

Chapter 9

The Benefits of Conscious Sensory Experience

So far the discussion has centred on the single conscious moment. The focus has been on the perception of static visual stimulus patterns, over which we can afford to linger. But consciousness is an ongoing, constantly changing experience. Each instant is interpreted in the light of what's gone before, and the limited information a single moment provides is usually only meaningful if it can be fitted into a larger pattern. The most important benefit of sensory consciousness must be that the process which creates the conscious moment also creates a very accurate though brief-lived record which can be related to following inputs. Sensory consciousness is above all a stream of consciousness.

For most stationary objects we need to put the products of several fixations together to get an adequate idea of the whole. The contribution of the very short-lived record must explain how it is that our conscious visual experience manages to be so much more informative than might be expected, given the limited amount of data that can be garnered towards the periphery of the retina. Each 'picture' is informed by what was learnt in previous fixations. Having identified something as a tree we remember it as a tree, though it's now just a blur on the far edge of vision. The record means that the products of several fixations can be put together to establish the wider-scale layout of the scene. In passing it may be noted that the evolution of a retina which concentrates most of its receptors and processing power in a small central area can only have been profitable in species which possessed such an ultra-short-term memory.

The very brief, very accurate record also serves, it may be inferred, to set up a slightly longer, not quite so accurate record which is maintained in existence even after the gaze shifts to a completely new focus. This enables the new set of data to be correlated with that which has just been absorbed. As you turn from, say, the building on your left to the building on your right you can acquire a very good idea of their spatial relationship, even though you can't get them both in view at the same time. You can also make a pretty good judgement as to their relative size, their relative distance from your viewpoint, and how their colours compare. These judgements aren't quite as precise as they might be if both buildings could be fitted into one retinal image, but they're a great deal more accurate than comparisons that can only be made in a longer-term memory. This sort of record makes it possible to build up schemata for very large objects which can't be studied with one set of overlapping fixations, and to add together the succession of views seen as one moves about.

Most of our sensory schemata are similarly distilled out of a succession of sensory inputs. In us humans a particularly impressive short-term memory is the one which serves audition, enabling us to take in lengthy sentences as meaningful units. It's pretty certain that an expansion of auditory short-term memory was important to the evolution of our capacity for complex language. Much simpler patterns can be extremely significant, however. A rustle in the undergrowth which is repeated at a nearer location, for instance, is a different matter from a series of rustles moving away. Being able to distinguish between the two patterns was no doubt the sort of benefit that first made a short-term aural memory, even one of small capacity, worth having.

This memory allows us not only to capture a pattern that unfolds over a brief period, but even to replay it, for further consideration. It's been termed echoic memory. When you haven't quite caught what's being said, or when what you're hearing doesn't seem to make sense and you need to reconsider your interpretation of the sounds, you can often rewind the tape, as it were, and play the sentence again. And with luck you decipher it on the second encounter.

Other sensory channels too have a short-term memory that allows patterns to be perceived, and which provide a replay facility. We can consider a series of tactile inputs as a whole, or interpret a moving touch. That enables us to learn about a whole object (not just an edge or two) by running our fingers over it, or to recognise a letter or other simple shape traced out on the skin. For those who are both blind and deaf it makes it possible to communicate by means of a finger alphabet.

We can also retain quite a lengthy flow of visual movement, and patterns of visual movement can tell us as much as stationary shape does, sometimes more. We can often identify a familiar person by the characteristic rhythm of their walk or gesture while they are still too far away for more than a basic, overall shape to be distinguished. Just how much information can be derived from visual movement alone was demonstrated by Gunnar Johansson. He attached twelve small lights to a collaborator, at shoulders, elbows and wrists, and at hips, knees and ankles. Then he took a ciné film of the illuminated figure in a completely darkened room, so that only the lights were visible. The film was shown to unprepared observers, with unexpectedly clearcut results. *During the opening scene, when the actor is sitting motionless in a chair, the observers are mystified because they see only a random collection of lights, not unlike a constellation. As soon as the actor rises and starts moving, however, the observers instantly perceive that the lights are attached to an otherwise invisible human being. They are able not only to differentiate between walking and jogging movements but also recognise small anomalies in the actor's behaviour, such as the simulation of a slight limp.* When two people, similarly kitted out with lights, performed a folk dance, the identification was again immediate. Further experiments showed that when the actors put emotion into their movements, the observers could recognise that too.

