Chapter 2

Reason and Consciousness

Reason is a tricky word, with several different meanings - creating the apparent paradox that only an entity possessed of reason can be accused of being unreasonable. I define it as the ability to abstract general ideas from experience, to deduce relationships, and to create new ones. Reason sniffs out similarities between apparently dissimilar phenomena and reveals differences between apparently similar ones, all without reward other than the pleasure of discovery. It also identifies chains of causation. It can then use the knowledge creatively.

This definition is designed to include reasoning processes and results which are purely concrete and practical. The idea of using a stone to crack a nut, for instance, entails bringing together a knowledge of the nature of stones with a knowledge of the nature of nuts, and appreciating the possibilities of a sufficiently sudden change in their spatial relationship. Modifying an object to create a more effective tool is a more advanced example of the same thing. The chimpanzee who first peeled a twig and used it to fish in an ant or termite nest had realised both that the relationship of the twig to its bark was something that could easily be altered, and that the resulting probe would fit into a hole in a termite mound. Whether the first successful termite extraction was the result of conscious planning or serendipitous accident we can only speculate. But if it was an accident it encountered a mind that could appreciate the possibilities of the twig for extracting food from inaccessible places, and the concept was passed on to other minds that were equally capable of making use of it.

At the other end of the scale the most advanced form of reason deals in abstractions, in causes and effects, forces and functions and influences, in the consistencies underlying perceived events. It discovers relationships that are not directly observable by the senses, but must be deduced. This is essentially an operation of the same sort as the analysis and planning that goes into the manufacture of tools. It differs in working with generalised ideas rather than with a knowledge of concrete objects, and in being dependent on language, but it still involves bringing together disparate concepts. It's logical, therefore, to suppose that the capacity for abstract reasoning has evolved through an elaboration of the apparatus which is responsible for concrete reasoning.

A lot of the mental activity that goes into the production of a scientific theory, or a new tool, or a work of art, is certainly done at unconscious levels. However, the final products clearly reach the consciousness of their progenitors, and can be transferred to other conscious minds. So if creativity and rational thought are to be explained sensory consciousness must be accounted for first. Indeed explaining consciousness is the major part of the challenge.

For a long time consciousness was pretty much a taboo subject in scientific circles. There was no way of studying it directly and objectively in other people, or of measuring it, and introspection was entirely out of fashion. All that has changed in recent years, however. It's clear now that consciousness performs several useful functions, so the evolutionary advantages it provides can be identified. There's also a good deal of evidence as to how sensory consciousness is produced.

On the other hand it must be admitted that while in one sense we can understand the mechanics by which an assemblage of nerve-cells produces this effect we call consciousness we may never comprehend just why it should do so. It may be something we have to take as given, in the same way we accept the existence of gravity, or the fact that electrons have a negative charge. As the philosopher David Chalmers put it, this is the hard problem - *explaining how any physical system, no matter how complex and well-organised, [could] give rise to experience at all.* On current evidence I have to agree with another philosopher, Colin McGinn, who suggested that it's a problem we're not equipped to solve.

Consciousness is another ambiguous word, of course. It's something which can be suspended by sleep or anaesthesia or concussion, and there's the complication that many things are potentially available to consciousness at any moment but only a selection of them actually occupies it. Before consciousness can be lost, however, or used, it must exist. All the evidence suggests that it's something which has evolved, and evolved gradually. And we need to identify several levels.

The most obvious is sensory consciousness - awareness of the external world through vision, audition and olfaction, and of internal sensations of hunger and satiation, pain, pleasure, warmth, and so on. This sort of consciousness is pretty certainly shared with a good many other species (but by no means all) though the precise range of senses included and the degree of definition within each sense must vary considerably. Life for a creature with only this sort of consciousness can be conceived of as a sequence of sensations and emotions, a wholly egocentric sort of experience.

The next level entails an awareness that these sensations and emotions are changeable, and reflect current circumstances. It implies a new sort of memory, so that consciousness is no longer a matter of living exclusively in the moment. And it introduces the possibility of conceiving of other individuals of the same species as beings much like the self, experiencing similar sensations and activations. We might call it *consciousness of self and others.* It's something that, at the level of the individual, develops with time. It depends on maturation processes in the brain as well as on accumulating experience. As Ulric Neisser put it: *babies know what they want but they do not know who wants it*.

The major evidence for the existence of this sense of self in other species is that it makes social interactions subtler. For instance, it enables a mother to manage her infants with a certain amount of insight, recognising their limited knowledge and abilities. This sort of behaviour can most notably be observed among great apes, elephants and dolphins. As in humans, however, there is considerable variation in the degree of insight shown.

