Hippocampus: Memory, Habit and Voluntary Movement

D. Gaffan

Phil. Trans. R. Soc. Lond. B 1985 **308**, 87-99
doi: 10.1098/rstb.1985.0012

Phil. Trans. R. Soc. Lond. B **308**, 87–99 (1985)
Printed in Great Britain

Hippocampus: memory, habit and voluntary movement

By D. Gaffan

Department of Experimental Psychology, South Parks Road, Oxford OX1 3UD, U.K.

A general method for studying monkeys' memories is to teach the animals memory-dependent performance rules: for example, to choose, out of two visual stimuli, the one that flashed last time the animal saw it. One may thus assess the animal's memory for any arbitrarily chosen event such as flashing even if the event itself has no intrinsic importance for the animal. The method also allows assessment of an animal's memory of the animal's own previous behaviour. The use of these methods has revealed a simple generalization about the function of the hippocampus in memory: hippocampal lesions impair memory of the voluntary movement that a stimulus previously elicited, but leave intact memory for relations between environmental events other than voluntary movements. The impairment in memory for voluntary movements produces deficits in exploration and in habit formation.

1. Memory tasks

Monkeys are adept at learning memory-dependent performance rules. A convenient method for teaching a monkey such a rule is to use a large population of visually discriminable stimulus objects. With the first pair of objects drawn from the population one proceeds to give one member of the pair one history, H1, and the other a different history, H2, in one or more 'acquisition trials' that are witnessed by the monkey. Subsequently the pair is presented for the animal's choice at a 'retention test'. Then that pair is discarded and the sequence of acquisition and retention is repeated with a new pair of objects. At retention tests it is always the object that has one of the histories, say H2, that is rewarded, and the other is not. Over a series of such tests the monkey will learn to choose between objects at retention tests according to their history. The accuracy of the choices made will therefore come to reflect the animal's ability to recall differentially H1 and H2, that is, to discriminate between the memories of H1 and H2. The stimulus populations used in the experiments reviewed below are either three-dimensional objects presented to the monkey by hand, or flat coloured patterns generated by a computer. The reward at retention tests is usually a peanut. The histories in question may be quite various, as table 1 exemplifies.

Some of the tasks in table 1 are well known and are given their conventional names in the table; for others a suggested name is given in quotation marks. The detailed methods of experiments with these tasks are described in the papers cited below in §2 and 3. For illustration in the table only the first four pairs of stimulus objects are shown for each task although in every case more retention tests than four would be required to teach a monkey the task. The two stimuli of each pair are identified by capital letters. The row of the table labelled 'retention test' shows the common retention test for all of the tasks, 'vs' indicating the permission to choose, '+' indicating the rewarded correct choice, and the order of the stimulus identifications indicating the left–right spatial position of the stimuli at the retention test. In the other rows of the table the acquisition trials for the different tasks are shown, separated by commas when there are more than one of them to be presented successively.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

Table 1. Some memory-dependent performance rules

task name	pair number 1	2	3	4
matching to baited samples	A^{+}	C^{+}	F^{+}	H^{+}
matching to unbaited samples	A	C	F	H
non-matching to randomly baited	B^{+}	D	E^{+}	G
'sensory recall (f)'	Af, Bx	Dx, Cf	Ex, Ff	Hf, Gx
'congruent recall'	A^{+}, B	D, C^{+}	E, F^{+}	H^{+}, G
object discrimination learning set	A^{+} vs B	D vs C^{+}	E vs F^{+}	H^{+} vs G
'incongruent recall'	A, B^{+}	D^{+}, C	E^{+}, F	H, G^{+}
retention test (all tasks)	A^{+} vs B	C^{+} vs D	E vs F^{+}	G vs H^{+}

See the text for explanation of the symbols.

The first three tasks, matching and non-matching, require the monkey to answer a simple question: which object has been recently presented and which has not? The object that is presented in acquisition is known as the sample. In cases where the sample at acquisition is baited with a reward, that is indicated by '+'.

Sensory recall requires an answer to a more complicated question: which stimulus was associated with which of two possible target events to be recalled? The target events in this case are 'f' which means that the stimulus flashes on and off at acquisition, and 'x' which means that the stimulus expands at acquisition. I have shown the version (f) where it is the memory of the stimulus having flashed that signals reward at retention, but naturally the opposite rule can also be learned. Though one or other of the memories of 'f' or 'x' signals reward, 'f' and 'x' themselves (not the memories of them) are both equally predictive of non-reward, since neither ever occurs at retention. This is an important difference between sensory recall and the bottom three tasks of the table, where acquisition events are themselves either rewarding or non-rewarding.

Congruent recall is so-called because the memory of an event predicts, in this case, that event, the events in question being reward and non-reward. In incongruent recall, on the other hand, the opposite relation is seen, since in that case the memory of reward predicts non-reward and the memory of non-reward predicts reward. In object discrimination learning set (o.d.l.s.) there is the same congruency between memory and expectation as exists in congruent recall. Furthermore, in o.d.l.s. the acquisition trial is operationally identical to the retention test.

