilly and Rudy, 2000). A conjunctive/configural representation is one that binds together (conjoins or configures) multiple elements into a novel unitary representation. This is consistent with the description of hippocampal function given above, based on the need to separate patterns to avoid interference. Indeed, it is clear that pattern separation and conjunctive representations are two sides of the same coin, and that both are caused by the use of *sparse* representations (having relatively few active neurons) that are a known property of the hippocampus (O'Reilly and McClelland, 1994; O'Reilly and Rudy, 2000). To summarize:

- Sparse hippocampal representations lead to pattern separation (to avoid interference) and conjunctive representations (to bind together features into a unitary representation).

One important application of the conjunctive representations idea has been to solve *nonlinear discrimination problems*. These

problems require conjunctive representations to solve because each of the individual stimuli is ambiguous (equally often rewarded and not rewarded). The negative patterning problem is a good example. It involves two stimuli, A and B (e.g., a light and a tone), which are associated with reward (indicated by +) or not (-). Three different trial types are trained: A+, B+, AB-. Thus, the conjunction of the two stimuli (AB-) must be treated differently from the two stimuli separately (A+, B+). A conjunctive representation that forms a novel encoding of the two stimuli together can facilitate this form of learning. Therefore, the fact that hippocampal damage impairs learning the negative patterning problem (Alvarado and Rudy, 1995; Rudy and Sutherland, 1995; McDonald et al., 1997) would appear to support the idea that the hippocampus employs pattern separated, conjunctive representations. However, it is now clear that a number of other nonlinear discrimination learning problems are unimpaired by hippocampal damage (Rudy and Sutherland, 1995). The next set of principles helps to make sense of these data so that they can be reconciled with our interference-based principles of conjunctive pattern separation, as discussed below.

Pattern Completion: Recalling a Conjunction

Pattern completion is required for recalling information from conjunctive hippocampal representations, yet it conflicts with the process of pattern separation that forms these representations in the first place (O'Reilly and McClelland, 1994). Pattern completion occurs when a partial input cue drives the hippocampus to complete an entire previously encoded set of features that were bound together in a conjunctive representation. For a given input pattern, a decision must be made to recognize it as a retrieval cue for a previous memory and perform pattern completion, or to perform pattern separation and store the input as a new memory. This decision is often difficult, given noisy inputs and degraded memories. The hippocampus implements this decision as the effect of a set of basic mechanisms operating on input patterns (O'Reilly and McClelland, 1994; Hasselmo and Wyble, 1997), and it does not always do what would seem to be the right thing to do from an omniscient perspective of knowing all the relevant task factors: this can complicate the involvement of the hippocampus in nonlinear discrimination problems.

Learning Mechanisms: Hebbian and Error-Driven

To more fully explain the roles of the hippocampus and neocortex, we need to understand how learning works in these systems (the basic principles just described do not depend on the detailed nature of the learning mechanisms; White, 1989). There are two basic mechanisms that have been discussed in the literature, Hebbian and error-driven learning (e.g., Marr, 1971; McNaughton and Morris, 1987; Gluck and Myers, 1993; Schmajuk and DiCarlo, 1992). Briefly, Hebbian learning (Hebb, 1949) works by increasing weights between coactive neurons (and usually decreasing weights when a receiver is active and the sender is not), which is a well-established property of biological synaptic modification mechanisms (e.g., Collingridge and Bliss, 1987). Hebbian learning

is useful for binding together features active at the same time (e.g., within the same episode), and has therefore been widely suggested as a hippocampal learning mechanism (e.g., Marr, 1971; McNaughton and Morris, 1987).

Error-driven learning works by adjusting weights to minimize the errors in a network's performance, with the best example of this being the *error backpropagation* algorithm (Rumelhart et al., 1986). Error-driven learning is sensitive to task demands in a way that Hebbian learning is not, and this makes it a much more capable form of learning for actually achieving a desired input/output mapping. Thus, it is natural to associate this form of learning with the kind of procedural or task-driven learning that the neocortex is often thought to specialize in (e.g., because amnesics with hippocampal damage have preserved procedural learning abilities). Although the backpropagation mechanism has been widely challenged as biologically implausible (e.g., Crick, 1989; Zipser and Andersen, 1988), a recent analysis showed that simple biologically based mechanisms can be used to implement this mechanism (O'Reilly, 1996), so that it is quite reasonable to assume that the cortex depends on this kind of learning.

