HIPPOCAMPUS 10:352–368 (2000)

Multiple Trace Theory of Human Memory: Computational, Neuroimaging, and Neuropsychological Results

L. Nadel,[1]* A. Samsonovich,[1] L. Ryan,[1]
and M. Moscovitch[2]

[1]*Department of Psychology, University of Arizona, Tucson, Arizona*
[2]*Department of Psychology, University of Toronto, and Rotman Research Institute, Toronto, Ontario, Canada*

ABSTRACT: Hippocampal-neocortical interactions in memory have typically been characterized within the "standard model" of memory consolidation. In this view, memory storage initially requires hippocampal linking of dispersed neocortical storage sites, but over time this need dissipates, and the hippocampal component is rendered unnecessary. This change in function over time is held to account for the retorgrade amnesia (RA) gradients often seen in patients with hippocampal damage. Recent evidence, however, calls this standard model into question, and we have recently proposed a new approach, the "multiple memory trace" (MMT) theory. In this view, hippocampal ensembles are always involved in storage and retrieval of episodic information, but semantic (gist) information can be established in neocortex, and will survive damage to the hippocampal system if enough time has elapsed. This approach accounts more readily for the very long RA gradients often observed in amnesia. We report the results of analytic and connectionist simulations that demonstrate the feasibility of MMT. We also report a neuroimaging study showing that retrieval of very remote (25-year-old) memories elicits as much activation in hippocampus as retrieval of quite recent memories. Finally, we report new data from the study of patients with temporal lobe damage, using more sensitive measures than previously the case, showing that deficits in both episodic and spatial detail can bed observed even for very remote memories. Overall, these findings indicate that the standard model of memory consolidation, which views the hippocampus as having only a temporary role in memory, is wrong. Instead, the data support the view that for episodic and spatial detail the hippocampal system is always necessary. *Hippocampus 2000;10:352–368.*

KEY WORDS: hippocampus; memory consolidation

INTRODUCTION

The establishment and maintenance of long-term episodic and semantic memory representations engage structures in both the hippocampal complex and neocortex. For 50 years a particular view of how hippocampal and neocortical structures interact has held sway, the current version of which goes something like this:

Information about an episode is encoded within both neocortical and hippocampal neural ensembles.

Within the neocortex, the various parts of a particular episode are represented in physically separated neural ensembles, since they involve different kinds of content (e.g., sights, sounds, smells).

This physical separation creates a "binding" problem with respect to subsequent retrieval of the memory for this episode.

The hippocampal complex plays a critical role in solving this binding problem, by providing a mechanism to link the physically separated neocortical fragments.

This view, articulated by Squire et al. (1984), and Teyler and DiScenna (1986), was given a computational interpretation by McClelland et al. (1995). In their view, the hippocampal component allowed for rapid learning that could then provide for slower learning within the neocortex. They argued that slow learning was essential within the neocortex, because rapid alterations in the synaptic weights comprising its memory stores would have the effect of eliminating previously stored information. In order to add new information without having this catastrophic effect, such new learning must occur in a slow, incremental fashion. In the hippocampus, by contrast, representations are formed in such a way that catastrophic interference of this sort is less likely. Here, similar representations are stored in ensembles that are chosen to be orthogonal to each other. This pattern-separation process has the effect of limiting interference, and it allows the hippocampus to rapidly acquire new memory traces without losing old ones, at least in the short run. Superimposed on this view of memory are at least two critical issues that remain to be resolved. First, there is the matter of the *kind* of memory one is talking about. Most relevant here is the presumed distinction

Grant sponsor: University of Arizona Cognitive Neuroscience Program and the Cognition and Neuroimaging Laboratories; Grant sponsor: McDonnell-Pew Cognitive Neuroscience Program; Grant sponsor: Flinn Foundation; Grant sponsor: State of Arizona Alzheimer's Disease Research Center; Grant sponsor: Center for the Study of Consciousness, University of Arizona; Grant sponsor: Natural Sciences and Engineering Research Council of Canada; Grant number: A8437.
*Correspondence to: L. Nadel, Department of Psychology, University of Arizona, Tucson, AZ.
Accepted for publication 1 May 2000

between episodic and semantic memory, first articulated by Tulving (1972). Second, there is the matter of how the interactions between hippocampal and neocortical structures play out over the lifetime of a memory. On these matters there are many ideas, but little consensus.

Within the traditional model (e.g., Squire, 1992; McClelland et al., 1995), episodic and semantic memory are treated as more or less equivalent with respect to their neural underpinnings. It is assumed that both forms of information initially require the involvement of hippocampal circuits, in the way noted above, but that over time both are stabilized within neocortical circuits in a manner that renders the hippocampal complex unnecessary for subsequent retrieval. This time interval, typically termed the "consolidation" period, is estimated to last anywhere from weeks to years, depending on the species and the kind of memory involved. We recently reviewed the rather extensive literature bearing on these questions, and have come to a somewhat different conclusion (Nadel and Moscovitch, 1997, 1998; Moscovitch and Nadel, 1998; Fujii et al., 2000). Put most succinctly, we have proposed, in contrast to standard consolidation theory, that the hippocampal complex is always involved in the storage and retrieval of episodic memories, independent of their age. On the other hand, and consistent with standard theory, we believe that long-term interactions between the hippocampal complex and neocortex can influence the stabilization of semantic memories within neocortex. Thus, our position emphasizes a critical distinction between episodic and semantic memory, both of which are influenced by the hippocampal complex, but only one of which becomes independent of this brain region in the course of consolidation (Tulving, 1972; Kinsbourne and Wood, 1975).

In what follows, we briefly review some of the literature that led us to this position. We then discuss a computational model of multiple trace theory (MTT), an analysis of remote memory for autobiographical and semantic memory in patients with damage to the hippocampal complex, some issues related to the scoring of remote memory retrieval, a case study of remote spatial memory, and, finally, data from a functional neuroimaging (fMRI) study. The overall picture presented by these data is one of agreement with the fundamental tenet of MTT, that the hippocampal complex is important for normal episodic memory function throughout the life of such memories.

REVIEW OF REMOTE MEMORY LITERATURE

In a recent review of the literature on the effects of temporal lobe lesions on memory, we examined all the published papers reporting patients with focal (nondegenerative) temporal-lobe lesions confirmed radiologically or by autopsy, and whose remote memory was tested formally. In just a few cases, we had to rely on observational reports of remote memory. The results of this review were encouragingly consistent with our model (Fujii et al., 2000).

We found that when the lesion was confined to the hippocampus proper (CA fields and dentate gyrus), there was either a short retrograde amnesia (RA), or almost complete sparing of remote memory, be it autobiographical or personal semantic, or whether it was of public events, personalities, or vocabulary. However, when the lesions went beyond the hippocampus proper, the temporal extent of remote memory loss was related to the size of the lesion, with autobiographical memory being affected most severely, memory for public events and personalities being affected substantially, and personal semantic memory being affected hardly at all. Thus, lesions of the hippocampal formation (hippocampus proper plus subiculum) were associated with autobiographical memory loss extending for about a decade (or more in some cases), with loss of memory for public events and personalities being affected just a little less in some patients and almost not at all in others. If the lesion included the remainder of the hippocampal complex, e.g., the entorhinal and perirhinal cortices, and parahippocampal gyrus, then memory loss extended as far back as 20–30 years, and sometimes longer, for autobiographical memory and, in most cases, for memory for public events and personalities, with involvement of personal semantic memory and vocabulary in some cases. When the lesion also included other parts of the temporal lobes, such as the inferotemporal cortex, loss of autobiographical episodes extended for the entire life. Only personal semantic memory was relatively spared.

Although the relation of extent of remote memory loss to lesion size is consistent with predictions of our model, it is important to note that we do not believe that size alone is the determining factor. More and more evidence is accumulating that each of the regions of the hippocampal complex makes separate contributions either to encoding, storage, or retrieval (Gabrieli et al., 1997; LePage et al., 1998, Schacter and Wagner, 1999), or is crucial to memory representations of a particular type, such as spatial, object, or objects in particular locations (Nadel, 1991; Aguirre et al., 1996; Owen et al., 1996; Epstein and Kanwisher, 1998; Epstein et al., 1999). Thus larger lesions implicate more areas whose contributions are needed to store or recover all the myriad aspects of memory that enter into the representations of a rich autobiographical memory, and even of detailed memories of public events and personalities. Simple semantics, whether personal or generic, do not depend as much on these medial temporal structures as on the neocortex, and so are relatively spared (Graham and Hodges, 1997; Graham, 1999; Hodges et al., 1998; Snowden et al., 1996).

Interpreting this evidence requires caution, taking into consideration the imprecise localization of the lesions, the not always satisfactory assessment of remote memory, the difficulty of distinguishing clearly among different types of memory, and the small number of cases with certain types of lesions, such as those confined exclusively to the hippocampus proper. Nonetheless, in its broad outline, the picture that emerges is closer to MTT as we formulated it than to the traditional model. Based on these analyses, we embarked on a program of research aimed at assessing the viability of MTT. This program included analytic and connectionist simulations of MTT, with the aim of establishing the plausibility of this approach in principle.

COMPUTATIONAL MODELS OF MTT

There are several reasons to pursue a computational model of a particular theoretical position. Such models force one to be explicit about the details of one's theory, and models can sometimes lead to novel understandings about the phenomenon under study. In any modeling enterprise one must balance biological plausibility against the reality of limited time and space on a computer. In seeking to model MTT, we did two things: we started with an analytic model describing MTT at a very general level, and then created a more specific, connectionist interpretation of this model, paying somewhat more attention to the constraints imposed by the actual organization of the neocortex and hippocampal complex, and their interactions.

The analytic model, which will be presented first, states the idea of memory improvement by multiple trace formation in mathematical terms in a very general form, under parsimonious assumptions, with virtually no biological specificity. This, however, does not mean that the model is biologically implausible: on the contrary, the assumptions of this model were chosen with the purpose of retaining biological plausibility. We see a model of this kind as a necessary first step toward understanding the application of MTT to the role of the hippocampus in remote memory and the phenomenon of retrograde amnesia.

AN ANALYTIC SOLUTION

Our general idealized model of multiple trace formation in the hippocampal-neocortical episodic memory system was based on the following assumptions:

1. A dynamical state of the system can be effectively characterized by an ordered, variable-length list of all episodic memories, together with their associated numbers of traces. Each episodic memory has a unique "time-stamp" referring to when it was created (i.e., its age), which can serve as a memory identifier.
2. All memories are created at random moments of time, independently of each other. On average, one memory is created per unit of time (this condition simply specifies our time units). A memory just created has exactly one trace.
3. Traces are considered as elementary units of the model. They have no internal freedoms: a trace either exists or does not exist. Traces can decay (i.e., disappear) and can replicate. By definition, each replica of a trace has the same time stamp as its ancestor. The actual time of creation of a replica is not essential for its further dynamics.
4. Trace decay can be viewed as an independent random process. All traces have the same lifetime, which reflects inherent trace degradation as well as interference with newly created traces. Trace lifetime is assumed to be large as compared to the characteristic time scale.
5. Trace replication is a nonhomogeneous point process, the intensity of which for a given memory is a function of memory age, the number of traces, and the total trace replication rate, which is assumed to be constant over time. Therefore, in general, at any moment of time, different memories may have different probabilities of creating an additional trace (see below for specific choices). In all cases, we assume that "dead" memories (i.e., memories having no traces left) do not replicate.
6. A dead memory cannot be retrieved. If a memory has one trace, then the probability of a failure in memory retrieval is a fixed parameter q that can be viewed as equivalent to a lesion of a particular size in the hippocampus, although q is not necessarily taken as zero for an intact system.
7. If a memory has more than one trace, then the effect of all of them on memory retrieval is additive: any trace may cause retrieval, and the reliability of each trace action is independent of the number of traces. Therefore, failure in retrieval only occurs when all traces of a given memory fail. This implies that the probability P of a successful retrieval for a memory with n traces is given by the formula (see Appendix for an alternative interpretation of this formula)

$$p = 1 - q^n. \tag{1}$$

8. The initial state of the system contains no memories.

Therefore, the expected number n of traces of an existing memory, and the probability P of its successful retrieval, are functions of time, memory age, and the three constant parameters of the model: the total rate of trace replication, the trace lifetime, and the extent of lesion q. Since q only effects P, and not n, and trace decay is taken as negligible, this is essentially a model with one free parameter: the replication rate. Therefore, the biological plausibility of the model is a function of the choice of replication rate and the plausibility of the multiple trace idea itself. Such a model cannot be biologically more specific than the choice that this one parameter allows.