Johansson's film provides two sorts of information. One is the rhythm of the movement. The other derives from the parameters within which the spatial relationships between the lights vary. The second ingredient is highly informative, but probably not absolutely necessary. Even the movement of a single light, I suspect, can be interpreted into a meaningful percept, or at least a reasonable hypothesis. Imagine unrelieved blackness on the cinema screen, implying a dark night. A single small light appears, bouncing up and down jerkily, and simultaneously swinging from side to side, and gradually grows a little larger. Any *aficionado* of classic horror films is surely going to suspect that this is a lantern being carried over rough ground, and will be prepared for the holder of the lantern to become visible as the light draws closer. A fairly similar erratic bouncy movement without the swing that indicates a flexible attachment between the light and its support might suggest a bicycle being ridden over a bumpy track. A light that wavered only a little as it approached would probably be interpreted as a bicycle on a smooth road - or perhaps as a flying saucer.

There are other circumstances where it may be possible to identify something by movement alone, although there's not so much as a hint of shape to be seen. When the leaves of a tree quiver and branches shake one can often guess roughly what sort of bird or mammal is causing the disturbance from the degree to which the branches bend, and from the rhythm of the movements, even though the thing itself remains invisible.

Finally, we have an accurate short-term memory for the feel of the movements we make ourselves. Immediately after making a gesture, or a sequence of gestures, we can recall the exact feel of doing it. When the movement results in stimulation to the vestibular organ the record includes that sensation as well. Patterns of rising, falling, rocking or twisting that are caused by external forces can also be replayed. Thus we come to develop schemata for the different styles of movement we can expect to be imposed on us when travelling in trains, boats, cars, escalators, or on different animals, and we have a schema for the feel of being on a swing. Having learnt to anticipate these effects we can adjust for them as they occur.

Echoic memory has a very short life, just long enough for several units from the even shorter-term memory which is the conscious moment to be put together, and appreciated as a whole. While it lasts it preserves the pattern of input from a sensory channel very efficiently. And the replay facility works even with sensory input that wasn't receiving attention at the point when it was registered. This is not only convenient in allowing you to fudge up a reply when the person you're ostensibly conversing with pauses expectantly, and you've been listening to the more interesting conversation going on behind you. It also means that attention can be switched to a stimulus pattern in retrospect, when it's only after the full pattern has been registered subconsciously that its attention-worthiness becomes apparent. (Very possibly that was how you came to be listening to the conversation behind you.) This indicates that incoming data can be preserved long enough for the pattern in the flow to

be registered at an unconscious level, regardless of whether attention is being paid or not, a valuable arrangement.

It's a necessary one too, since for stimulus patterns which manifest themselves over a period of time the matching-to-schema process can only begin after a certain amount of input has been registered. The individual units of a temporal pattern must be interpreted in the context of the whole pattern, and so decisions as to the identity of an individual unit may only be made after several units have been put together. A word containing an ambiguous consonant is heard in one way in one sentence, and as an alternative possibility in a different sentence - the choice made being the one that produces a more meaningful utterance. What we hear, when we listen to a familiar language, is whole words and whole phrases, rather than a succession of individual speech sounds.

The delay involved in creating the conscious experience is thus a positive benefit in listening to language and other extended series of inputs. It allows time for several units to be registered and fed through the processing system before the conscious experience emerges. This means that there's an opportunity for the interpretation of early phonemes to be influenced by what follows. So even with the very first phoneme of a sentence, perception can be influenced by context.

There's even more scope, of course, for the beginning of a sentence to influence how the later parts are interpreted. Languages take advantage of this to reduce the amount of perceptual work that is required in listening to them. There are always rules about acceptable combinations of sounds, and about the way words can be put together. Identifying one phoneme reduces the number of possibilities for the next one. And identifying a word reduces the number of potential following words, not only because of the grammatical rules but also because a sentence can be expected to convey some likely meaning. Experimental studies have shown that listeners develop a running hypothesis as a sentence unfolds. Under noisy conditions sentences that disobey accepted grammatical rules or express an improbable proposition, and so confound all hypotheses that the listener may have developed, are less likely to be heard accurately than more conventional sentences.