Although the association of Hebbian learning with the hippocampus and error-driven learning with the cortex is appealing in some ways, it turns out that both kinds of learning can play important roles in both systems (O'Reilly and Rudy, 2000; O'Reilly and Munakata, 2000; O'Reilly, 1998). Thus, the specific learning principles adopted here are that both forms of learning operate in both systems:

- Hebbian learning binds together co-occurring features (in the hippocampus) and generally learns about the co-occurrence statistics in the environment across many different patterns (in the neocortex).
- Error-driven learning shapes learning according to specific task demands (shifting the balance of pattern separation and completion in the hippocampus, and developing task-appropriate representations in the neocortex).

It is the existence of this task-driven learning that complicates the picture for nonlinear discrimination learning problems.

A Summary of Principles

The above principles can be summarized with the following three general statements of neocortical and hippocampal learning properties (O'Reilly and Rudy, 1999):

Learning rate

The cortical system typically learns slowly, while the hippocampal system typically learns rapidly.

Conjunctive bias

The cortical system has a bias towards integrating over specific instances to extract generalities. The hippocampal system is biased by its intrinsic sparseness to develop conjunctive representations of specific instances of environmental inputs. However, this conjunctive bias trades off with the countervailing process of pattern com-

pletion, so the hippocampus does not always develop new conjunctive representations (sometimes it completes to existing ones).

Learning mechanisms

Both the cortex and hippocampus use error-driven and Hebbian learning. The error-driven aspect responds to task demands, and will cause the network to learn to represent whatever is needed to achieve goals or ends. Thus, the cortex can overcome its bias and develop specific, conjunctive representations if the task demands this. Also, error-driven learning can shift the hippocampus from performing pattern separation to performing pattern completion, or vice-versa, as dictated by the task. Hebbian learning is constantly operating, and reinforcing the representations that are activated in the two systems.

These principles are focused on distinguishing the neocortex and hippocampus. We have also articulated a more complete set of principles that are largely common to both systems (O'Reilly, 1998; O'Reilly and Munakata, 2000). Models incorporating these principles have been extensively applied to a wide range of different cortical phenomena, including perception, language, and higher-level cognition. Next, we highlight the application of principles presented here to learning and memory phenomena involving both the cortex and hippocampus.

APPLICATIONS OF THE PRINCIPLES

The principles just developed have been applied to a number of different domains, as summarized below. In most cases, the same neural network model developed according to these principles (O'Reilly and Rudy, 2000) was used to simulate the empirical data, providing a compelling demonstration that the principles are sufficient to account for a wide range of findings.

Conjunctions and Nonlinear Discrimination Learning

We first apply the above principles to the puzzling pattern of hippocampal involvement in nonlinear discrimination learning problems. The general statement of the issue is that although we think the hippocampus is specialized for encoding conjunctive bindings of stimuli (and keeping these separated from each other to minimize interference), apparently direct tests of this idea in the form of nonlinear discrimination learning problems have not provided clear support. Specifically, rats with hippocampal lesions can learn a number of these nonlinear discrimination problems just like intact rats. The general explanation of these results according to the full set of principles outlined above is that:

- The explicit task demands present in a nonlinear discrimination learning problem cause the cortex alone (with a lesioned hippocampus) to learn the task via error-driven learning.




FIGURE 4. The model of O’Reilly and Rudy (2000), showing both cortical and hippocampal components. The cortex has 12 different input dimensions (sensory pathways), with 4 different values per dimension. These are represented separately in the elemental cortex (Elem). The higher level association cortex (Assoc) can form conjunctive representations of these elements, if demanded by the task. The interface to the hippocampus is via the entorhinal cortex, which contains a one-to-one mapping of the elemental, association, and output cortical representations. The hippocampus can reinstate a pattern of activity over the cortex via the EC.

- Nonlinear discrimination problems take many trials to learn even in intact animals, allowing the slow cortical learning to accumulate a solution.
- The absence of hippocampal learning speed advantages in normal rats, despite the more rapid hippocampal learning rate, can be explained by the fact that the hippocampus is engaging in pattern completion in these problems, instead of pattern separation.

We substantiated this verbal account by running computational neural network simulations that embodied the principles developed above (O’Reilly and Rudy, 2000) (Fig. 4). These simulations showed that in many (but not all) cases, removing the hippocampal component did not significantly impair learning performance on nonlinear discrimination learning problems, matching the empirical data. Figure 5 shows one specific example, where the negative patterning problem discussed previously (A+, B+, AB-) is impaired with hippocampal lesions, but performance is not impaired on a very similar *ambiguous feature* problem, AC+, B+, AB-, C-, (Gallagher and Holland, 1992). See O’Reilly and Rudy (2000) for a detailed discussion of the differential performance on these tasks.

