



Review

Hippocampal contributions to cortical plasticity

L. Nadel ^{a,*}, M. Moscovitch ^b^a *Department of Psychology, University of Arizona, Tucson, AZ 85721, USA*^b *Department of Psychology, Erindale Campus, University of Toronto, Mississauga, Ontario, Canada*

Accepted 23 March 1998

Abstract

The hippocampal complex and neocortex are both integrally important in memory function, in particular as regards memory for episodes and knowledge about the world that is derived from them. It is traditionally assumed that the role of the hippocampus is time-limited, after which retrieval of episodic memory depends only upon neocortical stores. A number of lines of evidence indicate that this traditional view is incorrect. We propose that the hippocampal complex is always necessary for retrieval of episodes and their contextual frame, and that hippocampal-neocortical interactions contribute instead to the extraction of semantic information to be stored in the neocortex. © 1988 Elsevier Science Ltd. All rights reserved.

Keywords: Cognitive maps; Hippocampus; Memory consolidation; Neocortex

1. Introduction

It has been clear for at least 40 years that the hippocampal complex and neocortex are interactively involved in memory function. While the exact nature of their interaction remains unclear, there is a current view that proposes the following: the hippocampal complex plays a critically important role in memory for some, but not all, of the life of a given memory. Initially essential, the hippocampal complex becomes less important with time, and ultimately only the neocortex is needed for accurate retrieval of information acquired in the past. During the period in which the hippocampal complex is essential, it interacts with the neocortex in ways that promote the ultimate independence of the neocortical memory store (e.g. Squire, 1992). Although widely accepted, this view is not in fact well supported by all relevant empirical data. This paper discusses an alternative view, asserting that the hippocampal complex is always important in reconstructing a specific episodic memory because it holds the traces representing the spatial scaffold within which the elements of experience are retrieved (Nadel and Moscovitch, 1997).

The current view speaks to our intuition that recent memories are labile while older ones are not, but it leaves many important questions unanswered. Consider the following list of mysteries:

- why are multiple brain systems engaged in memory?
- what form does a specific memory take in the brain?
- how is a memory of something we did yesterday or last month retrieved?
- to what extent is any memory a veridical replay of the past?
- how are memories influenced by subsequent experience?
- what happens to memories over long periods of time?

Attention has been paid to some of these questions, but the answers provided to date are rather vague. Thus, it seems clear, following Hebb's insights, the discovery of long-term potentiation (Bliss and Lomo, 1973), and subsequent studies linking LTP to behavioral change (e.g. Barnes, 1979; Morris, 1989), that memories involve alterations in synaptic efficacy, but little is known about how such synaptic plasticities occur across an ensemble of thousands of neurons so as to capture the richness of a single memory. Similarly, it is known that memories change over time, but we do not know how. It is not known if all aspects of memory change in the same way over time. Finally, it is thought that memories are

* Corresponding author.

reconstructed rather than replayed (e.g. Bartlett, 1932), but little is known about how this fundamental reconstructive process works, nor how veridical any given memory is. Some ingenious experiments have addressed these issues, but a satisfactory conceptual framework is lacking.

Matters are not much better when we consider the neural bases of memory. The neocortex and hippocampal complex are surely important, but various aspects of memory appear to require several other structures in the brain. It is said that there are two major kinds of memory systems, but there are indications that this is a considerable oversimplification (cf. Schacter and Tulving, 1994 and many of the chapters therein). Until recently, we were still talking as if there were memory traces (engrams) in the brain that could be pointed to if only we were smart enough, or had tools powerful enough.

This paper will focus on the special roles of, and interactions between, the hippocampal complex and neocortex. Little will be said about the memory functions of other brain structures and, in consequence, the focus will be on episodes and the knowledge about the world that we gain through these episodes. It is generally agreed that the hippocampal complex and neocortex are not essential for more ‘procedural’ or ‘reflexive’ types of memory.

2. An example

One sunny April afternoon you are walking on rue Tolbiac with your friend when you stroll past a shop selling many varieties of cheese, especially goat cheese. You stand at the window, commenting on the variety, and the names and prices of your favorites. The smells of the shop permeate the air, and the thought of buying a baguette from the bakery across the street runs through your mind.

Consider this quite average event, not highly charged or salient perhaps, but nonetheless memorable if you enjoy cheese in general, and goat cheese in particular. What might you remember about it sometime in the future, say, the next day when your spouse asks what you and your friend did on your walk yesterday?

You might remember:

where you walked	(rue Tolbiac)
when you walked	(mid-afternoon)
what the weather was like	(overcast—this is Paris after all)
who you were with	(a particular friend)

what you said	(the names and prices of the various cheeses)
what you saw	(the street, the shop, the cheese)
where you saw what you saw	(the exact location of the best cheese)
what you smelled	(the cheeses)
what you were thinking	(buying a baguette to go with some cheese)

But this is not all. You might also have noticed that many of the cheeses were called by similar names (cabacou, or crottin), and inferred that these names refer not to a specific cheese but rather to a class of cheeses. Further, you might have concluded that your friend was a great deal less interested in cheese than you were, given his reaction to the shop, and several similar interactions you had with him over the past month. These facts would also be part of the new knowledge you acquired as a result of this experience, although neither was a part of the episode in any simple way.