The ability to guess at the nature of another individual's perceptions and motivations also allows deceits to be perpetrated. Jane Goodall, who pioneered the study of wild chimpanzees, recounted how a bright young male in the troop she observed noticed a banana that was hidden from all the other chimpanzees present. Had he attempted to collect and eat it straight away he would pretty certainly have lost it to an older, dominant male. She received the impression that he carefully avoided looking any further in that direction until, the other bananas supplied to the troop being finished, he managed to initiate a general move away from the area. Then he sneaked back on his own and claimed the prize. Since then a good many clear examples of deceit have been recorded among great apes, along with quite a few suggestive observations in various monkey species. The subject is extensively explored in Byrne and Whiten's book *Machiavellian Intelligence.*

Deceit involves manipulating what another individual perceives, and rests on the ability to imagine another's viewpoint. The talent seems to coincide with the ability to imagine that objects might usefully be somewhat adapted, that current relationships might be altered, that roles other than the usual one might be played. In other words, it cohabits with the capacity for practical reasoning which can lead - in a species with appropriate manipulative ability - to the fashioning of simple tools. This sort of consciousness also begins to make language possible, since it is necessary to have a concept of another mind capable of receiving a message before there is any point in purposefully trying to frame one.

A third level of consciousness is consciousness of the self not just as a feeling, perceiving being, but as a thinking being. That means there is not only insight into the perceptions, emotions and motives of others, there is also the ability to take a step back from the action and reflect on these insights. Thoughts become detachable from the thinker, capable of taking the stage in their own right rather than just providing inspiration for action. The evolution of this mode of consciousness must have been closely interwoven with the development of language, which of course greatly facilitates such reflections. Furthermore, once thoughts can be contemplated as entities divorced from action they seem to cry out for expression and discussion. No doubt this was what motivated early humans to elaborate basic forms of language into structures governed by grammatical conventions.

Language in turn opened up the way for the evolution of a capacity for abstract thought, since abstract ideas need handles by which to grasp them. The thought must emerge, of course, before a word can be found for it, but it is difficult to get very far with an abstraction, let alone relate it to other abstractions, until there is a convenient means of storing it in memory and retrieving it again. Attaching it to the act of pronunciation and the sensory shape which is a word achieves this. Furthermore, a capacity for abstract thinking is generally only going to provide an evolutionary advantage if the results can be communicated to other members of the species. It's only likely to become established, therefore, in a species which already possesses a certain amount of language. The available evidence suggests that this sort of consciousness is pretty much confined to humans.

I shall argue that there is also a fourth, primitive sort of consciousness, merely a vague sense of current need and current purpose, which belongs at the beginning of this evolutionary progression - but that must wait for now.

What I have termed sensory consciousness appears to correspond to what Gerald Edelman and Giulio Tononi call primary consciousness and what Antonio Damasio calls core consciousness. Edelman and Tononi propose a 'higher order' consciousness which is *accompanied by a sense of self and an ability to explicitly construct past and future scenes.* Damasio writes of *extended consciousness* which has many levels and grades. It can create an elaborate sense of self, and provide awareness of a lived past and an anticipated future. The philosopher Daniel Dennett has similarly distinguished different levels of consciousness. All agree that language plays an important role in the most advanced form.

The second and third levels of consciousness must surely have evolved by elaborations of the mechanisms which produce conscious sensation. So before the question of abstract reason can be tackled it's necessary to consider what conscious sensation is, and how it might have come about. That's the main subject of this book.

We have this marvellous experience of a world out there, full of colours and shapes, textures and distances, rich in sounds, and enlivened with assorted odours. It's rather difficult to talk about, because we take it for granted that what we experience *is* the external world, in objective reality. Our sensory systems mesh together so well that the evidence offered by one is nearly always confirmed by the witness of another channel. Stretch your hand out towards what looks like a tangible object and the predicted tactile sensation occurs, nibble on something that looks and feels like an apple and it turns out to taste like one too, and provides the expected sense of nourishment. Our conscious sensory experience generally offers a very efficient guide to our environment, and enables us to fit our actions to it very effectively.

However, with increasing knowledge of nervous systems it has become apparent that the equation of sensory experience with reality isn't strictly correct. Our experience is not reality itself, but a model created in the brain out of the data which our sensory systems happen to be equipped to collect. (Other species collect somewhat different selections of data.) The greenness of leaves and grass is not integral to them, it's an effect created in the brain when light receptors in the eye register the wavelength of light that leaves and grass reflect. The sounds we hear exist only in brains. They are the transformation created by a marvellous neuronal mechanism when pressure waves travelling through the air wash up against the ear-drum. In short, conscious sensation models reality with what might be described as a set of coding systems. Similarly our pains and aches, feelings of comfort, sense of balance exist to provide a model of the current state of the self that will prompt appropriate behaviour. In 1973 Harry Jerison wrote: *A ….. modern way of thinking about the perceptual world of which either Man or other animals are conscious is as a construction of the nervous system designed to explain the sensory and motor information processed by the brain.*

