Brain and Intelligence in Vertebrates

Euan M. Macphail

Brain and intelligence in vertebrates

E. M. MACPHAIL
Senior Lecturer in Psychology,
University of York

Preface

The original impetus for writing this book lay in my early efforts to place my research work in a wider context. That work concerned the role of avian forebrain structures in learning, and led naturally to an interest in non-mammalian brains, and to following through the relevance of studies using non-mammals to the intelligence of mammals including, of course, man. I found that ideas concerning the functional organization of non-mammalian forebrains were changing rapidly, but that the relevant research was published in a wide range of journals, many of which were not widely read by psychologists. I also found that, although comparative studies of learning did appear in the standard psychological journals, there was no detailed recent survey of them available, and no attempt to relate them to neurological information. In this book I have tried to bring together neurological and behavioural findings relevant to intelligence in the various classes of vertebrates, in the hope of making information derived from diverse disciplines more readily accessible to all those with an interest in intelligence.

Although this book is intended as a contribution to comparative psychology, it should be noted that it does not deal with invertebrates nor (at least directly) with the evolution of either the brain or intelligence. My major concern has been to organize and describe the findings obtained from each group in as clear a way as I could, and to avoid speculations (such as those involved in surmising the course of evolution) for which little hard evidence was available. I have not, however, avoided speculation altogether; indeed, one of the unexpected outcomes of the hours spent in libraries involved in preparing my material has been that I have come to a very different view of the nature of intelligence in vertebrates from the one I held some years ago. My current view is, of course, represented throughout the book, but I hope that, although it may have coloured the presentation of some areas of research, it will not interfere with the utility of the survey to readers who wish to be given a sufficiently full account of the experimental results available to enable them to provide their own interpretations. I shall not be too downcast if my views are not adopted by all (or even many) who read this book, but I shall be disappointed if those who do not share those views find the book on that account less useful.

Two apologies are no doubt in order. First, to those who find some of the behavioural sections too complex. I can only say that theories of animal learning have never been easy to follow, and that, in order to grasp the essence of current controversies, we must go rather deep into theory. The second apology concerns errors which I have assuredly made, but have not yet detected, in presenting results from fields in which I am far from expert: here, I must hope that the virtue of having a single relatively consistent view outweighs the evident disadvantage of lack of direct research experience in many of the areas surveyed.

Many friends and colleagues have read drafts of the book and have given me comments, and, equally important, encouragement. Drafts of individual chapters were read by Peter Bailey, Bob Boakes, Anthony Dickinson, Charles Hulme, Bundy Mackintosh, and Neil Thomson; drafts of the entire manuscript were read by Jane Mitchell and Steve Reilly, my research associates at York, and by Richard Morris and Mick Rugg, and I am very grateful to all of them for their time and advice. Finally, Geoff Hall and Nick Mackintosh were both obliged to spend many hours discussing drafts with me, and made substantial contributions to the overall organization of the book: I am very glad to acknowledge my indebtedness to them both.

I should also like to express my gratitude to those friends who have borne the burden of typing for me: to Jenny Bailey, who typed almost all the final draft, and to June Edwards, Sue Medd, Jenny Debenham, and Jane Moor, who produced the references and indexes. My wife Kate typed all the early drafts, skilfully deflected children's demands on my time, and gave me constant support; my debt to her cannot be put into words.

York,   E. M. M.
August 1981

Acknowledgements

Contents

viii   Contents

1. The comparative approach to intelligence

Introduction

Throughout the history of Western thought there have been two contrasting interpretations of the behaviour of animals: some, like Descartes, have assumed that animals are qualitatively distinct from humans, that they are mere machines, devoid of reason or consciousness, having no mental life; others, Hume for example, have assumed a continuity between the mental processes of animals and man, arguing that the differences seen are essentially quantitative in nature. It is difficult to see what experimental observations might be relevant to the question of whether an animal does or does not enjoy consciousness, and the dispute between these two points of view was conducted with a minimum of reference to the facts of animal behaviour; this remained true until the late nineteenth century when, following the publication (in 1871) of Darwin's *The Descent of Man*, the impact upon the argument of the theory of evolution became clear. Just as the physical form of the body evolved, so must those capacities responsible for human mental life have evolved: the notion that there could be a sudden discontinuity between man and beast became much more difficult to sustain, and the search for evidence of continuity in mental life throughout evolution began. The fruits of that search constitute the field of enquiry now known as comparative psychology.