This example conveys two important things about episodic memory. First, it suggests that any episode memory is a melange of many bits and pieces. We argue below that this has both functional and structural implications. Second, it indicates that the knowledge acquired during an episode goes beyond the specific elements of that episode to include broader knowledge gained as a result of having had that experience. This poorly understood facet of episodic memory is central to any attempt to conceptualize the relation between the hippocampal complex and the neocortex.

3. Collecting the bits and pieces of episodic memory

If one looks in the dictionary for a definition of memory one finds the terms ‘recollection’ and ‘remembering’ prominently featured. What is interesting about these terms is the sense they convey of memory as a faculty dependent upon putting back together again the pieces of the past. Recently, Jacobs and Nadel (in press) extended the notion of ‘memory as reconstruction’ (Bartlett, 1932; and many others since) by arguing that the various parts of an experience are structurally disaggregated in the brain for purposes of storage, and must be re-aggregated when one attempts to retrieve an episode memory. This argument is based on two general facts: firstly, the consensually held view that the brain is organized to process, represent and store different kinds of information in structurally distinct systems (e.g. Nadel, 1992, 1994); and secondly, the evidence from a variety of syndromes resulting from brain dam-

age or disorder that individuals can lose access to just certain aspects of episode memories without losing access to the remainder.

If this perspective is valid, it is reasonable to conclude that retrieving an episode memory is necessarily a reconstructive act, and one that can, even under normal conditions, lead to only partial success. Thus, one can retrieve an episode memory for the ‘cheese shop stroll’ that lacks details about the smells, the sights, the presence of the friend—indeed many aspects of the episode can be lost and the episode still survive as a recognizable memory. Of course it is essential that some content survive, otherwise there would be no episode for the memory to be about. In this example, one could retrieve a memory about a cheese stroll (absent everything else), or a memory about a stroll with a friend (absent everything else) and so on. What this analysis shows is that while no particular content is essential, some content is necessary. There are, however, some features of episode memory that are not optional in this way, but instead are part of it by definition.

To start with, episodic memory is always about the past. When we engage in an act of remembrance it is with the clear understanding that what we are thinking about happened some time ago. It is not necessarily always this way—one plausible account of the intrusive ‘memory’ phenomena associated with post-traumatic stress disorder (PTSD) is that events from the past are re-experienced rather than referenced to the past and hence remembered (Nadel and Jacobs, 1996; Jacobs and Nadel, in press). The fact that this can happen at all suggests that the ‘pastness’ of a memory is not automatic but instead resides in some feature of the memory system itself. A theory of memory needs to account for how and why we know that something happened in the past. Because those few careful studies of our ability to date autobiographical memories (Linton, 1978; Wagenaar, 1986) suggest that temporal information is quite impoverished, we can suppose that the sense of pastness is not derived simply from a time-tag, but rather from something that differentiates a past episode from a present occurrence. Insofar as some interactive activation of hippocampal and neocortical ensembles characterizes both a retrieval of past events and the processing of current events, the genesis of this ‘pastness’ signal is not obvious.

Not only does a remembered episode occur in the past—it also occurs in a particular location in space. Episodes happen somewhere, and although our memory for all the details of the spatial setting may be spotty, this too is a necessary component of all episodic memories. In comparison to temporal tags, spatial tags are relatively effective as links to stored autobiographical memories (cf. Wagenaar, 1986).

This is not to say that recalling an episode without appropriate spatio-temporal contextual information is

impossible, because we are all familiar with instances of exactly that. It is not uncommon to recall an event from the past without a clear sense of where and when it happened. In experimental paradigms currently used to assess episodic memory, such cases would be considered ‘weak’ episodes and scored as defective. This is because information about where and when an event transpired is considered, by definition, to be an essential part of episodic memory.

Thus, an episodic memory is made up of parts, some mandatory, some optional, and each one likely retrieved from a different ‘knowledge module’ in the brain. This suggests several critical questions in pursuit of a theory of episodic memory:

- what are the modules, and which brain structures subservise each module?
- how is the information in each module represented?
- how do these representations change over time—either passively as a function of decay, or actively as a function of interactions with other representations in that module, previously or subsequently stored?
- how is a single episode, comprised as it is of dis-aggregated bits and pieces of information represented in disparate modules, re-aggregated in the process of retrieval?

