Proc. Natl. Acad. Sci. USA
Vol. 91, pp. 7041–7045, July 1994
Neurobiology

Memory consolidation and the medial temporal lobe: A simple network model

PABLO ALVAREZ*† AND LARRY R. SQUIRE*‡§

Departments of *Psychiatry and of §Neurosciences, University of California, San Diego, La Jolla, CA 92093; †Computational Neurobiology Laboratory, The Salk Institute, La Jolla, CA 92037; and ‡Veterans Affairs Medical Center, San Diego, CA 92161

Contributed by Larry R. Squire, April 6, 1994

ABSTRACT **Some forms of memory have been shown to depend on a system of medial temporal lobe structures that includes the hippocampus and the adjacent cortical areas (entorhinal, perirhinal, and parahippocampal cortex). The role of this system is only temporary, however, as indicated by the fact that, after damage to the medial temporal lobe, recent memories are impaired but very remote memories are intact. Here we review the evidence that the medial temporal lobe memory system is involved in a process of consolidation: memories are initially dependent on this system but gradually become established in other areas of the brain. We then review some of the ideas that have been proposed about the phenomenon of consolidation and suggest a synthesis of these views. Finally, we describe a simple neural network model that captures some key features of consolidation.**

The importance of the medial temporal lobe (MTL) for human memory was established in the 1950s, when surgical removal of this region was found to produce a profound and selective memory impairment (1). Subsequent work, using a model of human amnesia in the monkey (2, 3), identified the anatomical components of this brain system—i.e., the hippocampus (together with the dentate gyrus and subiculum) and the adjacent, anatomically related entorhinal, perirhinal, and parahippocampal cortices. Continuing studies in humans, monkeys, and rodents have illuminated how this system contributes to memory functions (for review, see ref. 4).

One important finding is that the MTL (and related structures in the diencephalic midline) is involved in a limited domain of learning and memory. These structures support the capacity for conscious recollections of facts and events (i.e., declarative memory) but are not necessary for various nonconscious (nondeclarative) forms of memory that are expressed through performance, including skills and habits, simple forms of conditioning, and the phenomenon of priming (5–7).

Retrograde Amnesia and Memory Consolidation

A second important finding (and the focus of this article) is that the role of the MTL is only temporary. This conclusion is based on the fact that damage to the MTL produces temporally graded retrograde amnesia. Memory for events that occurred a short time before the damage is impaired, but memory for remote events (events that occurred a long time before the damage) is spared. This phenomenon was recognized more than a century ago as central to understanding memory and the brain (8). Since that time, temporally graded retrograde amnesia has been observed repeatedly in rodents and humans using a variety of memory tasks and disruptive treatments (9–11). Observations of temporally graded retrograde amnesia led to the idea of memory consolidation: as time passes, the neural substrate of memory is gradually changed or reorganized in a way that makes memory resistant to disruption (for early versions of this idea, see refs. 12–15). Until recently the notion of consolidation could not be clearly linked to particular brain structures.

A link between consolidation and the MTL was suggested by the observation that the severely amnesic patient H.M. had good memory for the events of his early life (1). However, prospective studies in experimental animals with MTL damage were needed to demonstrate that this brain system is involved in memory consolidation. If consolidation does occur, one would expect memory for recent events to be significantly *worse* after MTL damage than memory for remote events (this argument is presented in detail in refs. 4 and 16). This point has now been established with several different species and tasks (17–20).

Thus, as time passes after learning, there must be gradual reorganization (consolidation) of memory storage, whereby memories that are initially dependent on the MTL eventually do not require this system. A more permanent memory that is independent of the MTL develops gradually, presumably in neocortex (2, 6, 21).

An alternative view is that consolidation occurs gradually but independently of the MTL (e.g., in neocortex). The MTL might be needed simply to permit retrieval until the consolidation process is completed. If this were the case, then premorbid memories lost as the result of MTL damage should recover with time as the memories become older and the neocortical consolidation process develops. Yet, this does not occur. In patients with MTL lesions, there is no evidence for recovery from retrograde amnesia as time passes. The MTL must therefore be actively participating in consolidation. One can begin with the idea that the neocortex is the permanent repository of long-term memory (2, 6, 21, 22). Memory consolidation would, therefore, involve an interaction between the MTL and neocortex.