The early comparative psychologists saw the conscious mind as the proper object of psychological enquiry: direct access was possible to only one mind (one's own), so that the properties of other minds (human or infrahuman) must be inferred. What rules, then, were to be used to infer from behavioural observations to mental events? An early proposal (Romanes 1882) was that one criterion for the possession of mind is the ability to learn; in Romanes' view, the evolution of mind was reflected in the evolution of learning capacity – 'The lower down we go in the animal kingdom, the less capacity do we find for changing adjustive movements in correspondence with changed conditions' (Romanes 1882, p. 8). Romanes' work, however, met criticism on two grounds. First, from Lloyd Morgan who, while sharing Romanes' general views on the primarily introspective nature of psychology, saw the dangers inherent in excessive anthropomorphism. Romanes, for example, had written: 'if we observe an ant or a bee apparently exhibiting sympathy or rage, we must either conclude that some psychological state resembling that of sympathy or rage is present, or else refuse to think about the subject at all' (Romanes 1882, p. 9). Not unnaturally the proposal that ants and bees might feel sympathy made many comparative psychologists uneasy, and this disquiet was crystallized in Lloyd Morgan's celebrated canon: 'In no case may we interpret an action as the outcome of the exercise

of a higher psychical faculty, if it can be interpreted as the outcome of the exercise of one which stands lower in the psychological scale' (Morgan 1894, p. 53). Although this principle still retains its importance today, and is particularly relevant to the analysis of problem-solving in animals, the issue to which it was originally directed, the inference of types of conscious processes in animals, is no longer of central importance to psychologists; the question, for example, of whether an animal (or even another human being) feels, say, pleasure or pain, has reverted once more to the philosophers.

Psychologists abandoned their efforts to provide an analysis of states of consciousness largely as a result of the attacks on introspectionism in the early years of this century, by biologists (Loeb, for example) and psychologists – Watson in particular. The essence of Watson's position is that there is no dichotomy between behaviour and consciousness, that thoughts and images, for example, are segments of behaviour on a par with limb movements, and to be studied in the same objective way. Now since that time, few comparative psychologists have concerned themselves with attempts to infer the mental life of animals; whether this is because most psychologists accept Watson's position, and see it as a resolution of the 'mind–body' problem, or because, although rejecting his psychological stance, they nevertheless agree that observations are not relevant to questions of consciousness, need not concern us. What is pertinent here is that since the early days of this century, comparative psychology has been concerned with the behavioural capacities of animals – with what they can and cannot *do*. Theoretical inferences have been made, not to the conscious experience of the animals, but to the types of internal processing required to produce the observed behavioural output.

A second criticism directed at Romanes was that his behavioural evidence was largely anecdotal in character, and so generally unreliable. The case against the uncritical use of anecdotal evidence was made out by Thorndike, in a style that is hardly likely to be bettered: 'In the first place, most of the books [on comparative psychology] do not give us a psychology, but rather an *eulogy,* of animals. They have all been about animal *intelligence,* never about animal *stupidity.* . . . Human folk are as a matter of fact eager to find intelligence in animals. Dogs get lost hundreds of times and no one ever notices it or sends an account of it to a scientific magazine. But let one find his way from Brooklyn to Yonkers and the fact immediately becomes a circulating anecdote' (Thorndike 1898, pp. 3-4).