4. Memory as a connection

Start with the last question first: many investigators (O’Keefe and Nadel, 1978; Squire et al., 1984; Nadel et al., 1985; Teyler and DiScenna, 1986; Moscovitch, 1995; Metcalfe and Jacobs, 1998) would agree that the hippocampus plays a kind of ‘binding’ role with respect to the elements of an episode spread out in different neocortical systems. That is, it provides a mechanism by

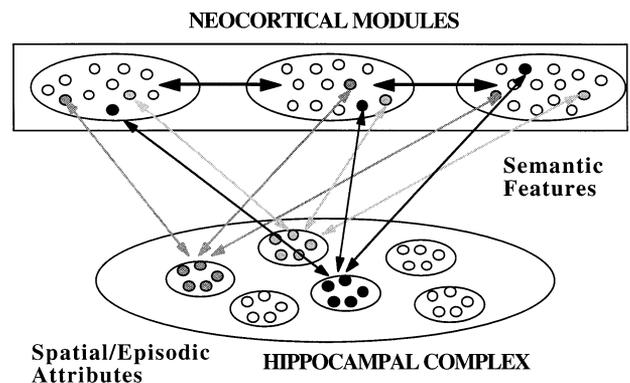


Fig. 1. The roles of neocortical modules and the hippocampal complex in the storage of episodic information. Representations of the elements of the episode are stored in dispersed neocortical modules. The elements of each episode jointly link to a hippocampal ensemble that contributes spatial information. This hippocampal trace serves as a means by which the dispersed elements comprising the content of an episode can be reknit together in an act of remembrance.

which the disaggregated bits of information comprising an episode can be kept in touch with one another. All would agree that the hippocampus is essential to this function for recent memories; its role in the retrieval of remote memories is a matter of considerable current debate (cf. Nadel and Moscovitch, 1997). Exactly how would the hippocampus play this role? Fig. 1 portrays one way of thinking about it.

Among the key points of this proposal are: (1) various aspects of an episode memory are represented and stored in dispersed neocortical modules, e.g. “concepts and categories, the look, the feel and the sound of things, the goodness or badness of objects” (O’Keefe and Nadel, 1978, p. 100); (2) each neocortical module interconnects with the hippocampal complex, such that the collection of representations comprising the features of an episode can create an ensemble within the hippocampal complex that encodes that episode¹; (3) this creation of a hippocampal ensemble (what O’Keefe and Nadel referred to as a ‘cognitive map’) occurs very rapidly, through the mechanism of long-term potentiation within the relevant hippocampal synaptic connections; (4) an act of episode retrieval can be accomplished in two ways: first, by activation of the relevant hippocampal ensemble, which then activates the dispersed, neocortically located, features; or by activating some subset of the dispersed neocortical feature-traces, which then activate the hippocampal ensemble.

These properties account for many of the basic facts of episodic memory as we know them, but they do not explain why spatial information should play such a central role. O’Keefe and Nadel (1978) first suggested that these spatial contextual attributes are contributed

¹ Exactly what is meant by ‘encodes’ in this context is left vague, so that we can skirt the issue of whether the information is ‘in’ the hippocampus or not. Either the information is in the hippocampus or the hippocampus ‘knows’ where to go in the neocortex to get it.

² O’Keefe and Nadel (1978) supposed that this spatial information was intrinsic to the hippocampal circuitry itself—recent results suggest that this might not be the case (e.g. O’Keefe and Burgess, 1996). An alternative possibility is that the spatial information is conveyed to the hippocampal circuits from an external source, perhaps the parietal neocortex. In either case, the hippocampus ‘contains’ this spatial information, and imposes it upon non-spatial contents derived from other neocortical sites.

³ This proposal is clear about the encoding of spatial relationships within an episode, but less clear with respect to temporal relations, hence the causal flow of events. It is reasonable to assume that in the process of reading an episode into (or out of) the hippocampus the temporal sequence information is accounted for—and likely that the frontal cortex is important in this regard. On the other hand, this proposal assumes that other relational information (such as the fact that one kind of goat cheese was softer than another) is not contained within the hippocampal trace, but instead within the relevant neocortical traces. In this sense the proposal differs markedly from the ‘general-purpose’ relational view espoused by Eichenbaum et al. (1994).

