

## Chapter 27

### A Tentative History of Consciousness

Consciousness, then, is the effect of neuronal mechanisms that bring a great deal of information together, making it worthwhile to collect a great deal of data from which to extract information. How these arrangements produce the subjective effect remains a mystery. (But this is not the sort of problem brains have evolved to cope with.)

Towards the beginning of this book I surmised that consciousness probably evolved in relation to one behaviour at a time, with each additional behaviour promoted to consciousness by the evolution of new neuronal connections. Specifically, I proposed that some of the machinery which underpins consciousness might first have emerged to cope with the challenge of directing locomotion, since this is where it could be most valuable. Genetically ordained rules can provide some guidance, and conditioned learning can extend it. But the geography of the environment into which an animal is born is essentially unpredictable, so something more is needed for knowing just where the best chance of food or shelter is, and the quickest way to get there from any starting point. And that's extremely valuable knowledge, the more so since locomotion, unlike other basic behaviours, is an extended activity which consumes a good deal of energy. I've assumed, therefore, that a cognitive map was the first sort of schema for which neuronal mechanisms evolved. For any species inhabiting a featured environment, equipped to travel, and living long enough to build and use a map, the ability to record the layout of the local environment and confine activity to the profitable areas must be a major benefit.

On the other hand stimulus patterns which represent food or conspecifics or signs of approaching danger can be adequately defined by hardwired mechanisms. Conspecifics are stimulus patterns that not only can but must be written in the genes, with sufficient detail that potential breeding partners can be reliably identified. The shape of danger is not so easily predicted, but experimentation is risky, so hardwired responses to broadly defined stimulus patterns provide the best way of avoiding harm. There's a great advantage, though, in learning to respond to cues that signal potential danger, however roughly such cues may be defined. Similarly, the stimulus pattern that represents food may be genetically determined, but the search for it is more efficient if the animal can learn to recognise signs that indicate its probable presence. All three behaviours can thus be managed pretty well by inbuilt stimulus-response programmes, but the efficiency of the latter two can be notably improved by some conditioned learning.

All three, moreover, simply involve responding appropriately to current sensory inputs. Learning where to look for food is another matter. If the territory covered is larger than that within immediate reach of the senses a medium-term memory is necessary to fit the different parts together, and to provide the means of establishing the overall pattern in long-term memory. There also has to be a means of accessing the records that are stored in long-term memory, so as to decide where to go.

Subjectively it's hard to imagine these essential requirements being fulfilled without at least a faint sense of place and a faint sense of motivation. Could the track that's now out of sight be correlated with the current surroundings, or a distant goal chosen and reached, by wholly unconscious means? Perhaps - there seems to be no way of finding a definite answer. But I think we must at least deduce that the necessary records are established by the sort of sustained and extended circulation of impulses that would be elaborated and expanded to produce consciousness. The business of guiding locomotion seems to be the likeliest place to look for the beginnings of the machinery that produces sensory consciousness, and for the earliest examples of such consciousness.

The hypothesis based on logic and subjective argument turns out to be capable of being fitted to the neurological evidence. Learning large-scale spatial patterns, getting any sort of information into the long-term store in such a way that it can be recalled, and remembering the goal of the current journey, all depend on the hippocampal formation. People who have suffered extensive hippocampal

damage can't do these things. For those of us with fully functional hippocampi, on the other hand, a conscious knowledge of the origin, path and aim of a current journey over known territory is constantly available, though for much of the time it doesn't occupy the focus of attention. There's a hazy, background awareness, against which other activities or thoughts about other things take place.

An equivalent of the hippocampal area seems to be present in all extant vertebrates, so it must be a very ancient part of vertebrate neuronal machinery. In simpler, smaller vertebrate brains it takes up a large proportion of the telencephalon, which implies that its function relates to a major part of what the telencephalon does. It's reasonable to suppose that it began by providing the medium-term memory that's needed to assemble a series of successively experienced scenes into a coherent record, creating a temporary cognitive map in association with the thalamus. As the telencephalon expanded the temporary maps could be put together and a longer-lasting, more extensive one assembled in the newly evolved area.

If these arrangements began to produce some sort of consciousness it may be imagined as a very vague consciousness, something like the minimal awareness of place and time that in us forms the background to events that are receiving focussed attention. Conceivably it might only operate when the animal was moving towards some goal, or about to do so.

The telencephalic map began, I've suggested, with a fair amount of input from the adjacent olfactory system and some minimal input about visual texture from the thalamus, which would also forward information about current actions. In addition to the olfactory and visual consciousness I think there was some sense of the need the locomotion was designed to fulfil - hunger, or unsatisfactory temperature, or, in the breeding season, hormonal arousal. It might be called a motivational or emotional consciousness.