O.d.l.s. is a familiar task that has been studied over many years in a number of laboratories. It is apparently simple. In a typical o.d.l.s. experiment all that happens is that a series of single simple two-choice visual discrimination problems is presented: first A and B are presented for a number of trials with A consistently rewarded, then C and D are similarly presented with C rewarded, and so on; and, not surprisingly, the rate of within-problem learning improves with practice. This simplicity is deceptive. The reason why all within-problem trials look the same in o.d.l.s. is, first, because each trial performs two functions that need to be separated logically, as an acquisition trial for subsequent retention tests and as a retention test for previous acquisition trials; and, second, because the peanut happens to be the recall target in the memory task as well as the inducement to perform the task correctly. Viewed in the context of monkeys' willingness to learn performance rules defined in terms of any arbitrarily chosen histories (table 1), the development of o.d.l.s. may most plausibly be assumed to reflect the acquisition of a memory-dependent performance rule: the monkey learns to choose whichever object is associated in memory with reward.

Serial visual reversal set is not shown in table 1 and differs from the tasks in the table in that only two stimuli are used throughout, in place of the large populations of stimuli used in the tasks described so far. So in reversal set, odd-numbered reversals will consist of repeated presentations of A vs B with A rewarded while even-numbered reversals will consist of A vs B with B rewarded. But the same performance rule as in o.d.l.s. is adequate for serial reversal set and, not surprisingly, monkeys transfer readily between these two tasks (Schrier 1966; Warren 1966).

Several implications of this type of approach may be noted at this stage and are taken up in greater detail in what follows. The first is an entirely open-minded attitude to the behavioural effects of memories. Formally speaking, one may say that a memory in the current treatment is nothing more than a property of the internal state elicited by retrieval cues at retention tests. If A and B are both paired with X and C and D are both paired with Y, then the memory of X is something that A and B but not C and D will elicit if they are subsequently presented. If the animal at retention tests discriminates or can be taught to discriminate between A and B and stimuli like them on the one hand, and C and D and stimuli like them on the other, then the hypothesis that the animal remembers X, or to be more precise that the animal can remember differentially X and Y, is justified since the sole purpose of speaking in this case of memories is to refer to the shared features of stimuli whose only systematic similarity to each other is their history. There is, for example, no further implication, unless explicitly stated, that the animal should behave as if expecting X to follow A. Second, it is essential to this approach that memories are clearly distinguished from habits. The phrase 'associative learning' is often used loosely to refer to the acquisition both of memories and of habits, since, in practice, the statement that an animal has associated A with X often means only that as the experimenter sees it, the animal has developed an appropriate response tendency to A which signals X. This usage must be eschewed if the relation of memory to behaviour is the question at issue, as below. Third, as the examples of o.d.l.s. and serial reversal learning make clear, the analysis of memory processes may have implications for understanding tasks that do not initially appear to be formally designed for the study of memory-dependent performance rules.

2. Learning and habit

In many theories of learning there is no scope for learning about memories. For example, Pavlov's principle of stimulus substitution is roughly speaking the principle that if A is associated in memory with B then the animal will react to A as if it were B. We have then on the one hand an associative memory mechanism that records the association of A with B, and, on the other hand, a performance rule according to which the output of memory (recall of B, in the example) elicits behaviour; and the performance rule is not itself a product of learning but a pre-existent and fixed condition of the manifestation of learning in behaviour. Figure 1*a* is an attempt to present these relations schematically.

The account above neglects Pavlov's treatment of 'trace conditioning'. However, as Revusky & Garcia (1970) have well argued, the non-associative process that bridged the delay in trace-conditioning was not explained as part of Pavlov's theory but was added on to the main body of theory by an unexplained assumption. The word 'trace' was used in the same way by Hull (1952), to refer to an unexplained short-lasting persistence of stimuli from one trial to the next. It is of historical interest that, during the era when memories appeared in theories only in this impoverished form, experiments on memory tasks other than discrimination

