THE ROYAL SOCIETY

doi 10.1098/rstb.2001.0945

Episodic-like memory in animals: psychological criteria, neural mechanisms and the value of episodic-like tasks to investigate animal models of neurodegenerative disease

Richard G. M. Morris

Department of Neuroscience, The University of Edinburgh, 1 George Square, Edinburgh EH8 9JZ, UK (r.g.m.morris@ed.ac.uk)

The question of whether any non-human species displays episodic memory is controversial. Associative accounts of animal learning recognize that behaviour can change in response to single events but this does not imply that animals need or are later able to recall representations of unique events at a different time and place. The lack of language is also relevant, being the usual medium for communicating about the world, but whether it is critical for the capacity to represent and recall events is a separate matter. One reason for suspecting that certain animals possess an episodic-like memory system is that a variety of learning and memory tasks have been developed that, even though they do not meet the strict criteria required for episodic memory, have an 'episodic-like' character. These include certain one-trial learning tasks, scene-specific discrimination learning, multiple reversal learning, delayed matching and non-matching tasks and, most recently, tasks demanding recollection of 'what, where and when' an event happened. Another reason is that the neuronal architecture of brain areas thought to be involved in episodic memory (including the hippocampal formation) are substantially similar in mammals and, arguably, all vertebrates. Third, our developing understanding of activity-dependent synaptic plasticity (which is a candidate neuronal mechanism for encoding memory traces) suggests that its expression reflects certain physiological characteristics that are ideal components of a neuronal episodic memory system. These include the apparently digital character of synaptic change at individual terminals and the variable persistence of potentiation accounted for by the synaptic tag hypothesis. A further value of studying episodic-like memory in animals is the opportunity it affords to model certain kinds of neurodegenerative disease that, in humans, affect episodic memory. An example is recent work on a transgenic mouse that over-expresses a mutation of human amyloid precursor protein (APP) that occurs in familial Alzheimer's disease, under the control of platelet derived (PD) growth factor promoter (the PDAPP mouse). A striking age- and amyloid plaque-related deficit is seen using a task in which the mice have to keep changing their memory representation of the world rather than learn a single fact.

Keywords: learning; one-trial learning; episodic-like memory; synaptic tagging; hippocampus; Alzheimer's disease

1. INTRODUCTION

The year 1983 witnessed the publication of two important books on learning and memory. One was Mackintosh's monograph *Conditioning and Associative Learning*, which ends with the following passage:

> 'It should not be forgotten that animals are probably not just machines for associating events. Their ability to represent different attributes of their environments, to respond in terms of spatial, and even of abstract relationships between events, to store and rehearse information for later use, are all important and little-understood capacities whose study requires the development of more sophisticated experimental arrangements than those of simple conditioning experiments.'
>
> (Mackintosh 1983, p. 277.)

The other was Tulving's monograph *Elements of Episodic Memory*, which began with the following (second) sentence:

> 'As far as we know, members of no other species possess quite the same ability to experience again now, in a different situation and perhaps a different form, happenings from the past, and know that the experience refers to an event that occurred at another time and in another place . . .'
>
> (Tulving 1983, p. 1.)

Published in the same year, the two passages echo each other's guarded views about the memory capabilities of animals, even to the rhetorical use of 'probably' by Mackintosh and 'quite' by Tulving. Mackintosh's book, a defence of associative conditioning as a theoretical framework for understanding animal learning, leaves the

reader in little doubt that he believes principles of conditioning will remain relevant even after the development of the 'more sophisticated experimental arrangements' that his open-minded attitude allows. Tulving also leaves the door open to a more inclusive view of memory in animals, but one is again left with the suspicion, as later articles were to spell out explicitly, e.g. Tulving & Markowitsch (1998), that episodic memory is not for the birds but for man.

Over the nearly 20 years since these books were published, there have been numerous developments in our understanding of the learning and memory capacities of humans and animals, and of the underlying physiological and pharmacological mechanisms that may mediate them. A now widely held view in neuroscience is that there are multiple 'types' of memory and these differ with respect to their psychological characteristics, the anatomical circuits involved and the underlying neural mechanisms of encoding, storage, consolidation and retrieval. Various taxonomic frameworks for thinking about these different forms of memory have been proposed, most having a common generic form (figure 1). All recognize a cardinal distinction between working (short-term) memory and long-term memory. Within the domain of long-term memory, the framework according to Tulving subsumes the further distinction between explicit and implicit memory (Tulving & Schacter 1990; sometimes referred to as 'declarative' and 'non-declarative' memory, Squire 1992). Procedural learning and the formation of perceptual representations are held to be instances of implicit memory; explicit memory encompasses both semantic (fact) and episodic (event) memory. The types of memory specified in this and other taxonomic frameworks are thought to map onto distinct anatomical circuits, although not necessarily in any simple or one-to-one manner.

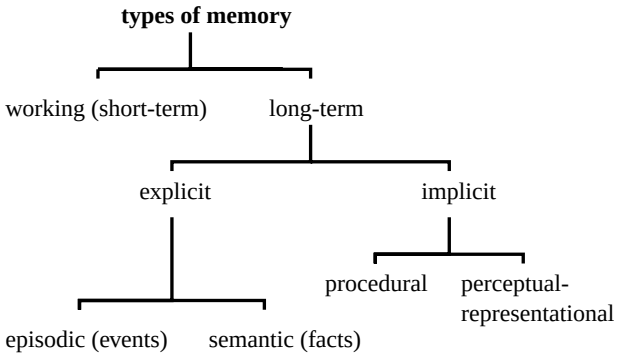

Figure 1. Taxonomy of memory according to Tulving. Working memory is distinguished from long-term memory at the top of the hierarchy. Long-term memory is first subdivided into implicit and explicit forms, with further subdivisions of these into procedural and perceptual representation (implicit), and episodic and semantic (explicit). Other similar taxonomies have been proposed (e.g. Squire 1992).

2. DO ANIMALS HAVE EPISODIC MEMORY?

The two key problems that confront us in thinking about how to apply Tulving's and other similar taxonomies to animals relate to: (i) the implicit/explicit distinction; and (ii) whether the concepts of semantic knowledge and episodic recollection are relevant to animal memory.

There are several grounds for caution. Assuming the mantle of Morgan (1894) a century before, Macphail (1982, 1998) has long held what he calls an agnostic, others a sceptical, view about the similarities between animal and human cognitive capacities. He argues that the possession of both language and consciousness are two critical differences between humans and animals. As consciousness is central to the distinction between implicit and explicit memory, and as language is the usual medium through which we describe our recollections, this view is also pertinent to episodic memory. Tulving's (1983) definition of it as a mental capacity that requires 'autonoetic' consciousness—a sense of the self—has also to be taken on board. Even if some animal species have forms of awareness that are analogous to particular states of consciousness in humans, Macphail's (1998) review of relevant literature presents cogent arguments for doubting that many (if any) species could possess something as sophisticated as autonoetic consciousness. Were we to accept his view, it would follow that animals cannot, by definition, have anything akin to human episodic memory. We would be left defending the position that all animal learning must be implicit.

There are, however, several grounds for exploring the possibility that some animals may have a form of memory that is explicit and 'episodic' in character, even if we can neither prove that animals are conscious nor be certain that they are not. This category of memory is what Clayton & Dickinson (1998) call 'episodic-like' in which an animal is thought to recall the 'what, where and when' of discrete events and can display this in its overt behaviour. To study this, they have developed an ingenious food-caching paradigm in which scrub jays cache perishable and non-perishable food items (such as worms and peanuts, respectively) and keep track of where they have stored each food type and how long ago. Their work is described more fully in a companion article (Clayton *et al.* 2001). However, before considering this and other experimental paradigms that are potentially 'episodic-like', there are various conceptual issues to consider that provide a basis for believing that the enterprise of studying episodic-like memory in animals is worthwhile.

First, is language as central to episodic memory as Macphail (1998) argues that it is to consciousness? To be sure, language is the usual medium of communication about the events of our lives, but having certain kinds of experiences and recalling them later is surely independent of their overt communication to others. It follows that a species that cannot communicate through language is not logically prevented from having private recollective experiences that may relate to the 'what, where and when' of past events. What is difficult to think about is the nature of this non-human experience—the issue often expressed as 'what might it be like to be a bat'? For, even if we do not communicate our past experience to others, our private recollections of objects and events do generally involve a linguistic code—Macphail's 'aboutness'. The experiential nature of recollection in any animal which lacks language—including, as Macphail discusses

at some length, very young infants—cannot be identified with any certainty.