Things that can be done without consciousness

All the evidence suggests that conscious experience is something that is created by a very substantial assemblage of nerve-cells (properly known as neurons). A great many species pretty certainly get along without it. In the simplest cases a sensory neuron is directly connected to the cell which controls a muscle, and activation of the sensory cell triggers a message which causes the muscle to contract. In the surf-clam, for instance, a light-sensitive neuron running around the edge of the body is illuminated when the shell is open and the animal is feeding. It connects to the muscle controlling the hinge of the shell, and if the amount of light reaching it is suddenly reduced - something which may indicate a predator passing overhead - the shell is quickly closed.

Behaviours of this sort are obviously fore-ordained by the architecture of the neuronal connections, and therefore determined by the genetic design of the species. We have a few such reactions ourselves. Sneezes and coughs are good examples. Another is the way one's leg jerks up in the air when the knee is tapped on a certain spot. That reflex is produced by a short connection from sensory cell to muscle which permits a rapid readjustment to muscle tensions when an unsteady footing threatens to throw us off balance. We notice these things happening, and sometimes we can make a conscious effort to suppress the action, but consciousness plays no part in its actual production. Moreover it's difficult to fake a reflex movement convincingly.

In many species the response to quite complex stimulus patterns, involving a combination of several features, may similarly be arranged through hardwired circuitry. It was the ethologist Niko Tinbergen who established the existence of these innate behaviours, experimenting with cardboard shapes and other models to discover just what it took to elicit a particular action. The evidence that the relationship between stimulus and response is genetically determined by the neuronal wiring is the fact that if the stimulus pattern contains the necessary elements it produces its effect, even though the normal context is missing. Infant herring gulls, for instance, gape hungrily when any sort of beak-shaped stimulus appears from above, and react most energetically to a shape that resembles an adult herring gull's beak - long, thin and yellow, with a red spot at the top. It doesn't matter if the rest of the gull is missing. All altricial birds (the sort that are helpless when they hatch and are fed by their parents in infancy) have some variation of this response, ensuring that they open their beaks to receive food when a parent turns up bearing it. No doubt they all react most strongly to the sort of beak proper to their species, but will also respond to a less than ideal stimulus, which means they don't starve if the parent's beak doesn't quite conform to the standard. (Once, hearing a squawk from under a bush while I was watering the garden, I found a newly fledged blackbird which, when I offered my little finger, gulped at that.)

In the breeding season a male stickleback builds a tube-shaped nest at the bottom of his patch of stream and tries to attract any pregnant female who passes to lay her eggs in it, where he will fertilise them, and then care for them. The sight of a rounded, silvery shape that approximates to the belly of a pregnant stickleback is sufficient to set off his courting display. He also defends his territory against rival males. Tinbergen showed that the male goes into its threat display - holding itself vertical with its tail uppermost and quivering violently - at the sight of any patch of red. In the streams which sticklebacks inhabit a patch of red is likely to be the breast of another male stickleback. In other surroundings strange things can happen. A fish that Tinbergen was keeping in an aquarium by a window went into a frenzied display at the sight of a passing Post Office van.

Often the stimuli which evoke these hardwired or reflex responses combine shape with movement. Day-old chickens crouch to the ground and freeze when a dark, roughly bird-like shape passes overhead if it has a short protrusion at the front and a longer one at the back. They don't do it if the long protrusion is at the front. It seems they are innately designed to try and hide from raptors, and not to waste time hiding from ducks and geese. Tinbergen identified many such hardwired responses, and many more have since been discovered by other ethologists.

Varied Types of Learning

In a great many species, however, many behaviours are obviously learnt. The detailed study of learning famously started with Pavlov, who was investigating digestive processes in dogs and measuring the rate at which they produced saliva when he noticed that his subjects salivated not only when their food actually appeared but also at the sight of the white-coated laboratory assistant who usually brought it. He then established that they reacted in this way to any sensory event which regularly occurred just before the appearance of food. Salivation is a genetically determined reaction to food, so he called it an unconditioned response, and the food an unconditioned stimulus. The bells, buzzers and other signals which the animals learnt to associate with food were conditioned stimuli, and the salivation which occurred in response to the conditioned stimulus was a conditioned response. To put it another way, a genetically programmed response comes to be evoked not only by the genetically ordained stimulus, but also by a stimulus pattern which has regularly been associated with it, and is likely to signal the imminence of the important event.

The ability to respond to a signal that indicates the probable nearness of something desirable is obviously useful. The ability to react to an advance warning of possible danger is even more valuable. This sort of learning is usually termed Pavlovian conditioning.