Quite a vague, undetailed account of the permanent features of terrain or seabed, along with olfactory guidance towards food or mating opportunities, should enable an animal to find its way around in a useful fashion more often than not. Once the goal area was reached the good old-fashioned hardwired forms of action guidance would take over. Food would be seized or breeding behaviour performed without any need for sensory or motivational consciousness to guide the action. Perhaps, though, a conscious sense of something resembling satisfaction might result, based on feedback about long-lasting internal states. Channelled to the hippocampus, it would reinforce the marking of the site on the map.

Insofar as current knowledge allows us to guess, a minimal, hippocampally based, travel-guiding consciousness might reasonably be attributed to frogs. The visual information which guides their prey-catching behaviours triggers extremely fast responses, leaving no time for the luxury of creating a conscious experience. Visual input which may represent a large descending foot or an approaching predator produces another rapid reaction. These mechanisms are clearly hardwired, and in neither case could a conscious sensation serve a useful purpose; for catching fast-moving insects or escaping from imminent danger speed is more crucial than seeing fine detail or planning the nature of the action.

The limited visual information which goes to the thalamus must be what guides locomotion. Internal factors, such as temperature or appetite, will prompt the decision to expend effort on moving about, and the urgency of different needs must be weighed. These judgements are no doubt made in the telencephalon, on the basis of messages from hypothalamus and elsewhere. The guess is that the correlation of internal needs with visual and perhaps olfactory inputs – all sampled over an extended period - involves a recurrent circulation of impulses sufficient to create a dim consciousness. Which seems appropriate for the sluggish, pause-punctuated form of travel generally practised by frogs.

## **Foreground consciousness**

For us humans the bright, fast-changing consciousness which paints the foreground of our subjective experience clearly depends on the neocortex. If a patch of neocortex is destroyed some part of the sensory experience is either lost, or becomes very confused. Without primary visual cortex

there is almost no conscious vision, while damage to one of the sectors a little further forward creates a gap in the patient's visual sensations. There may no longer be any perception of colour, or objects may appear to relocate suddenly from one place to another without any perception of the intervening movement.

Damage to a schema-area doesn't leave a blank where the percept ought to be, but renders the afflicted person incapable of making sense of a stimulus pattern that should be matched to that sort of schema, as in the case of Oliver Sacks' patient who 'mistook his wife for a hat'. When a bit of neocortex and its connections are knocked out an essential node in the circuit dealing with a certain sort of percept is removed, so the circulation of impulses which would allow a succession of inputs to be assembled into a meaningful perception is prevented. There would seem to be some sort of conscious experience but it's incomprehensible. The stimulus pattern cannot be linked to any associated information.

A conscious experience also depends on impulses reaching prefrontal cortex, as we saw in Chapter 8. This, we can guess, is necessary for the synchronisation of large-scale rhythms across activated pathways which enables the prolonged circulation of neuronal messages. Probably it's achieved by means of an interchange of messages between prefrontal and the basal ganglia, which between them record the known implications of the stimulus pattern.

This is relevant to several early experiments on conscious vision. If a visual stimulus pattern is undetectable when very briefly flashed by the tachistoscope, but becomes quite clearly visible after a series of such brief exposures, it must be because synapses along the pathway through sensory cortex have been sufficiently activated to become potentiated. As a result they become sensitive to even very brief inputs, and start firing strongly enough to forward a signal to prefrontal cortex. Prefrontal then initiates return messages, completes a circuit and establishes synchronisation. And George Sperling's experiment showing that one row of a very briefly displayed three-by-three array of letters can be read out from iconic memory after the exposure has ceased, but only one, implies that while signals from sensory cortex must reach prefrontal cortex to produce an awareness that there is something there, returning, top-down signals are necessary for the analysis that reveals what it is, and gives precision to the conscious experience.

We can also see how messages from sensory cortex can activate innate and conditioned responses before they reach prefrontal cortex, sometimes without getting that far, so that it's possible for attentional and emotional effects to occur even when sensory inputs don't reach consciousness. An innate or conditioned response to a sensory input makes it more likely that the message reporting that input will reach prefrontal cortex, but that result is by no means certain.

The situation must be somewhat different in the wholly inexperienced infant brain, where there are only genetically determined responses, no conditioning has yet occurred, and prefrontal cortex develops only slowly. If the hypothesis put forward earlier is correct a sensory input only causes an extended circulation of impulses in the neocortex if it evokes a genetically determined response, released by the basal ganglia. Any such action, even if it's only a saccade, serves to lift the inhibition exercised by the basal ganglia on the pathway from a midbrain sensory area to thalamus and thence to neocortex. That inhibition isn't complete, it's theorised, but without it the thalamic neurons can fire much more strongly, and a more distinct message goes to neocortex. This more powerful input produces a level of neocortical excitation sufficient to trigger a message back to the thalamic reticular nucleus, punching a hole over a 'driver' nucleus so that all available information about the stimulus that prompted the movement gets through.