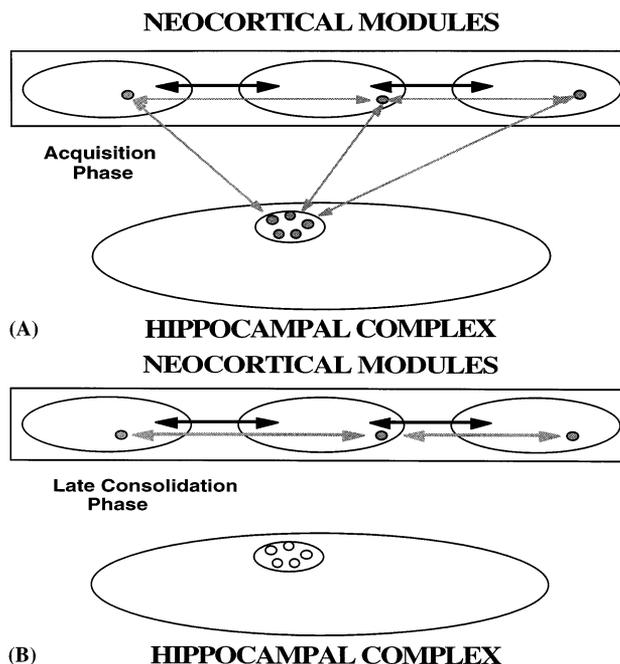


Fig. 2. (A) The acquisition of a memory within hippocampal and neocortical circuitry according to the ‘standard’ model of memory consolidation. Aspects of the memory are captured in dispersed neocortical modules. Each module communicates with the hippocampus, permitting the rapid establishment of a hippocampal ensemble trace. (The thick arrows linking the neocortical modules symbolize the neocortico-neocortical connectivity that links together the various modules.) The standard model assumes that direct neocortical–neocortical linkages are strengthened over time, perhaps during sleep (e.g. Wilson and McNaughton, 1994). The standard model assumes that as direct neocortical–neocortical connections emerge, neocortical–hippocampal connections weaken. (B) The final stage of memory consolidation within the standard model. Neocortical–neocortical connections have been firmly established, and neocortical–hippocampal connections have dissipated.

by the hippocampal circuitry itself. It seems likely that this attribution of a hippocampal role in temporal context was misplaced; however, the centrality of the hippocampal complex in representing spatial context is well established (cf. Nadel and Willner, 1980; Nadel et al., 1985; Nadel, 1991).

What does it mean to say that the hippocampal circuitry ‘contributes’ spatial contextual information? One possibility is that spatial information is represented in the internal circuitry of hippocampal ensembles², such that relations in space between aspects of a scene (distances and directions) are coded by the connections between the particular hippocampal elements that enter into the ensemble. This ensemble, then, acts as a kind of spatial scaffold, onto which the non-spatial elements of an episode are thrust. The combination gives rise to an episode, complete with spatial context and contents³.

It is in this sense that one can view the hippocampal role in memory as providing the basis for a connection among the various parts of a collection of elements. In so doing it enriches this collection with spatial at-

tributes that capture this essential aspect of any episode. An act of episodic remembering therefore must involve activity within both the neocortical memory stores and the hippocampal map/ensemble if it is to convey detailed contextual information; this squares with recent neuroimaging studies indicating widespread activation of hippocampal and cortical sites during retrieval of episodic memories (e.g. Fink et al., 1996; Fernández et al., 1998).

5. Memory dynamics over time

The above lays out a way of thinking about episodic memory, and its components, in the phase of acquisition of episodic information. What happens over time after the storage of an episode memory is the domain of memory consolidation. We have recently revisited the evidence held to support the standard model of memory consolidation (Nadel and Moscovitch, 1997; Moscovitch and Nadel, 1998). This model, to which we alluded at the outset, holds that the hippocampal complex is only ‘temporarily’ involved in memory storage, after which the neocortical system suffices for retrieval. The gist of the standard model is portrayed in Fig. 2(A) and (B).

We argued the following: given that current estimates of the length of retrograde amnesia after lesions to the hippocampal complex are of the order of 15–25 years, it makes little sense to think of the hippocampal complex as playing a temporary role in memory storage. By such an argument, consolidation of any individual memory might be taking place throughout the entire life of the organism. Instead, we proposed a new model, called ‘multiple-trace theory’ that explains the phenomenon of retrograde amnesia in another way. It

supposes that over time, as memories are re-activated and rehearsed, multiple memory traces are formed within the hippocampus. As each trace involves a widely dispersed hippocampal ensemble, the net effect is that older memories will, on average, involve ever greater extents of the hippocampal complex. Thus, it will take larger lesions to knock out older memories. The gist of this model is portrayed in Fig. 3.

Within multiple-trace theory the hippocampal complex plays a permanent role in the storage of certain aspects of an episodic memory, namely, the contextual elements discussed above. To the extent to which such information is preserved, the hippocampal complex will be involved in retrieval.

The evidence supporting this model comes from studies of retrograde amnesia in both humans and animals, and is summarized in Nadel and Moscovitch (1997). Even though multiple-trace theory rejects the claim of the standard model that the hippocampal complex only plays a temporary role in memory, it nonetheless accepts the view that there is an interesting and important interaction between the hippocampal system and neocortical systems in the time after memories are established. This interaction is at the center of the process by which semantic information is derived from episodic experience—perhaps the most important unsolved mystery of memory.

6. From episodic to semantic memory

Defining semantic knowledge is not simple. Part of the difficulty results from the historical fact that the term ‘semantic’ has been used quite differently by investigators concerned with language on the one hand and memory on the other. Linguists use the term to refer to meaning—the meaning of words, for example. It is contrasted to syntax, or meaning derived from the order with which words are used. Psychologists studying memory have used the term to stand for knowledge about the world that is divorced from the context in which it was acquired. This includes the learned meaning of words, but also includes facts about the world, knowledge about the features of specific objects, the names attached to faces of people we know, and so on. Indeed, just about any kind of memory that is neither episodic nor procedural qualifies as semantic.