The discovery inspired a great flood of experiments along similar lines. In addition to those where the desirable or aversive event simply followed the cue there were many in which an action had to be performed to obtain the desirable. Numerous pigeons learnt to peck at a key to obtain a pellet of food or a drop of water, and rats learnt to press bars for similar wages. This is known as operant conditioning, and the reward which encourages the action is termed a reinforcement. Animals also learnt to perform an action in order to avoid an undesirable experience - usually a mild electric shock, which is a highly unnatural sensory event but has the double advantage of not doing the animal any perceptible damage, and of being easy to produce repeatedly at a precisely determined strength.

Learning to perform a piece of behaviour so as to avoid a noxious experience means learning to recognise a signal which indicates that it is imminent - learning, for instance, that when a light shows it's necessary to press a bar to avoid receiving a shock. Passive avoidance can work on the same principle - the bar can safely be pressed when the light isn't showing. In yet another sort of experiment the animal must learn that a reward for pressing the bar is available only when a light is on. In these cases there are in effect two stimuli to provoke the response - the bar which must be pressed and the light which signals the timing.

These experiments led on to more complicated ones. A pigeon might learn that there would be a food pellet if it pecked at the key when a certain shape appeared in a little window next to the key, but nothing if it pecked when another shape was on show. Or perhaps there would be a reward when a light to the right of the key was on, but not when a light showed to the left. Or there might be a punishment for pecking at the 'wrong' time as well as reinforcement for the 'right' response. This is discrimination learning. As well as its interest in the context of learning it's a protocol that has proved very useful for discovering just what sensory information different species are equipped to collect.

As the evidence accumulated it became apparent that much discrimination learning can be explained in terms of pre-programmed behaviour systems which are capable of a certain amount of modification. A certain sort of stimulus prompts a certain response, but the specification of the stimulus isn't very precise. The reinforcement serves to indicate more exactly the sort of stimulus that is worth a response, thus enabling the individual animal to adapt its inbuilt behaviour patterns to the potentials of the environment into which it happens to be born. Harry Harlow speculated along these lines in the fifties, suggesting that perhaps *….conditioning does not produce new stimulus-response connections but …. operates instead to restrict, specify, and channelise stimulus-response potentialities already possessed by the organism.*

Newly hatched chicks, for example, peck at any very small, three-dimensional object below eye-level, a behaviour which points them towards picking up seed grains, but which may equally lead them to consume grains of sand. When an experimenter confronted some chicks with a pile of grain and a pile of sand they ate both in about equal amounts, and it took them some time to learn to distinguish food from non-food. After an hour with seeds and no sand, however, a distinct preference for seeds was established. But chicks left for an hour with sand and no seed didn't develop a taste for sand unless they were given an injection of high-calorie food direct to the stomach afterwards.

Under normal circumstances the chick's learning is guided by the mother hen, who may pick up seeds and drop them again in front of her brood, but the learning still takes time. Seeds come in many different shapes and colours, and there must be a good many different kinds which make acceptable food for chickens, so the advantages of the system are clear - the chick is born with a rough idea of what to peck at, but can adapt itself to whatever sorts of seed happen to be available in the locality.

The behaviour of pigeons in the standard experimental situation suggests an elaboration of a similar sort of basic pattern. The key in the experimental box is the only small round shape in evidence, and the nearest thing to a food-stimulus available. Indeed, it's the only conspicuous stimulus in sight, and of course the pigeon has been kept hungry in preparation for the experiment and is ready to respond to anything remotely appropriate. Actually acquiring food requires at least two pecks, one at the key and one at the food when it appears, but food and key are close together, so this is not a very great distortion of the normal way of things. And close observation showed that the pigeons treat the key as they treat the variety of reward for which they are working. When the bird is accustomed to receiving a food reinforcement it opens its bill just before contacting the key, and snaps it shut again as it raises its head. When it is used to working for a water reward its beak is almost closed as it hits the key, and remains there longer while the bird makes swallowing movements.

Moreover pigeons are very talented at learning to make all sorts of discriminations when pecking at a key to obtain food or drink, but they are pretty hopeless at learning to refrain from pecking at a key when they are only reinforced for abstention. In other words they find it very hard to suppress the normal food-getting action as a way of obtaining food.

In fact it has emerged that a species may be equipped with a large capacity for one sort of learning, and little or no talent for another sort. Thus pigeons readily learn to peck a key to obtain food or water, but not to avoid a shock - their inbuilt response to danger is to fly away. Similarly, a rat's natural reaction is to freeze when danger threatens, and rats are bad at learning to press a bar to avoid a shock (though they are slightly better than pigeons at adapting to such demands). Chaffinches will learn to press a key for food, or to perch in a particular place to hear some recorded chaffinch song, but it's very difficult to teach them to perch in a particular place for a food reinforcement, and impossible to persuade them to peck a key for song. A dog learns to take food when one auditory cue sounds, and to avoid it when it hears a different sort of sound, but doesn't, in this situation, make use of sound cues distinguished by coming from different directions. (Excellent guides to conditioned learning are to be found in the works of M. E. Bitterman, Robert Bolles, David Oakley and Paul Silverman.)