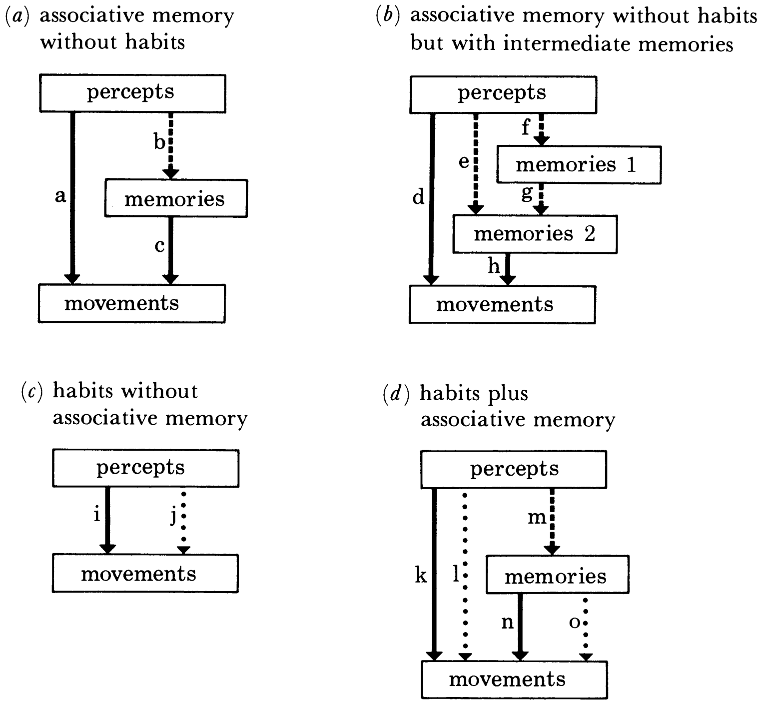

FIGURE 1. Schematic representations of four possible hypotheses of the relation between memories and movements. Solid lines indicate unlearned direct influences upon movement. Dotted lines indicate habits, namely, learned direct influences upon movement. Dashed lines indicate memory traces, that is, the ability of a retrieval cue to elicit a memory. Note that memories are treated as potentially able to do everything that percepts can do: act as a retrieval cue for a memory (line g), exert inflexible effects (line c), or elicit habits (line o).

learning required monkeys to remember only from one trial to the next. For example in delayed matching or non-matching, Nissen *et al.* (1948) and subsequently many others gave their monkeys only one sample to remember at a time, with an immediate retention test; this undemanding method may be contrasted with the lists of samples presented by Gaffan (1974), Mishkin (1978) and other modern experimentalists. Whether the lack of data held back the theories, or the lack of theory the experiments, is not clear; but at any rate it is worth stressing now that when monkeys are given lists of acquisition trials with different stimuli, their memory for events other than the reward, for example in delayed matching their memory of having seen an object, may be no less powerful than their memory for whether an object was rewarded or not (Gaffan 1976); and further that though the retention tests are presented next to the relevant acquisition trials in table 1 so that the rules should be clear to the reader, monkeys may not only perform but also learn such rules even when retention tests never follow immediately upon their objects' acquisition trials (Gaffan *et al.* 1984*c*).

With the minor exception of trace conditioning, therefore, the scheme presented in figure 1*a* is an adequate representation of the relation between associations and behaviour in Pavlovian theory. It is clearly inadequate as a model of monkeys' learning but the necessary extension of the simple model in figure 1*a* is not difficult to envisage and is presented in figure 1*b*. This model still appeals to no other process than associative memory in explaining learned modifications of behaviour, but it accommodates memories about memories. Consider, for example, matching or non-matching to sample. According to figure 1*b*, any stimulus that is associated in memory with food (line e) operates via a fixed translation rule (line h) to elicit approach. However there are, in addition, some other memories, 'memories 1' in the figure, that have no fixed unlearned effect (or only a weak fixed effect) on approach, but that can themselves act as a retrieval cue for the memory of food. Thus to take as an example the memory of having seen an object before or the memory of not having seen an object before, these memories will themselves become associated in memory with food, if it is the case that familiar or novel stimuli respectively are always rewarded with food. Then in, say, matching to sample, a sample at the retention test will elicit via line f a memory that via line g has been associated, over many similar retention tests with different objects, with food; and the sample at the retention test in matching will through this indirect process elicit approach.

Model 1*b* is a realistic proposal of at least one way that an animal could learn matching or non-matching. The idea that novelty or familiarity, as memory-retrieved properties of objects in general, are associated with food reward in matching or non-matching, receives strong support from experiments investigating the role of the sample bait at acquisition (Gaffan *et al.* 1984*d*): the sample when it appears at the acquisition trial is a novel object, and matching is therefore facilitated if the sample at acquisition, like the non-sample at the retention test, is not rewarded. Furthermore, this model is not restricted to a Pavlovian assumption of stimulus substitution. This type of account, where some associative memories have an unlearned effect upon behaviour via a fixed translation rule, includes also, for example, the inferential neo-Tolmanist aproach of Dickinson (1980, p. 115) and Mackintosh (1983, p. 111). According to this inferential hypothesis or model, unlearned logic compels a hungry rat to derive from knowledge of its hunger, and from the memory that approaching A has resulted in food, the practical inference that A should be approached, and some further unlearned process puts the practical inference into effect. In addition I may mention the version of this type of hypothesis that I myself favour.

I would suppose that just as certain percepts evoke unconditioned responses whose form is various since they are specialized to be appropriate to the individual stimuli in question, so also certain memories evoke unconditioned responses that are specially appropriate to those memories. These responses are not necessarily identical to the unconditioned responses evoked by the percepts that the memories are memories of: the memory of pain elicits responses that are not all identical with responses elicited by pain itself, and the memory of warmth elicits gentle pecking in chicks, which do not respond to warmth itself by pecking (Wasserman *et al.* 1975). This version of the hypothesis has the advantage over the Dickinson–Mackintosh version that the chicks need neither logical reasoning power nor the propositional knowledge that heat sources are likely to be hens and that the way to cajole a hen is to peck it gently. If memories are in this way seen as similar to percepts in their elicitation of unconditioned responses it becomes all the easier to see them as similar to percepts also in their potential elicitation of conditioned responses by a further association in memory. In the case of memory of peanuts in monkeys this hypothesis would however, as is often the case, be indistinguishable from the hypotheses either of stimulus substitution or of inference since all would agree that the effect is to elicit displacement of the object that elicits, either directly or via a mediating association, the memory of the food reward.