One way forward is to enquire whether there are any aspects of behaviour, other than communication using a representational code, that reflect an animal's knowledge of the world. Here I part company with Macphail for it is my view that non-human vertebrates do segment the world categorically into objects such as food, trees, nests, burrows and so forth. They lack words to describe these entities but it seems to me that they acquire 'factual' knowledge and can display through their behaviour that they 'know' what these and other types of objects are. That is, some animals do have 'aboutness' types of knowledge and they have the neural machinery to acquire it.

Second, is Tulving correct to assert that a particular form of self-consciousness is at the heart of episodic experience—the sense of an event happening to 'me'? A case can be made for seeing the selection pressure that might have led to episodic-like memory, and the brain structures that mediate it, being the need to deal with an important characteristic of the world that applies as much to animals as to humans—dealing with unpredictable events. Events are things that often happen only once and at times when ongoing behaviour is determined by other stimuli, motivational states or incentives. The world is also a dangerous place in which information about the availability of food and the proximity of predators tends not to arrive, as it sometimes does in the laboratory, in the form of ten predictable trials per day. It would therefore be adaptive for a species to evolve a mechanism for enabling its understanding of the world and its behaviour to benefit from encoding information about single events. This 'benefit' need not, as we shall see, necessarily require their explicit recollection—the learning process could be implicit. However, it could be advantageous to encode, store and later recollect events explicitly, and selection pressures must have existed for such memory processes to have evolved. The challenges are: (i) to identify these selection pressures; and (ii) to develop behavioural tasks that can only be carried out effectively were animals to possess a private episodic-like memory system.

This argument leads on to a third reason for being optimistic about identifying episodic-like memory in animals: the ingenuity of comparative and neuropsychologists in developing tasks for exploring animal cognition. When Macphail (1982) first outlined his claim that there were no differences in the learning capacities of vertebrates irrespective of brain size, and that learning by animals was radically different from human learning, he based his argument largely on conditioning tasks—such as habituation, classical conditioning and instrumental conditioning. He did, to be fair, also discuss more complex tasks, such as learning-set and conditional tasks, but the argument felt weaker partly because of frequent appeal to 'contextual' factors to explain away apparent species differences. Since that time, many new perceptual, learning, memory and motor tasks have been invented for animals that a sympathetic reader of Macphail's book at the time might have suspected to be beyond their capabilities. These include tests of mediated conditioning (Holland 1981, 1990), one-trial spatial memory (Foster *et al*. 1999; Steele & Morris 1999) and both scene and object-in-place memory (Gaffan 1994). These and other tasks may be amenable to solution in implicit associative terms, despite their procedural complexity. However, it is unlikely that all of them can be explained in such terms.

One example of this ingenuity is in the perceptual domain: Stoerig and Cowey's (1997) use of a 'commentary key' to reveal that monkeys with unilateral striate cortex lesions can report being unable to see an object that they can accurately reach towards. This is a particularly telling example because the experimental demonstration of 'blindsight' in primates is precisely the kind of phenomenon that a sceptic asserting the importance of language to experience might, until proved otherwise, have expected to be impossible. How could one ever get an animal to 'tell' us about his visual experience unless he had a representational code in which to convey this information? But this experiment suggests that we can.

To summarize this section, cogent arguments have been put forward in relation to the proper definition of episodic memory and the fact that we may never know whether any animal can ever be said to be conscious. Seeking true episodic memory in animals is therefore a forlorn exercise. However, there is value in taking a positive view of the possibility of identifying 'episodic-like' memory in animals: language may not be central to its expression; a sense of the self is not required; and the invention of new behavioural protocols should reveal more about how animals understand and remember their world.

3. DOES THE OCCURRENCE OF ONE-TRIAL LEARNING IMPLY THAT ANIMALS FORM EPISODIC-LIKE MEMORIES?

The next step of the argument concerns the issue of whether animals can learn in one trial and what the capacity to do so implies about the types of learning processes they possess.

Many forms of learning, in humans and in animals, require multiple 'trials' before a reliable change in behaviour is seen—the quintessential example being the learning of motor skills. Multiple trials are also a feature of most instances of associative and non-associative conditioning in animals. For example, in habituation (a non-associative waning of responsiveness to repetitive stimulation) and classical conditioning (an associative procedure in which an initially inconsequential stimulus acquires significance by virtue of being repeatedly paired with another stimulus), the overt behavioural change may only become apparent over repeated training trials. In these cases, animal learning theories speak of successive trials causing an incremental change in a parameter specifying the relationship between stimuli, between stimulus and response, or between responses and their outcomes. Formal models exist describing the circumstances in which associative strength changes and accumulates, the parameters affecting the rate of change (Rescorla & Wagner 1972; Mackintosh 1975; Pearce & Hall 1980).

It is important to appreciate that such accounts do not require that an animal encode, store or could ever retrieve information about the separate and unique events that gave rise to that gradual change. Animals have internal representations of stimuli and responses but, formally, the outcome of the learning process is a new

value of the parameter that specifies the strength of an association. The value of this parameter may be achieved by several routes and, for that reason, associative learning is said to have the characteristic of 'independence of path' (figure 2). What matters is the value of this parameter at the end of the learning process, like the bottom line of a bank statement, not how it got there.

Newer conditioning paradigms complicate the picture (such as occasion setting), but consideration of this simplest case is sufficient to explore the implications of one versus multitrial learning—for not all types of animal learning do require multiple trials and animals can learn to change their behaviour after a single learning trial. Examples of such rapid learning are poison avoidance, recognition memory, spatial learning and food caching. If a rat is given a single chance to drink saccharin-flavoured water and later made ill by being given a small dose of lithium chloride, it will avoid saccharin-flavoured water for a long time thereafter (Garcia & Koelling 1966). Similarly, if a rat explores one arm of a T-maze during an initial bout of exploration, it will display a strong tendency to explore the other arm of the T-maze later on—the phenomenon of spontaneous alternation (Halliday 1968). Are we to suppose that because learning has taken place so rapidly, an account in terms of implicit associative learning is inapplicable?

The answer to this is 'no'. The occurrence of one-trial learning is not a sufficient condition for asserting that a distinct episodic-like learning process must be engaged. In the case of poison avoidance learning, the value of learning in one trial is obvious—the animal might not survive were learning to take longer. The parameter specifying the associative relationship between a stimulus (in this case a foodstuff) and its consequence (illness) must increment very rapidly, but the animal is obliged to recall the events associated neither with eating nor with malaise for this learning to be successful. The same is true of humans. Indeed, we may enjoy telling each other stories about the restaurant where we had the oysters and what happened at home later that night, but such stories are what Weiskrantz (1997) calls a 'commentary' on the learning process rather than an integral part of it. For both animals and humans, rapid food-aversion learning and phenomena such as neophobia indicate that learning about what food to eat has evolved to be a conservative specialization (Rozin & Kalat, 1971). Nonetheless, poison avoidance is a curious instance of associative learning for it breaks many of the usual rules—such as the need for the stimulus and its consequence to be closely contiguous in time. Careful analysis has revealed, however, that it displays a number of the properties of other widely studied forms of associative learning and, in the view of most, a sufficient number to obviate the need to invoke any specialized learning process other than that of associative conditioning (Dickinson 1980). Similarly, although the argument will not be laid out in detail, implicit processes such as judgements of familiarity can mediate certain kinds of one-trial object recognition memory. That it also occurs in one trial is not in itself a basis for supposing that a distinct form of learning and memory is involved.