The theory explains how behavioural attention is converted into cerebral attention – or at least one of the mechanisms. It incidentally supplies another reason why the evolution of rotatable eyes would prove so seminal. The hardwired part of the machinery that governs saccades provides an easy and inexpensive way for the midbrain to tutor the neocortex.

As the infant neocortex learns about a stimulus pattern and its significance a pathway back to the basal ganglia is established. Instructions can be sent to lift the basal ganglia inhibition, and thereby, indirectly, the inhibition exercised by the thalamic reticular nucleus, even if no hardwired

response is performed. As this pathway is established the infant learns to direct attention more flexibly.

The supposition is that sensory or 'foreground' consciousness happens if the pathway to the thalamus which reports the performance of a hardwired response to a genetically determined stimulus pattern extends into a brain area that carries out a sophisticated analysis of the sensory input which prompted the action. In mammals this is of course the neocortex. The message indicates both the nature of the action and the position of the stimulus at which it's aimed. The neocortex takes over from there.

It's reasonable to suppose that the pathways were extended one by one, and that the first use to which the evolving sensory processing facilities of the telencephalon were put would be to add visual detail and precision to the cognitive map. A concomitant of this development, I propose, is that the cues which attracted approach or other responses were promoted into consciousness as actual sensory objects. No longer were they experienced merely in abstract, emotional fashion, as goals defined in terms of the need to be fulfilled, or threats to be retreated from. Now the actual shape or sound - to whatever extent it was revealed by the cortical sensory pathways - was consciously experienced. The account of the external world would continue, of course, to be pervaded by emotional and motivational significances. But it would be a more vivid and shapely world.

Other subcortical structures would also evolve links to the neocortex, or where links to the telencephalon already existed more value could be developed from them. Interconnections between cerebellum and neocortex, for example, would make it possible to learn in more detail than before which visual textures indicated solid barriers, which it was possible to push through, and what sort of input correlated with surfaces on which it was safe to put the feet. Pretty well all vertebrates must have some means of learning to distinguish impassable solidity from something flexible that can be pushed aside. Tetrapods also need to learn what sort of surface will bear their weight. The expanding sensory processing facilities of the neocortex must have provided mammals with a potential for more detailed learning about these subjects, once tactile inputs were reported to the neocortex - promoting another subject to consciousness.

This information would be marked on the map, but it's a slightly different thing from the map. Knowledge of landmarks and the overall geographical layout is acquired through vision and locomotion, with the gaze mostly focussed at some distance. Information about the ground and barriers is acquired through touch, and with a different convergence of the eyes. Moreover the feedbacks are mainly negative - the discovery that trying to push through dense bushes doesn't work, and that treading on the wrong sort of ground leads to alarming consequences. Such information must have come to be stored separately in the neocortex, available to attach to more than one location on the map. A neocortical area devoted to such matters has not, I think, been identified yet, but I suspect it was one of the earliest additions after the development of the new, elaborated cognitive map.

A similar exchange of messages with the amygdala would allow the stimulus patterns which prompted hormonal and visceral responses to be analysed in the neocortex. Later it would be a dialogue between amygdala and neocortex that made it possible for mothers and infants to recognise each other by visual and auditory means, as well as by olfaction. As maternal behaviour evolved the maternal amygdala must have become extra-sensitive to the olfactory signals given off by infants. The resulting excitations in the amygdala would be forwarded to neocortex, where complex visual and auditory discriminations could be made, enabling mother and infants to identify each other at a distance. And maternal hormones, as well as promoting suitable behaviour, came to have the sort of effect in the neocortex which is experienced as emotion.

It's sensory consciousness, I've suggested, that makes it possible for mothers and infants to maintain contact with each other while simultaneously pursuing other ends and dividing their attention. Mother's and infants' conscious perception of each other functions as a form of reward, though only sampled intermittently. Moreover, if the percept suddenly turns out to be unavailable

its absence functions very powerfully as an alarm and a source of distress, because there is a conscious sense of what ought to be there. The implication is that a well-developed sensory consciousness lies at the root of mammalhood, and must have been established before this means of nourishing the young evolved.

In some ways the challenges entailed in keeping mother and infants or the members of a social group together constitute a variation on those involved in travelling to a distant goal. The goal and the route to it have to be kept in mind despite any distractions along the way. The stimulus that has to be kept in mind by both mother and infant is a moving one, and the aim is to make sure it remains within view although it's not constantly observed. The hippocampal map must play an important role in keeping track of the not-constantly-attended figure or figures. And measuring the distance between mother and infant probably depends on much the same mechanisms that measure the distance to a visible goal.

To some extent it's easy to deduce where sensory consciousness is likely to be applied in other species. Many mammals chomp away at a genetically determined diet and presumably don't need to devote much neocortical space to the question of what to eat, so perhaps a consciousness of taste is unnecessary. Neo-cortical analysis is probably widely applied, though, to the cues that indicate where to find food, especially where it's the sort of food that is inconspicuous or widely scattered.