Sometimes, needless to say, the first hypothesis about a sentence is wrong. The last words arrive, and they just don't fit on to the supposed meaning of the first part of the message. Then it's necessary to adjust the hypothesis and re-check the identifications made of earlier words. This is another situation where the replay facility comes into operation, making it possible to match an ambiguous sequence of phonemes to a different schema, in the light of subsequent context.

Comparing memory and accumulating memory

As with the memory that produces the single conscious instant, the patterns perceived by means of the short-term memory are patterns of relationships. Music and spoken language demonstrate this. Few of us have perfect pitch, but most of us can learn to recognise a tune. Having done so, we can recognise it again in a different register, on another instrument, or in a different key; nor will a change of tempo confuse us, provided it's not too extreme. Indeed many of us, far from being confused by a change in key, timbre or tempo, take little notice of it. Clearly what we remember most readily is the relationship between the notes, not the absolute values. When Beethoven's fifth symphony begins with its famous *da da da BOM* what we essentially register is three notes of the same pitch and duration and a fourth that is lower, longer and louder.

Listening to language, too, is a matter of identifying relationships rather than absolutes. The records stored about speech-sounds are not precise specifications but comparisons. Vowels are characterised not by a precise shape but on a flexible scale of openness; their length is measured by comparing them to other sounds produced by the same speaker. As a result we manage to cope quite successfully with different accents and different speeds of delivery. We can understand both the very refined lady who says in clipped tones *Ai don't think ai quaike laike thet* and the Australian who drawls *Ah'll tike a plite 'v 'stike 'n chips*. When we encounter a new accent it takes a little while to understand it easily because we need a fair-sized sample of speech in order to make comparisons and work out the parameters of the sounds as they are used by this speaker. But once we have done that we are unfazed by the departures from whatever we happen to regard as normal pronunciation.

So, just as a visual object is a collection of features of various relative sizes, relative colours and relative distances in a certain spatial relationship, a tune is a sequence of notes measured in terms of higher and lower, shorter and longer, louder or softer. A movement, either seen, felt or performed, is a sequence of ups, downs, lefts, rights, nearers and furthers, along with the assorted compromises which constitute diagonals, all in relation to a fixed reference and combined with a sequence of variations in speed. Absolute values exist, obviously, but on the whole it takes training to register them accurately and remember them, whereas the patterns of relationships are learnt more easily. The schemata we build may well include limitations on absolute values, but the essence lies in the record of relative values.

Hardwired stimulus patterns can also be patterns extended over time. The flashes emitted by fireflies, for instance, are signals by which the females attract mates, and each species has its own distinctive pattern of flashing, so that where two or three different species inhabit the same area interbreeding is avoided. A comparable system exists among many species of frog. In the mating season male vocalisations summon the females, and the latter move only towards calls uttered at the correct pulse rate, distinguishing them both from the signals produced by other frog species, and from those produced by a male engaged in a disagreement with another male.

For the stickleback shape and colour is transformed into a mating signal with the addition of a special pattern of movement, and various other species use the same trick. The pigeon, for example, bobs along and bows repeatedly in his attempts to woo a female. Distinctive displays fulfil the same function as a change of male plumage in other species. In some birds the signals may be emitted by both sexes, and they work themselves up to the act with a sort of ritual dance.

Such extended signals imply something that could be termed a short-term memory. The efforts of the flashing firefly, the croaking frog and the posturing bird must work as a whole pattern, not just as a series of moments. But where the signal that must be registered is an innate one (as essential intraspecific signals always must be), the construction of a suitable memory is not too difficult. Imagine a species with a male mating call that goes *da da da BOM*. For the female to register it an auditory neuron which responded to *da* would be required, and another which responded to *BOM*. The *da* neuron would connect to an integrator neuron, which would fire only after receiving three inputs. The *da* integrator and the *BOM* neuron would in turn feed into a master neuron, and there would probably be inhibitory connections to prevent the excitation of this master neuron if the inputs arrived in the wrong order.