In order to assess MTT in this idealized version, we wanted to determine the average number of traces per memory n and the average probability P of successful retrieval of a memory as a function of the memory age and of the extent of lesions in the system. Further, we wanted to explore the effects of several different ways to approach the trace replication function (see assumption 5 above). When trace replication followed the assumption that all traces had an equal likelihood of replication, regardless of their age or numbers of traces per memory, the result was that the oldest memories formed a great number of traces, "running away" with the replication process: a singularity in the number of traces was observed at the maximal memory age (Figs. 1A, 2A). When, on the other hand, trace replication followed the assumption that all memories (rather than traces) had an equal likelihood of replication, the number of traces at the maximal age became finite; however, a cap was observed, again at the maximal age (Figs. 1B, 2B). These results led us to seek further modifications of the replication function that would produce more realistic effects.

Among possible choices for the replication function (assumption 5), we considered dependence of the specific replication rate on the age of the memory, and on the number of existing traces of

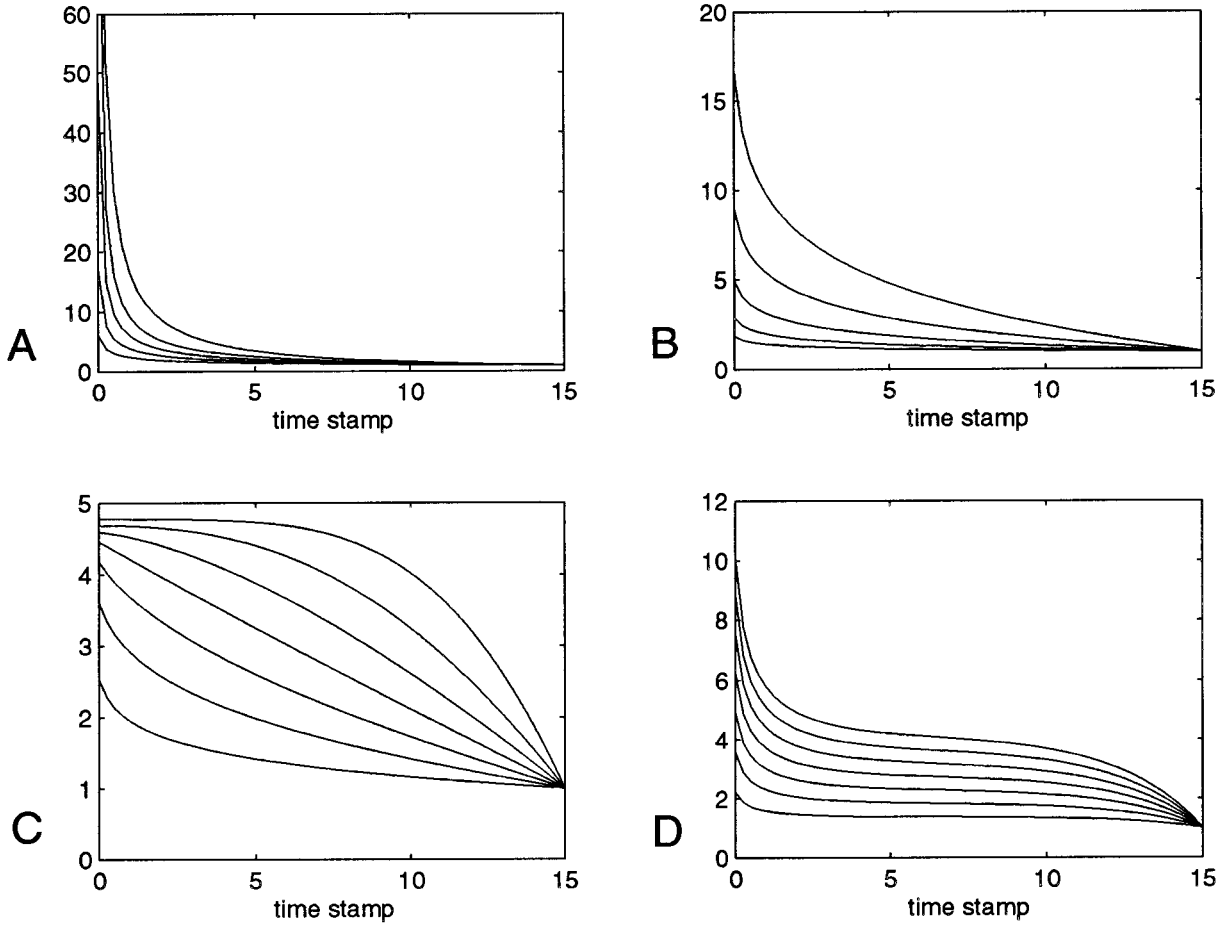

FIGURE 1. Average number of traces per memory as a function of the memory time stamp (i.e., current time minus memory age), obtained by numerical solution of the ordinary differential equation initial value problem for the probability distribution function (see Appendix) for the moment of time $t = 15$, with the inverse trace lifetime $\kappa = 0.015$, for various values of the total replication rate α. The cases A, B, C and D correspond to the 4 choices of the replication function, as described in Appendix. A: Equalized replication rate per trace. The values of α, from the bottom curve to the top curve, are: 0.6, 1.3, 2.5, 5, 10. B: Equalized replication rate per memory. The values of α, from the bottom curve to the top curve, are: 0.25, 0.5, 1, 2, 4. C: "Saturation". The replication rate for any memory linearly goes to zero at $m = 5$ traces: m is a model parameter. The values of α, from the bottom curve to the top curve, are: 0.5, 1, 1.5, 2, 2.5, 3, 3.5. D: "Recency". The replication rate exponentially decreases as a function of the memory age, with the exponent $1/\sigma$, where $\varsigma = 3$ is a model parameter. The values of α in this case are the same as in the case C.

a particular memory. The two cases that yielded the most interesting results were a model called "recency" (Figs. 1C, 2C), in which older memories were less likely to be replicated, and a model called "saturation" (Figs. 1D, 2D), in which the more traces in a particular memory, the less likely that memory would replicate. Both of these cases had the effect of limiting the number of traces of any particular memory (Fig. 2C,D), and both showed retrograde amnesia gradients on a short time scale (determined by model parameters), followed by a plateau in the probability of memory retrieval as a function of memory age (see Fig. 2C,D). As Figure 2C,D shows, with lesions of increasing size, retrieval probability decreases, and the overall shape of the functions approximates rather well what is typically observed in amnesic patients with medial temporal lobe damage.

The results from this analytic solution are satisfying, in that they demonstrate the formal plausibility of MTT, and enlightening, in that they suggest that limiting the number of replications per memory to a relatively small number actually produces the best fit to the empirical data. This is particularly interesting because one common objection to MTT is that the hippocampal complex would not have the capacity to store endless replications of old memories. This analysis shows that endless replications are neither necessary nor even advantageous in producing the kind of "retrograde amnesia" observed after real brain damage.

A CONNECTIONIST APPROACH

In order to see how these general principles might work in a more realistic neural network, we simulated a low-activity Hopfield-like network consisting of two components: a "hippocampus" (HC) and a "neocortex" (NC), each consisting of 1,000 neuronal

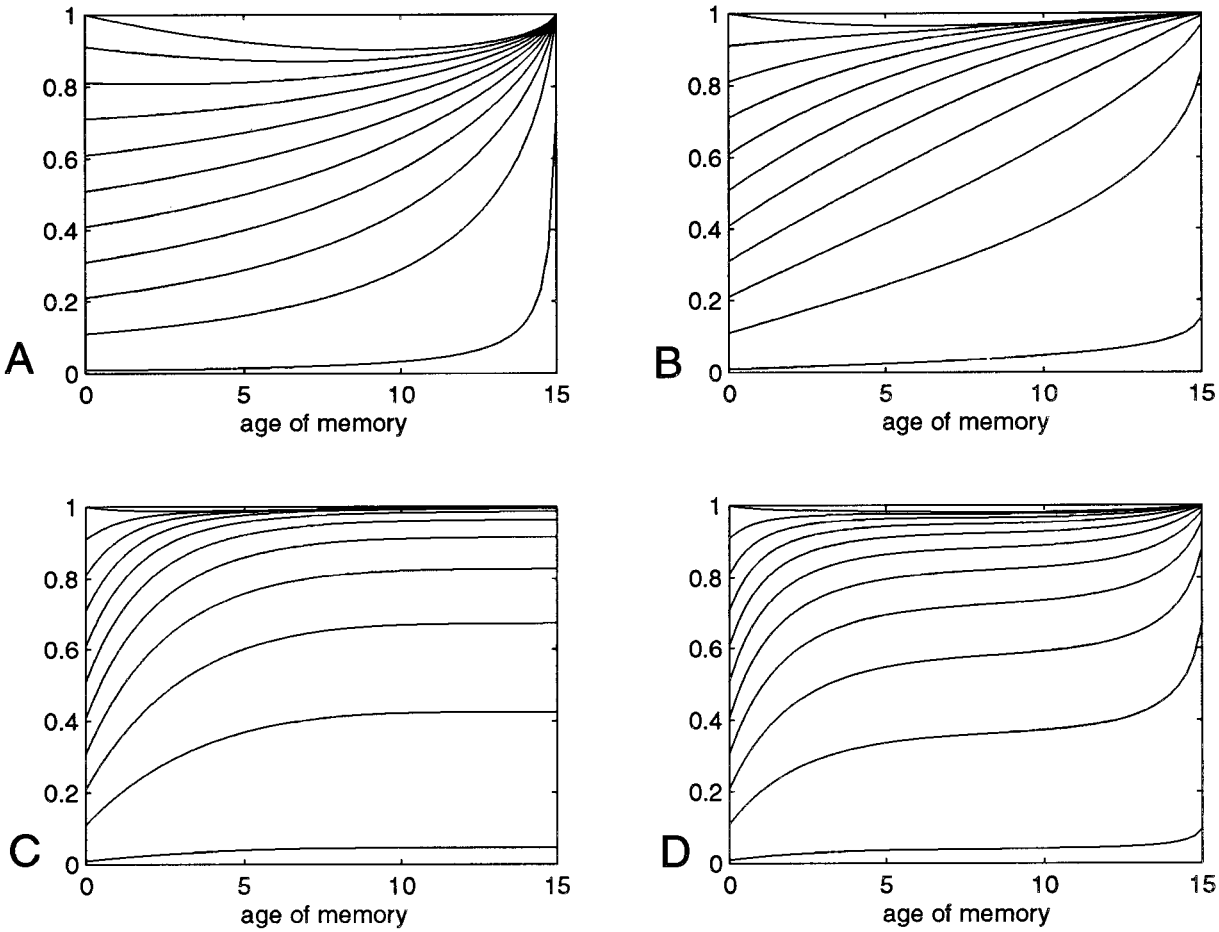

FIGURE 2. Average probability of recall as a function of memory age, obtained by numerical solution of the ordinary differential equation initial value problem for the probability distribution function (see Appendix) for the moment of time $t = 15$, with the inverse trace lifetime $\kappa = 0.015$, for various degree of lesion q. On each panel, form the top curve to the bottom curve, the parameter q takes the following values: 0, 0.09, 0.19, 0.29, 0.39, 0.49, 0.59, 0.69, 0.79, 0.89, and 0.99 (e.g., the last case corresponds to a 99% lesion to the hippocampus). The cases A, B, C and D that correspond to the 4 choices of the replication rate are the same as in Figure 1. The following values of the total replication rate α were selected. A: $\alpha = 10$. B: $\alpha = 4$. C and D: $\alpha = 3.5$.

units, using an asynchronous updating algorithm. A hard constraint in HC dynamics (two HC units active) and a soft constraint (Amit, 1989) keeping approximately 100 NC units active were used in these simulations. All-to-all synaptic connections in this model, which included NC-to-NC and HC-to-NC connections, were symmetric (bidirectional).