As a result of such arguments, there has since been fairly general agreement that the proper subject matter of comparative psychology is overt behaviour, and that the proper study of behaviour requires the use of objective methods of observation and experimentation. Agreement on these basic matters has not, however, resolved the question whether there is continuity in the processes controlling behaviour in animals and man. Thorndike's experiments led him to believe that the essence of all intelligent behaviour lay in the formation of associations, and that the differences in intellectual capacity between species reflected quantitative differences in the number, delicacy, complexity, permanence, and speed

of formation of associations. 'Amongst the minds of animals that of man leads, not as a demigod from another planet, but as a king from the same race' (Thorndike 1911, p. 294). Thorndike failed to obtain any evidence of the capacity to reason in monkeys, but argued that 'rational connections are, in their basic causation, like any others, the difference being in what is connected', so that 'the denial of reasoning need not mean, does not to my mind, any denial of continuity between animal and human mentality'. Maier and Schneirla, on the other hand, concluded from their analysis of experimental data that the ability to form associations and the ability to reason were distinct, and that the ability to reason required a well-developed cortex (found only in mammals).

One of our primary aims in assessing the evidence relevant to the intellectual capacities of various vertebrates will be to determine to what extent that evidence bears on the quantitative/qualitative or continuity/discontinuity controversy. The ultimate goal, however, is not simply to resolve that controversy, but rather to provide an explanatory account of intelligent behaviour in man and in animals; the goal, then, is to elucidate the nature of intelligence. The stance adopted by a particular psychologist on the continuity question generally reflects his view of the nature of intelligence. Most of those who believe that there are no major qualitative differences between vertebrate species argue that all intellectual acitvity proceeds according to the laws of associative learning and have, in this sense, a relatively unstructured view of intelligence: the only information-processing mechanisms involved in logical (as opposed to sensory, motor, or storage) processes are those required for the formation (and, perhaps, the dissolution) of associations. The opponents of continuity naturally tend to be somewhat more diverse in their views, but have in common the belief that intelligence is structured, that intelligent behaviour is the product of the interaction of a number of devices, each having distinct properties. As we shall see, these contrasting views of the nature of intelligence imply important differences over which types of experimental approach are likely to prove most fruitful.

Assessment of intelligence in animals

It is clearly important that, before attempting to compare the intellectual capacities of various species, we should have in mind a definition of intelligence. We shall not, however, attempt a formal definition here, and in case this should smack of evasion, some justification may be provided by citing the precedent set by Darwin, who wrote, in the chapter on Instinct in *On the Origin of Species:* 'I will not attempt any definition of instinct. It would be easy to show that several distinct mental actions are commonly embraced by this term; but every one understands what is meant, when it is said that instinct impels the cuckoo to migrate and to lay her eggs in other birds' nests' (Darwin 1866, p. 148). A similar view is taken here: we need not provide the definition for a word whose meaning is already well known. Moreover, to adopt any strict definition would run the risk of imposing some theoretical bias at the outset, in particular as regards the

question whether intelligence is some unitary capacity, or better seen as a complex of capacities, each of which might be quite independent of the others. It may be tempting to argue that intelligent behaviour must involve some complex behaviour, such as reasoning, and that, therefore, consideration of simpler types of behaviour is not relevant to the analysis of intelligence. The conclusion, however, supposes that reasoning involves mechanisms other than those involved in 'simple' learning, and this is an empirical matter, not one to be answered by appealing to the ways in which the word 'intelligence' is generally used. It may well be that 'complex' behaviour, such as reasoning, does involve mechanisms not engaged in 'simpler' behaviour—but this is an issue which can only be settled by evidence; to avoid ruling out such evidence, our notion of intelligence must at this stage be left open and flexible. In place, then, of a formal definition, it should be sufficient to indicate the sorts of tasks in which intelligence may play an important role: these include the entire range of learning tasks used by experimental psychologists, from habituation and elementary association formation to complex problem-solving, language analogues, and so on. Intelligence, in other words, is held to manifest itself in all those situations in which subjects are required to adapt to novel circumstances; a common feature of many such situations is that successful adaptation involves the detection of and appropriate response to regularities in the environment. This is the case as much with 'simple' association formation as with 'complex' concept formation: whether successful performance in all these situations involves the same mechanisms will remain, at this stage, an open question.