The master neuron might require several complete signals to be registered before the female began to move in the direction from which the sound came, and more calls would be needed for the approach to continue. If another male started to call and his call registered more strongly - whether because he had a louder voice or because he was nearer - she would divert towards him. And when a male and a female came together further signals would no doubt be exchanged, of a tactile or olfactory nature (no species can afford to be cavalier about mating matters) before the female accepted the male.

The neuronal wiring which registers such mating calls has in fact been worked out in several sorts of insect and various species of frog. It's invariably designed to require very precise timing of both the sounds and the pauses between them. If the rhythm is slightly wrong the integrator neuron doesn't fire, so there is little chance of a female being misled.

Such arrangements for registering extended signals exemplify an accumulating memory. The point to be stressed is that this sort of neuronal device works only for the genetically ordained pattern, and all other sequences of input pass it by. The memory that (within its capacity limits) can record any sequence of input the sensory channel is capable of receiving, and which is used to make comparisons among different patterns, is quite another matter. That sort of temporary storage arrangement clearly requires far more complicated neuronal engineering. And for mammals the neocortex provides this. Animals whose neocortex has been removed can make simple auditory

discriminations but cannot learn to distinguish changes in the temporal pattern of a tonal sequence. Damage to human auditory cortex may produce a similar deficiency.

Another piece of evidence that points to the same conclusion was provided by Sharpless and Jasper. They discovered that neurons very early in the cat's auditory pathway habituate to a given sound, but can be kicked back into full action if it changes in intensity. Neurons at the next level habituate similarly, but can respond to a change in pitch. Only in the neocortex does a repeated pattern of input produce habituation, which is relieved when the pattern varies.

We must conclude that the sustained activation of interlinked circuitry which creates the conscious moment persists to provide a longer-lasting record that preserves a succession of moments. This is what allows us to perceive very accurately an infinite variety of patterns of input that are extended over a brief period of time – to experience a stream of consciousness.

Echoic memories and sensory intelligence

The essential sensory schema - one that is developed without conscious effort - must deal in the sort of pattern that can be encompassed within the span of the accurate, replayable short-term memory. Larger echoic memories mean more complex schemata, and that in turn means a larger number of schemata, because more distinctions can be made. There must be considerable differences between the endowments of different species. We humans have an auditory echoic memory of particularly large capacity. Other species which produce long and variable vocalisations, such as various cetaceans and some songbirds, must be similarly equipped.

Monkeys and prosimians, meanwhile, must have especially large short-term memories for patterns of visual and performed movement, and for patterns of vestibular input. These provide the knowledge they need for their speedy and acrobatic travels through the trees, allowing them to match their jumps to the space that must be covered, maximise the propulsion to be gained from bendy branches, and tune their landings to the solidity or flexibility of the landing place. Youngsters must learn partly by experiment, partly by observing their elders, absorbing the rhythms and shapes of their movement patterns, and developing a vocabulary of what's possible. Similarly, those carnivores which pursue swift-running prey that jink and swerve must learn to time their final pounce exactly. We can deduce that they too have pretty good short-term memories with which to perceive the phraseology both of their own movement and of their prey's. The means of perceiving a fair-sized chunk of movement-pattern can be deduced to be essential to learning strategies of pursuit.

Then there are those species which are active mainly after dark, or live largely or entirely underground. It's likely that they have extensive short-term memories for whisker input - to which a great deal of their sensory cortex is devoted - since it's largely by means of this specialised tactile sense that they find their way around.

As well as major differences between species in the capacities of the various sensory memories, there must also be some variation between individuals of the same species. Amongst humans the guess must be that musicians, poets and actors are likely to have above-average auditory memories, painters above-average visual memories, and dancers particularly good memories for self-movement. Such artists can create patterns of a complexity which challenge the rest of us very slightly. They can also imitate patterns a little more easily than the rest of us. Learning a complex tune, or a series of ballet steps, or painting a realistic portrait, has to be that bit easier if slightly larger chunks of relationship can be perceived at one go.