The connectionist simulation study included 100 independent numerical epochs. Each epoch started with all synaptic weights set to the same value (zero for HC-NC connections, -0.005 for NC-NC connections due to the Amit soft constraint) and was performed with new values of all random variables. The total computer time was 63 h on an SGI Origin2000 supercomputer with 16 processors and 2.4 GB memory.

Each epoch included a sequence of 15 simulated experiences. During each experience, each network (HC and NC) was presented with a randomly generated pattern. All patterns were uncorrelated with each other. The total number of units in each pattern matched the normal activity level of each network (100 units for NC, and 2 units for HC). Learning occurred immediately upon experience in the NC-to-NC and in the NC-to-HC connections. Learning rules for the NC-to-NC connections were Hebbian, modified for a low activity network according to the soft constraint (Amit, 1989). Learning the rules for the HC-to-NC connections was associative: positive connections were created between all coactive units (simulated LTP), and erased between active HC and inactive NC units (simulated LTD).

Events of replay were interleaved into the sequence of pattern presentations, to simulate memory rehearsal episodes. Each of the 15 experiences was followed by 7 replay cycles. The total number of replays in each simulated epoch was 105. Each replay started with stimulation of a randomly selected, sparse (5-unit) fragment of a stored NC pattern. The fragment selection was purely random, except that we ensured that the selected fragment unambiguously determined the randomly selected stored pattern. Simultaneously, a random "noise pattern" consisting of one unit was presented to the HC. In addition, a small background static noise of the amplitude of 0.001 was added to the membrane potentials of all HC units. Up to 10 iterations of network updating were allowed in

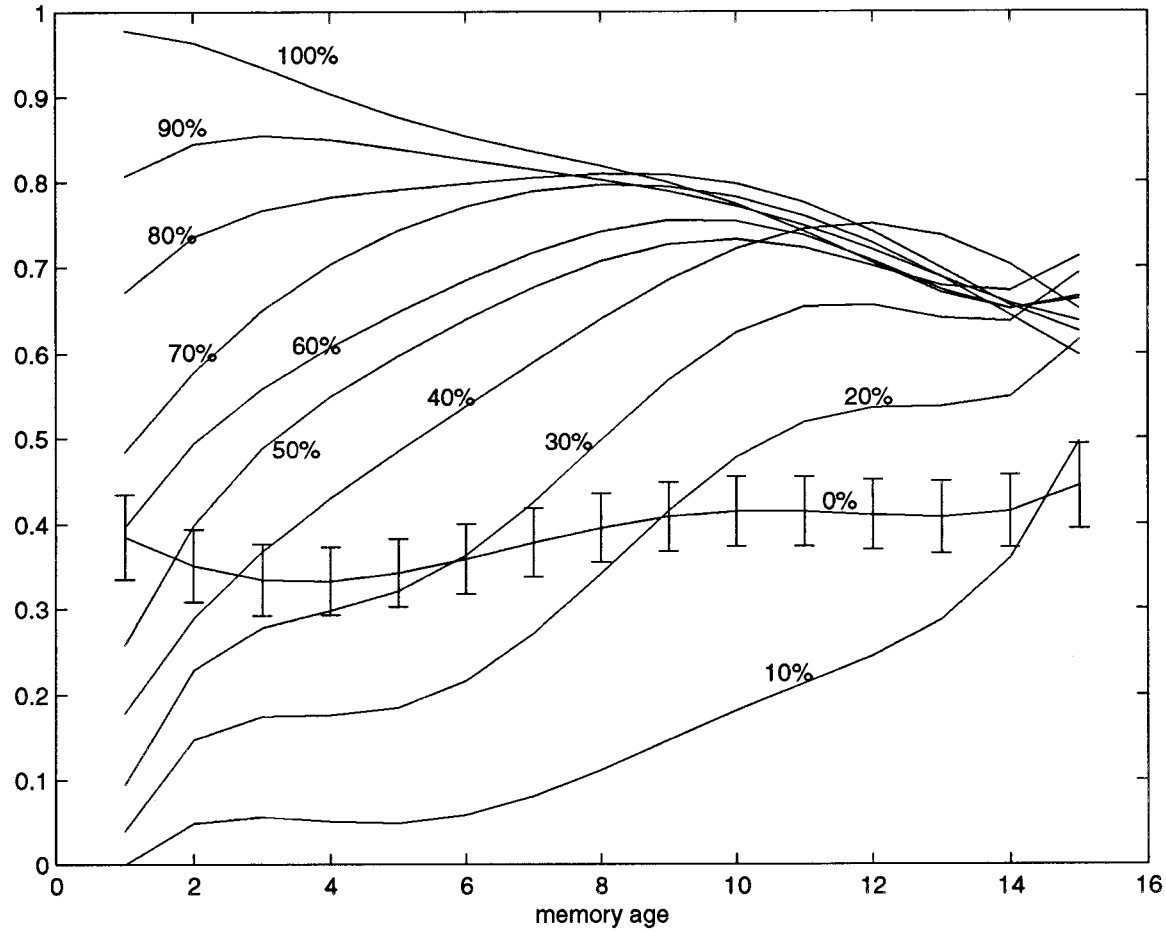

FIGURE 3. Average probability of recall as function of memory age, obtained by numerical simulation of the connectionist model (see text). The remaining fraction of the hippocampal component (i.e., 100% minus the degree of lesion *q*) is given in percents near each curve. The results were obtained by averaging over 100 independent simulation epochs and smoothening over a short range of memory ages. The error bars shown on the last curve (100% lesion) are approximately the same as for all the other curves. A qualitative similarity can be observed between this figure and Figure 2B. There are only 2 essential differences between these figures: (i) the normal forgetting in the connectionist model is higher and does not reach zero for the maximal memory age; (ii) There is a finite recall probability at a 100% lesion, which is due to the memory of the NC component and is substantially higher than for a 90% lesion.

each replay event. The normal scenario following presentation of a cue included subsequent reactivation of the associated HC pattern followed by retrieval of the entire NC pattern; then a partial alteration of the HC pattern occurred due to the noise, and the new HC unit(s) became associated with the active NC pattern via synaptic modification. Thus, multiple traces resulted from reactivation events. The NC-to-NC connections, however, were left unchanged during a replay event.

Each experimental epoch was followed by measurement of the system's performance in recall as a function of the lesioned fraction of all HC units. At this time, the system was subjected to a sequence of HC lesions. "Lesioned" units were randomly selected each time. Fractions of "lesioned" HC units constituted 0%, 10%, 20%, 30%, 40%, 50%, 60%, 70%, 80%, 90%, and 100%. The criterion for a successful recall was a Hamming distance of less than 20. There was no synaptic modification during each measurement cycle. The resultant retrieval curves were averaged over all 100 independent numerical epochs.

The results of this connectionist simulation are shown in Figure 3, where it can be seen that as "lesion" size increases, memory loss does as well, with a dependence on the age of the memory. The curves that correspond to lesions in the 50–90% range approximate what is seen in amnesic patients. It is of interest that in this model, a complete lesion actually has a smaller impact on retrieval than somewhat smaller lesions. This paradoxical result has actually been observed in some work with animal lesion models, and has been attributed to the potentially disruptive effects upon the neocortex of having a small and likely disorganized input from the hippocampus.

Overall, our analytical and numerical results and connectionist simulations demonstrate several things: first, that the core assumption of MTT, namely, that traces replicate with time, can indeed yield the kind of retrograde amnesia data observed in patients. Second, that limiting the number of such replications to a relatively small number works best of all. Of course, such simulations do not prove that this is what happens in the brain: they merely demon-

strate that it could work this way. In so doing, they encourage us to seek empirical data predicted by MTT. The remainder of the paper is concerned with describing several studies providing these data.

REMOTE MEMORY IN PEOPLE WITH UNILATERAL TEMPORAL LOBE EPILEPSY OR LOBECTOMY

To get a better handle on the effects of circumscribed lesions on remote memory, we examined 25 patients with unilateral temporal-lobe epilepsy either prior to surgery or following unilateral excision of the anterior temporal lobe (Viskontas et al., 2000). Of these, seizures originated in the right temporal lobe in 11 patients and in the left in 14 patients. Six of the right-sided and 9 of the left-sided patients were being assessed for resective surgery, and the remainder had already undergone it. Because the surgical procedure involved an en bloc resection, a substantial portion of the hippocampal complex was excised.

Previous studies that examined the effects of unilateral temporal lobectomy or temporal-lobe epilepsy (O'Connor et al., 1998) on remote memory found that memory for singular events, whether autobiographical or public, was impaired after left, but not right, temporal lobectomy (Barr et al., 1990). There was some evidence of a temporal gradient, with more recent memory for autobiographical incidents being somewhat more affected in the study by Barr et al. (1990) but not in O'Connor et al. (1998). In Barr et al. (1990), however, the onset of memory loss was calculated from the time of the surgery, though it is known that epilepsy may be of long duration, sometimes stretching back to infancy, and may affect encoding of new information over this entire time course. It is difficult to know, therefore, whether in Barr et al. (1990) the remote memory loss that was observed was truly retrograde, or whether it included anterograde effects caused by presurgical epilepsy. To account for the possibility that anterograde memory loss may contribute to remote memory deficits in our patient sample, we divided the groups into those with early (preadolescent, usually infantile) and late (after 18 years of age) onset seizures. We also looked at the extent of hippocampal sclerosis as determined by MRI to corroborate the longevity and extent of the lesion.

Each patient received the Autobiographical Memory Inventory (AMI) (Kopelman et al., 1989), a standardized test of remote memory that samples both autobiographical memories and personal semantics from three time periods: childhood, early adulthood, and recent, with finer distinctions between early, middle, and late childhood in the childhood period. In the autobiographical part of the test, one has to describe a particular personally experienced episode in detail, whereas personal semantics are tested by having the participant provide the names of friends, teachers, schools, home address, jobs, and so on. In addition to the AMI, we also examined our participants' memory for the pictures and names of famous personalities who became prominent during 5-year periods going back to the 1960s. Participants had to choose the name or face from among lures and to identify the individual by name (for the face) or by some other information that would distinguish that person from any other.

In comparison to age- and education-matched controls, people with unilateral temporal lobectomy showed a retrograde amnesia for autobiographical episodes that extended to childhood, and within that period extended to the very earliest part (see Fig. 4A,B). This effect was independent of the side of the lesion, indicating that both the right and left temporal lobes contribute to autobiographical memory. More importantly, there was no interaction between lesion onset and retrograde memory loss on the AMI. Though people with early-onset seizures had worse memory at all time periods, memory loss extended to very early childhood even in people with late-onset seizures. In both cases, the severity of retrograde memory loss was the same for recent and remote memories. These results indicate, in line with MTT but contrary to the traditional model, that both hippocampal complexes are needed for recovery of all remote, autobiographical episodic memories, no matter what their age.