Our use of the term intelligence is, then, intended to correspond reasonably well to its common, non-specialist, use: it should be noted, however, that it stands contrasted with another use current in the branch of experimental psychology known as artificial intelligence. The goal of workers in this field is to produce a detailed account of the information-processing that underlies behaviour, however simple that behaviour might seem on the surface. One area of particular concern is that of visual perception, and the problems that arise in attempting to devise a machine that can see emphasize the extreme complexity of perception. However, we shall not apply the word 'intelligent' to processes solely on the grounds that they involve complex information-processing: the meaning used here implies a generality and flexibility of application that may well not be true of at least some of the complex devices within the central nervous system. The processes of perception appear to involve something akin to problem-solving, but whether such processes enjoy any use outside perception is unclear, and until such use is shown, perceptual capacities will be assumed to be independent of intelligence. A further example may serve to emphasize the notion of general application implied by the present definition of intelligence. There is currently no evidence from laboratory experiments that amphibians can master complex learning tasks—but frogs do make use of sun-compass orientation: that is, they orient towards 'home' by observing the position of the sun, by calculating, from the use of an internal clock, some measure of the difference between that position

and the position occupied by the sun at home at that time of day, and finally, by calculating, from that measure, the direction in which home lies. Is this evidence of otherwise undetected intellectual capacity in frogs? The answer, of course, lies in the matter of definition: if by intelligence we had meant simply the capacity to carry out complex information-processing, then, of course, sun-compass orientation would have implied intelligence. On the other hand, as we wish to restrict our use of intelligence to more general capacities, sun-compass orientation does not of itself demonstrate intelligence, although it may stimulate efforts to discover other situations in which the relevant processing mechanisms may be brought into play.

As the preceding paragraph indicates, the decision to concentrate on mechanisms of general intelligence will lead to the exclusion from consideration of a number of capacities, such as navigation, imprinting, and poison-aversion learning, which are demonstrated in restricted contexts. These are capacities which have been the subjects of intensive investigation by ethologists and psychologists, and it may be worth considering whether there are grounds for supposing that such specific capacities might not be relevant to performance in the more general contexts in which we have declared an interest.

One possibility, which has been raised for at least some of these specific capacities (see, for example, Hoffman and Ratner's 1973 account of imprinting), is that they are not in fact as specific as they might seem but instances of the operation of general learning mechanisms responsible for virtually all behavioural adaptations. Now if this account is valid for any of the 'specific' capacities then we shall, of course, be considering the mechanisms responsible for the behaviour concerned when we discuss performance in standard learning tasks; no exclusion will in fact occur, and there would seem little value in pursuing the properties of those mechanisms in what might be misleadingly narrow contexts.

A second, very different, possibility is that there are specific capacities which, although having evolved to deal with restricted contexts, are available for at least some more general applications; in this case, interpretation of performance in conventional tasks should take account of the contribution to that performance that may have been made by (relatively) context-specific mechanisms. In its most exaggerated form, this view claims that 'general intelligence' is a chimera; that what we see in all examples of adaptive behaviour is the operation of mechanisms which evolved to cope with relatively specific environmental demands, but which have become to some degree emancipated so that they can be brought to bear on a wide range of problems (see, for example, Rozin 1976). According to this notion, it could be that by not examining 'specialized' forms of learning, we exclude from consideration those tasks which would exhibit the optimal operation of the capacities whose performance we are in fact exploring in our more general tasks.