Interactions Between the MTL and Neocortex: A Summary of Existing Views

Several proposals have been developed about memory consolidation and interactions between the MTL and neocortex. Many of these have focused specifically on the function of the hippocampus. [Proposals that focus on the role of the hippocampus in establishing a particular kind of representation of the information to be stored (23, 24), but do not address consolidation itself, are not considered here.] Marr (25) suggested that the hippocampus could act as a temporary memory store for the instantaneous storage of new data, a "simple memory," while the neocortex served as a permanent memory store. He proposed that neocortical inputs converge on hippocampal cells, which then form the representation of a memory. How the hippocampal representation

The publication costs of this article were defrayed in part by page charge payment. This article must therefore be hereby marked "*advertisement*" in accordance with 18 U.S.C. §1734 solely to indicate this fact.

Abbreviation: MTL, medial temporal lobe.

would interact with the neocortex to mediate consolidation was not discussed.

Later, it was proposed that the hippocampus is necessary for memory formation, but that its role is not to store memories. Instead, its role is to enable "chunking" (the formation of novel concepts and associations) to occur in the neocortex (26). Alternatively, it was suggested that the hippocampus does not store information, but rather acts as an "orienting system" that signals to neocortex the need to form a new representation (27, 28). Another proposal, based on studies of nonhuman primates (2, 29), is that the central representation of a stimulus object is stored in the higher-order cortical association areas, such as visual area TE. The role of the MTL is to "imprint" a memory or to serve as a "rehearsal" mechanism, strengthening the connections between cortical areas so that the same cells in visual area TE and in other areas would be activated if an experience (or a portion of it) were repeated. In this view, the MTL does not itself store any information but serves to imprint the memories in neocortex. This proposal did not address the time frame within which this process occurs or the idea of consolidation itself.

The idea that the MTL is actively involved in a gradual process of consolidation has been emphasized several times. In one view, information is embodied both in neocortical representations and in the interactions between the MTL and neocortex (16). The MTL is involved in maintaining the coherence of a memory (linking together the different cortical sites that together represent a whole memory) and is necessary for both memory storage and retrieval. Consolidation occurs gradually as a result of processes intrinsic to neocortex that require the operation of the MTL. Eventually, consolidation strengthens the connections between the different cortical sites enough to support the coherence of a memory without the help of the MTL. In another view, the MTL acts as a temporary memory store, and more permanent memory is stored in neocortex (30). This proposal suggests long-term potentiation (LTP) as the mechanism underlying rapid hippocampal plasticity. Cells in neocortex that represent an event are activated and, in turn, activate hippocampal cells. Cyclic activation of neocortical and hippocampal cells is needed to establish the memory trace initially. Repeated retrieval of an event over time will strengthen the connections between the neocortical cells that represent the event, eventually eliminating the need for the hippocampus to link them together.

Another proposal is that the hippocampus acts as a memory "index" (31). An event produces a pattern of activation in neocortex. The neocortical cells, in turn, activate a group of hippocampal cells that become linked together. These hippocampal cells then act as a retrieval index for the original pattern of cortical activation. This theory presupposes specific bidirectional connections from hippocampus to cortex. With respect to consolidation, it is supposed that the hippocampus might periodically reactivate the cortical representation, which would then change over some extended time course. Peter Milner (32) proposed a basic distinction between the operating principles of hippocampus and neocortex. The hippocampus acts as a temporary memory store because hippocampal synapses are "soft," capable of changing quickly, both during learning and forgetting. By contrast, neocortical synapses are "hard," change slowly in response to new stimuli, and also weaken slowly. The hippocampal synapses are used to reactivate the neocortical representation during recall. Consolidation occurs by repeated reactivation of the cortical representation, which eventually creates hard links in neocortex that can subserve memory even when the hippocampal synapses are lost (32).