The situation is different when it comes to other types of memories. Memory for the faces of famous personalities was impaired only in people with right temporal lobe seizures or lobectomy, and then only for more recent time periods. Memory for personal semantics was preserved in all participants in our study (see Fig. 4C), but only in right temporal lobectomy patients in O'Connor et al. (1998). Taken together with the results on autobiographical memory, these findings are consistent with MTT in the following way: when autobiographical episodes are reexperienced through retelling, they are unlikely to involve the same rich detail, both verbal and nonverbal, as occurred when they were initially experienced. As a result, damage to either medial temporal lobe will produce a long-lasting remote memory loss. In contrast, faces of famous personalities, by definition, are encountered more often but their representation is more dependent on the right hemisphere, and show the hemisphere-specific, temporally graded memory loss expected of such material. The information that constitutes personal semantic memory is repeated so often, almost constituting our official identity, that it suffers almost no loss following unilateral temporal lobectomy.

When medial temporal lobe lesions are more extensive, as they typically are following attacks of herpes encephalitis, remote memory also suffers. Thus, Kopelman et al. (1989) reported that in people with bilateral temporal lobe lesions, RA is severe, extensive, and ungraded for both autobiographical incidents and personal semantics, as well as for news events. When the lesion was unilateral, however, people with right temporal damage performed worse than those with left temporal damage in tests of autobiographical incidents and famous faces, but not for personal semantics or news events not containing a famous face. The left temporal group showed a gradient on the news events test.

Notwithstanding the consistency of these findings with MTT, some questions remain. In all the excision cases, and many of the encephalitis cases, the lateral surface of the temporal lobes and the anterior temporal pole were lesioned, leaving open the possibility that it is lesions to these structures, rather than to the hippocampal complex, that are responsible for the remote memory loss that was observed. Also, the extent of the deficit was unrelated to the extent

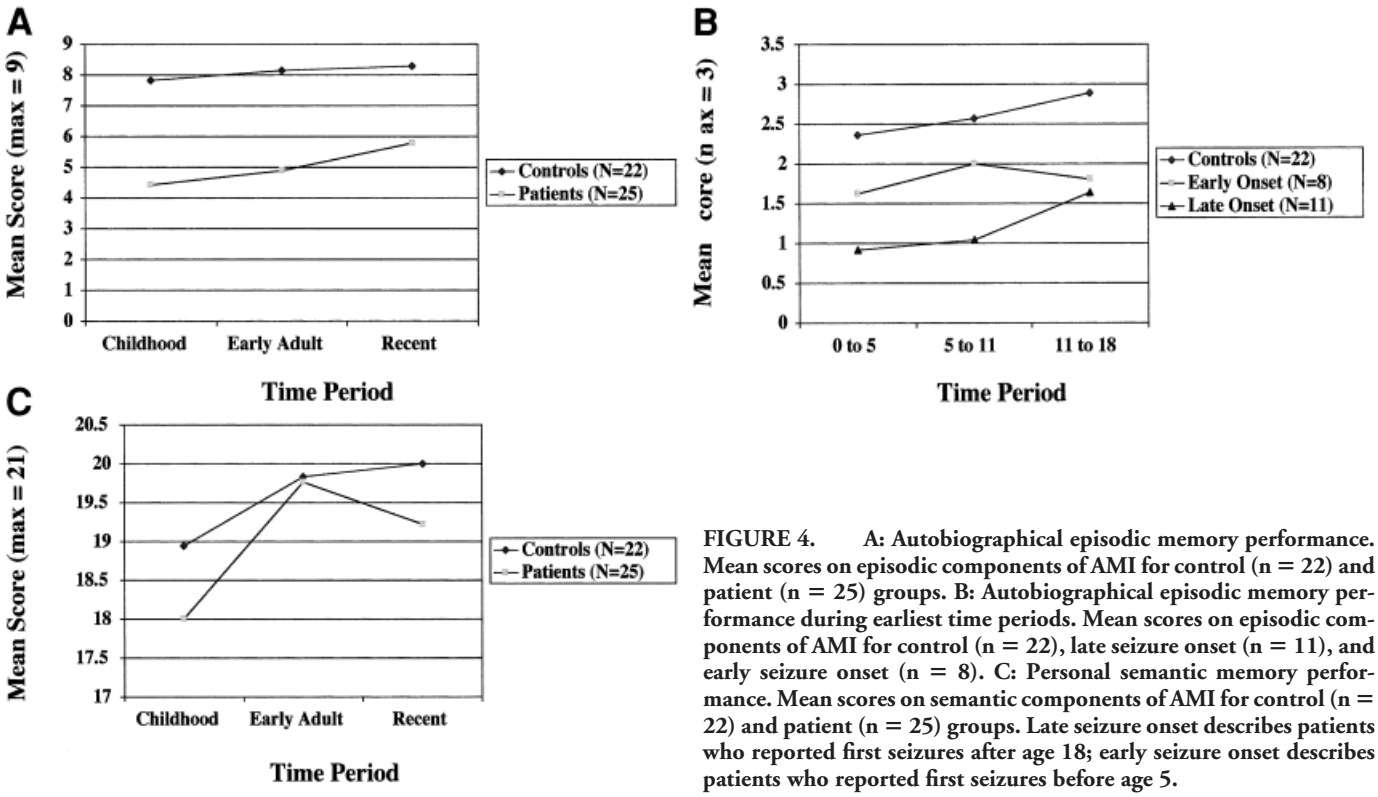

FIGURE 4. A: Autobiographical episodic memory performance. Mean scores on episodic components of AMI for control (n = 22) and patient (n = 25) groups. B: Autobiographical episodic memory performance during earliest time periods. Mean scores on episodic components of AMI for control (n = 22), late seizure onset (n = 11), and early seizure onset (n = 8). C: Personal semantic memory performance. Mean scores on semantic components of AMI for control (n = 22) and patient (n = 25) groups. Late seizure onset describes patients who reported first seizures after age 18; early seizure onset describes patients who reported first seizures before age 5.

of hippocampal sclerosis, which is a measure of lesion size. However, the evidence from Kopelman et al. (1989) suggests that more extensive bilateral lesions lead to greater remote memory loss. Further research on patients with hippocampal-amygdaloid removals that spare the lateral temporal cortex, and whose size of lesion is determined accurately by MRI or postmortem examination, should help to resolve these issues. Finally, there is a discrepancy among studies about the effects that side of lesion has on remote memory loss for autobiographical incidents, which may also be resolved once more accurate measures of tissue loss are taken into account, as well as time of seizure onset.

REMOTE MEMORIES FOR AUTOBIOGRAPHICAL EPISODES AND PUBLIC EVENTS: SCORING FOR NUMBER OF DETAILS

Almost all tests of episodic, autobiographical memory rely on a 3-point scoring technique, initially introduced by Crovitz and Schiffman (1974). The maximum number of points, three, is allotted if the time and place of the memory are specified as well as details of the event, two points if less detail is provided or the time and place is not specified, and one point if only general information is given. Apart from being somewhat subjective, this scoring procedure does not differentiate between memory reports that are very rich in detail from those that are just detailed enough to merit the maximum score.

We were alerted to the potential significance of this issue when comparing the performance of amnesic people with controls on an autobiographical and historical version of the Galton-Crovitz Cue Word Test (Moscovitch and Melo, 1997). In that test, participants are given a cue word, such as "letter" or "broken," and are required to describe a personal episode related to this word from their past in as much detail as they can. For the historical version, they are given a cue word such as "battle" or "king" and are required to describe a historical event that occurred before they were born. Although controls significantly outperformed the amnesics in both conditions, even with the traditional 3-point scoring method, the differences were rather small and did not reflect our impression that the amount of detail in the amnesics' recall was much less than in that of the controls in both conditions.

To corroborate our impression, we devised a simple new scoring technique in which we simply counted the number of details provided for each cue (Moscovitch et al., 1998, 1999). Initially, we also classified the details into various categories such as spatial, perceptual (visual, auditory), implication, temporal, emotional, and so on, but these proved to be difficult to sort reliably and did not provide additional information beyond the total number of details unsorted by type. Whereas with the old scoring technique there was only a 15–20% difference between amnesics and controls, on rescoring the data for number of details, we now found about a 50% difference (Moscovitch et al., 1998). We also noted that the hint of a temporal gradient evident with the old scoring technique was now eliminated. Because the study of Moscovitch and Melo (1997) was not designed to look at the temporal extent of