Whereas all episodic knowledge has certain defining features, there are no such obvious characteristics common to all examples of semantic knowledge. In some cases semantic knowledge can be acquired in a single trial (e.g. facts such as London is the capital of England). In other cases semantic knowledge only emerges after multiple presentations, so that the statistical properties of the input can be accurately divined (e.g. con-

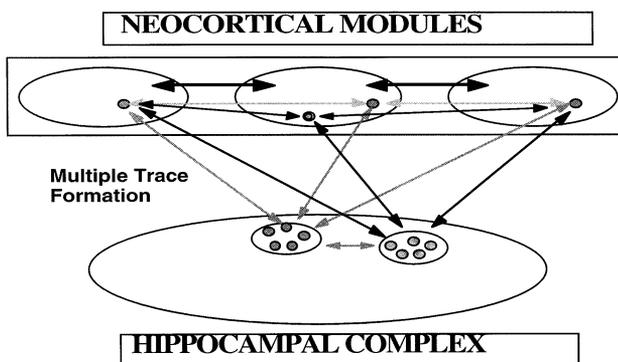


Fig. 3. Neocortical–hippocampal interactions within multiple-trace theory. Over time, a memory is re-represented within both hippocampus and neocortex. The new hippocampal trace involves a randomly chosen (hence largely non-overlapping) set of hippocampal neurons. The new neocortical trace involves many, but not all, of the same neocortical elements, resulting from the fact that the new trace of an old memory likely reflects some minor differences in the content of the memory.

cepts such as the presence of certain kinds of clouds indicating a high likelihood of a thunderstorm)⁴.

It is a general assumption that much semantic knowledge derives from a series of episodes in the course of which semantic structure is extracted from experience. This, however, cannot be the complete story, because there are at least two circumstances in which semantic knowledge can be acquired in the absence of a functioning episodic system—early infancy prior to the maturation of critical components of the hippocampal complex (Nadel and Zola-Morgan, 1984; Schacter and Moscovitch, 1984), and in neuropsychological cases of temporal lobe amnesia where episodic memory can be devastated but acquisition of semantic memory remains possible (e.g. Vargha-Khadem et al., 1997). It is instructive to consider these cases more closely.

It is well established that infants are incapable of forming lasting autobiographical episode memories until the age of 2 or so. There are various theories of this ‘infantile amnesia’, but one likely component of at least part of this syndrome is neurological immaturity. Parts of the hippocampus (e.g. the dentate gyrus) are still developing for several years after birth, and a reasonable guess is that the hippocampal complex is not functional until 18–24 months of age. This assertion is based on behavioral evidence indicating that the spatial learning dependent on the hippocampus is not in evidence in infants until that age (Mangan and Nadel, 1990; Newcombe et al., 1998). In the present context what is interesting about this phenomenon is that infants between 0 and 18 months of age are obviously capable of a great deal of semantic learning. The same has been shown to be true in cases of organic amnesia, whether observed in children or in adults. Vargha-Khadem et al. (1997) have carefully investigated a number of cases of children with organic amnesia resulting from hippocampal complex damage, and documented the extensive semantic learning capabilities of these subjects.

But exactly what kinds of semantic learning are possible in the absence of the hippocampal complex, and does this learning progress in the same fashion, and at the same pace, as it would in the presence of this system? In order to approach this question, we have to think about the nature of knowledge representations within the semantic memory system. What form do they take, and how are they formed?

⁴ To make matters more complicated, many investigators now separate both episodic and semantic memory into two classes, depending on whether or not the information is autobiographical. Most episodic information is autobiographical, though we certainly can store memories about events and episodes that we have read about, or have seen in the movies or on TV. Personal semantic information includes facts about our own lives (where we lived, went to school, etc.) as well as trait information (e.g. personality characteristics such as friendliness, hostility, and so on).

7. Varieties of semantic knowledge

Semantic knowledge and episodic knowledge share one very important feature: they are both examples of what Ryle called ‘knowing that’, as distinct from ‘knowing how’. What distinguishes episodic and semantic knowledge (and hence memory) is the fact that episodes are about specific events, and semantic knowledge cuts across such events. As noted, semantic knowledge comes in several varieties, such as the definitions of words, the nature of things such as dogs and cats, the meaning of actions such as eating, drinking and jumping, the properties of events such as going to class, attending a party, eating out at a restaurant, and the nature of complex abstract belief states and social systems. Is it reasonable to assume that all these kinds of knowledge are represented in the same way in the neocortex?

It is beyond the scope of this paper to tackle all, or even most, of the outstanding questions about semantic knowledge and how it is represented in the mind. For present purposes we concentrate on a comparison between knowledge about things, and knowledge about events, because this comparison will bring into focus the distinction between semantic and episodic knowledge. Consider the representation of a concept like ‘goat cheese’. Most would agree that concepts like this are represented by a network of elements that, in their turn, represent the various features of goat cheese such as its smell, texture, uses, color and so on. There is an ongoing debate about whether such concepts are represented in terms of exemplars, prototypes, or some combination of the two. We need not be concerned with this debate here, though the distinction between exemplars and prototypes is important in discussions of episodic memory.