Things must be more complicated for carnivores, at least if they hunt a variety of species, for there seems to be a need to learn how the different prey behave. Doubtless such hunters have a neocortical area which links the appearance of each prey species to its escape strategies, and the complementary strategies needed to capture such a meal. The particular taste which can be expected at the end of the hunt might also be expected to form part of the schema.

Something more is needed in adventurous omnivores, prepared to experiment cautiously with new possibilities, and equipped with some remarkable feedback connections that can derive learning from effects on the stomach which occur some time after the eating. The neocortical area which combines olfactory and gustatory information, turning it into the conscious experience of flavour, might be termed a schema area dedicated to the subject of food. It seems to be responsible for the remarkable conversion of a taste once found pleasant into something aversive when consumption has been followed by feelings of illness. In the gustatory cortex of rats the activity evoked by a sweet flavour began to look more like that evoked by a bitter one after it had become associated with 'a visceral malaise.' This neocortical sector no doubt evolved in conjunction with the development of an increasingly versatile digestive system.

Perhaps the most fascinating part of the neocortex is motor cortex, which provides the ability to shape actions with the aid of conscious feedback from vision, touch, and the senses reporting on balance, muscle stretch and joint position. In the lamprey there is a telencephalic area where artificial stimulation elicits movements. This ancient bit of forebrain has been greatly expanded in species which produce finely adapted movements shaped with the aid of experience - in other words those which go in for motor play in youth. The amount of such play indicates the amount of learning that must be done, and no doubt correlates with the size of the motor areas. We can deduce that where there is play there is a conscious experience of movement.

The expansion of motor cortex will have involved new links with the basal ganglia, and a considerable expansion of the latter. It will have coincided with the evolution of more complex musculatures that could be employed in more varied ways and of new neuronal pathways to carry instructions to them.

Such developments led to new uses for paws. No longer just things for walking on, forepaws variously became capable of patting in an investigatory way, clutching at food, clawing a way into it if necessary, or climbing trees - all activities which could benefit from learning just how movements should be adapted to fit different challenges. In some species jaws became more versatile, and different tactics could be adopted for tearing at different sorts of food.

The existence of motor cortex and the ability to learn how to use muscles most efficiently must have facilitated further evolutionary developments. Motor cortex and play were clearly well established long before the advent of hands, and extremities that could be described as fingers. We might guess that fingers were worth having because there was already a means of learning about the many varied possibilities for using them.

Perhaps, moreover, it was because primate fingers could readily acquire a variety of foods that were, in different ways, a little tricky to eat - fruit with inedible seeds or hairy skins, nuts that needed a lot of chewing - that some parts of the eating equipment also came under neocortical control. That meant that some of the muscles used in the consumption of food could be used as well for voluntary vocalisations, such as the lipsmacking and pant-hooting of chimpanzees. This laid one of the foundations for the development of speech and language.

As animals learnt to use their forepaws in new and varied ways, meanwhile, a neocortical area must have evolved that dealt in the objects to which such actions could be applied in a useful manner, and the sort of objects that should be left firmly alone. It probably began as a subdivision of the area that recorded which surfaces were safe to walk on. In primates, as forepaws turned into hands and objects began to be used as tools, the new schema-area would expand significantly, reaching its greatest extent, of course, in humans.

### **The interlocking webs**

The neocortex is well fitted for contributing to elaborate, self-maintaining circuits, of the sort believed to create conscious experience. Its major excitatory neurons, the pyramidal and stellate, have enormous numbers of input synapses, and the inhibitory interneurons which play such an important role are well endowed too. The thalamic nuclei which forward sensory information to the neocortex receive numerous projections back from it, as do other thalamic areas, and it's a general rule that any other subcortical area which sends axons to the neocortex also receives connections from it.

The largest proportion of the inputs to cortical neurons, however, comes from other cortical neurons, some local, some distant. Every stage of the visual pathway, for example, communicates directly with almost every other stage, and also receives inputs from the forebrain. These contributions synapse mostly in upper and lower layers, whereas incoming sensory information arrives mainly in the middle layers. In visual cortex the upper and lower layers play an important part in distinguishing small patches of one texture against a wider extent of a different sort, which no doubt indicates the sort of use that the horizontal connections within the cortex can be put to.

As we've seen, some parts of the neocortex are primarily concerned with processing a certain sort of sensory input, others with recording significant stimulus patterns, some with evaluating the input, and yet others with directing movement. But the whole is constructed like a three-dimensional net of interlinked spiderwebs, so that a small impact on a bit of one web can send tremors reverberating through other quite distant webs. Activation in a node that represents a familiar visual stimulus pattern can cause a subliminal level of excitation in the auditory node representing the sound associated with it, and in nodes representing other aspects or associations. The rate and pattern of neuronal firing quite early in the sensory pathways is modulated by attention and emotion, and influenced yet further during the development of a decision to act on the stimulus pattern being registered. Motor areas and motor neurons can show a small flutter of excitation in response to sensory inputs that merely suggest the idea of action. As we listen to speech there is a faint activation of the motor neurons that would be at work if the self was uttering those words. Tongue muscles can be faintly primed in anticipation of an expected phoneme, and the priming process is suspended if the expectation proves wrong.