What are the necessary features of episodic-like memory in animals? The key feature, in my view, is not

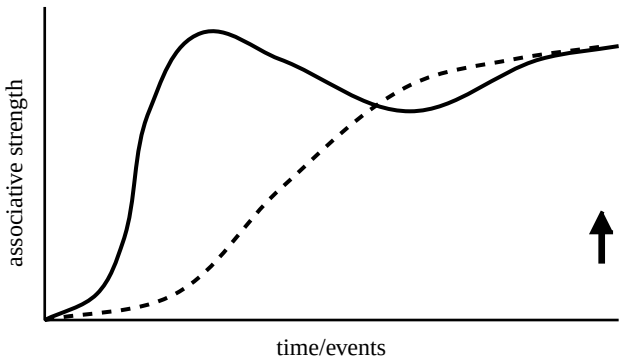

Figure 2. Independence of path assumption of classical associative learning theory. Experience (x-axis, time/events) results in increases or decreases of a single parameter, associative strength (y-axis). Probing 'memory' at any time point (arrow) involves reading off the value of associative strength without regard to the varied routes taken to the current value. Animals that learn according to the continuous line would display identical performance at the test point to those that learn by the dotted line.

just the occurrence of learning in a single trial but the evolution of some system of explicit encoding and recall. Whereas associative learning theories hold that multiple events have their effect on an animal's later behaviour by incrementing and decrementing a single parameter called associative strength, the possession of episodic-like memory has to do with being able to encode and recall specific features of such events. Not only might an animal be tested with respect to where something has happened, what happened and when, but the implicit/explicit distinction has also to be considered. The experimental test, as I see it, is as follows:

> We need to devise behavioural tasks that somehow distinguish between the changes in behaviour that occur because an animal remembers one or more prior events, and changes in behaviour that occur merely because these prior events have happened.
>
> (Adapted from Morris & Frey 1997, p. 1495.)

This is the challenge and, with the possible exception of the work of Clayton & Dickinson (1998) on food caching, few if any of the behavioural tasks we have at our disposal today yet meet it.

4. IN WHAT KINDS OF ANIMAL LEARNING AND MEMORY TASKS MIGHT AN EXPLICIT EPISODIC-LIKE MEMORY SYSTEM BECOME ENGAGED?

Numerous learning and memory tasks have been developed for animals that are 'episodic-like' in character. Certain tasks relate particularly to an animal's experience of the context in which learning takes place, others to the memory of prior events. Both are sufficiently different from classical associative learning tasks that they merit close attention. In both cases, a strictly implicit learning process may be sufficient for learning, once supplemented by abstract concepts such as recency or familiarity. However, whether it is always the most effective way of learning is another matter—explicit encoding and recall may sometimes be better.

The first group can be categorized as unusual discrimination learning tasks. These include tasks in which the solution of a series of problems depends upon, or at least could benefit from, memory of the scene against which the discrimination takes place. Prominent amongst these is a range of ingenious tasks developed by Gaffan prompted by a particular theoretical view of episodic memory. Working with monkeys, Gaffan & Harrison (1989) discovered that fornix lesions interfered with the acquisition of ambiguous object discrimination learning tasks when the ambiguity about which of two objects signalled reward could be resolved with respect to the direction an animal was facing in a complex scene. This finding led him to think about the experiential aspects of episodic recall, arguing that:

> 'When a human subject recalls specific information about some discrete, personal experience from the past, the recall process often involves a covert reconstruction of a spatially organised complex scene.'
>
> (Gaffan 1991, p. 262.)

Because securing propositional descriptions of the experiential aspects of recall in monkeys is not feasible, Gaffan and his colleagues went on to develop touch-screen based tasks in which a large variety of different scenes are projected onto a screen (Gaffan 1994). The role of these scenes is either to provide a set of stimuli to be discriminated in their own right, or to serve as backgrounds for other object discriminations that take place against them. A consistent finding is that such tasks are learned incredibly fast. They are also remembered very well and experimental damage to the hippocampus–fornix–mamillary bodies, a connected group of structures implicated in episodic memory in humans, causes impairments in learning (Gaffan & Parker 1996; Murray *et al.* 1998; Parker & Gaffan 1998; Gaffan & Parker 2000). Gaffan argues that this protocol is a model of episodic memory in the monkey.

The question arises of whether the monkeys are obliged to use explicit memory to encode and store the information necessary to solve these tasks. Both the faster speed of learning and the sensitivity to damage in particular brain areas differ from that seen for more conventional discrimination tasks, but we must also ask the information-processing question of whether the task meets Clayton & Dickinson's (1998) criteria for an episodic-like task, i.e. recall of 'what, where and when'. Leaving aside the issue of self-identity on the part of the monkeys, it is possible that the task involves no more than a combination of recognition and association. The monkey does not have to recall the scene during test trials—it is presented to him and it has only to be recognized. Moreover, it might be recalling which of two typographical 'objects' in the scene was associated with reward, using the background scene as a cue for recall, but as the scenes and objects associated with them are unique, the task can actually be solved through differential association with reward. A strictly implicit solution is therefore possible. In addition, one could argue, in keeping with Warrington and Weiskrantz's classical analysis of the effects of providing recall cues to amnesic persons, that the very provision of the background scenes during test trials would differentially aid the performance of lesioned monkeys (see McCarthy & Warrington 1990). Thus, while fornix-lesioned monkeys showed a deficit in this task, one wonders if that deficit would actually be larger if the distinctive scenes used during training were sometimes omitted during testing. Might normal monkeys be able to recall the scene with which individual typographical characters had been associated and, in doing so, use this explicit recall to help retrieve which of the two objects was associated with reward? Gaffan's scene-specific memory task seems to be neither an instance of memory for 'what, where and when' nor one that requires explicit recall. The task remains a pioneering and valuable approach to thinking about the experiential character of episodic memory in animals and could be adapted further.

The second group of tasks to be considered, this time studied using rodents, are those in which the animal has to keep changing what it should do in the light of what it has just done. A now classical paradigm for looking at this is Olton's 'radial maze' in which rats are required to search for single food items at the ends of each arm of the maze (Olton *et al.* 1979). The animal is confronted with the travelling salesman's problem of visiting all the places at which he can find food while minimizing the distance travelled. It is a task in which the animal must keep track of its own actions. Spontaneous alternation (Halliday 1968), delayed non-matching to place in a T-maze (Rawlins 1985; Aggleton *et al.* 2001) and the delayed matching to place (DMP) task in the water maze (Morris 1983; Steele & Morris 1999) are also examples of conceptually similar tasks that are so readily learned as to be almost spontaneously expressed (figure 3). Although each of these tasks is spatial, analogous non-spatial tasks also exist (Ennaceur & Delacour 1988).

Each of these tasks might be learned using an episodic-like memory system. In delayed non-matching to place in a T-maze, the animal might remember several aspects of what happened on the previous trial—where it went, what food it ate at the end of the maze arm, how long ago this eating happened. Unfortunately, there is an ostensibly more parsimonious solution. This is that the cues associated with turning left or turning right could acquire 'familiarity' through exposure on the sample trial and the animal need do no more than learn the rule, in the case of non-matching to place, of avoiding familiar cues when it confronts them on the choice trial. No explicit recall of prior events is then necessary—the task can be solved implicitly using familiarity.

A similar objection has been raised by Griffiths *et al.* (1999) to an episodic-like account of performance in the DMP task in the water maze (Steele & Morris 1999). In this, rats are trained to find a hidden escape platform that moves location between days. There are several trials each day, usually four, and the rats' behaviour very rapidly settles down to a pattern in which they show a long escape latency on trial one and a short latency on trial two and thereafter (figure 3*c*). In fact, learning takes place so rapidly that one wonders if the animal has to learn any 'rule' for solution at all. Griffiths *et al.* (1999) argue that this task could also be solved by familiarity—the animal has only to approach the set of cues that are most familiar, these being the ones associated with the last position in which the platform has been placed. This

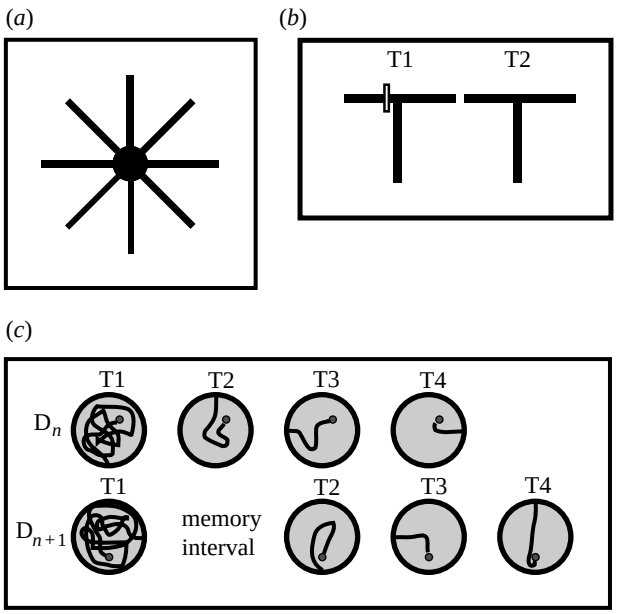

Figure 3. One-trial spatial learning tasks that require an animal to keep changing its memory representation. (*a*) The radial maze: many protocols are possible but the typical procedure involves the rat searching for food at the end of each maze arm. The best solution minimizes re-visits to arms from which the food has already been taken. Short periods of confinement in the central area prevent the use of turning algorithms that minimize memory demands. (*b*) T-maze delayed non-matching: on trial 1 (T1), the rat is required to go in one or the other direction from the stem (bottom) to secure food. There is a free choice on T2, but the rat must remember to go to the opposite arm to that just visited. The task has an episodic-like character and is exquisitely sensitive to hippocampal disruption, but can also be solved by familiarity. (*c*) Water maze delayed matching to place (DMP). Rats are given four trials per day to a platform that moves location between days (rows). Memory formed during T1 of each day is tested on T2 and subsequent trials. The memory delay interval can be manipulated. The task is exquisitely sensitive to hippocampal disruption. See § 4 for discussion of whether this task can be solved by familiarity.