The principal difficulty facing evaluation of the notion that general intellectual capacity is a composite of contributions from diverse specialized capacities is that it is vague; the proposal does not show how a given specific capacity might

contribute to performance in some more general context, nor how to assess the existence or otherwise of any such contribution. Taken at face value, the suggestion would appear to be that any comparative study of intelligence should consider all the behavioural capacities of the species considered — patterns of social behaviour, of rearing the young, of feeding, aggression, territoriality, and so on — but this without any indication of how such discussions would be relevant to the central issue. Such a procedure is clearly unsatisfactory, and in any case impractical. In the absence of specific proposals concerning possible contributions from specialized capacities, we shall be unable to consider such capacities further. This is not, of course, to dismiss the theoretical possibility of such contributions, and there is one general implication of the hypothesis that should be borne in mind. If we assume, reasonably enough, that not all species possess the same complement of specialized mechanisms, and that such mechanisms do contribute to a wide range of tasks, then we might expect to find qualitative differences between species in performance of standard tasks. Any discovery of qualitative differences should, then, cause us to look at differences between the specialized capacities of the species concerned, in case they might account for the qualitative contrast observed. Should that procedure also be successful, we should have both some evidence that the overall hypothesis might be valid, and, perhaps, indications of how to devise more general techniques for evaluating the contribution of specialized mechanisms to performance in standardized contexts.

The burden of the above argument is, then, not that 'general intelligence' should be assumed to be in some sense unitary, rather than a collection of specific capacities, but rather that, since we have no means of inferring from (apparently) context-specific capacities to corresponding general capacity, we should examine (apparently) general capacities initially, to see whether that analysis suggests factors which could be derived from context-specific capacities.

It is appropriate to emphasize at this point that by intelligence is meant, not the intelligent behaviour of an organism, but its capacity or potential for such behaviour; we cannot therefore conclude from the fact that a particular species is rarely if ever seen to perform intelligently in its natural life that it is of 'low' intelligence — its natural habitat may simply not provide an appropriate situation for the demonstration of its intelligence. Although it may seem to us that animals that are well adapted to a relatively monotonous environment, and that survive essentially by use of species-specific responses to restricted classes of stimuli, have no need of intelligence, and so would not have evolved such a superfluous capacity, the conclusion would be entirely speculative: efforts to determine the limitations of the animal's capacities would have to be made in 'unnatural' situations, to see how far novel demands might be met.

Our interest, then, is to be in the behaviour of animals in situations which make demands upon the subject that allow it to demonstrate the behavioural flexibility that is implied by intelligence. How are such situations to be designed? There are, unfortunately, no rules that enable us to say *a priori* what tasks are required: their design depends basically on the ingenuity of the experimenter, on

his intuitions concerning the organization of intelligent behaviour, and on the compatibility between the apparatus and his subjects. The experimenter's intuitions are of importance in the following way. If the task he designs makes very few demands on intelligence, and is, therefore, 'easy', then animals of a wide range of intellectual capacity might perform at a similar level in it – there would be, that is, a 'ceiling effect'. If, on the other hand, the task is too 'difficult', then only a few exceptionally intelligent creatures might succeed in it, the remainder appearing equal owing to a floor effect. Ideally, a number of tasks of varying levels of difficulty seem to be required: the problem is, how can the difficulty of a task be assessed in advance – what is it that makes one task more or less difficult than another? Task difficulty is determined by the organization of the intelligence of the species in question: it depends, that is, entirely on how the relevant information is stored and processed. Given a valid account of such information-processing we could of course design tasks that would make varying demands upon intelligence – but then at least one of the objects of the exercise is to obtain such an account, and, at present, *a priori* assessment of task difficulty depends on hypotheses that are far from commanding universal assent. The design of the tasks, then, will be guided by an experimenter's 'hunches' concerning intelligence; their success as assessors of intelligence is, on the other hand, an entirely empirical matter, although it is far from easy to set out criteria for success. Many problems in the interpretation of intelligence tests for animals are best considered individually, in the context of a particular test, but there are some general difficulties which may appropriately be considered here.