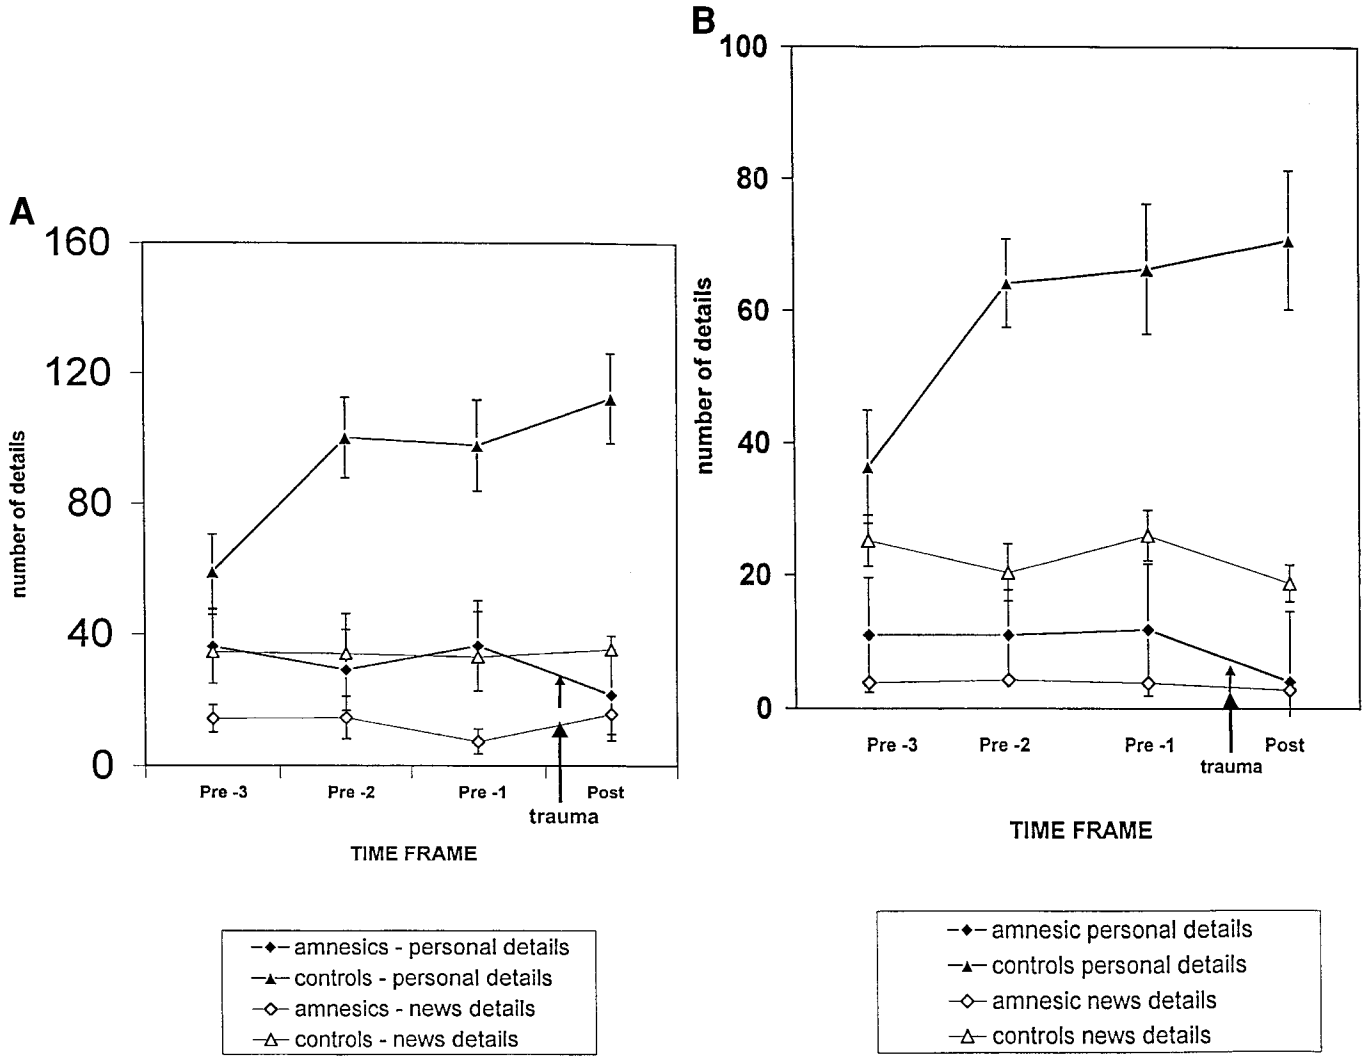

FIGURE 5. A: Total number of details freely recalled about personal episodes and news events that occurred before and after amnesia (trauma). Pre-3 refers to childhood, Pre-2 to adolescence, and Pre-1 to early adulthood. B: Total number of details recalled with prompting about personal episodes and news events that occurred before and after amnesia (trauma). Pre-3 refers to childhood, Pre-2 to adolescence, and Pre-1 to early adulthood.

remote memory loss, we devised a new study to examine this issue directly, and to test the predictions of the MTT.

In the new study, we had participants describe in as much detail as possible two personal memories and two public events from each of five time periods in their life: childhood, adolescence, early adulthood, middle age, and recent. We decided not to use the technique of Crovitz and Schiffman (1974) because recovering a memory in response to unusual cues may require the use of strategic retrieval strategies that could be impaired, though the memory itself may not be. Participants could simply choose any two memories they wished from each time period. To minimize the possibility that the problem was one of retrieval, we cued the participants extensively, and went so far as to provide them with a 75-item list (Levine, personal communication) of typical personal events and of public events from each time period. They could then choose those that were applicable to themselves. To control for the possibility that some people were more garrulous than others, we prompted them extensively to provide us with as many details as possible (Fig. 5A,B).

Despite having to choose only two events from each time period, and being cued and prompted extensively, the five amnesic people we tested (two with medial temporal lesions, and one each with anterior communicating artery aneurysm, diencephalic lesion, and early Alzheimer's disease), all showed remote memory loss for personal episodes back to early childhood (see Fig. 5). Though the amnesic people could often retrieve memories, they were impoverished in comparison to controls. When scored according to the old, 3-point procedure, a hint of a temporal gradient was observed, with older memories being better preserved than newer ones, However, with the new scoring procedure, there was no temporal gradient when we examined amnesic memory itself (see Fig. 6). The difference between amnesics' memory and that of controls was greater for recent than for remote memories, but that was due to the loss of detail

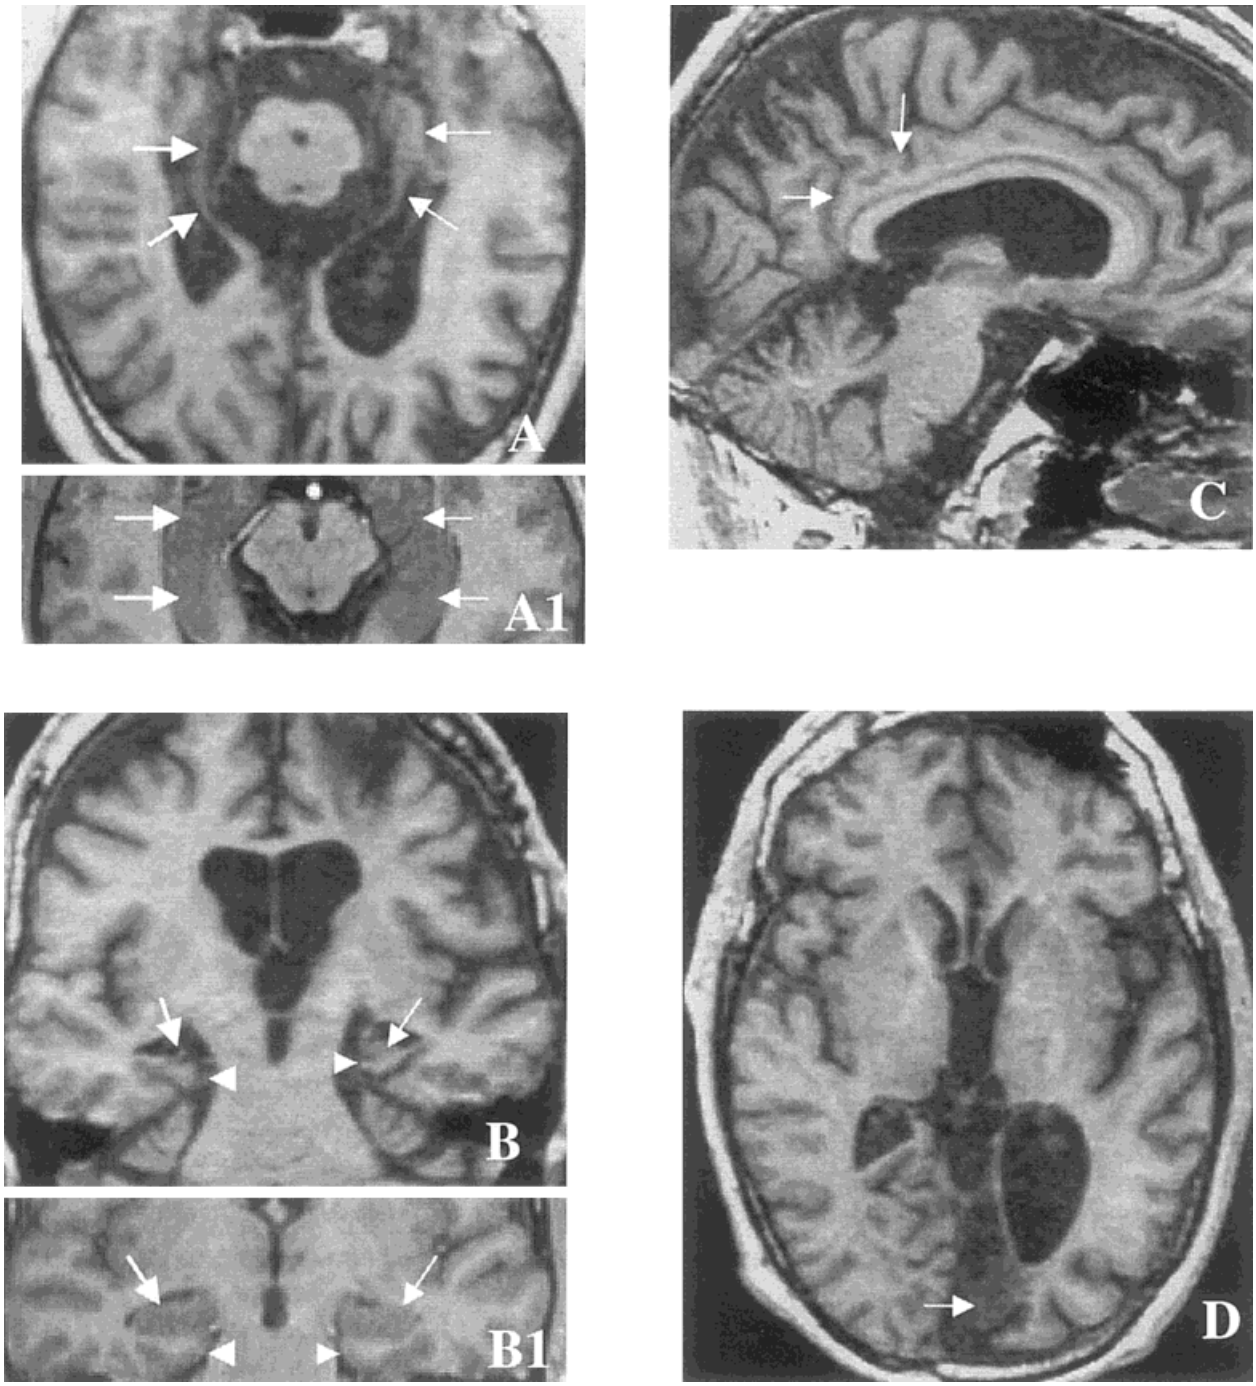

FIGURE 6. Magnetic resonance imaging (MRI) slices showing K.C.'s pathology. A: Axial view, showing severely atrophic right and left hippocampus (arrows). A1: Axial view, showing controls' intact right and left hippocampus (arrows). B: Coronal view, showing severe atrophy of the right and left hippocampus (arrows) and atrophy of the right and left parahippocampal gyri (arrowheads). B1: Coronal view, showing controls' intact right and left hippocampus (arrows) and right and left parahippocampal gyri (arrowheads). C: Midsagittal view, showing relatively preserved posterior cingulate (arrows). D: Axial view, showing left medial occipital infarction.

in the older memories of controls, rather than to the relative preservation of the more remote memories in amnesics.

These results suggest that as memories age, they lose richness of detail and may become more semanticized, as a number of investigators have suggested (Kinsbourne and Wood, 1975; Cermak and O'Connor, 1983). The results from the public events version of our study lend support to this view. For public events, little loss of detail was observed from recent to remote memories in control participants, suggesting that public events are somewhat semanticized from the beginning. Even for public events, amnesics provide fewer details than controls at each point in time, with no hint of a gradient. The flat gradient observed only for public events in con-

trols is also observed for personal events in amnesics, suggesting that what is retained of one's personal memories in amnesia has the character of memory for a public event: the memory is somewhat depersonalized and lacks the richness of detail that allows one truly to reexperience an episode.

The idea that what is lost in amnesia is not the gist of a memory but its details, no matter how old the memory is, will be explored below, in a discussion of spatial memory. At this point we should note that although the results are consistent with MTT, the amnesic individuals we investigated with temporal lobe lesions had damage extending beyond the hippocampal complex. At least two patients did not have damage to the hippocampal complex at all, but to other structures (diencephalon, basal forebrain) that are crucial for memory. In a study that examined the richness of memory in two people, one with lesions presumably confined to the hippocampal formation, and another with lesions to the hippocampal complex and inferior temporal lobes extending beyond it, Reed and Squire (1998) found the traditional temporal gradient for episodic memory, with only the most remote memories being spared in the person with hippocampal complex lesions, and no detectable loss in the person with lesions confined to the hippocampal formation. However, they used a new 3-point scale to score each of a number of categories of the reported memory, rather than counting details, and without the extensive cueing and prompting we used in our study. It is unclear, therefore, whether even the most remote memories of their amnesic people would have been shown to be impoverished had they used our procedure. At a more general level, it remains to be determined whether remote memory deficits extending to childhood are characteristic of all forms of amnesia or only those caused by medial temporal lobe lesions.