Finally in support of the generality of model 1*b*, the labelling of memories 1 and memories 2 in the figure is an aid to exposition when only one response is being considered, but a more accurate statement would be that memories can act as retrieval cues for further memories in a non-hierarchical fashion: if a monkey can learn, say, that the memory of warmth is associated with food then the animal can also presumably learn that the memory of food is associated with warmth.

However, in spite of all these strengths model 1*b* suffers from an obvious weakness: it cannot cope with a task such as incongruent recall. A situation where the memory of food is associated with no food, and the memory of no food with food, and where the animal's behaviour is directly elicited by the memories of food and no food, is every bit as confusing as it sounds for 1*b*. Some rear-guard action in defence of the model might have been possible if incongruent recall had been an exceptionally difficult task; unfortunately the monkeys do not find it confusing at all. In comparison with congruent recall, incongruent recall is neither a more difficult rule to learn nor a more difficult task to perform (see figure 2 and the acquisition data in experiments 3 and 4 of Gaffan (1979) and in experiment 1 of Gaffan *et al.* 1984*a*).

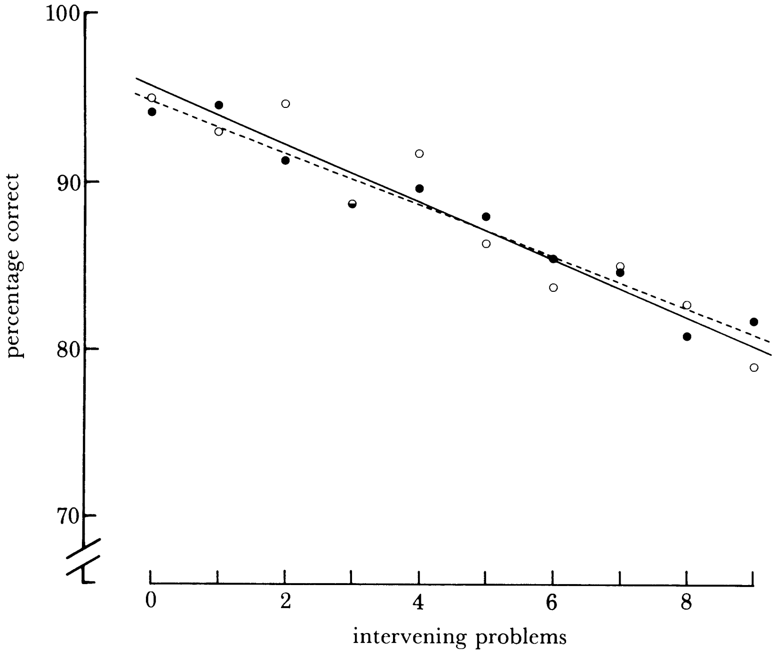

FIGURE 2. Performance of congruent recall (dashed line, filled circles) and of incongruent recall (solid line, open circles) by two groups of well-practised rhesus monkeys, containing eight and seven animals respectively. The groups are indistinguishable both in forgetting rate and in overall accuracy.

At this point it is worth stressing again the common feature of the hypotheses we have considered so far. They all rely for their explanation of learning upon the animal's acquisition of associative memories. Apart from the acquisition of associations in memory, the other mechanisms that they invoke to explain behaviour (inference, stimulus substitution, unconditioned reflexes to memories, and so on) are not themselves the product of learning; the only direct effect of learning that they allow is associative memory, and anything like a habit or a performance rule has to be explained in terms of associative memories and their sometimes complicated effects and interactions. But if 1*a* is, with the minor exception of the impoverished 'trace' concept, an adequate depiction of Pavlov's views on the relations we are interested in, then figure 1*c* is, with the same minor reservation, an adequate depiction of Hull's. Each of these extreme views is equally obviously inadequate to explain learning about memories, but it is nevertheless instructive to consider the hypothesis that habits unmediated by memories are one of the direct products of learning.

Consider next and finally, therefore, the composite model presented in figure 1*d*. This denies none of the mechanisms we have discussed. (The omission of the possibility of memories entering into further associations in memory is made for the sake of comprehensibility only.) Associative memories are formed and can have unconditioned effects; as well, habits are formed and can link either percepts or memories directly to behaviour. The main advantage of this model, from the evidence we have considered so far, is its ability to cope with incongruent recall. The habits envisaged in figure 1*d* are 'blind' habits in the sense that, unlike habits in the exclusively mnemonic models considered above, their performance does not require the mediation of a memory of the reward that established the habit. The habit of choosing stimuli that are associated in memory with non-reward in preference to stimuli associated in memory with reward can therefore proceed without any confusion or interference between the recall task and the motive for it.