To summarize, this work showed that it is essential to go beyond a simple conjunctive story and include a more complete set of principles in understanding hippocampal and cortical function. Because this more complete set of principles, implemented in an explicit computational model, accounts for the empirical data, these data provide support for these principles.

Rapid Incidental Conjunctive Learning Tasks

A consideration of the full set of principles suggests that another class of tasks might provide a much better measure of hippocampal learning compared to the nonlinear discrimination problems suggested by Sutherland and Rudy (1989). As we just saw, the very fact that nonlinear discrimination problems *require* conjunctive representations is what drives the cortex alone to be able to solve them via error-driven learning. Therefore, we suggest that *incidental* conjunctive learning tasks, where conjunctive representations are not forced by specific task demands, may provide a much better index of hippocampal function (O’Reilly and Rudy, 2000). Furthermore, the task should only allow for a relatively brief period of learning, which will emphasize the rapid learning of the hippocampus as compared to the slow learning of the cortex. Thus, we characterize these tasks as *rapid, incidental conjunctive learning tasks*.

There are several recent studies of tasks that fit the rapid, incidental conjunctive characterizations. In these tasks, subjects are exposed to a set of features in a particular configuration, and then the features are rearranged. Subjects are then tested to determine if they can detect the rearrangement. If the test indicates that the rearrangement was detected, then one can infer that the subject learned a conjunctive representation of the original configuration. The literature indicates that the incidental learning of stimulus conjunctions, unlike many nonlinear discrimination problems, *is* dependent on the hippocampus.

Perhaps the simplest demonstration comes from the study of the role of the hippocampal formation in exploratory behavior. Control rats and rats with damage to the dorsal hippocampus were repeatedly exposed to a set of objects that were arranged on a circular platform in a fixed configuration relative to a large and

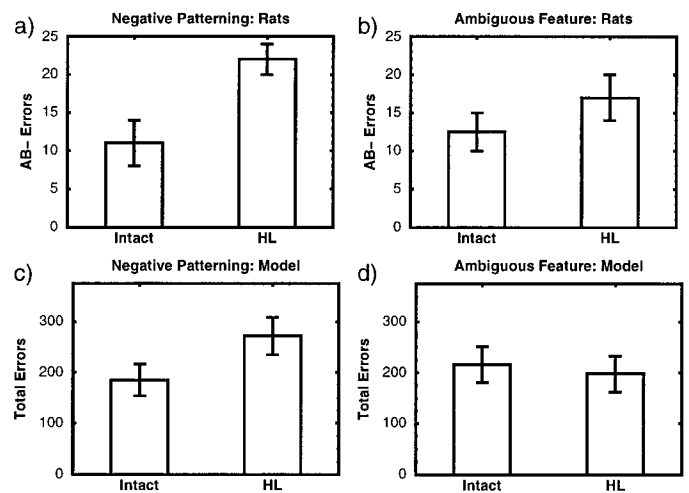


FIGURE 5. Results for the negative patterning (left) and ambiguous feature (AC+, B+, AB-, C-; right) problems. Top row shows data from rats from Alvarado and Rudy (1995), and the bottom row shows data from the model. Intact is intact rats/networks, and HL is rats/networks with hippocampal lesions. N = 40 different random initializations for the model. The hippocampally lesioned system is able to learn the problems, and all conditions require many trials (i.e. large number of errors). Negative patterning is differentially impaired with a hippocampal lesion. Data from O’Reilly and Rudy (2000).

distinct visual cue (Save et al., 1992). After the exploratory behavior of both sets of rats had habituated, the same objects were rearranged into a different configuration. This rearrangement reinstated exploratory behavior in the control rats but not in the rats with damage to the hippocampus. In a third phase of the study, a new object was introduced into the mix. This manipulation reinstated exploratory behavior in both sets of rats. This pattern of data suggests that both control rats and rats with damage to the hippocampus encode representations of the individual objects and can discriminate them from novel objects. However, only the control rats encoded the conjunctions necessary to represent the spatial arrangement of the objects, even though this was not in any way a requirement of the task. Several other studies of this general form have found similar results in rats (Honey et al., 1990, 1998; Honey and Good, 1993; Good and Bannerman, 1997; Hall and Honey, 1990). In humans, the well-established incidental context effects on memory (e.g., Godden and Baddeley, 1975) have been shown to be hippocampal-dependent (Mayes et al., 1992). Other hippocampal incidental conjunctive learning effects have also been demonstrated in humans (Chun and Phelps, 1999).