McClelland and colleagues (33–35) approached the idea of consolidation computationally. They began with many of the ideas already mentioned. The neocortex itself supports slow, gradual learning, and the hippocampus is necessary for fast, one-trial learning. Their key proposal is that the brain is organized this way for particular computational reasons. The hippocampus can acquire information quickly, but consolidation is slow so that the neocortex can change in a more gradual way, incorporating into its representations not only the elements of one experience but also the regularities of the environment that encompass many experiences. Events that are to be permanently stored in memory are gradually incorporated into an already existing framework, and as a result the framework is extended. The framework in neocortex can be modified gradually, but it would be unstable if it incorporated new information too quickly.

The proposals just summarized all include the idea that long-term memory storage in the neocortex depends in some way on the participation of the hippocampus (or the MTL). However, these proposals are often not explicit on several key points. Does the hippocampus specifically direct consolidation or does it enable it in some more diffuse, nonspecific way? What is required for consolidation to occur (overt rehearsal or endogenous activity within the hippocampus)? How does activity within the hippocampus cause disparate cortical areas to become linked together? What actually is stored in the hippocampus (an index or the representation of an event)?

In this article, we extend an earlier proposal from this laboratory (16), building especially on the two views just summarized (32, 34). The current proposal introduces a simple quantitative model that embodies the essential principles of consolidation and allows the relevant ideas to become more specific and concrete. The model shows that this proposal is internally consistent and provides a hypothesis about MTL function and consolidation that is specific enough to guide experimental work.

Memory Consolidation: A Proposal

Our proposal is that the MTL memory system (i.e., the hippocampus together with the adjacent entorhinal, perirhinal, and parahippocampal cortices, ref. 36) serves as a temporary memory store and that the neocortex is the permanent repository of long-term memory. Long-term memory is stored distributedly within the same higher-order cortical association areas that are specialized for processing and analyzing the particular kinds of information that are to be remembered. Each of the different specialized cortical areas contributes differently to memory storage, and all the areas participate together to store a whole memory. Thus, a major task for the formation and maintenance of long-term memory is the binding together of the geographically disparate areas that together constitute a complete memory.

During both learning and forgetting, plastic changes in the connections within the MTL and between the MTL and neocortex occur more rapidly than in the connections between different areas of neocortex. The MTL directs the recall of recent experiences by binding together the multiple cortical sites that constitute the representation. When a subset of the neurons within the neocortical representation are reactivated, they activate neurons within the MTL that are part of a strongly connected network unique to that stored event and that are able to revivify the complete neocortical representation. Consolidation occurs in a similar manner. Whenever the neocortical representation is reactivated, usually with the help of the MTL, functional connections are gradually established between the various constituents of the neocortical representation. In this way, the neocortex itself can eventually reconstruct the representation from partial cues, and the MTL is not required. Thus, with the passage of time after learning, the burden of long-term memory storage

is gradually assumed by neocortex, and the MTL memory system is always available for new learning.

We can summarize the present proposal in five statements. (*i*) The crucial event for the formation, maintenance, and retrieval of long-term declarative memory is an interaction between multiple, geographically separated areas of neocortex, and the structures of the MTL. (*ii*) The neocortex communicates with the MTL via reciprocal connections with entorhinal, perirhinal, and parahippocampal cortices. The latter two areas, in turn, are reciprocally connected to entorhinal cortex, which communicates with the hippocampus. (*iii*) Within the neocortex, the key event in consolidation is the gradual binding together of the multiple, geographically disparate cortical regions that together store the representation of a whole event. This gradual linking is the biological substrate of consolidation. (*iv*) The MTL learns quickly but has limited capacity. The neocortex learns slowly (i.e., disparate regions become bound together slowly) and has a large capacity. In both cases, learning proceeds according to the same simple (Hebbian) rules for changing synaptic strength. (*v*) Consolidation occurs when neural activity within the MTL coactivates separate regions of neocortex. These areas of neocortex are initially linked only weakly but become more strongly connected as a function of repeatedly being activated simultaneously by the MTL.