Consider the representation of an event, perhaps the ‘goat cheese stroll’ discussed above. Here too, it is generally agreed that an event is represented by a network of elements that, in their turn, represent the various bits and pieces that made up that event, as we have already seen. A consideration of exemplars and prototypes is important in this case, because it distinguishes between the representation of specific episodes (exemplars), replete with information about where and when, and that of generic ‘scripts’, which capture the

⁵ Episodic memory may itself be organized hierarchically much as semantic memory is. A specific episode may stand in the same relation to semantic knowledge of higher-order events such as life periods, occupations one has held, event scripts, and even stylized versions of the specific episode itself (all of which can be considered ‘personal semantics’ (Kopelman et al., 1989), as do unique exemplars to generic exemplars and to the superordinate semantic categories to which they belong. Access to a specific episode may depend on activating the semantic event hierarchy in which it is nested (cf. Conway, 1992; Hodges and McCarthy, 1995).

generalized features of all the ‘goat cheese strolls’ one might have taken, but lose the specific time and place of all of them. This ‘prototype’ is one kind of semanticized version of episode memories, and one presumes that such prototypes can be readily created if each occurrence of an object or event is stored in content-addressable networks. This guarantees that similar objects/events will be stored in roughly the same circuits, thereby leading quite directly to prototypes⁵. A single memory of an episode that has lost its connection to time and place is another sort of semanticized memory, and in this case there must be some mechanism for accurately re-aggregating the neocortical ensembles. However, are all neocortically based event memories deprived of their spatio-temporal specificity, or can there be neocortical event memories for specific episodes. The answer to this question will determine whether or not the hippocampal complex is always, or only temporarily, essential for episodic memory. We believe the former to be the case (Nadel and Moscovitch, in preparation).

8. How episodes give rise to concepts and categories

If we assume that every episode gives rise to both dispersed neocortical representations and a hippocampal ensemble that serves to bind them together in the ways described above, then there are several possible mechanisms by which the hippocampal episode trace could contribute to neocortical semantic representations:

- the hippocampal trace could serve as a basis for reactivating the collection of neocortical traces, thereby facilitating the creation (or strengthening) of links between these dispersed traces;
- the hippocampal trace could cause, or facilitate, the creation of a new ‘node’ that would have links to the neocortical elements involved in the original memory.

Both of these possibilities are open with regard to the issue of exemplar versus prototype in the representation of semantic event knowledge in the neocortex. One could argue that part of what gets established in neocortex over time is the spatio-temporal information that was originally contributed by the hippocampal complex. Or, one could argue, as we do, that only the semantic knowledge can be established in the neocortex, and the spatio-temporal framing always requires the contribution of the hippocampal complex⁶.

In this latter case, the contribution of the episodic system is then limited to facilitating the strengthening of links, or new nodes, among elements generically representing the features of things, actions, events, belief states and so on. However, if this is the case, why should any such episodic → semantic interaction be needed at all? Why could the semantic system not accomplish this learning by itself throughout life, much as it apparently does in early life, or in the presence of organic damage to the episodic system?

McClelland and colleagues have recently suggested an intriguing answer to this question (McClelland et al., 1992, 1995). They suppose that the episode system serves as a means by which multiple presentations of events can be fed to the semantic system, thereby allowing the latter to extract the statistical regularities slowly. One of the primary arguments in favor of this idea is that rapid extraction of such regularities (e.g. by the semantic system acting alone) could readily lead to considerable interference within previously stored concept representations. Thus, slow extraction permits both the acquisition of new conceptual information and the preservation of previously acquired concepts. In this view, hippocampal ensembles simply re-activate the pattern of neocortical activations, thereby increasing the strength of the connections among the parts of the neocortical pattern. This would have the effect of improving the linkages among elements within a semantic network. As McClelland et al. (1995) point out, there could be a developmental twist to this story: early in life when there is not a great deal of knowledge yet in the system, the incremental intervention of the episodic ‘teacher’ could be less important, since large-scale shifts in ‘connection weights’ within the existing (relatively impoverished) networks would be less troubling. As a corollary, plasticity functions within neocortical circuits could be ‘set’ at a relatively high value early in life, thereby allowing more rapid acquisition in the absence of the episodic ‘teaching’ input.

In contrast to the position spelled out here, McClelland et al. (1995) argued that all of the information contained in the episodic system can be ‘re-created’ within the semantic system, including the spatio-temporal contextual information contributed by the hippocampal complex. However, the various authors now have diverging views on this issue (cf. McClelland and Goddard, 1996; Samsonovitch and McNaughton, 1997), the latter agreeing with the position staked out here: namely, that spatial information is always dependent upon the hippocampal complex. While this is a crucial distinction between these positions with regard to the fate of episode-specific information, with respect to the more general implications of the episodic-se-

⁶ Part of the argument for this position resides in a distinction between the form of representation in the neocortex (similar objects/events processed in overlapping ensembles of neurons, thereby creating prototypes) and in hippocampus (where episodes are represented in multiple traces that employ orthogonally separated ensembles of neurons).

mantic interaction much agreement remains. That is, the episode ensemble trace acts to reactivate neocortical traces representing the various things/actions that comprised the episode. This reactivation permits the strengthening of links among these features, and hence learning about concepts, categories, and even the generic structure of event-types. In this way, episodic memory helps to consolidate semantic knowledge, and does so in a time-dependent fashion. This fact helps explain why following extensive lesions to the hippocampal complex, retrograde amnesia for personal or general semantic knowledge can be graded, while for episodic information it is not⁷.