We have shown that the same neural network model constructed according to our principles and tested on nonlinear discrimination learning problems as described above exhibits a clear hippocampal sensitivity in these rapid incidental conjunctive learning tasks (O'Reilly and Rudy, 2000).

Contextual Fear Conditioning

Evidence for the involvement of the hippocampal formation in the incidental learning of stimulus conjunctions has also emerged in the contextual fear conditioning literature. This evidence also provides a simple example of the widely discussed role of the hippocampus in spatial learning (e.g., O'Keefe and Nadel, 1978; McNaughton and Nadel, 1990). Rats with damage to the hippocampal formation do not express fear in a context or place where shock occurred, but will express fear towards an explicit cue (e.g., a tone) paired with shock (Kim and Fanselow, 1992; Phillips and LeDoux, 1994; but see Maren et al., 1997). Rudy and O'Reilly (1999) provided specific evidence that, in intact rats, the context representations are conjunctive in nature, which has been widely assumed (e.g., Fanselow, 1990; Kiernan and Westbrook, 1993; Rudy and Sutherland, 1994). For example, we compared the effects of preexposure to the conditioning context with the effects of preexposure to the separate features that made up the context. Only preexposure to the intact context facilitated contextual fear conditioning, suggesting that conjunctive representations across context features were necessary. We also showed that pattern completion of hippocampal conjunctive representations can lead to generalized fear conditioning.

We have simulated the incidental learning of conjunctive context representations in fear conditioning, using the same principles as described above (O'Reilly and Rudy, 2000). For example, Figure 6 shows the rat and model data for the separate vs. intact context features experiment from Rudy and O'Reilly (1999), with the model providing a specific prediction regarding the effects of hippocampal lesions, which has yet to be tested empirically.

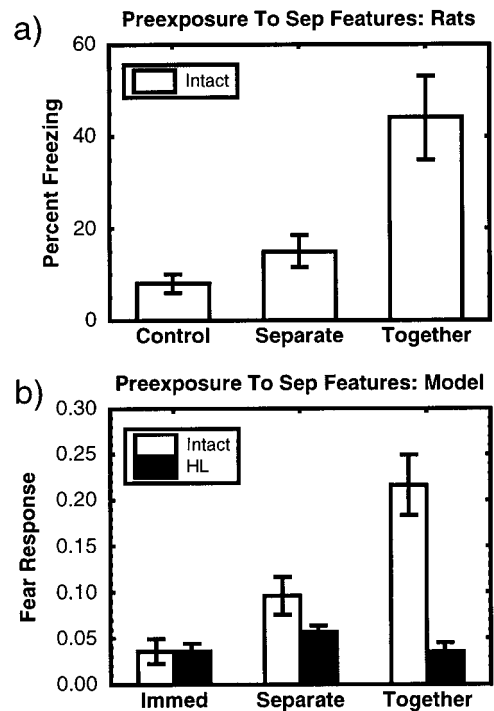


FIGURE 6. Effects of exposure to the features separately, compared to exposure to the entire context on level of fear response in (a) rats (data from Rudy and O'Reilly, 1999) and (b) the model. The immediate shock condition (Immed) is included as a control condition for the model. Intact rats and the intact model show a significant effect of being exposed to the entire context together, compared to the features separately, while the hippocampally lesioned model exhibits slightly more responding in the separate condition, possibly because of the greater overall number of training trials in this case. Simulation data from O'Reilly and Rudy (2000).

Transitivity and Flexibility

Whereas the previous examples concerned the learning of conjunctive representations, this next example is concerned with the flexible use of learned information. Several theorists have described memories encoded by the hippocampus as being flexible, meaning that 1) such memories can be applied inferentially in novel situations (Eichenbaum, 1992; O'Keefe and Nadel, 1978), or 2) they are available to multiple response systems (Squire, 1992). Although the term *flexibility* provides a useful description of certain behaviors, it does not provide a mechanistic understanding of how this flexibility arises from the properties of the hippocampus.

We have shown that hippocampal pattern completion plays an important role in producing this flexible behavior (O'Reilly and Rudy, 2000). Specifically, we showed that the *transitivity* studies of Bunsey and Eichenbaum (1996) and Dusek and Eichenbaum (1997) can be simulated using the same model as in all of the previous examples, with pattern completion playing a key role. Interestingly, the model shows that the training parameters employed in these studies interact significantly with the pattern completion mechanism to produce the observed transitivity effects. We are able to make a number of novel empirical predictions that are

inconsistent with a simple logical-reasoning mechanism by manipulating these factors (O'Reilly and Rudy, 2000).