Obviously an essential requirement for operant conditioning and discrimination learning is a feedbeck mechanism which allows the results of an action to be related to the stimulus which evoked the action. These experiments in animal learning suggest that feedback channels tend to be specialised. A certain sort of stimulus pattern elicits a certain action, and feedback is provided only about the type of result that is likely to follow such an action. Similarly, only the sort of stimulus that is likely to be relevant will function as a cue.

Discrimination learning is thus often a matter of honing an innately determined (or unconditioned) stimulus to a sharper definition. Or perhaps the genetic recipe provides the specification for an ideal stimulus, but in the absence of the ideal an animal will respond to the best approximation available, and thus learn what makes an acceptable substitute and what doesn't.

Many species are not only pretty efficient at learning how to obtain rewards but also play the odds intelligently when rewards arrive only intermittently. They behave, in fact, with the logic that a human might show in similar circumstances. Those reinforced on every trial learn fast, but give up responding fairly rapidly if the reward ceases to arrive. Animals reinforced on a fixed ratio schedule (for every tenth action, say), or at fixed intervals, learn more slowly but go on trying longer after reinforcement ceases. Animals whose reward arrives on a totally irregular basis learn slowly but then keep working hard all the time, and persist for even longer with no reward.

Early experimenters tended to assume that the neocortex - that thick, rumpled blanket of grey matter which forms so conspicuous a component of the mammalian brain and is especially large in *Homo sapiens* - functioned as a store for all the learnt connections that constitute conditioned learning. But it turned out that Pavlovian conditioning is possible in some of the lowliest invertebrates. Furthermore rats whose neocortex was removed could still learn very effectively, and in some tasks progressed faster than normal animals. Such rats may also retain the ability to perform responses they learnt with a whole brain. This is not to say that the absence of the neocortex doesn't restrict what may be learnt; it does, very considerably. But the limitations seem to be attributable not so much to deficiencies in basic learning ability as to the fact that the most sophisticated parts of the sensory systems are lost, and lower parts of the brain have less capacity for analysing the sensory input so as to make the more difficult discriminations.

Beyond Behaviourism

The first scientists to conduct this sort of experiment were known as Behaviourists. Some of them seemed to think that all learning could be explained as conditioning produced by simple reinforcements (though the theory, paradoxically, doesn't allow for the existence of theories, and could logically only be tenable if they were persuaded to it by bribery). The improbable idea was laid to rest by Edward Tolman, using a variation on a standard experiment. Rats put in a maze soon learn to find their way to food or water, and the assumption was that it was the prospect of the reward which motivated the animal to learn the most efficient path through the maze. Tolman set one group of rats to learn a fairly complicated maze in the conventional way, while another group ran in the same maze without finding any food. After ten sessions a prize was provided for the previously unrewarded rats as well, and they immediately began to perform slightly better than those which had been rewarded all along. From the twelfth day onwards - as soon as they had had the chance to establish where the food was - they logged slightly fewer errors on the way to the goal. In a further experiment the most direct route to the food was blocked off, once it had been learnt; once again the animals which had initially explored the maze without reward found the best new route to the food more quickly than those which had always had food to aim for.

Thus it was demonstrated that rats study and learn practical geography for its own sake. Not only do they do it in the absence of reinforcement, but they do it more efficiently that way. It's as if the presence of food tends to distract atten-tion from the business of exploration. (An introspective human may not find this surprising - we too tend to learn more of the geography of a strange city while wandering around it for pleasure than when hungrily searching for a restaurant.)

What a rat acquired when it learnt how to find its way through the maze, said Tolman, was a cognitive map - not just a memory of how to get from A to B but an idea of landmarks and spatial relationships. The hypothesis was tested with a simple, cross-shaped maze. A rat would be put in at one end of the cross, and food would consistently be provided at the end of the arm to the rat's right. When the rat had mastered this situation and consistently turned right the procedure was changed and it was put into the maze at the opposite end. For half the rats in the experiment the food was still in the same place, which meant they now had to turn left, and for the other half the reward was now in the opposite arm, so that they still had to turn right. If maze learning was a matter of conditioned responses the latter group should perform more efficiently in the new circumstances. If there was a cognitive map then the rats whose reward was still in the same place - identifiable by features of the laboratory, visible beyond the walls of the maze - should do better. It turned out that the group whose food was still in the same place got to it most rapidly, supporting the cognitive map theory.