However, mention of the motive and the reward calls attention to what the model is lacking, namely a mechanism of reinforcement. For Hull or Thorndike, reinforcement had an unalterable effect of strengthening the S–R conjunction which preceded it, and if that hypothesis was adopted then figure 1*d* would have a difficulty of its own in incongruent recall, because the peanut at acquisition would strengthen the habit of displacing that stimulus object under which it was found and would therefore detract from correct performance at the subsequent retention test. What is needed therefore is some more subtle hypothesis of reinforcement if figure 1*d* is to be viable. This is clearly not an insurmountable obstacle but before tackling it the evidence from fornix transection should be examined.

3. Impaired habit formation after fornix transection

In clinical neuropsychology several well-known and ancient lines of evidence suggest that the hippocampal system (hippocampus, fornix and mammillary bodies) is specially important for memory. In experimental neuropsychology and neurophysiology much effort over the last 30 years has gone into the attempt to provide firm evidence for the role of the hippocampus in memory and learning. Anyone who has dipped into this literature will not need to be told that in the space available here I do not intend to try to review or summarize it. However, there are some predecessors in this field of study whose influence I would particularly like to acknowledge. The remarkable correlation between hippocampal electrical activity and voluntary movement was observed and documented by Vanderwolf (1971) and Black (1975). Vanderwolf pointed out that one way to define voluntary movement is as movement whose elicitation is modified by its consequences, and that a connection might therefore be found between his observations and the behavioural evidence of some kind of learning impairment after hippocampal lesions. The study of fornix transection in monkeys was initiated by Mahut (1972). The fornix of a monkey is an accessible and clearly demarcated structure surrounded by ventricle. It is possible to transect the fornix reliably, in every animal in a group, and completely. If extra-fornical damage is produced there is no difficulty in describing the extent of that damage since the adjacent structures are clearly differentiated from the fornix, and though extra-fornical damage needs to be only slight and asymptomatic it can, if desired, be replicated in a control group. If psychologists cannot make sense of the behavioural effects of fornix transection then we may as well abandon the hypothesis that anatomically discrete structures have discrete functions that can be delineated by behavioural experimentation, since the fornix is the clearest possible example of an anatomically discrete structure and it is also by good fortune a pathway that can be cleanly and completely transected. Finally the origin of many of the ideas expressed below can be found in a lucid and strongly argued paper by Douglas (1967).

A lesion that impairs performance of a memory task may do so either by impairing the memories that are necessary for performance, or by impairing the execution of the performance rule that translates those memories into correct choices. With one important exception to be discussed below, that of incongruent recall, every task in table 1 has been shown to be performed normally by fornix-transected monkeys when they are given ample opportunity for practice at it (Gaffan *et al.* 1984*a*, *c*, *d*). The same applies to visual reversal set (Mahut 1972; Gaffan & Harrison 1984*b*) and to visual–visual memories and their reversal in a complicated paradigm to test acquisition of object–object associations (Gaffan & Bolton 1983; Gaffan *et al.* 1984*a*, experiment 3). These results suggest, therefore, that memories of the type required by these tasks are not impaired by fornix transection, and that in those instances where an impairment has been observed in tasks of this type (when, as we shall see in a moment, a new procedure was introduced after initial acquisition of the task, and performance in the new procedure was assessed without extensive practice in it) the memories were available to the monkeys but the performance rule was impaired. Since, according to model 1*d*, memory-dependent habits are similar in principle to perceptually based habits, the results from these memory tasks may imply a general impairment of some kind in the formation or modification of habits. As the following discussion shows, the evidence is consistent with that implication.

In matching or non-matching, for example, fornix-transected monkeys can show the same high levels of performance as normal monkeys (Gaffan *et al.* 1984*c*, *d*). When an impairment has been observed in these tasks it was when an established performance rule was required to be modified in response to a change in procedure (Gaffan 1974; Gaffan *et al.* 1984*c*). In the first experiment on matching and fornix transection (tasks 1 and 2 of Gaffan 1974) I began by adopting the procedure which as noted above (section 2) had been normal since the report by Nissen *et al.* (1948): the monkeys were initially trained with all retention tests following immediately upon their sample's acquisition trial. To my disappointment, the fornix-transected animals learned delayed matching just as fast as the controls. Then to test the power of the animals' memories I began introducing some very long retention intervals, and a deficit was immediately apparent in the fornix-transected group. At the time it seemed obvious that the deficiency was in long-term memory as opposed to immediate memory. It was not until several years and hypotheses later that my colleagues and I considered the possibility that the cause of this deficit and of a similar deficit in non-matching (Owen & Butler 1981) might have been the change from the conditions of original learning rather than the retention interval as such. We therefore taught non-matching to naive monkeys with a variety of retention intervals from the beginning of training: there was now no sign of an impairment in the fornix-transected group at any retention interval, although an impairment could be produced by some subsequent changes of procedure (Gaffan *et al.* 1984*c*).