Dual-Process Memory Models

The dual mechanisms of the neocortex and hippocampus provide a natural fit with dual-process models of recognition memory (Jacoby et al., 1997; Aggleton and Shaw, 1996; Aggleton and Brown, 1999; Vargha-Khadem et al., 1997; Holdstock et al., 2000; O'Reilly et al., 1998). These models hold that recognition can be subserved by two different processes, a *recollection* process and a *familiarity* process. Recollection involves the recall of specific episodic details about the item, and thus fits well with the hippocampal principles developed here. Indeed, we have simulated distinctive aspects of recollection using essentially the same model (O'Reilly et al., 1998). Familiarity is a nonspecific sense that the item has been seen recently; we argue that this can be subserved by the small weight changes produced by slow cortical learning. Current simulation work has shown that a simple cortical model can account for a number of distinctive properties of the familiarity signal (Norman et al., 2000).

One specific and somewhat counterintuitive prediction of our principles has recently been confirmed empirically in experiments on a patient with selective hippocampal damage (Holdstock et al., 2000). This patient showed intact recognition memory for studied items compared to similar lures when tested in a two-alternative forced-choice procedure (2AFC), but was significantly impaired relative to controls for the same kinds of stimuli using a single item yes-no (YN) procedure. We argue that because the cortex uses overlapping distributed representations, the strong similarity of the lures to the studied items produces a strong familiarity signal for these lures (as a function of this overlap). When tested in a YN procedure, this strong familiarity of the lures produces a large number of false alarms, as was observed in the patient. However, because the studied item has a small but reliably stronger familiarity signal than the similar lure, this strength difference can be detected in the 2AFC version, resulting in normal recognition performance in this condition. The normal controls, in contrast, have an intact hippocampus which performs pattern separation and is able to distinguish the studied items and similar lures, regardless of the testing format.

Retrograde Amnesia

Several lines of empirical evidence suggest that there is a *retrograde gradient* for memory loss as a function of hippocampal damage, with the *most recent* memories being the most severely affected, while *older* memories are relatively intact (e.g., Squire, 1992; Winocur, 1990; Kim and Fanselow, 1992; Zola-Morgan and Squire, 1990). Theoretically, this phenomenon can be understood in terms of the cortex gradually acquiring hippocampal information (e.g., McClelland et al., 1995; Alvarez and Squire, 1994). However, this account has been called into question recently, both from failures to replicate the retrograde findings (Sutherland et al., 2000), and reinterpretations of the existing findings in ways that do not require that the cortex acquires information from the hippocampus (e.g., the "multiple hippocampal trace" theory of Nadel and Moscovitch, 1997).

Our computational principles suggest that to the extent there are opportunities for the hippocampus to reactivate cortical patterns of activity, the consequent cortical learning will necessarily produce a consolidation-like effect. We were able to fit a number of different retrograde amnesia gradients using these principles (McClelland et al., 1995).

COMPARISON WITH OTHER APPROACHES

A number of other approaches to understanding cortical and hippocampal function share important similarities with our approach, including, for example, the use of Hebbian learning and pattern separation (e.g., Hasselmo and Wyble, 1997; McNaughton and Nadel, 1990; Touretzky and Redish, 1996; Burgess and O'Keefe, 1996; Wu et al., 1996; Treves and Rolls, 1994; Moll and Miikkulainen, 1997; Alvarez and Squire, 1994). These other approaches all offer other important principles, many of which would be complementary to those discussed here, so that it would be possible to add them to a larger, more complete model.

Perhaps the largest area of disagreement is in terms of the relative independence of the cortical learning mechanisms from the hippocampus. There are several computationally explicit models proposing that the neocortex is incapable of powerful learning without the help of the hippocampus (Gluck and Myers, 1993; Schmajuk and DiCarlo, 1992; Rolls, 1990), and other more general theoretical views that express a similar notion of limited cortical learning with hippocampal damage (Glisky et al., 1986; Squire, 1992; Cohen and Eichenbaum, 1993; Wickelgren, 1979; Sutherland and Rudy, 1989). In contrast, our principles hold that the cortex alone is a highly capable learning system, that can for example learn complex conjunctive representations in the service of nonlinear discrimination learning problems.