This supports the idea that the neocortex is organised essentially in relation to action, thanks to those inputs from midbrain and basal ganglia via thalamus - a reasonable proposition, since the essential function of neocortex is to decide where action is needed, and if so, what sort. The result is a very effective filing system. As I've emphasised, mammalian sensory systems have to deal with an

enormous variety of sensory inputs, with the potential number of classifiable stimulus patterns limited only by the amount of detail that can be registered. Inputs must be organised into schemata flexible enough to fit many variations on a theme. In contrast, the range of possible movements is limited, even in the most agile of animals, and the movement required in response to a stimulus pattern often has to be very precise, honed by much practice. The records in motor memory are thus much more closely defined than those in perceptual memory. Using the former to help deal with the latter is a canny strategy.

The extensive connections of the neocortex provide an enormous potential for large and complex circuits, with many possibilities for recycling a current back to points it has passed before. I envisage foreground consciousness as a flow of activation snaking through the neocortex along many bifurcating pathways in a constantly evolving pattern, fuelled by inputs from lower brain areas. Many sectors of pathway are covered repeatedly as a perception evolves, but there may be modifications to the circuit at every pass. Maintenance of the current is facilitated by a temporary co-ordination of the larger-scale rhythms in all the areas involved. Well-synchronised and strongly activated pathways equate with the focus of attention, while less synchronised rhythms elsewhere provide the unattended and sometimes unconscious background to the main sensory experience.

The reiteration of the conversations between cortical sensory areas and schema-areas means that information originating from many different sources can be assembled into a coherent whole, even if the processing of input takes a little longer in some pathways than in others - as in the magnocellular, parvocellular and koniocellular pathways from retina to visual cortex. Consciousness, I suggested earlier, is the answer to the philosophers' 'binding problem'. It's what makes it practicable to have so many and such extended sensory processing pathways. Their contributions can be brought together over the considerable number of milliseconds during which the conscious moment is created. Thus quite complex correlations can be discovered.

On a slightly longer timescale, the short-term memory provided by the sustained circulation of a large proportion of the impulses is substantial enough to reveal the pattern in a succession of inputs, and to allow components of the pattern to be evaluated in the context of the whole. The bigger the circuit, the longer the circulation is likely to be kept going. And the larger the chunk of data that can be fitted into one perception, the greater the potential variety of percepts.

The reiterated circulation of impulses around the extended circuits also serves to reduce the danger that the elaborate processing which produces fine discriminations might sometimes lead to unreliable or confused results. In the case of moving or evolving stimulus patterns the interpretation of each component can be influenced by the context of the whole. In the case of static ones the information extracted from the data undergoes multiple rechecking, and the interpretation can be refined or altered before action is undertaken. There's time, too, for correlating the input of different senses.

For emerging mammals, operating in the dark, it was no doubt quite practical to delay action while triple-checking such information as was available, and very worthwhile. There were probably no nocturnal predators to threaten them, and when some did evolve they would have to depend heavily on olfaction and hearing to locate their prey. A prey animal could afford to freeze until it had sorted out where the threat was located, and which direction was best for escape.

In short, consciousness functions, as Giulio Tononi and others have proposed, to correlate all sorts of disparate information, information which may be acquired and processed over a period of time. The correlations have to be established through experience, of course, which is why we can make a reasonable guess as to just which sorts of information a species is equipped to correlate in this way by observing what sort of exploration and experimentation the young undertake. In humans, meanwhile, the expansion of the final stages in the cortical processing machinery, at the very front of the frontal lobe, has extended the reach of consciousness and made abstract thought possible.

The major divisions of the neocortex are determined by axons arriving from the thalamus and perhaps other subcortical areas. The structure built on these foundations is to a large extent self-

shaping. Neurons are produced in exuberant quantity, and those that find no employment die away. And the nodes of the neocortical webs are heavily shaped by experience, thanks to the plasticity of the synaptic connections. The evidence suggests that in the newborn infant a great deal of the neocortex is activated by any sensory input, and the interconnections are gradually pared down over years.

This may be true even of early stages of the sensory pathways. Some people enjoy a condition known as synaesthesia, in which some sensory inputs produce, as well as the usual experience, an additional one. For instance words or numbers or tunes may regularly evoke sensations of particular colours. In some of these individuals this idiosyncratic form of perception is lost as they grow up, in others it persists. One theory is that synaesthesia may be the norm in infancy. Perhaps the excitations from sensory inputs initially spread widely, but then the excitatory synapses that multiply profusely in the first year of life are pruned and inhibitory synapses are potentiated. The infant brain somehow discovers which links in the network produce useful predictions and which don't, and the synaesthetic effects disappear. But maybe some individuals find such reward in a world with extra colour, or other associations, that some links escape the editing process.