This deficit in the ability to modify learned memory-based habits is seen also in perceptual habits. Fornix-transected monkeys were slow to reverse their performance in a task I shall call 'f.r.-dro reversal', where fixed-ratio and 'dro' response requirements were alternated between sessions: in the same apparatus and with the same stimuli every day, repeated manual contact (f.r.) with a stimulus was rewarded on one day and withdrawal of the hand from it (dro) was rewarded on the next (Gaffan & Harrison 1984*b*, experiment 2). In search of a contrast with serial visual reversal set which as noted above is performed normally by monkeys with lesions of the hippocampal system, we taught our monkeys a serial reversal set in this f.r.-dro reversal task. After several reversals the normal animals all reached a pre-established criterion of rapid reversal learning. None of the fornix-transected animals reached the criterion, and at the end of the experiment their within-reversal learning was still slower than the normal monkeys' even though by then they had had substantially more practice than the controls at the serial reversal task.

The modification in the animal's learned movements required from day to day in serial f.r.-dro reversal is quite different from that required in serial visual reversal learning. In serial visual reversals (as in o.d.l.s. or congruent recall) a memory is available to guide the animal's displacement of one stimulus object rather than the other. That memory, namely the memory of the food reward that sometimes one and sometimes the other object recalls, is a valid guide throughout the task. Objectively speaking, the animal's response in displacing an object switches each day from one object to the other, but as we have seen above the memory-dependent habit is not changed, only the association in memory (recall the strict distinction between habit and memory that was drawn above). But there is no similar memory to mediate between stimulus and response in the f.r.-dro reversal, since an association in memory between the stimulus features of the task and the reward outcomes is not by itself in f.r.-dro a valid guide for the animal's response. The habit itself, the tendency to make either an f.r. or a dro response in the context of the task, must change. A similar analysis applies to reversals of a simultaneous spatial discrimination, when an animal learns to go left one day and to go right the next. Fornix transection impairs performance of this task (Mahut 1972), and since the impairment is not the result of an impaired ability to discriminate spatial location as a perceptual cue (Gaffan *et al.* 1984*a*, experiment 6) it is attributable to a defect in the control of the locomotor movement, comparable to the defect in the control of manual movements that is revealed in the f.r.-dro reversal set.

Even within the paradigm of visual reversal learning it is possible to distinguish the case of habit reversal from the case of memory reversal. Within the very first visual discrimination an animal learns, say A positive vs B, A is, of course, no less valid a guide for the animal's response than the memory of reward is, and the animal might therefore acquire a perceptual habit of approaching A as well as a memory of A's association with reward. Subsequent reversal of that discrimination would then require a change in habit. In conformity with this suggestion, fornix transection did produce an impairment in the first post-operative visual reversal (Gaffan & Harrison 1984*b*, experiment 1).

In the examples discussed so far fornix transection impaired the ability to change an established habit. Some further observations suggest that this impairment can be subsumed under a more general defect, namely impaired learning ability when one habit is to be formed in one set of circumstances and a different habit is to be formed in a different set of circumstances that is similar to the first and therefore liable to be confused with it. In reversal learning the two circumstances follow each other, but in learning a conditional response they are intermingled in time and the formation of the two habits proceeds simultaneously with neither necessarily taking precedence. An example of such conditional response learning is a task where the animal must go left in the presence of one set of visual stimuli and go right in the presence of a different set. This task was learned abnormally slowly by fornix-transected monkeys (Gaffan *et al.* 1984*a*, experiment 5). We have observed a similar impairment in recent unpublished experiments with visual conditional learning of the f.r. and dro responses.

In reversal learning the contrast between visual reversal set and response reversal set shows that the lesion effect is concerned with some aspect of the learned control of the animal's own movements, since if there had been a general impairment in reversal as such there would have been a deterioration in the rate of reversal of memories in the visual reversal set. The same point can be made about impaired conditional response learning. Configural discriminations, which involve conditional relations between stimulus properties rather than between stimuli and the animal's own movements, were learned normally. There was no impairment in learning an auditory–visual configural task where an auditory cue signalled which of two visual stimuli was rewarded, and there was also no impairment in learning a visual–spatial configural task (Gaffan *et al.* 1984*a*, experiments 6 and 7). Thus the impairment in conditional response learning is not a general impairment in learning about conditional relations.

Finally, consider one more example of impaired habit formation, in this case of a memory-dependent habit. The explanation offered above for normal animals' high levels of performance in incongruent recall was that they formed a memory-dependent habit. This habit requires some reversal learning in the sense that monkeys that are experimentally naive show some tendency to perform according to the performance rule of congruent recall (Gaffan 1979, experiment 3); it also requires a kind of conditional response to be learned, since memories similar to and liable to be confused with those that control incongruent recall performance control something more like congruent recall performance in the animal's daily experience of extra-experimental feeding. As might therefore be expected, fornix transection severely disrupted incongruent recall performance (Gaffan *et al.* 1984*a*, experiment 1).