Empirically, the data that appear to support the limited cortical learning view tend to be based on larger lesions of the medial temporal lobe. With the advent of selective lesion techniques in rats and monkeys, and the study of people with highly selective hippocampal lesions, it is becoming clear that the cortex is capable of quite substantial learning on its own. Perhaps the most dramatic evidence comes from a group of human amnesics who suffered bilateral selective hippocampal damage at relatively young ages (Vargha-Khadem et al., 1997). Despite having significantly impaired hippocampal function (as was supported by brain scans and very poor performance on recall tests), these individuals had acquired normal or nearly normal levels of cognitive functioning in language and semantic knowledge, and had normal or nearly normal IQs. Although it is difficult to completely rule out the idea that this preserved semantic learning is the result of residual hippocampal functioning (as advocated by Squire and Zola, 1998), this seems somewhat implausible in the face of the patients' significant recall impairments and the brain scan evidence.

One important conclusion from this line of reasoning is that the cortical regions surrounding the hippocampus in the medial temporal lobes are particularly important for many kinds of learning and memory. We suggest that this is because of a significant con-

vergence of other cortical association areas in these regions (O'Reilly and Rudy, 2000; Mishkin et al., 1997, 1998).

CONCLUSIONS

We have shown that a small set of computationally motivated principles can account for a wide range of empirical findings regarding the differential properties of the neocortex and hippocampus in learning and memory. In addition, these principles make a large number of empirical predictions that will be tested in future research.

REFERENCES

- Aggleton JP, Brown MW. 1999. Episodic memory, amnesia, and the hippocampal-anterior thalamic axis. *Behav Brain Sci* 22:425–490.
- Aggleton JP, Shaw C. 1996. Amnesia and recognition memory: a re-analysis of psychometric data. *Neuropsychologia* 34:51.
- Alvarado MC, Rudy JW. 1995. A comparison of kainic acid plus colchicine and ibotenic acid induced hippocampal formation damage on four configural tasks in rats. *Behav Neurosci* 109:1052–1062.
- Alvarez P, Squire LR. 1994. Memory consolidation and the medial temporal lobe: a simple network model. *Proc Natl Acad Sci USA* 91:7041–7045.
- Barnes JM, Underwood BJ. 1959. Fate of first-list associations in transfer theory. *J Exp Psychol* 58:97–105.
- Bunsey M, Eichenbaum H. 1996. Conservation of hippocampal memory function in rats and humans. *Nature* 379:255.
- Burgess N, O'Keefe J. 1996. Neuronal computations underlying the firing of place cells and their role in navigation. *Hippocampus* 6:749–762.
- Chun MM, Phelps EA. 1999. Memory deficits for implicit contextual information in amnesic subjects with hippocampal damage. *Nat Neurosci* 2:844–847.
- Cohen NJ, Eichenbaum H. 1993. *Memory, amnesia, and the hippocampal system*. Cambridge, MA: MIT Press.
- Collingridge GL, Bliss TVP. 1987. NMDA receptors—their role in long-term potentiation. *Trends Neurosci* 10:288–293.
- Crick FHC. 1989. The recent excitement about neural networks. *Nature* 337:129–132.
- Dusek JA, Eichenbaum H. 1997. The hippocampus and memory for orderly stimulus relations. *Proc Natl Acad Sci USA* 94:7109–7114.
- Eichenbaum H. 1992. The hippocampal system and declarative memory in animals. *J Cogn Neurosci* 4:217–231.
- Fanselow MS. 1990. Factors governing one-trial contextual conditioning. *Anim Learn Behav* 18:264–270.
- French RM. 1992. Semi-distributed representations and catastrophic forgetting in connectionist networks. *Connect Sci* 4:365–377.