A Simple Model of Memory Consolidation

To make these ideas more specific, we next describe a simple connectionist model that implements the basic ideas just outlined. Information is first stored in a fast-learning "MTL" area, which then participates in gradually strengthening slower-changing connections in a geographically distributed "neocortical" network.

The model consists of two "cortical" areas that are reciprocally interconnected with one MTL area. Each of the two cortical areas consists of eight units, whereas the MTL area consists of four units (Fig. 1). Each unit represents a simplified neuron. Along the lines discussed by Rumelhart *et al.* (37), the activity a_i of each unit i is calculated as follows:

$$a_i = \delta a_i + \Sigma a_j w_{i,j} + \varepsilon, \quad [1]$$

where δ determines the decay of activation over time, a_j is the activation level of unit j, $w_{i,j}$ is the strength of the connection from unit j to unit i, and ε is a noise term. At each time step, activations are first updated synchronously in the cortical areas and then updated synchronously in the MTL area. All inputs are excitatory, and inhibition is simulated using a winner-take-all arrangement. That is, units are arranged in competitive groups of four, such that the most active unit in each group retains its current level of activation, whereas the activations of the other three units become zero.

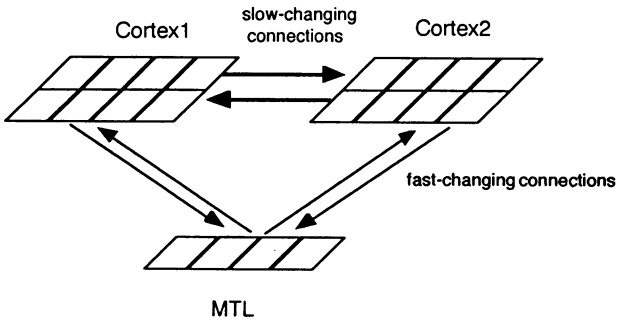

FIG. 1. Schematic diagram of the model. Areas cortex1 and cortex2 represent association neocortex. Each unit in each of the areas (4 in MTL and 8 in cortex1 and cortex2) is reciprocally connected to each unit in the other areas. There are no connections within areas, only a form of winner-take-all inhibition. A key feature of the model is that connections to and from the MTL area (thin lines) change much faster than connections between the two cortical areas (thick lines).

The strengths of the connections change after each time step according to a modified Hebbian, competitive learning mechanism (38, 39). Specifically, the strength of the connection from unit j to unit i changes only if i is active. If unit j is making a large contribution to the activity of unit i (i.e., its activity is greater than the mean of the activities of all the units connected to i), the connection from j to i is strengthened. If unit j makes only a small contribution to the firing of unit i (i.e., its activity level is smaller than the mean of all the units connected to i), the connection from j to i is weakened. The learning rule is as follows:

$$\Delta w_{i,j} = \lambda a_i (a_j - \bar{a}), \quad [2]$$

where $w_{i,j}$ is the change in the strength of the connection, $\bar{a}$ is the mean activation level of all units projecting to unit i, and λ is the learning rate.

Forgetting is simulated by reducing all connection strengths in proportion to their current strength (producing exponential forgetting). Strong connections decline faster than weaker connections. The forgetting rule is

$$\Delta w_{i,j} = -\rho w_{i,j}, \quad [3]$$

where ρ is the forgetting rate.

Unit activations and connection weights are both constrained by the limits (0, 1). A key feature of the model is that the connections to and from the MTL area change more rapidly than the connections between the cortical areas. That is, both the learning rate (λ) and the forgetting rate (ρ) are greater by an order of magnitude for the connections between cortex and MTL than for the "cortico–cortical" connections (see Fig. 2 for values). Thus, when the network is presented with a stimulus (the pattern AB, for example), changes in the cortico–MTL connections are sufficient after only a few presentations to allow for recall of the complete pattern (AB) when only part of it (i.e., A, the part of the pattern corresponding to cortex1, or B, the part of the pattern corresponding to cortex2) is presented to the network. Many more presentations of the stimulus are needed for the cortico–cortical connections to change enough to support such recall on their own.