objection cannot be definitively dismissed, but I would make two points:

First, it is not clear that there are any overt cues associated with the most recent goal location, in or around the testing arena, that are more or less familiar than any others. The extra-maze cues provide the basis for the animal forming a representation or 'map' of space. This map is learned rapidly and remains stable over the days or weeks of testing. All that changes is the entry into the map of where the hidden platform has last been located. This location is unmarked by local cues, it is a 'place' within the animal's map that has to be recalled rather than merely recognized at the start of a trial. To object that performance could rely exclusively on the recognition of what is familiar and what is not also requires us to believe that the rat can recognize, from the starting point, the familiarity of cues that may only be visible from the goal location. Foster *et al.* (1999) have recently modelled performance in this and a standard reference memory task in the water maze using a temporal difference learning algorithm. Their work revealed that an associative solution to both tasks was possible, but the solution of the DMP task relied upon: (i) the formation of a map of space using hippocampal 'place-cells' (O'Keefe 1976); and (ii) the ability of the animal to recall goal locations in the map from start locations at any point in the maze. The model therefore contains a cryptic form of episodic-like recall. Interestingly, Foster *et al.* (1999) show that the standard reference memory task of finding an escape platform that stays in a single location does not require a map of space, nor recall of the goal location. The animal need learn only to execute appropriate approach behaviour at locations that are individuated by place cell firing but these cells are not necessarily part of a map.

Second, performance of the DMP task is exquisitely sensitive to disruption of *N*-methyl-D-aspartic acid (NMDA) receptor-dependent synaptic plasticity (Steele & Morris 1999). Performance of the DMP task is disrupted completely by hippocampal lesions, even at short memory delay intervals, arguably because recall of information from the animal's map of space cannot occur without synaptic transmission in this brain area. Accordingly, to examine the neural mechanisms of episodic-like memory encoding independently of performance, we need a way of disrupting it in a manner that does not, or at least may not, affect retrieval. Work on the underlying neurobiological mechanisms of learning and memory in animals is currently at a transition point. We are moving from studies searching for process dissociations on the basis of discrete lesions which permanently damage brain tissue through to studies that use reversible, pharmacological manipulations (Izquierdo & Medina 1998; Riedel *et al.* 1999; McGaugh 2000). Our developing understanding of glutamate neurobiology has given us a range of tools with which we can manipulate function within discrete brain areas in highly selective ways to investigate learning mechanisms (Danysz *et al.* 1995). Intra-hippocampal infusion of the NMDA receptor antagonist D-2-amino-5-phosphopentanoic acid (D-AP5) is one way of achieving this. Steele & Morris (1999) found that doing this had no effect on performance at a short memory delay of 15 s between trials 1 and 2 but disrupted performance at the longer interval of 2 h (figure 4). Our preferred interpretation of this finding is that the animal encodes information about the act of escaping from the water and where this happened at the end of each trial using a hippocampal NMDA receptor-dependent mechanism. Navigational performance on the succeeding trial would be guided by recall, not recognition, of this context-event association encoded and stored during the preceding trial.

Unfortunately, the water maze DMP task shares various limitations with Gaffan's scene-memory task. One is that only a single kind of event can ever happen at the end of each trial—escape from the water—and for this reason is an inadequate analytic tool to look at 'what, where and when'. The same limitation applies to the paired comparison object recognition memory task of Ennaceur (Ennaceur & Delacour 1988). What we need are tasks in which several different kinds of events can occur. This is precisely the problem to which Clayton & Dickinson (1998) so successfully turned their attention in their studies of memory for food caches by scrub jays.

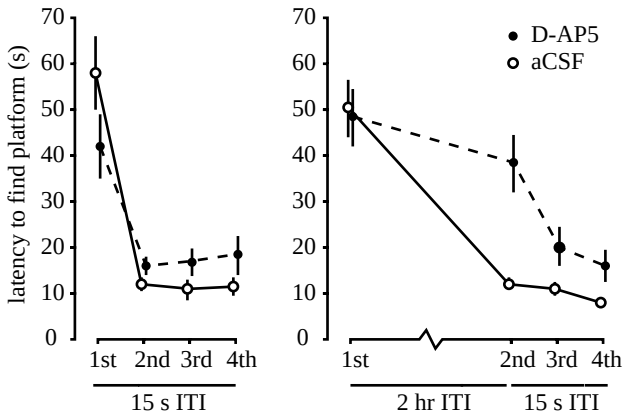

Figure 4. Intra-hippocampal microinfusion of the NMDA antagonist D-AP5 causes a delay-dependent impairment of one-trial spatial memory. Rats implanted with bilateral infusion cannulae aimed at the dorsal hippocampus were trained for on the DMP task to asymptote. They were then tested with D-AP5 or artificial cerebrospinal fluid (aCSF) (vehicle) infusions over a series of days, with these treatments given within-subject on alternate days. The inter-trial memory delay interval (ITI) between trials 1 and 2 was varied between 15 s (short-term memory) and 2 h (long-term memory) with all other trials at 15 s. Averaged across days, there was an impairment on the second trial of the day at the long memory delay. Data from Steele & Morris (1999, experiment 3).

5. ARE THERE ANY CONSTRAINTS ON THE ANATOMICAL SYSTEMS INVOLVED IN EPISODIC-LIKE MEMORY?

A separate reason for suspecting that mammals, and perhaps all vertebrates have an episodic-like memory system is because of similarities in the structure of particular areas of their brains. Mammals contain all the cortical areas in the human brain that have been implicated in episodic memory from clinical and functional imaging studies on humans—although some comparative anatomists question the development of the frontal cortex in rodents (Preuss 1995). At least one of these structures, the hippocampus, has an evolutionary well-conserved structure. It contains similar cell types and apparently similar connectional architecture in humans and non-human primates—and its structure in rodents is strikingly similar albeit with about 10 times fewer cells (Amaral & Witter 1989; Amaral & Insausti 1990) and more extensive inter-hemispheric connections. Parsimony dictates consideration of the possibility that the same algorithm is being carried out in human and non-human brains in this and other relevant brain structures. Even if the information upon which it is operating is different by virtue of the human possession of language, there are no obvious cellular, connectional or biochemical differences to which we might relate the supposed differences in conscious awareness of animals and humans.

Several neural network models of episodic-like memory have been proposed (Marr 1971; McNaughton & Morris 1987; Gluck & Granger 1993; Granger *et al.* 1996; Rolls & Treves 1998; Redish 1999). Despite differences of emphasis, these generally focus on certain distinctive features of hippocampal excitatory circuitry that may enable particular types of associations to be formed. The architecture of the dentate gyrus, with a much larger number of cells than the layer II neurons of the entorhinal cortex that are afferent to it, has a structure that is ideal for orthogonalizing inputs. The architecture of area CA3 is remarkable, with each cell receiving about 3500 perforant path connections that bypass the dentate gyrus, each receiving around 10 000 feedback connections from other CA3 neurons (the longitudinal-association pathway) and, most intriguing of all, a very limited number of mossy fibre inputs from the dentate gyrus. In the rat brain, mossy fibre axons of individual dentate gyrus cells make a total of 15 ± 2 connections in CA3 as they traverse through the thousands of cells that lie in a single lamella plane, with each CA3 cell receiving about 50 mossy fibre terminals. No one working on the hippocampus can look at this unique neuronal architecture without wondering what algorithm it computes. Other brain areas have architectures that immediately suggest their function, such as the coincidence-detecting 'delay lines' of the nucleus laminaris used for prey localization by the barn owl (Konishi 1986). Yet CA3 remains a mystery. Rolls & Treves (1998), in proposing that CA3 is an autoassociative type of memory device, have made the interesting suggestion that its architecture is ideal for controlling the conditions of memory encoding, via mossy fibre activation, somewhat separately from the determinants of cued recall. Lisman (1999) has proposed a somewhat heretical model consisting of reciprocal interconnections between the dentate gyrus and CA3. Synaptic potentiation in area CA3 of the hippocampus is thought to encode traces relevant to sequences of events occurring within episodes, and maintain sequence order despite the essentially passive nature of trace storage at synapses of the longitudinal-association pathway. During retrieval, CA3 neurons are, speculatively, held to project information retrieved in response to items earlier in a sequence back to the mossy cells of the dentate gyrus. There, pattern completion corrects recall errors, the corrected recall pattern is then projected forward to CA3 to retrieve the next item in the sequence, and so on. Lisman's (1999) model also allows a role for context to bias the firing of CA3 cells (achieved by the direct perforant path input to CA3), with recoding of the hippocampal representation back into a neocortical form accomplished by area CA1.