- Gallagher M, Holland PC. 1992. Preserved configural learning and spatial learning impairment in rats with hippocampal damage. *Hippocampus* 2:81–88.
- Glisky EL, Schacter DL, Tulving E. 1986. Computer learning by memory-impaired patients: acquisition and retention of complex knowledge. *Neuropsychologia* 24:313–328.
- Gluck MA, Myers CE. 1993. Hippocampal mediation of stimulus representation: a computational theory. *Hippocampus* 3:491–516.
- Godden DR, Baddeley AD. 1975. Context-dependent memory in two natural environments: on land and under water. *Br J Psychol* 66:325–331.
- Good M, Bannerman D. 1997. Differential effects of ibotenic acid lesions of the hippocampus and blockade of n-methyl-d-aspartate receptor-dependent long-term potentiation on contextual processing in rats. *Behav Neurosci* 111:1171.
- Hall G, Honey RC. 1990. Context-specific conditioning in the conditioned-emotional-response procedure. *J Exp Psychol Anim Behav Proc* 16:271–278.
- Hasselmo ME, 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* 67:1–27.
- Hebb DO. 1949. *The organization of behavior*. New York: Wiley.
- Holdstock JS, Mayes AR, Roberts N, Cezayirli E, Isaac CL, O'Reilly RC, Norman KA. 2000. Memory dissociations following human hippocampal damage. *Hippocampus*. In press.
- Honey RC, Good M. 1993. Selective hippocampal lesions abolish the contextual specificity of latent inhibition and conditioning. *Behav Neurosci* 107:23–33.
- Honey RC, Willis A, Hall G. 1990. Context specificity in pigeon autoshaping. *Learn Motiv* 21:125–136.
- Honey RC, Watt A, Good M. 1998. Hippocampal lesions disrupt an associative mismatch process. *J Neurosci* 18:2226.
- Jacoby LL, Yonelinas AP, Jennings JM. 1997. The relation between conscious and unconscious (automatic) influences: a declaration of independence. In: Cohen JD, Schooler JW, editors. *Scientific approaches to consciousness*. Mahway, NJ: Lawrence Erlbaum Associates. p 13–47.
- Kiernan MJ, Westbrook RF. 1993. Effects of exposure to a to-be-shocked environment upon the rat's freezing response: evidence for facilitation, latent inhibition, and perceptual learning. *Q J Psychol* 46:271–288.
- Kim JJ, Fanselow MS. 1992. Modality-specific retrograde amnesia of fear. *Science* 256:675–677.
- Kortge CA. 1993. Episodic memory in connectionist networks. In: *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum. p 764–771.
- Maren S, Aharonov G, Fanselow MS. 1997. Neurotoxic lesions of the dorsal hippocampus and pavlovian fear conditioning. *Behav Brain Res* 88:261–274.
- Marr D. 1971. Simple memory: a theory for archicortex. *Philos Trans R Soc Lond [Biol]* 262:23–81.
- Mayes AR, MacDonald C, Donlan L, Pears J. 1992. Amnesics have a disproportionately severe memory deficit for interactive context. *Q J Exp Psychol* 45:265–297.
- McClelland JL, McNaughton BL, O'Reilly RC. 1995. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol Rev* 102:419–457.
- McCloskey M, Cohen NJ. 1989. Catastrophic interference in connectionist networks: the sequential learning problem. In: Bower GH, editor. *The psychology of learning and motivation*, volume 24. San Diego: Academic Press, Inc. p 109–164.
- McDonald RJ, Murphy RA, Guarraci FA, Gortler JR, White H, Baker AG. 1997. Systemic comparison of the effects of hippocampal and fornix-fimbria lesions on the acquisition of three configural discriminations. *Hippocampus* 7:371–388.
- McNaughton BL, Morris RGM. 1987. Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends Neurosci* 10:408–415.
- McNaughton BL, Nadel L. 1990. Hebb-marr networks and the neurobiological representation of action in space. In: Gluck MA, Rumelhart DE, editors. *Neuroscience and connectionist theory*. Hillsdale, NJ: Lawrence Erlbaum Associates. p 1–63.
- McRae K, Hetherington PA. 1993. Catastrophic interference is eliminated in pretrained networks. In: *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum. p 723–728.