The task of the network was a simple one: to reconstruct from incomplete cues two patterns that had been previously trained. Each pattern consisted of two active units in each of the cortical areas (one unit in each competitive group of four). To simplify the task, the patterns used were orthogonal to each other—i.e., nonoverlapping. Training involved presenting a pattern twice to the network. For each presentation, external inputs to the four cortical units that were to represent the trained pattern were first set to 1. Activation then cycled through the network for three time steps. At each time step, activations were updated in cortex and MTL, and connection strengths were updated accordingly. The network was then tested by presenting only half of the pattern (i.e., the half corresponding to activity in one of the two cortical areas). For the test, activation was again allowed to cycle through the network for three time steps, and the final pattern of activity in the cortical areas was compared with the complete pattern. The error in recalling the pattern was calculated as

$$\text{error} = \Sigma (a_i - p_i)^2, \quad [4]$$

where p_i is the activation level of unit i in the original, complete pattern. To obtain a measure of the network's

overall performance (total error), the network was tested on each of the four possible half patterns, and the error was summed across the four test patterns.

As expected from the model's architecture, the learning that occurs in the cortico–MTL connections can support good recall of the original patterns after only a few presentations. That is, memories can be stored in the strengths of the connections between the cortical and MTL areas. At this initial stage the memory cannot be reconstructed if the network is "lesioned" by inactivating the MTL.

However, activity in the MTL area can reactivate the stored patterns in cortex and allow for the strengthening of the corresponding cortico–cortical connections. Random activity in the MTL area (i.e., activation of a randomly selected MTL unit) was used to simulate the processes that drive consolidation. When random activity is allowed to occur repeatedly in the MTL, the cortico–cortical connections can over time become sufficiently strengthened so that patterns that have been learned can become independent of the MTL. That is, the network can reconstruct the complete pattern AB from the presentation of only A (the part of the pattern corresponding to cortex1) or B (the part of the pattern corresponding to cortex2), even if the network has been "lesioned" by disconnecting the MTL area from the cortical areas. In this way, the model can simulate the retrograde amnesia gradient observed after MTL lesions.

Fig. 2 shows the result of testing the network in the manner just described. Patterns were presented twice to train the network. After the training, various amounts of time were allowed for consolidation. For each unit of consolidation time, the activation of one randomly selected MTL unit was set to 1 (i.e., its maximal value). Activation was then allowed to cycle through the network for three time steps. The network was tested by presenting all four half patterns as described before Eq. 4, both in the "normal" state and in a "lesioned" state in which the connections to and from the MTL were inactive. Fig. 2 shows the error in reconstructing the patterns based on how much time had passed between