Are such differences genetically determined? Probably they are, to a great extent. One of the benefits accruing to a social species is that all members of a group can profit, in many cases, from the special talents of a few, so many varied talents can flourish, and this is of course particularly true of *Homo sapiens*. But learning and practice are required to make the most of an individual's genetically determined potential. And since we generally enjoy practising the things we are good at more than struggling with the activities in which we cannot compete successfully, the bedrock of natural talent is difficult to distinguish from the effects of practice.

In any case, there are all sorts of difficulties about measuring these sensory short-term memories. They hinge on the vital question: what sort of units of measurement do we use? Some of the earliest experiments, concerned with the auditory memory, were done by the nineteenth-century psychologist Wilhelm Wundt. He used the simplest and most elementary of stimuli, the clicks of a metronome, running at around one per second. A subject would listen to a series of clicks, then, after a brief pause, a second series, which might be of the same length, or one click longer, or one shorter. It transpired that a series of eleven clicks could readily be distinguished from a series of ten. Counting wasn't necessary - it was just immediately perceptible as different. One might conclude that the short-term memory is sufficient for recording a minimum of ten or eleven separate stimuli; or perhaps that it can hold at least twenty odd seconds' worth of input. However, Wundt found that it's possible to impose a subjective pattern on the clicks, hearing them as groups of two, or groups of four. Doing this increases the size of the stimulus-trains which can be discriminated as similar or different quite substantially.

Using echoic memory to register a string of assorted stimuli is more demanding. Experimental subjects asked to repeat an arbitrary sequence of letters usually manage around seven, with the variation plus or minus two. But it's possible that reciting the list interferes with the echo.

However, if the subjects are asked to repeat an arbitrary string of words the score is still somewhere around seven, even though the number of speech-sounds is now noticeably greater. Moreover if the words make up a meaningful sentence the length of the input that can be perfectly reproduced increases considerably. It seems that if the stimulus patterns can be readily identified - matched to an existing schema - the amount of information that can be preserved in echoic memory is very much expanded. The count is the same for words as for single letters, and if the words are meaningfully linked it can be made in concepts rather than stimulus patterns. Perhaps echoic memory must be measured in matching-to-schema units. The conclusion to be drawn may be that the real capacity of echoic memory is always much the same, and its apparent capacity is determined by which stage of the sensory pathway is involved, whether it's the early stage that identifies speech-sounds, a later one that identifies words, or the complex of brain areas that makes sense of whole sentences.

The experience of learning to hear a foreign language provides a natural experiment in the use of echoic memory. First we have to build up schemata for the speech-sounds the new language employs and our native tongue doesn't, before we can begin to hear words properly. And only when we can hear the words clearly do we begin to decipher whole sentences, rather than guessing at their general drift. If it takes all your concentration to identify the unfamiliar speech-sounds then the useful capacity of echoic memory may be one phoneme. If you can identify the odd word, but only with effort, then the capacity is one word. When you attempt to replay a heard utterance many of the sounds remain ill-defined, and only odd words or fragments of words turn out to have been captured. This implies that when there are no schemata to which to refer the input the capacity of echoic memory is very small. Only when all the words are familiar can echoic memory encompass a normal-length sentence. The fact that there *is* a normal maximum length for sentences must say something about the typical echoic memory capacity in humans.

Measuring the span of time or input across which similarities attract notice might be one way of assessing echoic memory. How wide can the gap between two stimulus patterns be if a match is still to be immediately noticeable? For auditory memory a good clue may be provided by poets' use of rhyme. Clearly rhyme only works if the second occurrence of the sound arrives while the first is still echoing in the mind. Line 27 of a poem may rhyme with line 8 but nobody will notice. Rhyming couplets, on the other hand, are certainly unmissable. Rhymes that are two pentameters apart are pretty salient. Those separated by three pentameters, as in some sonnet forms, attract notice more delicately. Leaving a longer interval is rare, and the effect tends to be more of a hint than an attention-getter. Once the rhyme has been registered, however, it may recur five, perhaps even six lines away and still be noticed. But by then the repetition is not wholly dependent on echoic memory for recognition, since the first occurrence has given the sound a toehold in a slightly longer-term memory.