- Mishkin M, Suzuki W, Gadian DG, Vargha-Khadem F. 1997. Hierarchical organization of cognitive memory. *Philos Trans R Soc Lond [Biol]* 352:1461–1467.
- Mishkin M, Vargha-Khadem F, Gadian DG. 1998. Amnesia and the organization of the hippocampal system. *Hippocampus* 8:212–216.
- Moll M, Miikkulainen R. 1997. Convergence-zone episodic memory: analysis and simulations. *Neural Networks* 10:1017.
- Nadel L, Moscovitch M. 1997. Memory consolidation, retrograde amnesia and the hippocampal complex. *Curr Opin Neurobiol* 7:217.
- Norman KA, O'Reilly RC, Huber DE. 2000. Modeling neocortical contributions to recognition memory. Poster presented at The Cognitive Neuroscience Meeting, San Francisco, 2000.
- O'Keefe J, Nadel L. 1978. *The hippocampus as a cognitive map*. Oxford: Oxford University Press.
- O'Reilly RC. 1996. Biologically plausible error-driven learning using local activation differences: the generalized recirculation algorithm. *Neural Comput* 8:895–938.
- O'Reilly RC. 1998. Six principles for biologically-based computational models of cortical cognition. *Trends Cogn Sci* 2:455–462.
- O'Reilly RC, McClelland JL. 1994. Hippocampal conjunctive encoding, storage, and recall: avoiding a tradeoff. *Hippocampus* 4:661–682.
- O'Reilly RC, Munakata Y. 2000. *Computational explorations in cognitive neuroscience: understanding the mind by simulating the brain*. Cambridge, MA: MIT Press.
- O'Reilly RC, Rudy JW. 2000. Conjunctive representations in learning and memory: principles of cortical and hippocampal function. *Psychol Rev*. In press.
- O'Reilly RC, Norman KA, McClelland JL. 1998. A hippocampal model of recognition memory. In: Jordan MI, Kearns MJ, Solla SA, editors. *Advances in neural information processing systems*, volume 10. Cambridge, MA: MIT Press. p 73–79.
- Phillips RG, LeDoux JE. 1994. Lesions of the dorsal hippocampal formation interfere with background but not foreground contextual fear conditioning. *Learn Mem* 1:34–44.
- Rolls ET. 1990. Principles underlying the representation and storage of information in neuronal networks in the primate hippocampus and cerebral cortex. In: Zornetzer SF, Davis JL, Lau C, editors. *An introduction to neural and electronic networks*. San Diego: Academic Press. p 73–90.
- Rudy JW, O'Reilly RC. 1999. Contextual fear conditioning, conjunctive representations, pattern completion, and the hippocampus. *Behav Neurosci* 113:867–880.
- Rudy JW, Sutherland RJ. 1994. The memory coherence problem, configural associations, and the hippocampal system. In: Schacter DL, Tulving E, editors. *Memory systems 1994*. Cambridge, MA: MIT Press. p 119–146.
- Rudy JW, Sutherland RW. 1995. Configural association theory and the hippocampal formation: an appraisal and reconfiguration. *Hippocampus* 5:375–389.
- Rumelhart DE, Hinton GE, Williams RJ. 1986. Learning representations by back-propagating errors. *Nature* 323:533–536.
- Save E, Poucet B, Foreman N, Buhot N. 1992. Object exploration and reactions to spatial and nonspatial changes in hooded rats following damage to parietal cortex or hippocampal formation. *Behav Neurosci* 106:447–456.
- Schmajuk NA, DiCarlo JJ. 1992. Stimulus configuration, classical conditioning, and hippocampal function. *Psychol Rev* 99:268–305.
- Sherry DF, Schacter DL. 1987. The evolution of multiple memory systems. *Psychol Rev* 94:439–454.
- Slovan SA, Rumelhart DE. 1992. Reducing interference in distributed memories through episodic gating. In: Healy A, Kosslyn S, Shiffrin R, editors. *Essays in honor of W. K. Estes*. Hillsdale, NJ: Erlbaum. p 227–248.
- Squire LR. 1992. Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. *Psychol Rev* 99:195–231.
- Squire LR, Zola SM. 1998. Episodic memory, semantic memory, and amnesia. *Hippocampus* 8:205–211.
- Sutherland RJ, Rudy JW. 1989. Configural association theory: the role of the hippocampal formation in learning, memory, and amnesia. *Psychobiology* 17:129–144.
- Sutherland RJ, Weisend MP, Mumby D, Astur RS, Hanlon FM, Koerner A, Thomas MJ. 2000. Retrograde amnesia after hippocampal damage: recent vs. remote memories in several tasks. *Hippocampus*. In press.
- Touretzky DS, Redish AD. 1996. A theory of rodent navigation based on interacting representations of space. *Hippocampus* 6:247–270.
- Treves A, Rolls ET. 1994. A computational analysis of the role of the hippocampus in memory. *Hippocampus* 4:374–392.
- Vargha-Khadem F, Gadian DG, Watkins KE, Connelly A, Van Paesschen W, Mishkin M. 1997. Differential effects of early hippocampal pathology on episodic and semantic memory. *Science* 277:376.
- White H. 1989. Learning in artificial neural networks: a statistical perspective. *Neural Comput* 1:425–464.
- Wickelgren WA. 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.
- Winocur G. 1990. Anterograde and retrograde amnesia in rats with dorsal hippocampal or dorsomedial thalamic lesions. *Behav Brain Res* 38:145–154.
- Wu X, Baxter RA, Levy WB. 1996. Context codes and the effect of noisy learning on a simplified hippocampal CA3 model. *Biol Cybernet* 74:159.
- Zipser D, Andersen RA. 1988. A backpropagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature* 331:679–684.
- Zola-Morgan S, Squire LR. 1990. The primate hippocampal formation: evidence for a time-limited role in memory storage. *Science* 250:288–290.