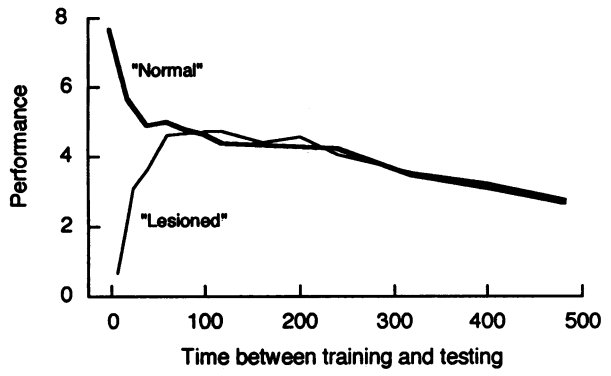

FIG. 2. Performance of the model in a retrograde amnesia experiment. The model is trained with two patterns in cortex and tested at different intervals after training for its ability to recreate the patterns based on partial input. The abscissa is a measure of the time (in consolidation units, see text) allowed for consolidation. The ordinate is a measure of the network's performance, measured as 8 − total error (see text). The intact network (thick line) shows a normal forgetting curve. The "lesioned" network (thin line) had its MTL area disconnected from the cortical areas immediately before testing. Its performance is qualitatively similar to that of amnesic patients and animals with MTL damage: memory for recently learned material is worse than memory for material learned long before the lesion. The performance data shown are averaged over 50 separate simulations. Parameters used were: $\lambda = 0.1$ for MTL, $\lambda = 0.002$ for cortex; $\rho = 0.04$ for MTL, $\rho = 0.0008$ for cortex; $\delta = 0.7$; ε was uniformly distributed between -0.05 and $+0.05$. All connection strengths were initialized to uniformly distributed random values from 0.0 to 0.2.

training and the lesion (each unit of time corresponds to three time steps of activity). "Normal" networks can reconstruct recently learned patterns better than ones learned more remotely. The effect of the "lesion" is similar to that observed in experimental studies: performance is better for patterns learned at a remote time period than for patterns learned more recently.

Thus, the model is capable of producing behavior similar to that observed in lesion experiments, and it exhibits a phenomenon like consolidation. The central idea is that the MTL quickly learns to bind together different portions of a cortical representation, and consolidation is due to slower learning in neocortex that is driven by the MTL representation.

Issues Raised by the Model

Making this proposal specific enough to be implemented in a simple, quantitative model addresses some of the questions about consolidation posed earlier. In the model, the MTL directs consolidation in a specific manner. Consolidation is initiated by an endogenous event—random activity within the MTL. This activity excites the ensembles of neocortical cells that together form a representation of the original event. As a result, the cells in different neocortical areas are active concurrently, and the connections between them slowly become stronger. Eventually links are formed that are independent of the MTL.

The model also addresses the nature and extent of initial information storage in MTL. On the one hand, the MTL has been proposed to have only a nonspecific modulatory role in memory formation (26, 28, 40); on the other hand, the MTL could store information that specifies which neocortical cells form the memory representation (31, 32). The model proposed here falls within the second category, in that specific information about the cortical patterns that have been learned is stored in the connections between neocortex and the MTL.

Within this second category of explanation, a further distinction can be made. On the one hand, the MTL may simply store an "index" of the neocortical cells or areas that need to be reactivated to reconstruct a memory (31). On the other hand, the MTL could initially contain the memory representation itself (34, 41). The model described here is an example of the index view, in that the role of the MTL is to point to and activate the relevant neocortical cells. However, in the model, the MTL can also be considered to store the memory because it is capable of completely reconstructing it. Accordingly, in the model, the distinction between indexing and storage is unnecessary.

Simply eliminating this distinction does not address, however, an important question underlying it: how much information can the MTL hold? It seems reasonable to suppose that it does not hold all of the details of all unconsolidated memories. One answer to this question comes from the proposal (21) that representations of events consist of the simultaneous activation of neuronal ensembles in multiple cortical areas and that this activation is coordinated by a hierarchy of convergence zones. If the MTL is viewed as the highest level of the convergence zone hierarchy, then it may store (or index) only some of the higher-order features of a memory, while lower-order details are the responsibility of neocortical convergence zones.

Another question concerns the nature of the connections between MTL and neocortex. Neuroanatomical studies show that the connections between MTL and neocortex, and within the MTL, are not organized in a precise topographic manner (42). How could these connections be specific enough to reconstruct the neocortical representation of a memory from a partial input? In particular, how does activation of the MTL by neocortex find its way back to a specific group of cells in neocortex? In the model, the connections are

not topographic: initially, there are weak connections between all cortical cells and all MTL cells. Specific connections are formed during initial learning in a manner similar to that proposed by Halgren (30). Some MTL cells will initially have slightly stronger reciprocal connections with the cortical cells that constitute the learned patterns, and these cells will receive more input than other MTL cells. Then, through continued activity, the connections between the cortical cells in the pattern and these MTL cells become even stronger. In the model, there is no need for *a priori* specificity: it is established during the initial learning of the patterns. Selection occurs as the result of continuing activity through a recurrent network.