The benefits to be gained from a cognitive map are obviously substantial, and the energy expended on exploration pays good dividends. The animal which has learnt its way around its environment knows the quickest way to the nearest shelter if danger threatens, where to find water or a safe place to sleep, the easiest route to a likely place for food. But what prompts it to explore in the first place, in the absence of any immediate concrete reward? Building an internal model of the environment must constitute a different sort of learning from the conditioned learning which merely serves to modify the application of instinctive behaviours.

This is not to say that conditioned learning and model-building are opposed. On the contrary, they work together. Successful models incorporate the fruits of conditioned learning. It may be conditioned learning that links food to the place where it's frequently found, and causes the site to be marked on the map.

Learning and sensory abilities

The evolution of learning abilities must have been closely intertwined with the evolution of sensory abilities. A species whose behaviour is entirely innate, composed exclusively of genetically determined responses to genetically determined stimuli, needs only such sensory apparatus as is required to register those stimuli. Any additional sensory equipment would be superfluous, indeed it might make the animal function less efficiently. But an organism capable of Pavlovian conditioning must of course be capable of registering the cue, or 'conditioned' stimulus as well as the unconditioned or innate one. Learning to make discriminations entails collecting enough sensory information to detect some difference between the relevant stimulus patterns. And where there is a capacity for building a cognitive map almost any new type of sensory receptor or new way of extracting information from sensory input may prove to be useful.

We humans have evolved the means of collecting enormous amounts of sensory data. How do we manage to make sense of it all, translating quite varied patterns of sensory input into consistent and meaningful percepts? For instance, we learn to recognise an animal as an ongoing entity, regardless of its current pose or the angle from which we happen to see it. We recognise a tune, regardless of whether it's played on the piccolo or the double bass, or performed by a double orchestra and massed choirs. We don't need to hear it in the familiar register; nor are we fazed by unusually fast or slow performances, unless the change is so extreme as to alter the whole mood of the piece.

All the evidence indicates that what we record is essentially a pattern of relationships. A tune is a pattern of tonal and rhythmic relationships between notes. A cow is a shape made up of components in spatial relationships some of which change, within certain parameters, and some of which don't. The distance between shoulder and knee or knee and hoof is always pretty much the same, but that between shoulder and hoof or nose and shoulder may alter, within certain limits. Thus we recognise a cow with equal ease whether it is standing up, walking, or lying down, and from almost any angle. We also recognise a distortion, exceeding the parameters of normal spatial relationships, which indicates that the animal is damaged, or dead. In addition, we come to recognise the way the spatial relationships alter as the cow moves. And we store the information that the cow is bigger than a rabbit, smaller than an elephant.

The bundle of sensory information that enables us to recognise the same essential pattern of relationships in many different manifestations can be termed a *schema* (borrowing a usage coined by Frederick Bartlett and adopted by Jean Piaget). The schemata created in different sensory channels become organised, of course, into multisensory schemata. We remember that cows moo, and give milk, that their short hair has a certain sort of feel, and we learn that it's important to distinguish a cow from a bull. Recording all this information must require a good deal of space, and we can deduce that this is what the complex circuitry of the neocortex is needed for.

Relationships, spatial ones, are of course the very essence of maps. Making a cognitive map means learning to recognise those features of the landscape which are conspicuous enough to function as landmarks, and establishing how they relate to each other. The positions of less conspicuous but more vital things - food sources, water, shelter - can then be related to the pattern of landmarks.

Three things make the creation of the map a particularly challenging business. One is that the landmarks currently visible must be related to those no longer within sight. The second is that the map must be capable of being rotated - the landmarks and the relationships must be recognisable from any angle, and be usable as guides to various goals. The third is that an animal needs to keep track of its own frequently changing position on the map.

The benefits of a cognitive map are huge. Its possessor is not limited to reacting to whatever is within reach of its sensory systems, or else wandering about at random until it encounters something that prompts a response. It can set off for some distant place where particularly tasty food is often to be found, or for a remembered spot of comfortable shade, or a secure place to sleep. In other words it can apply the knowledge gained through past experience to take itself to areas beyond the current range of its senses.

There are other means of directing travel to useful goals that are out of sensory reach. Several species of insect, such as bees and ants, have been shown to be pretty good at getting around in an organised fashion. But the general consensus is that they do it by a means more like conditioned learning. They learn to head towards a landmark, and then to change course towards the next landmark, and they can record how far they have travelled from home. But they don't appear to recognise an asymmetrical landmark from different sides, or to have an overall picture of the landscape which would allow them to take shortcuts.