9. Conclusion

The hippocampal episode memory system makes a very important contribution to neocortical plasticity. It provides the linkages through which dispersed neocortical traces can be co-activated in a cohesive, collective fashion. In so doing it contributes the spatial and contextual information central to any specific episode memory. At the same time, and in the course of this re-activation of neocortical traces, it facilitates the strengthening of links among these dispersed traces. As a consequence, concepts, categories and a great deal of generic knowledge results. This process can be called consolidation, though it represents a somewhat different sense of the term than that implied in the 'standard' model. Present evidence suggests that such consolidation does not include establishment of a full episodic trace in neocortex—the hippocampal complex is always needed both for its encoding and its subsequent retrieval.

Acknowledgements

We thank Mary Peterson and Richard Morris for

⁷ Although retrograde amnesia for episodes is very extensive when damage to the hippocampal complex is severe, in many cases some episode-like memories survive. What we do not know is whether these episode-like memories are equivalent to those observed in normal subjects. Even if they are, the possibility remains that for both normals and amnesics these memories have lost the spatial contextual information necessarily provided by the hippocampal system. Recent animal model work supports the latter supposition (Kubie et al., submitted) in showing that although some information survives hippocampal lesions in rats when the lesions are made a long time after learning a spatial task (14 weeks in this case), the kind of information that survives is different than the kind of spatial information normally available within the hippocampal system. The authors suggest that while the hippocampal system provides a map-like representation of spatial information, neocortical systems can only support a form of spatial information they refer to as 'vector-field'. Should this proposal turn out to be correct, it might explain why some amnesics can demonstrate remote spatial memories.

their insightful comments on earlier drafts, and Rich Zemel for his discussion of some of the issues. We are indebted to Endel Tulving for his innovative work on episodic memory. It was only after we finished the manuscript that we became aware of how much it bears the marks of his influence. L.N. also wishes to acknowledge the contributions of W.J. Jacobs to his thinking on many of these issues. L.N. was supported by the McDonnell Foundation, and M.M. by the Natural Sciences and Engineering Research Council of Canada.

References

- Barnes, C.A., 1979. Memory deficits associated with senescence: A neurophysiological and behavioral study in the rat. *J. Comparative Physiol. Psychol.* 93, 74–104.
- Bartlett, F.C., 1932. *Remembering: A Study in Experimental and Social Psychology*. Cambridge University Press, Cambridge, UK.
- Bliss, T.V.P., Lomo, T., 1973. Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J. Physiol.* 232, 331–356.
- Conway, M.A., 1992. A structural model of autobiographical memory. In: Conway, M.A., Rubin, D.C., Spinnler, H., Wagnaar, W.A. (Eds.), *Theoretical Perspectives on Autobiographical Memory*. Kluwer Academic Publishers, Dordrecht, pp. 167–194.
- Eichenbaum, H., Otto, T., Cohen, N.J., 1994. Two functional components of the hippocampal memory system. *Behav. Brain Sci.* 17, 449–518.
- Fernández, G., Weyerts, H., Schrader-Bölsche, M., Tendolkar, I., Smid, H.G.O.M., Tempelmann, C., Hinrichs, H., Scheich, H., Elger, C.E., Mangun, G.R., Heinze, H.-J., 1998. Successful verbal encoding into episodic memory engages the posterior hippocampus: A parametrically analyzed functional magnetic resonance imaging study. *J. Neurosci.* 18, 1841–1847.
- Fink, G.R., Markowitsch, H.J., Reinkemeier, M., Bruckbauer, T., Kessler, J., Heiss, W.-D., 1996. Cerebral representation of one's own past: neural networks involved in autobiographical memory. *J. Neurosci.* 16, 4275–4282.
- Hodges, J.R., McCarthy, R.A., 1995. Loss of remote memory: A cognitive neuropsychological perspective. *Curr. Opin. Neurobiol.* 5, 178–183.
- Jacobs, W.J., Nadel, L., in press. *Neurobiology of reconstructed memory*. *Psychol., Public Policy, Law*.
- Kopelman, M.D., Wilson, B.A., Baddeley, A.D., 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, J.L., Sutherland, R.J., Muller, R.U., submitted. Hippocampal lesions produce a temporally-graded retrograde amnesia on a dry version of the Morris swimming task. *J. Neuroscience*.
- Linton, M., 1978. Real world memory after six years: an in vivo study of very long-term memory. In: Gruneberg, M.M., Morris, P.E., Sykes, R.N. (Eds.), *Practical Aspects of Memory*. Academic Press, Orlando, pp. 69–76.
- Mangan, P., Nadel, L., 1990. Development of spatial memory in the human infant. *Psychonomic Soc. Abstr.* 31, 35–36.
- McClelland, J.L., Goddard, N.H., 1996. Considerations arising from a complementary learning systems perspective on hippocampus and neocortex. *Hippocampus* 6, 654–665.
- McClelland, J.L., McNaughton, B.L., O'Reilly, R., Nadel, L., 1992. Complementary roles of hippocampus and neocortex in learning and memory. *Soc. Neurosci. Abstr.* 18, 1216.