Relating the inputs of different senses

A further important benefit of sensory consciousness lies in enabling us to link the inputs of different senses. It's something we can do regardless of whether there is anything to be gained by it - not just when one input functions as feedback concerning the effect of responding to another slightly previous one. We can work out, for instance, that the song must come from the bird with the moving beak, or that the delicate scent goes with the shape and colour which is the rose. Multisensory schemata can be created simply for the interest of it. In many species associations between stimulus patterns registered by different senses only seem to be learnt if there is some profit to be gained, by way of predicting the appearance of a desirable or undesirable. Our ability to correlate different sensory inputs must depend on the existence of neocortical pathways for the senses concerned, on connections between those pathways, and on a potential for sustaining a circulation of impulses around those connections.

But it is of course the intersensory correlations which relate directly to the wellbeing of the self that are most important. With appropriate neuronal connections an animal can learn that an input in one sensory channel is frequently followed by a desirable or a noxious input to another sense. There's reason to think that the short-term record which equates with sensory consciousness extends the possibilities for discovering such predictive cues. Humans can learn to blink at just the right moment to avoid a small puff of air to the eye without being aware of either the puff of air or the blink, but only if the signal which warns of the unwelcome stimulus lasts long enough to overlap with it. If there is a brief gap between the warning cue and the puff of air only those people who notice both sensory events learn to make the avoidance response. This suggests that the process which produces the conscious sensory experience also sets up the record which allows the cue to be related to the unconditioned stimulus.

There's one form of learning, meanwhile, which clearly involves a conscious feedback, as I noted earlier. From the subjective perspective it's obvious that we humans can plan actions on visual objects because both movement and visual object are conscious experiences. We can tell whether the movement was successful by means of conscious vision and conscious touch. If it isn't, we perceive how it needs to be adjusted with the aid of conscious vision. We learn most, in fact, through our failures, and it's the picture-in-the-brain which provides a means of measuring the degree of failure. If the hand didn't quite reach the ball, or didn't quite achieve a lasting grasp, what went wrong? Here the classical arrangement is reversed. Instead of a contact sense providing feedback about the effects of responding to a stimulus registered by a distance sense, a distance sense provides feedback on an attempt at making a contact.

Other distance senses too can provide information about consequences, as when, for instance, an infant bangs two objects together and discovers they make an interesting noise - and promptly proceeds to do it again. Most importantly, it's the conscious auditory experience which enables an infant to compare the sounds she herself makes with the speech she hears around her, and gradually to adapt her vocal productions to the same style. We can conclude, in short, that sensory consciousness is what turns distance senses into sources of feedback.

Modifying motor behaviour through learning, meanwhile, requires not only feedback from distance senses about the results of an action. That needs to be correlated with information about muscle deployment, joint position and tactile inputs. The senses which register how limbs and joints are currently disposed (proprioception) must be long-established. But where all varieties of movement are fore-ordained by the nature of the muscles and the neurons that control them the proprioceptive data only needs to be applied with equal automaticity, to keep track of the position from which any new movement must start. Species whose actions are organised in this way begin life already capable of very efficient movement, or need only the briefest of 'running-in' periods for neurons and muscles.

Promoting proprioceptive information into the arena of consciousness must have been a vital factor in turning species with a predetermined repertoire of very effective hardwired actions into animals capable of more flexible, adaptable movements, which can be honed to suit a wide range of

tasks. This was the development that led to species where the young start out with wobbly, uncertain movements, and must learn to perfect them, developing their motor potential by practice and experimentation. The repertoire of movements is no longer limited, and the animal can learn to shape its movements to suit varied goals.

First, though, a young animal may have to learn how to use the feedback. The role of the picture-in-the-brain may possibly be relevant to one of the most enigmatic aspects of vision. A cause of much early pondering was the fact that the image formed on the retina is upsidedown, and reversed left to right. The maps in visual cortex also maintain this relationship to reality. Does this matter? From a modern perspective the answer might be assumed to be no. Neither the orientation of the retinal image nor of the brain's retinal maps would seem to be important to the complex processing of the information they provide, or to the parietal calculations about the location of the objects registered.