There are a number of questions that the model does not address (e.g., exactly what kind of representation is established in the MTL and exactly how connections are changed and neurons are activated). Nevertheless, the model is still useful in identifying and defining these and other questions. As an example, consider the question of when consolidation occurs. Is consolidation occurring constantly, which would imply that old memories in neocortex are constantly being revived as part of normal brain activity? Or is there, in fact, a special state during which most consolidation occurs? In the first scenario, an explanation is needed for why the consolidation process appears not to intrude on consciousness. In the second scenario, this state and the characteristics that make it optimal for consolidation need to be identified. Recent proposals include the idea that consolidation occurs during rapid-eye-movement (REM) sleep (43) or during hippocampal sharp waves (i.e., during slow-wave sleep or quiet alertness, 44).

Summary and Conclusions

Declarative memories depend on the MTL for a limited period of time after learning. We propose that the MTL performs the function of binding together the different neocortical portions of a memory representation. The MTL stores this information in a rapid, labile manner, whereas changes in the neocortex are slower and longer lasting. The information in MTL directs the slow changes in neocortex that underlie consolidation. We show that a simple network model based on this proposal behaves in a way consistent with experimental observation. The model has limitations that include its small size, the fact that patterns do not have distributed representations, and the unrealistically high density of connections, especially between neocortical areas. It has the advantages of being simple and of having clearly defined properties. As a consequence, it can address several current questions about consolidation, and it provides a reference point for further experimental investigation and a starting point for improved models.

This work was supported by the Medical Research Service of the Department of Veterans Affairs, National Institutes of Health Grant NS19063, The Office of Naval Research, the McKnight Foundation, and the McDonnell-Pew Center for Cognitive Neuroscience. Dr. Kyle Cave, now at Vanderbilt University, contributed to the early development of the model. We are also grateful to Drs. Stuart Zola-Morgan, Jay McClelland, and Terrence Sejnowski for their comments.