- McClelland, J.L., McNaughton, B.L., O'Reilly, R.C., 1995. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol. Rev.* 102, 419–457.
- Metcalf, J., Jacobs, W.J., 1998. Emotional memory: the effects of stress on 'cool' and 'hot' memory systems. In: Medin, D.L. (Ed.), *The Psychology of Learning and Motivation: Advances in Research and Theory*, vol. 35. Academic Press, San Diego, CA.
- Morris, R.G.M., 1989. Does synaptic plasticity play a role in information storage in the vertebrate brain? In: Morris, R.G.M. (Ed.), *Parallel Distributed Processing: Implications for Psychology and Neurobiology*. Oxford University Press, Oxford, pp. 248–285.
- Moscovitch, M., 1995. Recovered consciousness: a hypothesis concerning modularity and episodic memory. *J. Clin. Exp. Neuropsychol.* 17, 276–290.
- Moscovitch, M., Nadel, L., 1998. Consolidation revisited: in defense of the multiple-trace model. *Curr. Opin. Neurobiol.*, in press.
- Nadel, L., 1991. The hippocampus and space revisited. *Hippocampus* 1, 221–229.
- Nadel, L., 1992. Multiple memory systems: what and why. *J. Cogn. Neurosci.* 4, 179–188.
- Nadel, L., 1994. Multiple memory systems: what and why. An update. In: Schacter, D., Tulving, E. (Eds.), *Memory Systems 1994*. MIT Press, Cambridge, MA, pp. 39–63.
- Nadel, L., Jacobs, W.J., 1996. The role of the hippocampus in PTSD, panic and phobia. In: Kato, N. (Ed.), *Hippocampus: Functions and Clinical Relevance*. Elsevier Science, Amsterdam.
- Nadel, L., Moscovitch, M., 1997. Consolidation, retrograde amnesia and the hippocampal formation. *Curr. Opin. Neurobiol.* 7, 217–227.
- Nadel, L., Moscovitch, M., in preparation. Consolidation and Memory. *Handbook of Neuropsychology*, 2nd ed. Volume on Memory Disorders.
- Nadel, L., Willner, J., 1980. Context and conditioning: a place for space. *Physiol. Psychol.* 8, 218–228.
- Nadel, L., Zola-Morgan, S., 1984. Infantile amnesia: a neurobiological perspective. In: Moscovitch, M. (Ed.), *Infant Memory*. Plenum Press, New York, pp. 145–172.
- Nadel, L., Willner, J., Kurz, E.M., 1985. Cognitive maps and environmental context. In: Balsam, P., Tomie, A. (Eds.), *Context and Learning*. Erlbaum, Hillsdale, NJ, pp. 385–406.
- Newcombe, N., Huttenlocher, J., Drummey, A.B., Wiley, J.G., 1998. The development of spatial location coding: place learning and dead reckoning in the second and third years. *Cognitive Dev.* 13, 185–201.
- O'Keefe, J., Burgess, N., 1996. Geometric determinants of the place fields of hippocampal neurons. *Nature* 381, 425–428.
- O'Keefe, J., Nadel, L., 1978. *The Hippocampus as a Cognitive Map*. Oxford University Press, London.
- Samsonovitch, A., McNaughton, B.L., 1997. Path integration and cognitive mapping in a continuous attractor neural network model. *J. Neurosci.* 17, 5900–5920.
- Schacter, D.L., Moscovitch, M., 1984. Infants, amnesics, and dissociable memory systems. In: Moscovitch, M. (Ed.), *Infant Memory*. Plenum Press, New York, pp. 173–216.
- Schacter, D.L., Tulving, E. (Eds.), 1994. *Memory Systems 1994*. MIT Press, Cambridge, MA.
- Squire, L.R., 1992. Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. *Psychol. Rev.* 99, 195–231.
- Squire, L.R., Cohen, N.J., Nadel, L., 1984. The medial temporal region and memory consolidation: a new hypothesis. In: Weingartner, H., Parker, E. (Eds.), *Memory Consolidation*. Erlbaum, Hillsdale, NJ, pp. 185–210.
- Teyler, T.J., DiScenna, P., 1986. The hippocampal memory indexing theory. *Behav. Neurosci.* 100, 147–154.
- Vargha-Khadem, F., Gadian, D.G., Watkins, K.E., Connelly, A., Van Paesschen, W., Mishkin, M., 1997. Differential effects of early hippocampal pathology on episodic and semantic memory. *Science* 277, 376–380.
- Wagenaar, W.A., 1986. My memory: A study of autobiographical memory over six years. *Cogn. Psychol.* 18, 225–252.
- Wilson, M.A., McNaughton, B.L., 1994. Reactivation of hippocampal ensemble memories during sleep. *Science* 265, 676–679.