This navigational system is more extravagant in terms of the energy spent on locomotion, but must require much less brain space than a cognitive map, and much less time spent on learning. In the species in which it's been studied it's augmented by the ability to register the polarisation of light, so that even on a cloudy day the sun provides an aid to orientation. Ants also leave a trail of chemical marker which they can follow to return to the nest, and to guide future excursions. The zigzag flight of bees and wasps, meanwhile, probably gives them a wider view of a landmark than a more direct approach would, as well as providing motion parallax, so that the landmark can stand out from its background.

A hypothesis

Conditioned learning, as defined above, requires a relatively simple form of memory. When innate responses are modified by the effect of reinforcements all that's needed is to be able to suppress the response to the unprofitable variants. For building a cognitive map a good deal more is necessary. Features currently in view must be related to features no longer present, which means remembering where one has just come from, and what lies behind one's back. There must also be some record of how much time and effort has been expended on getting from A to B to C. These are the threads that connect the place where one currently is to the place one has visited before and to which one now thinks of going. And most importantly, to take advantage of a cognitive map it's necessary to be able to exercise recall - to refer to areas on the map other than those within sensory reach. Memory must become a store out of which records can be retrieved, not just matched up against current input.

Something else is implied too. An animal whose lifestyle is dictated by hardwired stimulus-response reflexes doesn't get to make choices. Its repertory of responses is accompanied by a hardwired means of deciding which takes prior-ity when more than one potential stimulus presents itself. When two stimuli coincide the response will be made to the more urgent one (escaping being eaten, for instance), or if that criterion is not applicable, to the one that is most urgent in the context of the animal's current needs. If the animal is confronted with two stimuli of the same sort it will respond to whichever is exciting its sensory receptors most strongly. Sometimes this will be the nearest one, and if it's a food stimulus the response will result in its disappearance, and the animal will be free to go and deal with the next nearest. Where innately determined stimulus-response patterns can be modified by conditioned learning the methods of deciding between stimulus patterns must be much the same, except that the weighting given to the various possibilities is influenced by past experience.

For the species with a cognitive map the same priorities must be applied, but new possibilities open up. Using its capacity for recall an animal can make a prediction as to where it will find a satisfactory supply of food, and on that foundation make a plan. If memory provides the prospect of really tasty food on the other side of the hill it may be worth passing a less exciting meal on the way, suppressing the response to a present stimulus in favour of saving up the appetite for something better. The means of making predictions and plans leads on to the possibility of making choices. As more and more sensory information was collected, the power to make choices will have become increasingly significant.

Our hardwired responses are produced without conscious choice, and in many species a good deal of behaviour clearly depends on genetically determined neuronal wiring and can be assumed not to involve consciousness. Pavlovian conditioning does not necessarily require consciousness either, for humans can be conditioned to produce certain responses without ever noticing the conditioned signal, the unconditioned stimulus or the response. For many species this is probably also true of operant conditioning. Discrimination learning poses a more difficult question, but the best guess is that it depends on just how complex the stimulus patterns to be distinguished are.

However, there are several reasons for supposing that some form of consciousness is necessary for making and using a cognitive map. Firstly, how can the current location be related to a place that is now out of sight? The answer I find easiest to conceive is that the scene currently occupied can be related to those recently passed through because one is a conscious perception and the others are conscious recollections. It's difficult to imagine unconscious recollection, and equally difficult to imagine totally unconscious planning. An animal which chooses to head for somewhere beyond the reach of its senses surely has some sort of conscious idea - even if it's only the haziest possible awareness of the goal, and of the need which it's expected to answer. This, it seems reasonable to suppose, is what keeps the animal on its path, carrying out the plan, and resisting the distractions that crop up along the way. There's no way of proving this, of course. But I think we can be sure that when a vertebrate behaves in this purposeful manner it must at least be using a more primitive version of neuronal mechanisms that in many species have been expanded to produce sensory consciousness.

How might the evolution of the capacity for firstly conditioned learning and then schema-building and consciousness have come about? Paul Rozin suggested that new adaptive specialisations appear at first in particular contexts. Each innovation evolves as an answer to a particular problem, as it were, and *tends to manifest itself only in the narrow set of circumstances in which that problem arises*. This must certainly be true of the capacity for conditioned learning, which clearly can only happen where there is both a sensory mechanism to register the consequence of some particular sort of response and a channel to transmit a message about that consequence to the site where responses are triggered. The neuronal connections which convey different sorts of feedback must have come into being one at a time, resulting in the very specific learning abilities of different species.

It's reasonable to guess that schema-building and consciousness evolved in the same way, functioning only in one context to begin with, and then in some lines of descent being extended to guide other sorts of behaviour. Each extension to consciousness is likely to have required the addition of some new wiring in the brain. If so, then the nature of the activities in which a species can make use of consciousness must be determined by the particular range of connections its brain has acquired. Thus there may be species which, say, use consciousness only for building a cognitive map, can apply some discrimination learning to their eating habits but only Pavlovian conditioning to the business of escaping from danger, and which govern their relations with conspecifics entirely by means of genetically determined responses.