REMOTE SPATIAL MEMORY AND COGNITIVE MAPS

Most studies on remote memory in humans are concerned with memory for events or people. Only a handful deal with spatial memory. This is surprising in view of the important position that spatial memory holds in theoretical and empirical research on animal models of hippocampal function (O'Keefe and Nadel, 1978; Nadel, 1991). It is probably not an exaggeration to say that the most commonly used tests of memory in rodents are spatial, and that even theories that do not ascribe an exclusively spatial function to the hippocampus concede that it has a special role in creating, representing, or operating on cognitive maps of the environment. Studies on anterograde spatial memory have been reported in monkeys (Parkinson et al., 1988) and humans (e.g., Bohbot et al., 1998), with the predictable result that it is impaired following hippocampal lesions, though it remains a matter of debate whether spatial memory is affected more than any other following such lesions (cf. Henke et al., 1999).

There are far fewer studies of remote spatial memory, even in rodents. Of the studies that have been published, most report a severe and long-lasting retrograde amnesia with no temporal gradient (Murray and Bussey, in press), consistent with MTT. Studies of remote spatial memory in humans, however, show the opposite, as reflected in the preservation of cognitive maps of homes, neighborhoods, and cities. There is anecdotal evidence that the amnesic patient H.M. could remember and navigate in the neighborhood in which he dwelt prior to his surgery and could draw an accurate map of the floor plan of his house. In a more formal test of spatial memory, Beatty et al. (1987) corroborated that amnesic people with bilateral medial temporal lobe lesions could draw, as accurately as their relatives, maps of floor plans of the homes in which they dwelt long before. They were impaired, however, in their drawings of the floor plans of dwellings they lived in more recently. Similarly, Teng and Squire (1999) reported that remote spatial memory of a large-scale environment, a district near San Francisco, was preserved in a person with bilateral medial and inferior temporal lobe lesions, even though he had not lived in or visited that neighborhood for decades.

Considering the importance of spatial memory for theories of hippocampal function, of remote memory for theories of consolidation, and the dearth of studies of remote spatial memory in humans, we (Rosenbaum et al., submitted) undertook to investigate spatial memory in K.C., a 47-year-old, severely amnesic person with extensive bilateral hippocampal damage (see Fig. 6), though the damage also included other regions (Kohler et al., 1998). On tests other than spatial memory, K.C. shows a retrograde amnesia extending for his entire life, without a temporal gradient, for autobiographical episodes, and a more restricted retrograde amnesia encompassing the 5–10 years prior to his trauma, for famous people, public events, and vocabulary (Westmacott and Moscovitch, 2000).

We knew from observing him informally that K.C. was familiar with the neighborhood in which he lived since he was 10, and continues to live in today. He goes unaccompanied on daily walks in the neighborhood and can direct people to different locations in it confidently and without hesitation. Nonetheless, we thought it was important to document these impressions with more formal tests, and to determine whether indeed he relies on internally represented cognitive maps of the environment or perhaps navigates by some other means, such as memorized routes or landmarks which are less dependent on the hippocampus. We also wanted to test his knowledge of world, national, and provincial geography to see if there was any relation between that and remote spatial memory for his neighborhood.

We administered the following tests to K.C. and two controls, his mother and brother, who were equally familiar with the neighborhood: 1) drawing a sketch map of his neighborhood, 2) drawing vector maps in which participants had to indicate with an arrow drawn on a blank piece of paper, the direction and proportional (drawn-to-scale) distance between landmarks in his neighborhood, 3) blocked-route problem solving in which participants had to indicate the route one would take from one location to another if the most convenient or typical route were blocked, 4) distance judgments in which participants had to give verbal estimates of distances between points in the neighborhood, 5) proximity judgments in which participants were shown sets of three

photos of different locations in the neighborhood and had to indicate which two were closest to the third, and 6) sequencing landmarks along a route, in which participants were presented with randomly arranged photographs of landmarks along a route and had to order the photographs in the sequence they would appear as one traveled from one point to another along that route.

K.C. performed as well as controls in all of these tests. Indeed, he was perfect on the blocked-route test, which is considered a prototypical cognitive map test. On the vector-drawing and distance estimation tests, his deviation error both in terms of angle (about 10 degrees) and distance (about 0.3 Km) was comparable to controls. The only hint that K.C. may have had some deficit in spatial memory came from a close examination of his sketch map. Although his map was accurate in drawing the streets in correct spatial relation to one another, he drew fewer landmarks than the controls.

That K.C. had particular problems with landmarks was confirmed unequivocally in a test of landmark recognition and identification. In that test, participants viewed 96 photographs. Half the photographs were of familiar landmarks (e.g., a school) and salient houses (e.g., located at street corners, belonging to friends) from his neighborhood, and the other half were photographs taken in another neighborhood. Half the lures resembled the landmarks and houses in K.C.'s neighborhood, and the other half were distinctly different. After choosing photographs which participants recognized as belonging to their neighborhood, they had to describe the location of the landmarks they recognized as familiar in relation to other landmarks in the neighborhood and the vantage point from where the photograph was taken.

Whereas controls scored 45/48 and 46/48, KC scored only 15/48. The few errors that controls made were always on photographs that resembled those in their neighborhood, but K.C. chose similar and highly dissimilar lures equally. Finally, controls were able to identify the location of almost all the photographs they recognized (43/45 and 46/46), whereas KC could do so for only 7 of the photographs which were of the most salient landmarks in his neighborhood. These results suggest that although K.C. has a good memory of the general layout of his neighborhood, and some of the major landmarks, his memory of specific details was very faulty. One can truly wonder whether he managed to retain any knowledge of incidental, spatial details despite having lived in the neighborhood for over 20 years prior to his accident and for 17 years since then.

The results of the spatial geography test were consistent with the results of the neighborhood spatial memory test. For the geography test we modified the Fargo Map Test for Toronto (Toronto Map Test) of Beatty (1988). Participants first had to identify gross geographical features (oceans, continents, countries) on an outline map of the world and of North America, and then proceeded to identify more specific features (cities) on outline maps of Canada and Ontario. Each map contained numbered dots, approximately two thirds of which corresponded to geographical features listed on a separate sheet of paper. Participants were instructed to match each number with its correct geographical feature on the list. K.C. scored normally when locating the gross features but progressively more poorly as identification became more specific.

It should be noted that K.C. was grossly impaired in acquiring new spatial memories. He could not draw a floor plan of the library in which he has worked for the last 2 years, nor learn a route from the psychology laboratory to the washroom in 12 trials. He also performed at the same level as H.M. on Smith and Milner's tabletop test of spatial memory in which he was asked to replace, in their proper location on a board, items he had studied 5 min earlier.

The picture that emerges from the results of these tests of remote spatial memory, the most extensive by far undertaken in a single amnesic individual, is that memory for spatial location, as memory for autobiographical episodes, public events, and famous personalities, suffers in the details. The spatial representation or cognitive map that K.C. has of his neighborhood is akin to the personal semantic memory he has of his life, without the specific details of the episodes constituting that life.

There are a number of questions that remain unresolved. We do not know whether the cognitive map K.C. has of his neighborhood is similar fundamentally to normal cognitive maps, though his normal performance on tests designed to assess topographical knowledge suggests that it is. One possibility is that instead of cognitive maps, K.C. uses heading vectors from salient landmarks to navigate in his neighborhood (for a discussion of the use of such heading vectors by rats with hippocampal lesions, see Pearce et al., 1998; Kubie et al., 1999), though it is unlikely that such vectors can support the kind of knowledge K.C. demonstrated on the various tests. We also cannot say with certainty that his spatial deficit is due to hippocampal complex lesions, because his damage extends beyond the medial temporal lobes. These other areas, however, are not known to be involved in memory for landmarks. Nonetheless, it will be important to examine people with more circumscribed lesions to resolve this issue.

In addition, it will be important to examine memory for spatial location in tests that are similar to those of autobiographical events. What distinguishes an episodic memory from other types is that the event occurred only once (though it may have been rehearsed later). We know of no such tests of spatial memory, though they are crucial if we are to compare one type of remote memory loss to another. In an ongoing study (Booker et al., unpublished findings), we are asking normal people to describe locations they have visited only once and comparing these to descriptions of locations they have visited many times. If one can recollect such locations (e.g., vacations to different places) from different times in one's life, one could then ask whether remote memory for once-visited spatial locations following hippocampal damage is preserved in amnesia, as the traditional view would have it, or whether it is impaired, as MTT predicts.

If the hippocampus is not needed to retain the broad topographical knowledge K.C. has of his neighborhood, what brain areas are crucial? One candidate is the posterior cingulate gyrus. It is relatively preserved in K.C. and has been implicated in topographical memory loss in a number of cases (Aguirre and D'Esposito, 1999).

Although many issues are still unresolved, the similarity in the extent of remote memory loss across various domains is striking. Though realizing the need for caution, we venture to conclude that the hippocampal complex is needed to retain detailed knowledge of all types of memory (of autobiographical episodes, of public

events and personalities, and of cognitive maps), no matter the age of the memory.

NEUROIMAGING AND REMOTE MEMORY

A recent functional MRI (fMRI) study we conducted provides additional evidence in support of MTT. As discussed earlier, one of the clearest differences between classical consolidation theory and MTT concerns the recollection of remote memories. Consolidation theory suggests that the role of the hippocampus in recollection is time-limited, although the period of consolidation in which the hippocampus is still required for recollection is variable, with estimates ranging from 3–15 years in humans (Rempel-Clower et al., 1996). Nevertheless, recalling memories of very old life events should not require the hippocampus, but instead should depend exclusively on neocortical structures. The MTT model, on the other hand, posits that successful retrieval of event-related information, especially detailed contextual information such as time, place, temporal sequence, and emotional content and perceptual features or vividness, should require the hippocampus, regardless of the age of the memory.

We explored this issue in an fMRI study in which normal older subjects were asked to retrieve remote memories of various ages and kinds (Ryan et al., 2000). If the role of the hippocampus ceases after memories are consolidated into the neocortex, retrieval of remote memories should show little or no activation in the hippocampus, and considerable activation in neocortical regions. If, on the other hand, the hippocampus remains important in retrieval even of remote memories, activation in this area should still be observed when subjects think about events from the distant past.

Seven subjects between ages 56–72 participated in the study. While undergoing fMRI, subjects were asked to recall the details of 10 recent events (occurring within 2 years of the study) such as "a rent holiday" or "a family dinner" and 10 remote events (occurring at least 20 years prior to the study) such as "learning to drive" or "your wedding day." Most of the remote events happened when subjects were in their twenties or early thirties, so for some subjects, the events had occurred as much as 45 years ago.

Two control conditions were also included in the design. After recollecting each event, subjects were shown the word "relax" and were instructed to stop thinking about the event and to focus their mind on relaxing. A second, more engaging control condition was also added. Subjects were periodically shown sets of five high-cloze sentences with the last word missing, such as "The cat was chased by the _____." Subjects were required to complete each sentence with an appropriate word, and were informed that they would be given a memory test for the sentences afterwards. This provided a control task that required the retrieval of semantic rather than autobiographical information, and it also provided a comparison between the retrieval of autobiographical information and the encoding of novel information within the scanning session.

Images were acquired on a 1.5T GE full-body echo speed magnet, using single-shot spiral acquisition (Glover and Lee, 1995). Sections (5 mm, Noskio) were placed obliquely, perpendicular to the long axis of the hippocampus, in order to minimize partial-voluming of the hippocampus on each section, covering the whole brain.

Briefly, images were reconstructed, reregistered to correct for movement, and normalized to allow the data to be combined across subjects, using correlation analysis (AFNI; Cox, 1996). This allowed us to compare the fMRI signal during recent events, remote events, rest and the sentence completion task.