Area CA1 of the hippocampus also has an architecture whose complexity is often underestimated. Most writers about hippocampal neuroanatomy emphasize the Schaffer-collateral inputs from CA3 which synapse *en passant* as they course through the transverse plane. However, recent research indicates that CA1 may also receive a direct, topographic excitatory input from layer III of the entorhinal cortex. The existence of this pathway is highly controversial (see the articles in *Hippocampus* 1994) but it raises the possibility that information might be recalled from the cortex to the hippocampus by a route that bypasses the orthogonalizing dentate gyrus and associative machinery of CA3. This route may enable representations retrieved from memory to be integrated with other information that has been newly processed. Perhaps their intersection is the basis of context-event associations?

6. ARE THERE ANY CONSTRAINTS ON THE PHYSIOLOGICAL MECHANISMS INVOLVED IN EPISODIC-LIKE MEMORY?

Given current emphasis on the role of the hippocampus in explicit memory, it is also worth considering whether the physiological properties of the activity-dependent synaptic plasticity that has been intensively studied in this brain area relate in any way to episodic-like memory. One well-known idea is that long-term potentiation (LTP) is a suitable model of the synaptic changes that might be involved in the formation of memory traces. Proponents of this hypothesis have often noted that three properties of LTP—input specificity, associativity and persistence—are suitable properties of a physiological memory mechanism. Evidence in favour of the generic 'Synaptic Plasticity and Memory' hypothesis is gradually getting stronger as various objections raised by early critics are addressed by ever more sophisticated experiments (for a review, see Martin *et al.* 2000). For the purposes of the present discussion, however, these basic properties are as suited to a strictly implicit learning machine as to any other.

What may be less well appreciated is that more recent physiological studies have revealed additional properties of LTP (and its companion long-term depression, LTD) that are ideal for mediating a particular form of memory mechanism. These additional properties include metaplasticity, silent synapses, the apparently digital nature of the synaptic change at individual terminals, and synaptic tagging. Two of these properties seem particularly pertinent to episodic-like memory.

First, when LTP is studied in whole animals *in vivo*, as in the classical work of both Bliss & Lomo (1973) and McNaughton *et al.* (1978), successive tetanizations cause an incremental change in the size of evoked field potentials. This graded increase has sometimes been likened to learning which, so the analogy goes, also increments gradually. But is this really what is happening at the level of individual synapses? To study this, Petersen *et al.* (1998) used brain slices *in vitro* and explored the effects of minimal stimulation until, in the limit, they would have been activating only single fibres. In the Schaffer collateral input to CA1, it is known that each fibre makes, on average, one *en passage* synaptic bouton per pyramidal cell as it passes through area CA1. It follows that many fibres were activated at which it was possible to examine the consequences of attempting to induce LTP at single synapses. Using both extra- and intracellular recording in neurons of the target CA1 region of the hippocampus, Petersen *et al.* (1998) observed that the gradual incremental increase in the extracellular field potential that followed successive tetanizations was accompanied by step-like changes at the level of single fibres. Single-fibre potentials jumped from their baseline level to their maximal level at varying times but, once potentiated, could be potentiated no further by later bouts of high frequency stimulation. The implication is that, at the level of individual synapses, LTP in the hippocampus is a digital change in synaptic strength—from zero (or weak) to strong. As activity patterns in CA3 project to CA1 by axons that enable any given CA3 cell to project onto a particular CA1 cell at (on average) only one synaptic terminal, such a system has the potential to change the strength of individual synaptic weights to the maximum allowable level following single events. That is, patterns of neural activity in CA3 could be associated with patterns of firing in area CA1 and, with a single pairing, realize maximal synaptic change at many of the thousands of individual synapses that connect neurons that were firing.

A second of the newer properties of LTP relates to whether synaptic potentiation necessarily results in the formation of a lasting memory trace or is merely creating the potential for creating a lasting trace. Episodic memory occurs continuously and automatically—we cannot decide to turn off what Frey and I have previously described as the 'automatic recording of attended experience' (Morris & Frey 1997). Unpredictable events happen to us and, whatever our separate intentions or concerns at the time such events happen, we necessarily make some record of attended events even if that command of attention has been involuntary. Our theoretical supposition is that this is what LTP in the hippocampus is doing all the time—automatically recording traces of experience. The Frey and Morris hypothesis further suggests that very few of the synaptic changes that occur as part of this ongoing record are preserved. The vast majority decay to baseline.

If this idea is on the right lines, we must then ask what determines whether synaptic potentiation decays, or that it persists and becomes the basis for lasting memory traces in the brain? The recent research revealing additional properties of LTP suggests that the past and future history of postsynaptic neuronal activation is the critical determinant. Importantly, whether a trace lasts a short or a long time does not have to be determined at the time that the trace is first formed—persistence can be influenced heterosynaptically by other activity patterns impinging on the target neuron over a window of time (of about an hour). This cellular property of LTP could be very relevant to episodic-like memory because it provides a potential mechanism for distinguishing between creating the potential for lasting memory traces and ensuring that a lasting trace is actually encoded.

The physiological studies first suggesting this possibility were reported by Frey & Morris (1997, 1998*a*). Using hippocampal brain slices *in vitro*, they found that application of a weak tetanus (21 pulses, 100 Hz) to one stratum radiatum pathway that ordinarily induced long-term potentiation decaying over 4 h (early LTP) could be made to induce a more persistent LTP (lasting > 8 h). This happened if stimulation with the weak tetanus was closely followed (1998*a*) or preceded (1997) by the application of three strong tetani (3×100 pulses, 100 Hz) to a separate input that successfully induced late LTP. Our explanation for this result centred on the synthesis of plasticity proteins in the target neurons. On the basis of data secured using the protein synthesis inhibitors anisomycin and emetine, Frey and Morris suggested that 'synaptic tags' set at the postsynaptic terminals of the weakly stimulated pathway sequester plasticity related proteins whose synthesis is induced in response to strong tetanization. These proteins, probably synthesized in the cell body, are presumed to travel in a non-targeted manner via transport mechanisms in the dendrite. Their function is to stabilize or 'consolidate' the otherwise transient synaptic potentiation; elegant studies in culture by

Martin *et al.* (1997) have established the mechanism at the single cell level. Put together (see figure 5), the input specificity of transient or lasting synaptic change is determined by the pattern of glutamatergic synaptic activation that, in addition to causing transient potentiation, also sets synaptic tags. The persistence of such changes is, however, determined by the history of activation of the neuron. This history governs the intersection between the availability of plasticity proteins and synaptic tags that have not yet reset to baseline.