In conclusion then fornix transection impairs habit formation both in line l and in line o of figure 1*d*. It does not completely abolish it: as the examples illustrated have shown, fornix transected animals form habits more slowly, but they do form them. The contrast is with sensory memories where the memories are acquired without impairment at a normal rate of acquisition. The deficit in habit formation is especially evident when there is the possibility of confusion between conflicting habits appropriate to similar circumstances, as in reversal or conditional learning. These two features, the survival of some capacity for habit formation, and the special difficulty in discriminating habits that are liable to confusion, may not be theoretically distinct since it is presumably in the face of such liability that efficient acquisition of appropriate habits is most necessary. But the idea of a perceptual stimulus eliciting a voluntary movement via a habit requires careful definition. Fornix transection does not produce a general difficulty in acquisition of reversed or conditional memories, actions or percepts, as is shown by the examples above of reversal in sensory memories and of learned performance rules that involve conditional relations between percepts; the difficulty observed is therefore something specially to do with the animal's own learned movements as opposed to events in the environment; and yet there is no impairment in the smooth execution of movements, nor in postural adjustment, nor in the sequencing and control, presumably via sensory feedback, of a complex behaviour such as the emission of an f.r. response. We may think therefore of a habit as a mechanism for conveying a vague instruction ('leave it alone', 'go left', 'hammer it') from the visual or auditory modalities, in which the cue for the habit is recognized, to a motor command that has its own complex routines to determine in detail just how a dro response, a locomotor response or an f.r. response (respectively) is constituted in terms of muscular contraction and sensory feedback. The function of the fornix, then, is to facilitate the acquisition of excitatory links which carry vague instructions from the temporal lobe where perceptual recognition takes place to the thalamus where movements are initiated. But the facilitation of the acquisition of habits is precisely what is meant by reinforcement, of which as noted at the end of the previous section a hypothesis is required. Not surprisingly then, further examination of effects of fornix transection may supply such a hypothesis.

4. Reinforcement, exploration and response memory

Exploration in animals is by no means a matter of running around aimlessly when there is nothing better to do. It is a memory-dependent performance rule, though presumably an innate one. Exploration is directed towards alternatives that are novel in the sense that they have not recently been explored, and a memory of past behaviour is therefore required in exploration to direct present behaviour towards novel choices. Much research upon exploration is performed with rats and is concerned with their locomotor choices between spatially defined alternatives, but exploration can be similarly demonstrated in monkeys' choices between visually defined alternatives. Naive monkeys were tested in a paradigm identical with matching or non-matching to sample except that both objects at the retention test were rewarded. Since the sample at acquisition was also rewarded this was a pleasant task in which every displaced object revealed a peanut. Normal monkeys showed a tendency towards spontaneous non-matching. That tendency was abolished by fornix transection (Gaffan *et al.* 1984*c*).

Consider therefore the hypothesis that since exploration is directed by memory, and since damage to the hippocampal system impairs exploration in animals and memory in men, the direct effect of the fornix transection in the experiment above might be to impair memory in some way and thus to cause indirectly the observable decrement in exploration. If this hypothesis is true then the memory impairment in question must affect some aspect of memory that is more important for spontaneous non-matching than for trained non-matching, since the latter is unimpaired. The obvious possibility was that the crucial aspect of memory might be the animal's memory for its own responses to stimuli as opposed to the sensory familiarity of the stimuli in themselves, since the function of exploration is to direct the animal's movements. This possibility was easy to test. A monkey will displace an object that covers a food well, but will leave alone an object that is just behind the food well, even though the latter object will be seen if the food well is baited and a peanut is retrieved from it. To return to the terminology of §1, we have here a difference between two histories. If an object is placed on top of a baited well then its history will include the fact that the monkey displaced it. If an object is placed instead just behind a baited well then its history will be similar to that of the first object except that it will not have been displaced. So one can set up a memory task to distinguish between these two histories, either by rewarding, at subsequent retention tests, objects that have been displaced and not rewarding objects that have not been displaced, or vice versa. The task may be called respectively 'push-matching' or 'push-non-matching'. Both of these tasks were impaired by fornix transection and this impairment was not alleviated by practice (Gaffan *et al.* 1984*c*, *d*). There was thus a clear contrast between memory for responses and memory for the various aspects of environmental events that are involved in the unimpaired tasks of table 1.

The distinction between memory and habit is once again crucially important. When a stimulus elicits a memory of the animal's previous response to that stimulus it is difficult not to think in terms of an 'S–R association', but the connotation of that term is entirely inappropriate to the present context. In the terms of the conclusion of § 3 above, the memory of a previous response is not the issuing of a vague instruction but instead the memory of a vague instruction having been obeyed. It is only if this distinction is clear that one can ask the interesting question: what is the relation between the memories of previous responses to itself that a stimulus evokes, and the habit that that stimulus elicits?

A concise answer to that question is that reinforcement is the converse of exploration: a habit is what is left after exploration has ceased. Moreover, having defined exploration as the choice of a novel response, which is already a memory-dependent performance, one can further propose that at a higher level exploration itself is controlled by memories. The form of control would be that, other things being equal, if a stimulus is familiar, and evokes a memory of consistent previous responses to itself, and also evokes no memory of any recent surprises, then that stimulus suppresses exploration in favour of the opposite performance rule to exploration, namely habit, the production of the response that the stimulus recalls.