1. Scoville, W. B. & Milner, B. (1957) *J. Neurol. Neurosurg. Psychiatr.* **20**, 11–21.
2. Mishkin, M. (1982) *Phil. Trans. R. Soc. London B* **298**, 85–92.
3. Squire, L. R. & Zola-Morgan, S. (1983) in *The Physiological Basis of Memory*, ed. Deutsch, J. A. (Academic, New York), pp. 199–268.
4. Squire, L. R. (1992) *Psychol. Rev.* **99**, 195–231.
5. Schacter, D. L. (1987) *J. Exp. Psychol. Learn. Mem. Cogn.* **13**, 501–518.
6. Squire, L. R. (1987) *Memory and Brain* (Oxford Univ. Press, New York).
7. Schacter, D. & Tulving, E. (1994) *Memory Systems 1994* (MIT Press, Cambridge, MA).
8. Ribot, T. (1881) *Les Maladies de la Memoire [English translation: Diseases of Memory]* (Appleton-Century-Crofts, New York).
9. McGaugh, J. L. & Gold, P. E. (1976) in *Neural Mechanisms of Learning and Memory*, eds. Rosenzweig, M. R. & Bennett, E. L. (MIT Press, Cambridge, MA).
10. Squire, L. R., Slater, P. C. & Chace, P. M. (1975) *Science* **187**, 77–79.
11. Squire, L. R. & Spanis, C. W. (1984) *Behav. Neurosci.* **98**, 345–348.
12. Muller, G. E. & Pilzecker, A. (1900) *Z. Psychol. Erganzungsband* **1**, 1.
13. Burnham, W. H. (1903) *Am. J. Psychol.* **14**, 382–396.
14. Glickman, S. (1960) *Psychol. Bull.* **53**, 322.
15. McGaugh, J. L. & Herz, M. J. (1972) *Memory Consolidation* (Albion, San Francisco).
16. Squire, L. R., Cohen, N. J. & Nadel, L. (1984) in *Memory Consolidation*, eds. Weingartner, E. & Parker, E. (Lawrence Erlbaum Associates, Hillsdale, NJ), pp. 185–210.
17. Zola-Morgan, S. & Squire, L. R. (1990) *Science* **250**, 288–290.
18. Winocur, G. (1990) *Behav. Brain Res.* **38**, 145–154.
19. Kim, J. J. & Fanselow, M. S. (1992) *Science* **256**, 675–677.
20. Cho, Y. H., Beracochea, D. & Jaffard, R. (1993) *J. Neurosci.* **13**, 1759–1766.
21. Damasio, A. R. (1989) *Cognition* **33**, 25–62.
22. Singer, W. (1990) *Concepts Neurosci.* **1**, 1–26.
23. Rolls, E. (1990) in *An Introduction to Neural and Electronic Networks*, eds. Zornetzer, S. F., Davis, J. L. & Lau, C. (Academic, San Diego), pp. 73–90.
24. Gluck, M. A. & Myers, C. E. (1993) *Hippocampus* **3**, 491–516.
25. Marr, D. (1971) *Phil. Trans. R. Soc. London, Series B* **262**, 23–81.
26. Wickelgren, W. A. (1979) *Psychol. Rev.* **86**, 44–60.
27. Grossberg, S. (1976) *Biol. Cybernet.* **23**, 121–136.
28. Carpenter, G. A. & Grossberg, S. (1987) in *The Adaptive Brain I: Cognition, Learning, Reinforcement, and Rhythm*, ed. Grossberg, S. (Elsevier, Amsterdam), pp. 238–286.
29. Mishkin, M. (1990) in *Vision, Memory and the Temporal Lobe*, eds. Iwai, E. & Mishkin, M. (Elsevier, New York), pp. 7–15.
30. Halgren, E. (1984) in *The Neuropsychology of Memory*, eds. Squire, L. R. & Butters, N. (Guilford, New York), pp. 165–182.
31. Teyler, T. J. & Discenna, P. (1986) *Behav. Neurosci.* **100**, 147–154.
32. Milner, P. (1989) *Neuropsychologia* **27**, 23–30.
33. McClelland, J. L., McNaughton, B. L., O'Reilly, R. & Nadel, L. (1992) *Soc. Neurosci.* **18**, 1216 (abstr.).
34. McClelland, J. L., McNaughton, B. L. & O'Reilly, R. C. (1994) *Technical Report PDP.CNS.94.1* (Carnegie Mellon Univ., Pittsburgh, PA).
35. McClelland, J. L. (1994) *Rev. Neurol. (Paris)*, in press.
36. Squire, L. R. & Zola-Morgan, M. (1991) *Science* **253**, 1380–1386.
37. Rumelhart, D. E., Hinton, G. E. & McClelland, J. L. (1986) in *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, eds. Rumelhart, D. E. & McClelland, J. L. (MIT Press, Cambridge, MA), pp. 45–76.
38. Hebb, D. O. (1949) *The Organization of Behavior* (Wiley, New York).
39. Rumelhart, D. E. & Zipser, D. (1986) in *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, eds. Rumelhart, D. E. & McClelland, J. L. (MIT Press, Cambridge, MA), pp. 151–193.
40. McClelland, J. L., & Rumelhart, D. E. (1986) in *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, eds. McClelland, J. L. & Rumelhart, D. E. (MIT Press, Cambridge, MA), pp. 503–528.
41. McNaughton, B. L. & Morris, R. G. M. (1987) *Trends Neuro-Sci.* **10**, 408–415.
42. Squire, L. R., Shimamura, A. P., and Amaral, D. G. (1989) in *Neural Models of Plasticity*, eds. Byrne, J. & Berry, W. (Academic, San Diego), pp. 208–239.
43. Winson, J. (1985) *Brain and Psyche* (Doubleday, Garden City, NY).
44. Buzsaki, G., Horvath, Z., Urioste, R., Hetke, J. & Wise, K. (1992) *Science* **256**, 1025–1027.