An important footnote here is that such a self-shaping structure is only likely to evolve in species where infants are nourished and protected by their parents, and don't initially need to use their brains for the vital purposes of finding food and avoiding danger. That provides an opportunity for educating the neocortex through experimentation and play. A potential for extending the shaping process over a longer period arose where parental protection came to last beyond the feeding, and especially where a social lifestyle added to the security.

### **Olfaction as a probable pioneer among senses**

The shape of some of the most ancient fossilised fish skulls suggests that olfactory bulbs took up a good deal of space, implying that there was already the means of registering a fair variety of odorants. It looks, therefore, as if the means of collecting a quite extensive range of information may have been acquired earlier in olfaction than in the other distance senses, which seems quite probable. Odours can be intrinsically informative, so olfactory powers could be expanded simply by the evolution of an increasingly large variety of receptors. In extant olfactory specialists such as rats the number of different types is around a thousand, in dogs it's greater, and even we humans possess four or five hundred - in strong contrast to the limited repertoire of photoreceptors or mechanoreceptors possessed by any one species. The size of the olfactory bulb probably reflects the number of different receptor types, so we can deduce that many early fishes had a decent selection. Early tetrapods are unlikely to have been worse off.

With such a variety of receptors there would be less need for the correlations and comparisons required to extract better data from the inputs of just a few types of visual receptors. The earliest odorant receptors are likely to have been tuned to the things a species most needed to know about - conspecifics and food. Receptors responsive to odorants regularly associated with food, rather than emitted by the food itself, were probably then added. All this input was worth mining for subtler information, however, and the three-layered olfactory cortex, at the front of the brain, close to the olfactory apparatus and the olfactory bulbs, was well developed long before the neocortex appeared. When that new bit of telencephalon was added it probably elaborated on techniques that had first been applied to odours.

An olfactory receptor is activated when an odorant molecule keys in to a docking site. For some receptors there's only one odorant that fits, for others there may be several bearing the right sort of projection. To complicate things, some odorants can activate more than one type of receptor, but with different degrees of efficiency. Consequently some odours may be securely differentiated only by the varying proportions in which the same set of receptors is activated. In these cases the input levels of different types of receptors must be compared.

To add to the challenge a significant aspect of olfaction in tetrapods is the way the odorants are drawn across the olfactory epithelium with the intake of air. Encounters between odorant molecules

and the receptors they fit must be somewhat random, on the timescale relevant to neuronal excitations. Moreover some odorant molecules engage with their receptors more rapidly than others. In rodents the different types of receptors seem to be located along the length of the epithelium in such a way as to allow for the latter factor, but it seems likely that chance still plays a role in the timing of excitations. The solution is to collect up the input over a brief period. This must be the only way to identify an odour that activates several types of receptor to varying degrees, and to establish at the same time just how intense it is.

The axons of the receptor neurons run to the olfactory bulbs, terminating in structures called glomeruli, each of which receives inputs from only one type of receptor. The overall pattern of activation develops over the course of an inhalation, reflecting both the temporal pattern of receptor activation and the increasing amount of inter-glomerulus inhibition which is triggered as the levels of excitation grow stronger. The initial pattern of responses may be common to several different odorants but the final pattern usually seems to be unique to the odour currently being inhaled.

So in effect the inputs over the whole intake of breath are compared, and treated as a unit. Tests indicate that a rat can recognise a familiar and easy odour before a whole sniff's worth of input has had time to register, but dealing with a more complex odour or an unfamiliar one takes longer. For discriminating between similar odours it's necessary to wait for the later-registering components, and weigh the relative contributions. In short the olfactory system uses a chunking system, very much like that of the cortical part of the visual system, in which the inputs gathered over a short space of time are compared. There is in effect an ultra-short-term memory, which can winnow more reliable information out of the input than would otherwise be possible. Some learning contributes to the process, for there is a degree of plasticity in the synapses of the excitatory and inhibitory neurons which weave among the glomeruli.

The evolving pattern of responses in the olfactory bulb is conveyed to olfactory cortex, the largest sector of which is piriform cortex. Here the terminals of the glomerular axons may spread themselves out over wide areas, and a piriform pyramidal cell may receive inputs from glomeruli connected to a few or many different receptor types. Imaging experiments with humans show that cells grouped together in the piriform cortex tend to deal with odours that are subjectively rated as similar. Moreover a piriform cell which is excited by a combination of odorants may not respond at all to just one of the components on its own, which correlates with the fact that quite different conscious experiences are produced by such inputs. On the other hand the neuron may be quite tolerant if a minor component of a complex odour is missing - which echoes our ability to recognise a whole visual stimulus pattern when part of it is concealed.