We may now return to the question that was posed at the end of § 2. It can now be seen that the real problem of the cognitive model of habit formation was not its appeal to associative memories but more specifically its appeal to memory of reward. The hypothesis that habit is to be explained by unconditioned responses to memories has no difficulty in allowing for the habitual choice of stimuli associated with non-reward in incongruent recall, so long as the memories that in turn elicit the performance of the habit are not themselves memories of reward but, as envisaged above, memories of the previous performances of the habit.

The fact that the cognitive model of reinforcement appeals to memory of the reward that established the habit is no coincidence. In the everyday language which cognitive psychology adopts, memory is in the paradigm case the ability to report verbally upon the past, and intentional or voluntary action implies an ability to justify the action verbally. In the case of a reinforcement as used to induce learning of a task, that justification would of course be stated in terms of the past delivery of the reinforcer. To make the argument from the other point of view, a cognitive psychologist would have every reason to object to the hypothesis formulated above if it were offered as a piece of cognitive psychology, that is as an explanation of habitual action in terms of knowledge. A man may even perform a habit to find out what he habitually does in a certain circumstance, because he cannot remember what he does; more generally, it would be difficult to claim that every habitual action is preceded by a memory of itself in the cognitive sense of memory as conscious awareness of the past. The technical sense of memory defined above (§ 1) is, of course, another matter. The study of physiological mechanisms of behaviour leads one by gradual steps away from the natural-language categories in which questions about behaviour are most easily formulated in the initial stages of that study.

REFERENCES

Black, A. H. 1975 Hippocampal electrical activity and behavior. In *The hippocampus*, vol. 1 (ed. R. L. Isaacson & K. H. Pribram), pp. 129–167. New York: Plenum.

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

Douglas, R. J. 1967 The hippocampus and behavior. *Psychol. Bull.* **67**, 416–442.

Gaffan, D. 1974 Recognition impaired and association intact in the memory of monkeys after transection of the fornix. *J. comp. physiol. Psychol.* **86**, 1100–1109.

Gaffan, D. 1976 Recognition memory in animals. In *Recall and recognition* (ed. J. Brown), pp. 229–242. London: Wiley.

Gaffan, D. 1979 Acquisition and forgetting in monkeys' memory of informational object-reward associations. *Learning Motiv.* **10**, 419–444.

Gaffan, D. & Bolton, J. 1983 Learning of object–object associations by monkeys. *Q. Jl exp. Psychol.* **35 B**, 149–155.

Gaffan, D., Saunders, R. C., Gaffan, E. A., Harrison, S., Shields, C. & Owen, M. J. 1984*a* Effects of fornix transection upon associative memory in monkeys: role of the hippocampus in learned action. *Q. Jl exp. Psychol.* **36 B**, 173–221.

Gaffan, D. & Harrison, S. 1984*b* Reversal learning by fornix-transected monkeys. *Q. Jl exp. Psychol.* **36 B**, 223–234.

Gaffan, D., Gaffan, E. A. & Harrison, S. 1984*c* Effects of fornix transection upon spontaneous and trained non-matching by monkeys. *Q. Jl exp. Psychol.* **36 B**. (In the press.)

Gaffan, D., Shields, C. & Harrison, S. 1984*d* Delayed matching by fornix-transected monkeys: the sample, the push and the bait. *Q. Jl exp. Psychol.* **36 B**. (In the press.)

Hull, C. L. 1952 *A behavior system*. New Haven: Yale University Press.

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

Mahut, H. 1972 A selective spatial deficit in monkeys after transection of the fornix. *Neuropsychologia* **10**, 65–74.

Mishkin, M. 1978 Memory in monkeys severely impaired by combined but not by separate removal of amygdala and hippocampus. *Nature, Lond.* **273**, 297–298.

Nissen, H. W., Blum, J. S. & Blum, R. A. 1948 Analysis of matching behavior in chimpanzee. *J. comp. physiol. Psychol.* **41**, 62–74.

Owen, M. J. & Butler, S. R. 1981 Amnesia after transection of the fornix in monkeys: long-term memory impaired, short-term memory intact. *Behav. Brain Res.* **3**, 115–123.

Revusky, S. & Garcia, J. 1970 Learned associations over long delays. In *The psychology of learning and motivation*, vol. 4 (ed. G. H. Bower), pp. 1–84. New York: Academic Press.

Schrier, A. M. 1966 Transfer by macaque monkeys between learning-set and repeated-reversal tasks. *Percept. Motor Skills* **23**, 787–792.

Vanderwolf, C. H. 1971 Limbic–diencephalic mechanisms of voluntary movement. *Psychol. Rev.* **78**, 83–113.

Warren, J. M. 1966 Reversal learning and the formation of learning sets by cats and rhesus monkeys. *J. comp. physiol. Psychol.* **61**, 421–428.

Wasserman, E. A., Hunter, N. B., Gutowski, K. A. & Bader, S. A. 1975 Autoshaping chicks with heat reinforcement: the role of stimulus–reinforcer and response–reinforcer relations. *J. exp. Psychol. Anim. Behav. Processes* **1**, 158–169.