At one level of analysis, the synaptic tagging hypothesis of the persistence of LTP is no more than a possible solution to the conceptual problem of targeting plasticity proteins to dendritic sites at which they are needed. However, this molecular machinery also endows the memory system of which it is a part with an almost magical property. This is to allow the memory system to 'keep its options' open with respect to the persistence of memory traces until and if other events occur that might influence whether or not a permanent trace should be created. The 'decision' to make a long-term memory does not have to be made at the time that the event to be remembered actually happens—a key contribution to the decision can occur beforehand or afterwards. The time window of decision making is determined, in part, by the kinetics of the synthesis and intracellular distribution of plasticity proteins. Such a property is ideal within that subset of the episodic memory system that automatically encodes attended experience. To speculate, we might think of the hippocampal formation as mediating this automatic component of episodic memory and the frontal lobes as the substrate of the intentional, executive parts of the system. In the intentional part, the focus of attention is on a limited body of new information that has to be learned—such as a list of paired associates in an experiment on cued recall. In contrast, in the automatic component of episodic memory, the subject has much less control over what gets temporarily recorded. The hippocampus records the incidentals of life as they happen to us on a minute-by-minute, hour-by-hour basis—we cannot stop ourselves from doing this. However, the mechanism of encoding is one that results in erasure of a very large part of this record of events within a matter of hours unless other events happen that could stabilize memory traces selectively. Behavioural experiments to examine this speculation from the physiological properties of LTP are underway.

7. DOES EPISODIC-LIKE MEMORY IN ANIMALS BECOME DYSFUNCTIONAL IN MODELS OF HUMAN DISEASE AND, IF SO, CAN WE USE SUCH MODELS TO DEVELOP NEW THERAPIES?

One value of working with animals is to develop models of human or animal diseases with a view to understanding them better and developing appropriate therapeutic strategies. Certain brain diseases, such as Alzheimer's disease, are well known to affect episodic memory during an early stage of its insidious progression. The importance of using appropriate behavioural tests of learning and memory is illustrated by recent work using transgenic mice engineered to be 'models' of aspects of the disease.

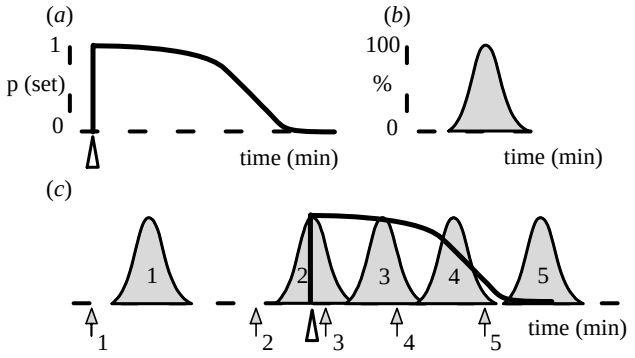

Figure 5. The synaptic tagging hypothesis of long-term potentiation (LTP). (*a*) Synaptic tag. In addition to inducing a transient form of synaptic potentiation (lasting at most a few hours), induction of LTP is also thought to set a 'synaptic tag' at activated terminals, which will reset after a few hours. (*b*) Plasticity proteins. Relevant gene activation, probably through heterosynaptic activation, sets in train the synthesis and intracellular distribution of plasticity proteins. These travel diffusely from their site of synthesis, whether the cell body or at local sites of dendritic synthesis. (*c*) Dynamics of tag–protein interactions. Tag–protein interactions are responsible for stabilizing or consolidating the otherwise transient potentiation. These interactions can occur when the availability of proteins intersects temporally with the activation of synaptic tags (positions 2, 3 and 4 in the diagram). Based on Frey & Morris (1998*b*).

It has long been known that mice over-expressing human mutant amyloid precursor protein (hAPP) show learning deficits. However, a puzzle that has eluded explanation is the apparent lack of relationship between these deficits and the progressive β-amyloid plaque formation that these mice sometimes display. In the standard reference version of the water maze in which the animals search for a hidden escape platform in a fixed location in space (Morris *et al.* 1982), hAPP mice are impaired throughout the life span—before and after amyloid plaque deposition (D'Hooge *et al.* 1996; Nalbantoglu *et al.* 1997; Moechars *et al.* 1999). Puzzled by the apparent lack of relationship to this striking age-related neuropathological change, Chen *et al.* (2000) wondered whether the water maze reference memory task was the appropriate behavioural assay. They decided to explore whether PDAPP mice, in which a human mutation of amyloid precursor protein is over-expressed under the control of a platelet derived growth factor promoter (Games *et al.* 1995), show age- and amyloid plaque-related learning or memory deficits using a more 'episodic-like' training protocol.

Accordingly, a new water maze training procedure was developed in which the mice had to keep changing their memory representation of the environment in much the same way as in the DMP task. A preliminary study using a version of the DMP task adapted for mice was unsuccessful, partly because performance was so poor in the PDAPP group that they failed to learn to return to the most recent location of the platform. To improve performance, yet retain the important feature that the memory representation of where escape is possible has to be updated frequently, a serial reversal procedure was used. A series of separate spatial learning problems was given to the mice,

each of which had to be learned to a pre-set criterion of performance. When the mice reached this criterion, whether transgenic or control, they were then switched to learning a new platform location. Reversal learning, as usually analysed, involves repeatedly switching the reward assignments of two cues with all other cues being irrelevant. This new procedure is analogous excepting that the significance and spatial relationship between the extra-maze cues—the animal's map of space—remains constant; all that changes is the location of the hidden escape platform within this map. Associative learning theory can offer a satisfactory account of varying rates of learning in such a procedure (often couched in terms of varying attention to relevant and irrelevant cues—see Mackintosh 1983, p. 252), but the procedure is 'episodic-like' in the following sense. The animal cannot treat the extra-maze cues as irrelevant and so ignore them because they provide the basis for its successful navigation. Nor can the animal use only its map of space to recall the current platform location. Instead, and as in the DMP task, the animal must either keep updating its map of space by ensuring that only one possible location of the platform can be recalled, or it must use recency as an additional cue for distinguishing the multiple long-term traces of where the platform might be located.

The results revealed that heterozygous PDAPP mice displayed both an age-independent and an age-related deficit in learning. Consistent with previous findings, the age-independent effect was apparent in learning the first of the series of spatial problems. This was to be expected because the training procedure, with the sole exception of training to criterion, is identical to a standard reference memory protocol as previously studied by other groups. However, as training continued on successive problems, an age-related deficit became apparent. Young PDAPP mice became successively better at learning to a point where they were indistinguishable from littermate controls; middle-aged and old PDAPP were even more impaired in an age-related manner (figure 6). This learning impairment was apparent in both a cross-sectional study (in which mice were first trained at different ages) and a longitudinal study (in which mice, first trained when they were young, were re-tested throughout the life span). It also correlated with β-amyloid (Aβ) plaque burden, with age itself removed as a covariate (see Chen *et al.* 2000). Training on other tasks, including navigation to a visible escape platform and a test of object recognition memory, revealed a further degree of selectivity to these animals' learning deficit.

These findings have an important practical implication. They indicate that the over-expression of Aβ and/or the frank deposition of Aβ plaques are associated with disturbed cognitive function and, importantly, suggest that some but not all tests of learning and memory are suitable behavioural assays of the progressive cognitive deficits associated with Alzheimer's disease-type pathologies. As disturbances of episodic memory are an early symptom of the disease, there is value in creating animal models of aspects of the disease using 'episodic-like' tests. In addition, two recent reports (Janus *et al.* 2000; Morgan *et al.* 2000) indicate that immunization against human Aβ in closely related hAPP mice can not only protect against the neuropathological deposition of amyloid plaques, as first shown by (Schenk *et al.* 1999), but can also prevent this age-related learning deficit from developing. It is noteworthy that variants of the 'episodic-like' DMP task and this serial reversal procedure were used in these successful vaccination experiments.