Piriform responses are modified by learning, and fine discriminations can be acquired when odours are predictive of noxious inputs or potential rewards. When two odours of very similar composition were regularly presented to rats, one indicating a reward on the right, the other on the left, the pattern of responses in piriform cortex became more differentiated as the animals learnt the trick of gaining the reward. When circumstances changed and there was nothing to be gained by distinguishing between the odours the piriform responses to both became similar again.

Piriform cortex, it seems, is where odours are classified by significance, rather than by molecular structure, and where significance can be determined by experience. Although it only has three layers rather than six it thus operates on the same essential principle as the association or schema areas of neocortex. Combinations of inputs are classified according to the action that needs to be taken in response to them. Neocortex performs a similar job, with the difference that it's able to give the data more intensive processing before the decision on significance is reached.

There's no way of telling how much of this olfactory processing power was present in distant ancestral species, but it seems reasonable to hypothesise that olfaction was the sense in which an ultra-short-term comparing memory first evolved, and the first to acquire the means of making subtle distinctions between similar patterns of inputs.

It also looks likely to have been the first sense to benefit from a somewhat longer preservation of data. Odours can provide useful contextual information, and probably that was their first use - indicating that it was worth slowing down and turning frequently because food or potential breeding partners were around. The next development would be an ability to pursue an odour to its source by comparing its strength in different places and turning when the signal faded, much as male moths do when searching for the female emitting a sex pheromone. A simple mechanism which prompts a reversal of direction whenever the olfactory input ceases can produce a reasonably efficient search pattern. But clearly the pursuit can be much more efficient if a record of the olfactory input can be maintained for a while, and the search can be continued even if the scent has been temporarily blown away by wind or current.

It should be facilitated even more if the intermittent olfactory inputs can be related to a map of the current environment. Probably the hippocampus evolved close to the olfactory processing system because olfactory information made a major contribution to the earliest form of map. The first maps perhaps received only a limited amount of visual information from the thalamus, just sufficient to relate the animal's position to stationary surroundings. Such maps may have been painted quite richly, however, with odours. The colour would shift, of course, as currents of water or air deflected the stream of odour molecules, which is why some visual information would be needed as an anchor. Keeping a record of the sites of recent olfactory inputs, so that odours could still be pursued if temporarily lost, may have been one of the most valuable functions of the early hippocampus. As vertebrates acquired a hippocampal consciousness odours may have played a big part in it.

### **Consciousness in non-mammalian species**

Most mammalian species must apply consciousness over a narrower range of subjects than we do. Where the neocortex is small we may guess that foreground consciousness is quite limited. Indeed the richness of sensory consciousness must correlate with the complexity of the data that is collected, and the length of the period over which the data can be assembled.

What of non-mammals? Conceivably a hippocampal, geographical consciousness is common to many if not all vertebrates. It may be a talent closely intertwined with the evolution of the backbone and more efficient locomotion, and perhaps a significant factor in the success of vertebrates at the expense of most non-vertebrate chordates.

If the theories about sensory consciousness are correct there's no reason why it shouldn't also happen in brains that are structured quite differently from ours. It might be expected wherever the circulation of impulses around the sensory processing arrangements can be maintained after the sensory input ceases for long enough to allow multiple inputs to be analysed and correlated, and information to be sifted from 'noise'.

In reptiles the telencephalon – rather than mushrooming as the mammalian one did – has been extended with a cluster of nuclei which bulge inwards, into a ventricle, forming what's known as the dorsal ventricular ridge (DVR). In line with this expansion the thalamus grew and increased in complexity, and so did the tectum. Whether that led to sufficiently elaborate circuitry to produce some degree of 'foreground' consciousness we can only guess, but some crocodylians seem pretty intelligent.

Birds have expanded the tectum further, and turned the DVR into a collection of nuclei that probably does a job roughly similar to that of the neocortex, and achieves some of the same ends. Clearly birds not only have well developed cognitive maps, but some at least do things for which sensory consciousness seems to be necessary, such as giving parental care to mobile youngsters. Quite a few of the larger-brained species - parrots and crows, for example - show signs of a considerable and flexible intelligence. It seems reasonable to deduce that there is foreground consciousness here, the amount varying from species to species, as it may be supposed to do in mammals.

The larger cartilaginous fish, meanwhile, have brain-body ratios high enough to put them in the same camp with smaller-brained mammals, and in some bony fish the telencephalon is quite large.

It's certainly conceivable that skilled hunters, such as tuna and the larger sharks, might have evolved a form of intelligence similar in principle to that of some tetrapods, though employing differently arranged neuronal machinery. Looking outside our phylum, the octopus brain is composed of two interconnected lobes which could well provide sustainable circuitry - as J. Z. Young proposed. The behaviour of some of the larger species certainly suggests a high level of curiosity, which in turn implies consciousness.