This list of possibilities is not random. A conspecific is very much a genetically determined stimulus pattern. Its appearance, vocalisations and odours are dictated by the same set of genes which shapes the sensory systems that must register them. Consequently appearance, sounds and smell can be entirely predictable. A purely hardwired system must be very adequate in species which don't distinguish among individual conspecifics.

For avoiding danger a hardwired system can be similarly effective. In fact it's likely to be the best bet, since the hardwired response must be pretty efficient if generations have survived to pass it on. Experimenting with new ways of dealing with threats is likely to be risky. On the other hand, learning to react to the stimulus which signals the possible approach of danger, so that evasive action can be taken sooner, is clearly a very useful talent. There are obvious advantages, too, in being able to improve on the hardwired indication of what's edible.

But geography, at the local level, is unpredictable. The only sort of rules about pathfinding that can be written into the genes are relatively generalised ones. Quite a few of the more elementary aquatic species work on the principle of heading up towards light when food is needed and staying in the dark at other times, or heading one way in the morning, the other in the evening. Some rules that work on a global scale guide the migrations of birds, turtles, and other species equipped with a magnetic sense, or with the ability to register the changing pattern of the stars - rules like head north in the spring and south when the days grow shorter. But there is no way the genes can encode the particular pattern of landscape into which an animal will be born. So a capacity for the most complex form of learning would be most valuable in the context of getting about.

The guess, then, is that in the evolutionary history that led to *Homo sapiens* the capacity for conditioned learning probably evolved first in relation to the form of behaviour where it was most useful, and gradually spread to other areas of life. And the capacity for creating a conscious sensory model of the world grew in the same way. Each extension enabled the individual to adapt to the environment in a more thoroughly flexible way. This, said Kenneth Craik, an influential psychologist and philosopher, *may be one of the functions of consciousness - to permit a greater 'elasticity' and flexibility and unity of response*.

However, the adaptability of the individual mustn't be confused with the adaptability of the species. Bacteria may have a very limited and fixed repertoire of behaviour, but they reproduce at furious speed when circumstances are favourable, and have plenty of opportunity to produce genetic adaptations to changing circumstances. The capacity for conditioned learning is only profitable, and therefore only exists, in species which live long enough to derive some benefit from having learnt something. And the construction of a cognitive map requires an investment of time and energy that is only worthwhile if the individual's potential lifespan allows the rewards to be reaped. The prospects for the survival of a species' genes may not be so very different, therefore, whether the strategy is one of short existences, huge numbers, high wastage-rates and frequent genetic mutations, or whether it is one of long lives, more modest numbers of individuals, and an investment in learning which increases the chances of individual survival.

We should note, too, that while the capacity for constructing really sophisticated schemata, in other words the capacity for advanced reasoning, increases adaptability yet further, it's a high-risk strategy for survival. The capacity to reason is also the ability to produce ideas that are bad for the prospects of the individual's genes, and in some cases may threaten the whole species - to decide that copulation is wicked, for instance, that suffering is inherently better than pleasure, that there is a god who is pleased when we murder fellow humans, or, as some misguided and rapidly deceased individual apparently once did, that it would be a good idea to eat nothing but carrots. Reason can undermine the guidance provided by our genes. The bacterial strategy for passing genes on to many gener-ations of descendants may well be sounder than ours.

Recapitulation

Reason can be defined as an ability to perceive relationships of all sorts, and to juggle with them and create new relationships - whether by making a simple modification to a twig to turn it into a tool, or by constructing an abstract idea.

Rational thought requires consciousness, and consciousness will be the subject of much of this book. It can be subdivided into at least three levels.

The most obvious is sensory consciousness, out of which the more complex levels must have evolved. The last to evolve is closely intertwined with the invention of language.

Even in humans not all behaviours require consciousness, and some learned responses can be performed without it.

Learning has been studied extensively, in various species. Some forms, which answer nicely to the term conditioning, concern when to produce a certain response.

Much conditioning seems to entail refinement of the stimulus pattern to which a hardwired response is produced.

Learning the geography of the local environment is a much more complicated business, but extremely valuable. I hypothesise that it was for this application that some of the machinery of consciousness first evolved.

A map is a pattern of essential relationships, and can be regarded as a form of schema, to adopt a term used by Piaget and Bartlett.

A capacity for building schemata was gradually extended to other subjects, in addition to the layout of the environment.

The essence of a sensory schema is that it records the essential features of stimulus patterns and their relationships, allowing unimportant variations to be discounted.