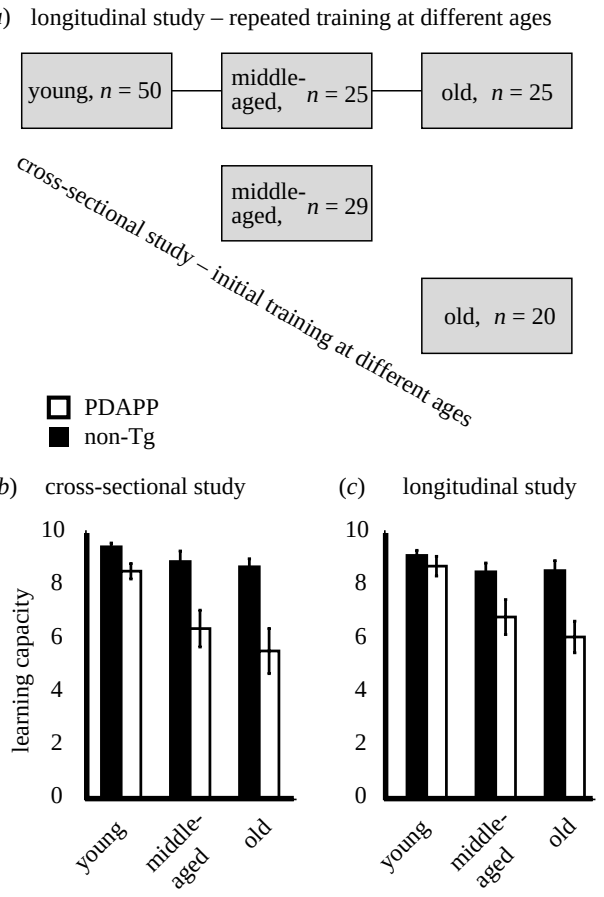

Figure 6. Use of serial-reversal to monitor the age- and plaque-related decline in 'learning capacity' in transgenic mice. (*a*) Longitudinal study—repeated training at different ages. Chen *et al.* (2000) used both a longitudinal and a cross-sectional design to examine the capacity of heterozygous PDAPP mice and their non-transgenic littermate controls to learn a series of spatial problems. (*b*) Cross-sectional study. The results showed that the 'learning capacity' (the total number of problems that could be acquired in a ten-day period) declined in an age-related manner in the PDAPP mice. This decline in learning also correlated with the age-related deposition of amyloid plaques. (*c*) Longitudinal study.

8. CONCLUSION

The primary aim of this paper has been to point to reasons for thinking that mammals may possess an 'episodic-like' memory system even though Tulving's formal definition of episodic memory puts the concept outside the realm of experimental study in non-human species. In developing episodic-like memory tests for animals, I have suggested that the current emphasis on one-trial learning may be misplaced; many one-trial tasks are ambiguous with respect to the demands they place on implicit versus explicit memory processes. The

focus now needs to be on developing behavioural protocols that capture the essentially 'explicit' nature of recall. An episodic-like system must, to be sure, have the capability of encoding unique events and a newly discovered property of LTP suggests that activity-dependent synaptic plasticity at the level of the individual synapse may be digital. However, the encoding mechanism might also be one in which the creation of a transient but automatic record of attended events does not preclude the selective stabilization of memory traces at other times through heterosynaptic interactions (Frey & Morris 1998*b*; Bailey *et al*. 2000).

In contrast to encoding and storage, episodic recall is primarily about 'mental time travel', reconstructing the past at the point of retrieval and knowing that it is the past. The behavioural criteria for identifying this mental state in animals remain ill-defined but they must be distinguished from the control of behaviour by parameters that merely accumulate value (associative strength, familiarity, etc.) independently of the path taken to that value. There must also be some preservation of the 'what, where and when' of events. Few of the tasks that have been developed to model different types of memory meet these strict definitional criteria. However, certain tasks, such as serial reversal learning, one-trial scene and object-in-place memory, delayed matching and non-matching to place are exquisitely sensitive to hippocampal lesions and to the selective disruption of hippocampal synaptic plasticity. While the argument is circular, it seems likely that these tasks are mediated by the same neural networks that subserve episodic recall in humans—including the hippocampal formation. Pending the development of analytically less ambiguous tasks, their use to investigate animal models of neurodegenerative disease is valuable pragmatically, as they could provide an effective way of evaluating new therapeutic strategies.

There remains at least one key issue where the argument of this paper differs sharply from that of Macphail's (1998) discussion of the evolution of consciousness. We have seen that Macphail asserts that animals have only an implicit learning system, that the development of language by humans is central to our having consciousness, and that animals, lacking language, cannot be conscious as we are. It follows from this argument that they cannot have anything like episodic recall. However, he also suggests that, in humans, the hippocampal formation is one of a group of structures that mediate explicit recall. The problem with this last part of the argument is that it leaves the hippocampal formation in mammals as what he calls a pre-adaptation or, to quote, a 'forerunner of the system used by humans to gain conscious access to memories' (Macphail 1998, p. 173). This is unsatisfactory because there must have been a selective advantage to the possession of this 'forerunner' (*sic*) long before the advent of hominids and the emergence of language. For Macphail, the mammalian hippocampus seems to be left in a kind of no man's land between having nothing to do with the primitive but stable implicit learning mechanisms shared by all vertebrates, but everything to do with the conscious, lateralized, language-based system used by humans. In contrast, the view taken here is that the hippocampus is a system for the automatic recording of attended experience that enables the encoding, storage and private recollection of experience in a form that would be advantageous to an animal but cannot yet be communicated to another. The challenge is to develop new analytically powerful behavioural tasks to study this memory system.

I am grateful to many colleagues and friends for discussing the ideas contained in this paper including Nicola Clayton, Anthony Dickinson, J. Uwe Frey, Susan Healy, Stephen Martin, Edvard Moser, Gernot Riedel and Emma Wood. The experimental work described includes studies conducted with J. Uwe Frey, Robert Steele, Guiquan Chen and colleagues at Elan Pharmaceuticals (Karen Chen, Dora Games and Stephen Freedman). This work was supported by grants from the Cunningham Trust, Medical Research Council, the Human Frontiers Science Programme and the European Union Framework V.

REFERENCES

Aggleton, J. P. & Pearce, J. M. 2001 Neural systems underlying episodic memory: insights from animal research. *Phil. Trans. R. Soc. Lond.* B **356**, 1467–1482.

Amaral, D. G. & Witter, M. P. 1989 The three dimensional organization of the hippocampal formation: a review of anatomical data. *Neuroscience* **31**, 571–591.

Amaral, D. G. & Insausti, R. 1990 The hippocampal formation. In *The human nervous system* (ed. G. Paxinos), pp. 711–755. San Diego: Academic Press.

Bailey, C. H., Giusetto, M., Huang, Y. Y., Hawkins, R. D. & Kandel, E. R. 2000 Is heterosynaptic modulation essential for stabilizing Hebbian synaptic plasticity and memory? *Nature Neurosci.* **1**, 11–20.

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. (Lond.)* **232**, 331–356.

Chen, G. (and 10 others) 2000 A learning deficit related to age and β-amyloid plaques in a mouse model of Alzheimer's disease. *Nature* **408**, 975–979.

Clayton, N. S. & Dickinson, A. 1998 What, where and when: episodic-like memory during cache recovery by scrub jays. *Nature* **395**, 272–274.

Clayton, N. S., Griffiths, D. P., Emery, N. J. & Dickinson, A. 2001 Elements of episodic-like memory in animals. *Phil. Trans. R. Soc. Lond.* B **356**, 1483–1491.

D'Hooge, R., Nagels, G., Westland, C. E., Muck, L. & De Deyn, P. P. 1996 Spatial learning deficit in mice expressing human 751-amino acid beta-amyloid precursor protein. *Neuroreport* **7**, 2807–2811.

Danysz, W., Zajaczkowski, W. & Parsons, C. G. 1995 Modulation of learning processes by ionotropic glutamate receptor ligands. *Behav. Pharmacol.* **6**, 455–474.

Dickinson, A. 1980 *Contemporary animal learning theory.* Cambridge University Press.

Ennaceur, A. & Delacour, J. 1988 A new one-trial test for neurobiological studies of memory in rats. 1. Behavioral data. *Behav. Brain Res.* **31**, 47–59.

Foster, D. J., Morris, R. G. M. & Dayan, P. 1999 A model of hippocampally-dependent navigation using the temporal difference learning rule. *Hippocampus* **10**, 1–16.

Frey, U. & Morris, R. G. M. 1997 Synaptic tagging and long-term potentiation. *Nature* **385**, 533–536.

Frey, U. & Morris, R. G. M. 1998*a* Weak before strong: dissociating synaptic-tagging and plasticity-factor accounts of late-LTP. *Neuropharmacology* **37**, 545–552.

Frey, U. & Morris, R. G. M. 1998*b* Synaptic tagging: implications for late maintenance of hippocampal long-term potentiation. *Trends Neurosci.* **21**, 181–188.

Gaffan, D. 1991 Spatial organization of episodic memory. *Hippocampus* **1**, 262–264.

Gaffan, D. 1994 Scene-specific memory for objects: a model of episodic memory impairment in monkeys with fornix transection. *J. Cogn. Neurosci.* **6**, 305–320.

Gaffan, D. & Harrison, S. 1989 Place memory and scene memory: effects of fornix transection in the monkey. *Exp. Brain Res.* **74**, 202–212.

Gaffan, D. & Parker, A. 1996 Interaction of perirhinal cortex with the fornix-fimbria: memory for objects and 'object-in-place memory. *J. Neurosci.* **16**, 5864–5869.