Whatever sensory consciousness other species possess it must vary, of course, according to the sort of sensory information they are equipped to collect, as well as the amount of processing it gets. The olfactory world of a mouse, for instance, must be much richer than ours, the visual one much poorer, drawn with much less detail. How else might the sensory worlds constructed by other species diverge from ours?

Some aspects of visual experience seem to present only limited possibilities for variation. The perception of edges and points and angles is likely to differ only to the extent that the range of orientation-sensitive neurons in primary visual cortex might differ. In this respect many species might have a less detailed perception of the visual world than we do, but the essential ingredients might be expected to be the same. Distinctions between orientations may be less subtle, perhaps, and lines, edges and angles may collapse into textures more readily, especially for those species without foveas. Indeed it's very possible that for some species textures form the whole story.

Colour is another matter. Language works on the assumption of shared perceptions, but there's no way of being sure that we don't all experience colours differently - not just those of us who lack one or other wavelength-sensitive photoreceptor, or have alternative versions of the red receptor. (It's ironic that when philosophers discuss 'qualia' they so often instance the 'redness of red', given that the uniformity of that sensory experience is particularly in question.) No chance, then, of guessing what the colour qualia of another species might be like. If there are species with conscious vision which can detect ultra-violet light but not red, is their colour experience different from ours, or is the same rainbow palette shifted a notch, as it were, to be evoked by the lower range of wavelengths?

Then there are those species which have sensory powers we don't possess, such as sonar. This, like vision and whiskers, is a sense involving active exploration, so cetaceans and bats could be expected to be endowed with a sonar consciousness. Is it quite different from visual consciousness, or is it fused with it, so that these animals 'see' objects detected by sonar in much the same way as those detected by their eyes? The latter seems quite probable, for two blind people who have been using tactile devices as a substitute for vision over many years report that they have come to obtain some degree of visual sensation thereby. The essential function of both vision and sonar is locating things in space. It seems reasonable to suppose that the nature of the conscious sensory experience might be dictated by the use to which the sensory information is put, as much as by the sort of receptor which has provided it. Occipital cortex may perhaps be the cortical sector concerned with depicting space and spatial relationships, rather than simply the visual area.

Pursuing this line of thought, it's conceivable that colour might be a sensation applied to whichever sensory input undergoes the greatest amount of neo-cortical processing, with the subtlest results. Maybe olfactory specialists such as dogs and rats smell in colour, rather than seeing in colour. A fascinating idea, which perhaps philosophers will like to have fun with.

### **Could a computer experience consciousness?**

This has become a popular question in some quarters. It may be a futile one, for if computers did develop consciousness how would we know? I assume that my fellow humans experience the world in much the same way I do because you seem to be entities constructed along the same principles as me, behaving with similar motivations. I am happy, too, to deduce that other species have some form of consciousness if their brains look something like ours and their behaviour suggests it. But these factors don't apply to computers. The only possibility for checking on whether a

computer has consciousness would seem to be to ask it. And how would we know it had understood the question?

But we might ask why a computer would need consciousness. Electronic circuits seem to be somewhat more consistent than neuronal ones, and the input comes in decisive form, so there's no necessity for a comparing memory to clarify it. And nourishment is provided by an outside force, so emotions and motivations and a means of weighing up internal needs are all unnecessary.

The most promising line of research, still at an early stage, would seem to lie with the development of robots which have to learn how to interpret input, how to act, how to get their battery recharged, and how to communicate with each other. But it will still be difficult to determine whether they experience anything parallel to this curious effect produced by numerous interlinked neuronal circuits.

On the other hand it might be argued that in a sense computers already have consciousness. The monitor screen constitutes their form of it, revealing the results of the umpteen calculations going on out of sight as a result of the sensory input received through the keyboard. The difference, it might be said, is that computers use their consciousness to communicate with us, whereas we use ours first and foremost to communicate with ourselves.

## **Recapitulation**

A hazy, hippocampally-based form of consciousness probably evolved first, supporting the construction and use of a cognitive map – first a temporary map relating the animal to its immediate surroundings, then a more permanent and extensive one.

Fullscale, vivid sensory consciousness as we know it must depend on the neocortex. (But something comparable may well be created by other brain areas in some non-mammals.)

Consciousness involves a maintained circulation of impulses through complex circuitry.

It allows sensory evidence to be assembled into meaningful wholes even when the inputs arrive at slightly different times, or when simultaneous inputs are processed at slightly different rates.

The first application of the sensory-processing machinery of the neocortex probably lay in adding detail to the cognitive map, and to cues for approach or avoidance.

As a result, we can hypothesise, such cues were promoted from a vague motivational consciousness to something like the sort of sensory experience that we know.

Subtle ways of combining the evidence from different senses became possible.

The potential for subtle distinctions opened the way to new behaviours, such as maternal care of mobile young.

The olfactory telencephalon was well developed long before neocortex appeared, and olfaction was perhaps the first sensory system to process inputs with a very short-term comparing memory.