All 7 subjects showed activation in the hippocampus during recollection; 5 subjects showed bilateral activation, 1 subject showed significant activation only in the left hippocampus, and 1 subject showed significant activation only in the right hippocampus. The degree of hippocampal activation was significantly greater for recollection of events than either the relax condition or the sentence completion task. Most importantly, the magnitude of activation did not differ depending on the age of the recalled event: the hippocampus was equally active during the recollection of remote events as it was for recent events, despite the fact that the remote events occurred 20–45 years prior to the scanning session.

The results of this study are absolutely clear: the hippocampus is as activated by retrieval of very remote memories as it is by retrieval of recent events. These results are inconsistent with the standard model of memory consolidation, which posits that memories become independent of the hippocampus as they age. Instead, the findings strongly support the MTT model, within which autobiographical or episodic recollection always depends on the hippocampus. We suppose that when some event from the past is reactivated, a new memory trace is created within the hippocampus (and possibly an old one is strengthened), and this proliferation of traces renders older memories less susceptible to disruption from brain damage than more recent ones. This mechanism provides a way of understanding the retrograde amnesia gradients sometimes reported after damage in the hippocampus.

Our findings, however, are open to alternative interpretations. One might argue, for example, that a normally functioning hippocampus is involved not only in retrieval of prior events, but in the ongoing encoding of new events. The hippocampal activation described here might therefore be associated with the encoding of a very salient novel event, namely, the experience of the subject recalling prior events within the MRI scanner, i.e., the reencoding of an old event in a new context. If this were the case, we would still expect at least differences in neocortical activation between recent and remote memories, but no such differences were evident. It is also important to note in this regard that hippocampal activation was greater during the recollection of autobiographical events than during sentence completion. Though the sentence completion task did not include explicit instructions to the subjects to remember the materials, subjects were nevertheless very good at recalling the sentences afterwards, indicating that they had engaged in encoding while completing the sentences. This finding suggests that activation within the hippocampus occurred, at least in part, because of the retrieval of autobiographical events and not solely because of new encoding during the scanning experience.

Despite the compelling evidence, we concede it is unlikely that fMRI studies alone can provide a definitive answer to the question

of whether the role of the hippocampus in memory is time-limited, and for this reason we sought converging evidence from other domains of inquiry. Taken together with the studies described earlier detailing the persistent deficits in amnesics and in temporal lobectomy patients for episodic material, these findings provide consensus that the hippocampus is not merely a temporary memory system whose role in recovering old memories is superfluous once consolidation is complete. Instead, it appears that the hippocampus is essential for the long-term availability of particular aspects of memory, namely, episodic detail for all memories, no matter how old they are.

CONCLUSIONS

The emergence of new techniques such as computational modeling and neuroimaging, along with finer-grained analysis of neuropsychological deficits, is providing the possibility of clearer understanding of the separate yet interacting roles of the neocortex and hippocampal complex in episodic and semantic memory. We argued some years ago that the standard models in this domain were oversimplified, and not capable of accounting for the details of the data. In the intervening years this has become even more clear. In this paper we provide further reasons for seeking new ways to think about memory function.

Our computational models, though far from biological realism, nonetheless demonstrate, in a formalistic way, that "memories" within a system can replicate, and that such replication can lead to results rather like naturally observed retrograde amnesia, when the system is subjected to what amounts to a "lesion." What is more, these models suggest that by placing constraints on the extent of trace replication, one can more accurately simulate the empirical results.

The data from our neuroimaging study demonstrate that activation in the hippocampal complex and neocortex is not affected by the age of a retrieved episodic memory. Memories as old as 20–30 years elicited hippocampal activations as intense as those seen with memories retrieved from the recent past. These results are most consistent with the idea that in the intact human brain, the hippocampal complex is involved in the retrieval of both remote and recent episodic memories, as predicted by MTT.

But, given this, how can one explain the apparently normal remote memory capacities reported in many studies over the years? The answer, in our view, is that the normalcy of remote memory in these cases was more apparent than real. Careful testing of the detailed capacity of remote memory in these cases makes it clear that neither spatial nor episodic memory is truly normal. Rather, our data suggest that what is relatively normal is semantic memory, though this too may show a temporal gradient when detailed knowledge is sought. Episodic memory, on the other hand, is always impaired. Thus, these data collectively indicate that, contrary to standard consolidation theory, the hippocampal complex plays a continuing role in episodic memory, no matter the age of the particular memory being retrieved.

Further study of the precise role of the hippocampal complex in remote memory storage and retrieval should illuminate the exact nature of the relationship between episodic and semantic memory. In so doing, this program of research will help elucidate the nature of human memory as a whole.

Acknowledgments

L.N., A.S., and L.R. were supported by the Cognitive Neuroscience Program and the Cognition and Neuroimaging Laboratories, both at the University of Arizona. Funding from the McDonnell-Pew Cognitive Neuroscience Program, the Flinn Foundation, and the State of Arizona through the Alzheimer's Disease Research Center is gratefully acknowledged. A.S. was also supported by a grant from the Center for the Study of Consciousness at the University of Arizona. M.M. was supported by grant A8437 from the Natural Sciences and Engineering Research Council of Canada.

REFERENCES

Aguirre GK, D'Esposito M. 1999. Topographical disorientation: a synthesis and taxonomy. Brain 122:1613–1628.

Aguirre G, Detre J, Alsop DC, D'Esposito M. 1996. The parahippocampus subserves topographical learning in man. Cereb Cortex 6:823–829.

Amit DJ. 1989. Modeling brain function: the world of attractor neural networks. New York: Cambridge University Press.

Barr WB, Goldberg E, Wasserstein J, Novelly RA. 1990. Retrograde amnesia following unilateral temporal lobectomy. Neuropsychologia 28: 243–255.

Beatty WW. 1988. The Fargo Map Test: a standardized method for assessing remote memory for visuospatial information. J Clin Psychol 44:61–67.

Beatty WW, Salmon DP, Bernstein N, Butters N. 1987. Remote memory in a patient with amnesia due to hypoxia. Psychol Med 17:657–665.

Bohbot VD, Kalina M, Stepoankova K, Spackova N, Petrides M, Nadel L. 1998. Spatial memory deficits in patients with lesions to the right hippocampus and to the right parahippocampal cortex. Neuropsychologia 36:1217–1238.

Cermak LS, O'Connor M. 1983. The anterograde and retrograde retrieval ability of a patient with amnesia due to encephalitis. Neuropsychologia 21:213–234.

Cox RW. 1996. Software for analysis and visualization of functional magnetic resonance neuroimages. Comput Biomed Res 29:162–173.

Crovitz HF, Schiffman H. 1974. Frequency of episodic memories as a function of their age. Bull Psychonom Soc 4:517–518.

Epstein R, Kanwisher N. 1998. A cortical representation of the local visual environment. Nature 392:598–601.

Epstein R, Harris A, Stanley D, Kanwisher N. 1999. The parahippocampal place area: recognition, navigation, or encoding? Neuron 23:115–125.

Fujii T, Moscovitch M, Nadel L. 2000. In press. Memory consolidation, retrograde amnesia, and the temporal lobe. In: Cermak L, editor. Handbook of neuropsychology, 2nd edition, volume 4. Amsterdam: Elsevier Science B.V.

Gabrieli JDE, Brewer JB, Descond JE, Glover GH. 1997. Separate neural bases of two fundamental memory processes in the human medial temporal lobe. Science 276:264–266.

Glover GH, Lee AT. 1995. Motion artifacts in fMRI: comparison of 2DFT with PR and spiral scan methods. Magn Reson Med 33:624–635.

Graham KS. 1999. Semantgic dementia: a challenge to the multiple-trace theory? Trends Cogn Sci 3:85–87.
Graham KS, Hodges JR. 1997. Differentiating the roles of the hippocampal complex and neocortex in long-term memory storage: evidence from the study of semantic dementia and Alzheimer's disease. Neuropsychology 11:77–89.
Henke K, Kroll NEA, Behniea H, Amaral DG, Miller MB, Rafal R, Gazzaniga MS. 1999. Memory lost and regained following bilateral hippocampal damage. J Cogn Neurosci 11:682–697.
Hodges JR, Gerrard P, Patterson K. 1998. Semantic dementia. In: Kertesz A, Munoz DG, editors. Pick's disease and Pick's complex. New York: John Wiley Co. p 83–104.
Kinsbourne M, Wood F. 1975. Short-term memory processes and the amnestic syndrome. In: Deutsch J, Deutsch D, editors. Short-term memory. New York: Academic Press. p 258–291.
Kohler S, Moscovitch M, Winocur G, Houle S, McIntosh R. 1998. Networks of domain-specific and general regions involved in episodic memory for spatial location and object identity. Neuropsychologia 36:129–142.
Kopelman M, Wilson BA, Baddeley AD. 1989. The autobiographical memory interview: a new assessment of autobiographical and personal semantic memory in amnesic patients. J Clin Exp Neuropsychol 11: 724–744.
Kubie JL, Sutherland RJ, Muller RU. 1999. Hippocampal lesions produce a temporally graded retrograde amnesia on a dry version of the Morris swimming task. Psychobiology 27:313–330.
LePage M, Habib R, Tulving E. 1998. Hippocampal PET activations of memory encoding and retrieval: the HIPER model. Hippocampus 8:313–322.
Levine B. 1999. A list of potential personal experiences from childhood to old age. Personal communication.
McClelland JL, McNaughton BL, O'Reilly RC. 1995. Why are there complementary learning systems in the hippocampus ande neocortex: insights from the successes and failures of connectionist models of learning and memory. Psychol Rev 102:419–457.
Moscovitch M, Melo B. 1997. Strategic retrieval and the frontal lobes: evidence from confabulation and amnesia. Neuropsychologia 35: 1017–1034.
Moscovitch M, Nadel L. 1998. Consolidation and the hippocampal complex revisited: in defense of the multiple-trace model. Curr Opin Neurobiol 8:297–300.
Moscovitch M, Yaschyshyn T, Melo B, Ziegler M. 1998. Remote episodic and semantic memory in amnesia. Paper presented at the Brain, Behavior and Cognitive Science Meeting, Ottawa, Canada.
Moscovitch M, Yaschyshyn T, Ziegler M, Nadel L. 1999. Remote episodic memory and retrograde amnesia: was Endel Tulving right all along? In: Tulving E, editor. Memory, consciousness and the brain: the Tallinn Conference. New York: Psychology Press. p 331–345.
Murray EA, Bussey TJ. In press. Consolidation and the medial temporal lobe revisited: methodological considerations. Hippocampus.
Nadel L. 1991. The hippocampus and space revisited. Hippocampus 1:221–229.
Nadel L, Moscovitch M. 1997. Memory consolidation, retrograde amnesia and the hippocampal complex. Curr Opin Neurobiol 7:217–227.
Nadel L, Moscovitch M. 1998. Hippocampal contributions to cortical plasticity. Neuropharmacology 37:431–439.
O'Connor M, Marin M, Verfaellie M, Greenblatt D, Doherty R, Cahn G, Schomer D. 1998. Performance of temporal lobectomy patients on tests of remote memory. Paper presented at the Annual Meeting of the American Epilepsy Society.
O'Keefe J, Nadel L. 1978. The hippocampus as a cognitive map. Oxford: Clarendon Press.
Owen A, Milner B, Petrides M, Evans C. 1996. A specific role for the right parahippocampal gyrus in the retrieval of object location: a positron emission tomography study. J Cogn Neurosci 8:588–602.
Parkinson JK, Murray EA, Mishkin M. 1988. A selective mnemonic role for the hippocampus in monkeys: memory for the location of objects. J Neurosci 8:4159–4167.
Pearce JM, Roberts ADL, Good M. 1998. Hippocampal lesions disrupt navigation based on cognitive maps but not heading vectors. Nature 396:75–77.
Reed JM, Squire LR. 1998. Retrograde amnesia for facts and events: findings from four new cases. J Neurosci 18:3943–3954.
Rempel-Clower NL, Zola SM, Squire LR, Amaral DG. 1996. Three cases of enduring memory impairment after bilateral damage limited to the hippocampal formation. J Neurosci 16:5233–5255.
Rosenbaum RS, Priselac S, Kohler S, Black S, Gao S, Nadel L, Moscovitch M. Submitted. Studies of remote spatial memory in an amnesic person with extensive bilateral hippocampal lesions.
Ryan L, Nadel L, Keil T, Putnam K, Schayer D, Troward T, Moscovitch M. 2000. Hippocampal activation during retrieval of remote memories. Neuro Image 11:5396.
Schacter DL, Wagner AD. 1999. Medial temporal lobe activations in fMRI and PET studies of episodic encoding and retrieval. Hippocampus 9:7–24.
Snowden JS, Griffiths HL, Neary D. 1996. Semantic-episode memory interactions in semantic dementia: implications for retrograde memory function. Cogn Neuropsychol 13:1101–1137.
Squire LR. 1992. Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. Psychol Rev 99:195–231.
Squire LR, Cohen NJ, Nadel L. 1984. The medial temporal region and memory consolidation: a new hypothesis. In: Weingartner H, Parker E, editors. Memory consolidation. Hillsdale, NJ: Erlbaum. p 185–210.
Teng E, Squire LR. 1999. Memory for places learned long ago is intact after hippocampal damage. Nature 400:675–677.
Teyler TJ, DiScenna P. 1986. The hippocampal memory indexing theory. Behav Neurosci 100:147–154.
Tulving E. 1972. Episodic and semantic memory. In: Tulving E, Donaldson W, editors. Organization of memory. New York: Academic Press. p 381–403.
Viskontas IV, McAndrews MP, Moscovitch M. In press. Remote episodic memory deficits in patients with unilateral temporal lobe epilepsy and excisions. J Neurosci.
Westmacott R, Moscovitch M. In press. Autobiographical, episodic and semantic memory in semantic dementia and amnesia: a test of consolidation theory. Neurocase.