Suppose that a task has been devised which does obtain different rates of learning for various species of animals – may we assume that the species are now ranked in intelligence? Clearly, we cannot. Obvious alternative explanations are that the task merely distinguishes between their sensory capacities (as, for example, in the acquisition of a visual pattern discrimination), or between their motor skills (as, for example, in learning to fit one object into another), or reflects differences in motivation or incentive – some species may find the reward offered more attractive than others. It seems obvious that such 'contextual variables' (Bitterman 1965a) would indeed generate species differences in many learning situations – how could we rule out the possibility that one of them was responsible for a given difference in any situation? One proposal for overcoming this problem is known as the method of 'systematic variation', and was advocated by Bitterman (1965a): this technique involves the use of varying levels of a contextual variable in a particular task. Suppose, for example, that pigeons have acquired a particular discrimination more rapidly than goldfish, and we are concerned to know whether this is in fact because the pigeons are more highly motivated. We should then run groups of pigeons and goldfish at different levels of motivation (by, for example, varying the duration of food deprivation) on the same discrimination: if neither group shows much variation in performance with variations in motivation, and if there is no sign of an overlap between the groups at any point, then it is not unreasonable to conclude that the observed difference

is not brought about by differences in motivation. At this point, of course, the method would have to be applied to other confounded contextual variables, which should in turn be varied. Moreover, variations in the level of one variable should be tested at various levels of other variables, in case there might be inter-actions between the effects of such variables; for example, variations in size of reward might have little effect at one level of deprivation, but a marked effect at another level. Even then, a conclusion as to absence of effects of contextual variables would be unsafe, as there is no finite catalogue of such variables. The method of systematic variation, taken at its face value, clearly makes heavy demands on both time and subjects, and cannot guarantee, however extended the series of variations, that there is not still some relevant but untested variable present in a situation.

This conclusion is not, however, as damning as might at first sight appear; the conclusion is, after all, only that, by using this technique, it is not possible to exclude all logically possible artefacts. On the other hand, the technique can be used to rule out any apparently plausible source of artefact—there is, indeed, no alternative, as one cannot 'equate' the sensations of members of different species. Once again, the implication is that we must wait upon results, and then con-sider what interpretations are plausible rather than all those that are logically possible.

As the method of systematic variation is so time-consuming, experimenters have attempted to devise tasks in which contextual variables may seem unlikely to play an important role in generating species differences. One popular technique has been to use situations in which the measure is not the absolute number of trials or errors in the acquisition of a task, but some relative measure that depends on a baseline score made by the species. For example, an animal might be taught to choose stimulus A rather than stimulus B in a discrimination task, the number of errors being recorded; the reward value of the two stimuli is now reversed, and the number of errors made in the course of learning to choose B rather than A is recorded, and expressed as a percentage of the errors made in acquisition. It might be supposed that such a relative score (the Reversal Index of Rajalakshmi and Jeeves 1965) would be less liable to contamination by contextual variables than would raw acquisition data, the grounds for this supposition being, of course, that the contextual variables, present presumably in both acquisition and reversal, should cancel themselves out. However, this is evidently not a watertight con-clusion: for example, reversal is a different task from acquisition—it may, there-fore, employ a mechanism that is not employed in acquisition, and that mechanism might be sensitive to contextual variables that are of little significance in acquisi-tion; alternatively, it might be simply that reversal is more difficult than acquisi-tion (or vice versa), and that the effects of some contextual variable vary with task difficulty. This technique, then, can give no guarantee of success in pro-viding an uncontaminated measure of intelligence, but this is not to deny its pos-sible utility: if a test using a 'relative' score provides results that appear to form a meaningful pattern, lending itself to plausible theoretical analysis, then it will

have been successful at least to the extent that further empirical analysis, employing, no doubt, the method of systematic variation, will be justified.

There is one additional important problem that requires comment. The situations used by experimental psychologists expose animals to environments very different from those that would be encountered in the wild; they require unnatural responses (key-pecking, lever-pressing) to unnatural stimuli (monochromatic light, electric shocks). Poor performance by an animal might reflect, not a low level of intelligence, but rather that the situation is not one that allows expression of the adaptive capacities available in more natural environments. Experimental psychologists have become increasingly aware of this problem, and are very conscious of the need for compatibility between the behavioural demands of their situations and the species-specific behaviour patterns of their subjects. For example, it is extremely difficult to train pigeons to peck a lit key to escape or avoid trains of electric shocks (Hineline and Rachlin 1969); but this probably reflects the fact that key-pecking is, in pigeons, a most unlikely response to either fear or pain. It is not, in fact, difficult to train pigeons to avoid shocks by moving from one compartment to another in a shuttle-box (Macphail 1968) or by pressing a treadle with their feet (Foree and LoLordo 1970).