Gaffan, D. & Parker, A. 2000 Mediodorsal thalamic function in scene memory in rhesus monkeys. *Brain* **123**, 816–827.

Games, D. (and 33 others) 1995 Alzheimer-type neuropathology in transgenic mice overexpressing V717F β-amyloid precursor protein. *Nature* **373**, 523–526.

Garcia, J. & Koelling, R. A. 1966 Relation of cue to consequence in avoidance learning. *Psychonomic Sci.* **4**, 123–124.

Gluck, M. A. & Granger, R. 1993 Computational models of the neural bases of learning and memory. *A. Rev. Neurosci.* **16**, 667–706.

Granger, R., Wiebe, S. P., Taketani, M. & Lynch, G. 1996 Distinct memory circuits composing the hippocampal region. *Hippocampus* **6**, 567–578.

Griffiths, D., Dickinson, A. & Clayton, N. 1999 Episodic memory: what can animals remember about their past? *Trends Cogn. Sci.* **3**, 74–80.

Halliday, M. S. 1968 Exploratory behaviour. In *Analysis of behavioural change* (ed. L. Weiskrantz), pp. 107–126. New York: Harper and Row.

Holland, P. C. 1981 Acquisition of representation-mediated conditioned food aversions. *Learning and Motivation* **12**, 1–18.

Holland, P. C. 1990 Event representation in Pavlovian conditioning: image and action. *Cognition* **37**, 105–131.

Izquierdo, I. & Medina, J. H. 1998 On brain lesions, the milkman and Sigmunda. *Trends Neurosci.* **21**, 423–426.

Janus, C. (and 16 others) 2000 Aβ peptide immunization reduces behavioural impairment and plaques in a model of Alzheimer's disease. *Nature* **408**, 979–982.

Konishi, M. 1986 Centrally synthesised maps of sensory space. *Trends Neurosci.* **9**, 163–168.

Lisman, J. E. 1999 Relating hippocampal circuitry to function: recall of memory sequences by reciprocal dentate-CA3 interactions. *Neuron* **22**, 233–242.

McCarthy, R. A. & Warrington, E. A. 1990 *Cognitive neuropsychology.* San Diego: Academic Press.

McGaugh, J. L. 2000 Memory—a century of consolidation. *Science* **287**, 248–251.

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

McNaughton, B. L., Douglas, R. M. & Goddard, G. V. 1978 Synaptic enhancement in fascia dentata: cooperativity among coactive afferents. *Brain Res.* **157**, 277–293.

Mackintosh, N. J. 1975 A theory of attention: variations in the associability of stimuli with reinforcement. *Psychol. Rev.* **82**, 276–298.

Mackintosh, N. J. 1983 *Conditioning and associative learning.* Oxford: Clarendon Press.

Macphail, E. M. 1982 *Brain and intelligence in vertebrates*. Oxford: Clarendon Press.

Macphail, E. M. 1998 *The evolution of consciousness*. Oxford University Press.

Marr, D. 1971 Simple memory: a theory for archicortex. *Phil. Trans. R. Soc. Lond.* B **262**, 23–81.

Martin, K. C., Casadio, A., Zhu, H. Y. E., Rose, J. C., Chen, M., Bailey, C. H. & Kandel, E. R. 1997 Synapse-specific, long-term facilitation of aplysia sensory to motor synapses: a function for local protein synthesis in memory storage. *Cell* **91**, 927–938.

Martin, S. J., Grimwood, P. D. & Morris, R. G. M. 2000 Synaptic plasticity and memory: an evaluation of the hypothesis. *Ann. Rev. Neurosci.* **23**, 649–711.

Moechars, D. (and 12 others) 1999 Early phenotypic changes in transgenic mice that overexpress different mutants of amyloid precursor protein in brain. *J. Biol. Chem.* **274**, 6483–6492.

Morgan, C. L. 1894 *An introduction to comparative psychology.* London: Walter Scott.

Morgan, D. G. (and 14 others) 2000 Aβ peptide vaccination prevents memory loss in an animal model of Alzheimer's disease. *Nature* **482**, 982–986.

Morris, R. G. M. 1983 An attempt to dissociate 'spatial mapping' and 'working memory' theories of hippocampal function. In *Neurobiology of the hippocampus* (ed. W. Seifert), pp. 405–432. London: Academic Press.

Morris, R. G. M. & Frey, U. 1997 Hippocampal synaptic plasticity: role in spatial learning or the automatic recording of attended experience? *Phil. Trans. R. Soc. Lond.* B **352**, 1489–1503.

Morris, R. G. M., Garrud, P., Rawlins, J. N. P. & O'Keefe, J. 1982 Place navigation impaired in rats with hippocampal lesions. *Nature* **297**, 681–683.

Murray, E. A., Baxter, M. G. & Gaffan, D. 1998 Monkeys with rhinal cortex damage or neurotoxic hippocampal lesions are impaired on spatial scene learning and object reversals. *Behav. Neurosci.* **112**, 1291–1303.

Nalbantoglu, J. (and 10 others) 1997 Impaired learning and LTP in mice expressing the carboxy terminus of the Alzheimer amyloid precursor protein. *Nature* **387**, 500–505.

O'Keefe, J. 1976 Place units in the hippocampus of the freely moving rat. *Exp. Neurol.* **51**, 78–109.

Olton, D. S., Becker, J. T. & Handelmann, G. E. 1979 Hippocampus, space, and memory. *Brain Behav. Sci.* **2**, 313–365.

Parker, A. & Gaffan, D. 1998 Interaction of frontal and perirhinal cortices in visual object recognition memory in monkeys. *Eur. J. Neurosci.* **10**, 3044–3057.

Pearce, J. M. & Hall, G. 1980 A model for Pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychol. Rev.* **87**, 532–552.

Petersen, C. C., Malenka, R. C., Nicoll, R. A. & Hopfield, J. J. 1998 All-or-none potentiation at CA3-CA1 synapses. *Proc. Natl Acad. Sci. USA* **95**, 4732–4737.

Preuss, T. 1995 The argument from animals to humans in cognitive neuroscience. In *The cognitive neurosciences*, lst edn (ed. M. Gazzaniga), pp. 1227–1241. Cambridge, MA: MIT Press.

Rawlins, J. N. P. 1985 Associations across time: the hippocampus as a temporary memory store. *Behav. Brain Sci.* **8**, 479–497.

Redish, A. D. 1999 *Beyond the cognitive map: from place cells to episodic memory.* Cambridge, MA: MIT Press.

Rescorla, R. A. & Wagner, A. R. 1972 A theory of Pavlovian conditioning: the effectiveness of of reinforcement and nonreinforcement. In *Classical conditioning. II. Current research and theory* (ed. A. H. Black & W. F. Prokasy). New York: Appleton-Century-Crofts.

Riedel, G., Micheau, J., Lam, A. G. M., Roloff, E. L., Martin, S. J., Bridge, H., De Hoz, L., Poeschel, B., McCulloch, J. & Morris, R. G. M. 1999 Reversible neural inactivation reveals hippocampal participation in several memory processes. *Nature Neurosci.* **2**, 898–905.

Rolls, E. T. & Treves, A. 1998 *Neural networks and brain function*. Oxford University Press.

Rozin, P. & Kalat, J. W. 1971 Specific hungers and poisoning as adaptive specializations of learning. *Psychol. Rev.* **78**, 459–486.

Schenk, D. (and 24 others) 1999 Immunization with amyloid-beta attenuates Alzheimer-disease-like pathology in the PDAPP mouse. *Nature* **400**, 173–177.

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

Steele, R. J. & Morris, R. G. M. 1999 Delay-dependent impairment of a matching to place task with chronic and intrahippocampal infusion of the NMDA antagonist D-AP5. *Hippocampus* **9**, 118–136.

Stoerig, P. & Cowey, A. 1997 Blindsight in man and monkey. *Brain* **120**, 535–559.

Tulving, E. 1983 *Elements of episodic memory.* Oxford: Clarendon Press.

Tulving, E. & Schacter, D. L. 1990 Priming and human memory systems. *Science* **247**, 301–306.

Tulving, E. & Markowitsch, H. J. 1998 Episodic and declarative memory: role of hippocampus. *Hippocampus* **8**, 198–204.

Weiskrantz, L. 1997 *Consciousness lost and found.* Oxford University Press.