APPENDIX: ANALYTICAL MODEL OF MULTIPLE TRACE FORMATION

The dynamical variables of the model are: i) the variable-length list of time stamps $\{\tau_i : I = 1..L\}$ of all once created memories, and ii) the numbers of traces $\{n(\tau_i)\}$ associated with each listed time stamp. The total number of memories $L = |\{\tau_i\}|$, as well as each particular number of traces $n(\tau_i)$, are finite nonnegative integers. Thus, a dynamical state of the system is given as a set of tuples: $\{(\tau_i, n_i): I = 1..L\}$. In other terms, a dynamical state at a time t can be given by an integer-valued, discontinuous function $n(\tau)$ of one nonnegative variable. This function is everywhere zero, except a finite set of points $\{\tau_i : I = 1..L\}$.

Under the rather general and oversimplified assumptions 1–8 of the model (see above, Analytic Solution), the evolution of $n(\tau)$ over time t is given by two intensities: the intensity λ^+ of trace formation ($n \rightarrow n + 1$) and the intensity λ^- of trace elimination ($n \rightarrow n - 1$). These intensities are introduced as functions of time t, the

memory time stamp τ, and the number of traces n(τ, t), to which the contribution is made. The problem is formulated as follows:

$$\begin{cases} \lambda^{+}(\tau, t) = \delta^{0}(t-\tau) + \text{sign}(n(\tau, t))\,\dfrac{\alpha\rho(n(\tau, t), t-\tau)}{Z([n], t)}, \\ \lambda^{-}(\tau, t) = \kappa n(\tau, t), \\ n(\tau, 0) \equiv 0. \end{cases} \quad (2)$$

Here α is the total replication rate, which is constant in this model (assumption 5). The source function $\delta^0(t-\tau)$, which is 1 for $t = \tau$ and 0 otherwise, provides the constant, unitary rate of new memory creation (assumption 2). The factor sign(n) ensures that dead memories do not replicate (assumption 5). The specific replication rate function $\rho(n, t-\tau)$ will be specified below. The constant forgetting rate κ (assumption 4) is a small parameter that can be interpreted as the total trace formation rate ($\alpha + 1$) times the probability that a newly created trace will destroy a given trace by interference (a small constant), plus the relatively small constant rate of natural trace degradation. The last line in (2) is the initial condition for n (assumption 8). Finally, the normalizing factor Z in (2) is

$$Z([n], t) = \sum_{\tau>0} \text{sign}(n(\tau, t))\rho(n(\tau, t), t-\tau), \quad (3)$$

where the sum is actually taken over the set of all current memories $\{\tau_i: I = 1..L\}$. We consider 4 choices of the replication function ρ, corresponding to the 4 versions of the model, as described in the main text: "equalized over traces" (A), "equalizd over memories" (B), "saturation" (C), and "recency" (D).

$$\begin{aligned} &\text{A: } \rho(n, t-\tau) = n; \\ &\text{B: } \rho(n, t-\tau) = 1; \\ &\text{C: } \rho(n, t-\tau) = m - n; \\ &\text{D: } \rho(n, t-\tau) = exp\left(-\frac{t-\tau}{\sigma}\right). \end{aligned} \quad (4)$$

Here m and σ are additional parameters of the model (positive constants).

In a "mean field" approximation, the number of traces n is replaced by the expected mean number of traces μ per memory with a time stamp τ (the average $\langle \ldots \rangle$ is taken over the statistical ensemble of identical systems):

$$\begin{aligned} n(\tau, t) &\rightarrow \mu(\tau, t) = \langle n(\tau, t) \rangle, \\ f(n) &\rightarrow f(\mu) \approx \langle f(n) \rangle. \end{aligned} \quad (5)$$

As a result, the problem (2)–(4) is reduced to the following initial value problem for the kinetic equation on μ:

$$\begin{cases} \dfrac{\partial}{\partial t}\mu(\tau, t) + \kappa\mu(\tau, t) = \alpha\theta(t-\tau)\,\dfrac{\rho(\mu, \tau, t)}{Z([\mu], t)} + \delta(t-\tau); \\ \mu(\tau, 0) = 0. \end{cases} \quad (6)$$

Here θ is the Heaviside step function, δ is the Dirac delta-function (the derivative of θ), and the normalizing factor is

$$Z([\mu], t) = \int_0^t d\tau\rho(\mu(\tau, t), t-\tau). \quad (7)$$

The problem (6)–(7) for the average number of traces per existing memory μ can be solved analytically. The solution for the case A is:

$$\mu(\tau, t) = \theta(t-\tau)\exp \times \left\{ \frac{\alpha}{1+\alpha}\log\left(\frac{\sinh\frac{\kappa t}{2}}{\sinh\frac{\kappa\tau}{2}}\right) - \frac{\kappa(2+\alpha)}{2(1+\alpha)}(t-\tau) \right\}, \quad (8)$$

where sinh is the hyperbolic sinus. The solution for the case B is:

$$\mu(\tau, t) = \theta(t-\tau)\{e^{-\kappa(t-\tau)} + \alpha e^{-\kappa t}[\text{Ei}(\kappa t) - \text{Ei}(\kappa\tau)]\}, \quad (9)$$

where Ei is the integral exponent. For the case C we have a solution in quadratures:

$$\mu(\tau, t) = \alpha e^{-\beta b}\int_a^b \frac{dx}{x + ce^{-\beta x}}\exp\left[\beta x - \frac{\alpha}{m}\varphi(x, b, c, \beta)\right] + \exp\left[\kappa(\tau - t) - \frac{\alpha}{m}\varphi(a, b, c, \beta)\right]; \quad (10)$$

where

$$\begin{aligned} \varphi(a, b, c, \beta) &= \int_a^b \frac{dx}{x + ce^{-\beta x}}; \\ a = m\tau - \frac{1+\alpha}{\kappa};\ b &= mt - \frac{1+\alpha}{\kappa}; \\ c = \frac{1+\alpha}{\kappa}\exp\left(-\frac{1+\alpha}{m}\right);\ \beta &= \frac{\kappa}{m}. \end{aligned} \quad (11)$$

Finally, for the case D we have the following solution in quadratures:

$$\mu(\tau, t) = e^{\kappa(\tau-t)} + e^{\tau/\sigma-\kappa t}\alpha\int_{e^{\tau/\sigma}}^{e^{t/\sigma}} dx\,\frac{x^{\kappa\sigma-1}}{x-1}. \quad (12)$$

The probability of recall P as a function of the amount of lesion q is given by (1), (5):

$$p(\mu, q) = 1 - q^{\mu}, \quad (13)$$

with μ given by (8)–(12).

Comparison of the results (8)–(12) with the curves obtained numerically based on a more rigorous approach (Fig. 1, and see below) shows good quantitative agreement everywhere, except the singularity region, where a good qualitative agreement is nevertheless observed. Especially good agreement is observed for the cases C and D, where the maximal discrepancy is at a level of few percent or below that level. Therefore, the results (8)–(12) can be accepted as good analytical approximations for the average number of traces

per memory in this model. However, due to the essential nonlinearity of (1), the results (8)–(13) do not provide everywhere accurate data for the retrieval probability P, in all cases (A–D).

Therefore, we shall consider a more rigorous approach, based on description of model dynamics in terms of a probability distribution function $f(\tau, t) = \{f_n(\tau, t)\}$, which is defined as the probability that an existing memory τ will have n traces at a time t. Assuming self-averaging of Z (3) at a large number of memories $L \sim t$,

$$Z(t) = \sum_{i=1..L} \mathrm{sign}(n_i)\rho(n_i, t - \tau_i) \cong \langle Z(t)\rangle$$

$$= \sum_{n=1}^{M} \int_0^t d\tau \rho(n(\tau, t), t - \tau) f_n(\tau, t),$$

$$t \gg 1, \qquad (14)$$

where M is the maximal possible number of traces per memoty, the original problem (2)–(4) can be reformulated as follows:

$$\begin{cases} \dfrac{\partial}{\partial t} f = Wf,\ t > \tau; \\ f_n|_{t=\tau+0} = \delta_{n,1};\ f_n|_{t<\tau} = \delta_{n,0}; \end{cases} \qquad (15)$$

where the matrix W is given by

$$W_{i,j}(\tau, t) = (\delta_{i-1,j} - \delta_{i,j}) a_j(\tau, t) + (\delta_{i+1,j} - \delta_{i,j})\kappa j;$$

$$a_j(\tau, t) = \frac{\alpha}{Z(t)} \rho(j, \tau, t)(1 - \delta_{j,0})(1 - \delta_{j,m}). \qquad (16)$$

Here δ is the Kronecker delta, and the expression for Z (3) should be replaced by the second line of (14). With a discretization of the time stamp τ, the problem (14)–(16) can be further reduced to an initial value problem for a system of ordinary differential equations, the numerical solution of which produces the results represented in Figures 1 and 2.

To finish up with the analytical model, we provide an alternative derivation of the formula (1), based on the following general assumption. We assume that the system has M possible trace locations (in some abstract space), each of which can be occupied by only one trace. We further assume that the fraction q of these abstract locations is destroyed by a lesion, and the rest remain intact. Now, given the actual number of traces n of a memory before lesion, what are the chances P(n, q) that at least one of these traces will remain in the system after a random lesion? This is a combinatorial problem, and the answer is

$$p(n, q) = 1 - \frac{\Gamma(M + 1 - n)\Gamma(qM + 1)}{\Gamma(M + 1)\Gamma(qM + 1 - n)}, \qquad (17)$$

where Γ is the gamma-function. Assuming that M is large as compared to n, we arrive at (1).