This problem may be exaggerated: the situations preferred by psychologists have become popular precisely because the species with which they are used perform efficiently in them, and do not require extensive (and tedious) training. Provided that animals are not grossly maladapted to the test situation, it does not seem that the objection has much force, except perhaps as a caveat to be borne in mind where subjects fail to perform up to expectation. Moreover, there is no alternative, unless the comparative analysis of intelligence is to be abandoned: it has already been argued that passive observations of animals in the wild do not necessarily reveal the full extent of their capacities — unnatural demands, therefore, must be made. Provided that the situations used are carefully designed, and that a variety of situations are used, there is no reason why invaluable information should not be gained in the laboratory.

Classification of learning tasks

In a number of the chapters that follow, experimental results are considered under one of four headings — habituation, classical conditioning, instrumental learning, and complex learning. It should be emphasized here that these labels refer to the nature of the procedures used by the experimenter; the allocation of results to separate headings is not intended to imply that the different procedures tap distinct learning mechanisms — that, in other words, they are different 'types of learning'. They may, of course, be different types of learning, and indeed one of our major interests will be to see whether some species succeed in giving evidence of learning in one procedure, but not in another — a finding which would give support to the view that different mechanisms are involved in the various procedures.

Habituation

The term habituation refers to the process whereby the magnitude of an unconditioned (unlearned) response (e.g. a change in heart-rate) to an originally novel stimulus declines with repeated presentations of the stimulus. One way in which habituation may differ critically from (and be 'simpler' than) the other procedures is in requiring no (or at least no obvious) association-formation: the subject has, as it were, only to register that a particular stimulus has occurred previously in order to perform adequately. It makes sense, then, to begin the sections on learning by considering habituation first, as it has some claim to be a 'simple' form of learning. Our interest will be not only in whether habituation occurs in given sets of species, but also in whether there is evidence for comparability of mechanisms underlying the phenomenon in the various species in which it is observed.

Classical and instrumental conditioning

Our major concern in the sections on these two types of conditioning will be to determine whether each type is obtained in the class of animals under review. In recent years, the distinction between classical and instrumental learning has attracted much theoretical interest, and we shall consider at this point some of the problems which apply generally to the distinction, independent of its investigation in any particular species.

In classical, or Pavlovian conditioning, the occurrence of a reinforcer (an unconditioned stimulus, or UCS, normally of biological significance, that reliably elicits a particular response, the unconditioned response or UCR) is contingent upon a prior occurrence of some arbitrary stimulus (the conditioned stimulus or CS); the CS, then, acts as a signal that the UCS is imminent, and nothing that the animal does can affect the occurrence of the UCS. That the animal has detected the contingency between CS and UCS is shown by the occurrence of some conditioned response (CR) to the presentation of the CS alone; the CR generally, but not invariably, closely resembles the UCR, and appropriate controls are, of course, required to show that the CR is not obtained in the absence of the CS–UCS contingency.

In instrumental conditioning, the occurrence of reinforcers is contingent upon the production of some specified response by the animal; the detection of that contingency is demonstrated by a change in the spontaneous rate of occurrence of that response (by an increase, where the reinforcer is a reward, or by a decrease, where punishment is used). As with classical conditioning, controls are necessary to show that the rate changes obtained are indeed due to the contingency imposed. Successful acquisition of instrumental conditioning is prima facie evidence of the capacity to form associations involving a response as one of the terms in the association.

The major complexity that arises from these two superficially simple procedures is the possibility that learning in a classical procedure might be mediated by a covert instrumental contingency, and vice versa for an apparently instrumental procedure. Consider a typical classical procedure using an appetitive re-