However, towards the end of the nineteenth century a classic experiment was performed to discover what happened if the image on the retina was the 'right' way up. G. M. Stratton devised a pair of goggles with prisms instead of lenses, which inverted the pattern of light entering the eye. Heroically, he wore these all day and every day for a period of weeks. And he saw, of course, an upsidedown world, not to mention an inverted self, with hands that descended into view from above. With practice he managed to adapt his movements to this new perspective. But then, after several days, he found that he was looking at the world the right way up again, despite the goggles. His picture-in-the-brain had changed to match the reality. Furthermore, when he finally left the prisms off his view of the world turned upsidedown again, and he needed a further period of readjustment.

The experiment has been repeated since, sometimes with up/down inversion, sometimes right/left, with varying results. Not everyone's conscious experience adapts. When it does it seems likely to be the result of relating visual input to the eye movements and arm movements that achieve desired goals. This implies an intriguing possibility - that a considerable element of early learning may go into the functioning of the 'where' pathway, just as into that of the 'what' pathway. Very young infants concentrate on very small bits of the visual scene - one fixation's worth at a time - and don't seem to deal in larger spatial patterns, so if they lack a proper sense of orientation it probably doesn't matter. The earliest eye movements and arm movements are stereotyped enough to indicate they are hardwired; it can be deduced that they are controlled by subcortical circuits, and can be performed without reference to the conscious 'picture'. Maybe the infant has to learn how eye movements change the visual input. And perhaps, as he looks at an attractive object and experiences proprioceptive feedback from the arm that is flung out towards it, he learns to construct his picture-in-the-brain in the correct orientation. In a species which uses consciousness to bring many strands of sensory evidence together, working out just how to do that may conceivably be the first step.

A fair amount of evidence supports the proposition that always seemed likely: as far as we mammals are concerned the sort of fullscale sensory consciousness we experience depends on the neocortex. (The neocortex can only work, of course, with the extensive support provided by its subcortical foundations.) The neocortex is a massively complex network, even in small mammals, and especially so in primates. There are numerous parallel pathways and complex interconnections among them. That makes it a very efficient machine for making correlations - enormous numbers of correlations. It's also very thoroughly connected to all the other areas of the brain. There are complex and lengthy pathways around which the neuronal impulses can circulate.

The sustained currents of impulses that create consciousness must be necessary to knit all the diverse activity together. At the same time modifications can be made to the circuitry at each pass, making the correlations increasingly delicate and more accurate. The benefits of these arrangements in extracting subtle and reliable information from the sensory data generally outweigh the disadvantage of extending the interval that elapses before a decision to act is made.

The process which creates the conscious moment, I've suggested, can be regarded as an ultra-short-term memory. The record it creates is incorporated into a more capacious short-term memory,

which allows a succession of sensory inputs to be assembled into a single percept, and equates with the stream of consciousness. That in turn allows the more definitive aspects of stimulus patterns to be preserved, as experience accumulates, in longer-term memory.

The information is stored in the later stages of the neocortical sensory pathways, in association areas of the cortex, or schema-areas. What's recorded are those aspects of the pattern which have attracted attention, which are generally those which will prove useful for recognising it again, or for recognising something of similar significance. Further encounters with similar stimulus patterns will, however, be able to modify the schema.

The record, reciprocally, helps to shape later conscious experience, as we saw in chapter 7, *Learning to See*. A two-dimensional arrangement of straight lines looks like a cube because we are familiar with cubes. An Ames room looks like a conventional set of walls joined at right angles because that's what we've learnt to expect. In short, sensory consciousness functions to bring together large amounts of sensory information, both information that comes from all sorts of different sensory channels and information that is spread out over time.

Harry Jerison's account sounds even better now than it did in 1973. *I regard the mind and conscious experience as constructions of nervous systems to handle the overwhelming amount of information that they process.* Or as Francis Crick and Christof Koch put it, with reference to vision, *the biological usefulness of visual consciousness in humans is to produce the best current interpretation of the visual scene in the light of past experience and to make this interpretation directly available - for a sufficient amount of time - to the parts of the brain that plan possible voluntary outputs of one sort or another.*

Operating without sensory consciousness

A modern animal that probably doesn't have much conscious vision is the frog. The study of this species' visual abilities began with a seminal paper by J. Y. Lettvin and colleagues in 1959, *What the frog's eye tells the frog's brain*.

Lettvin and co. found that the frog's eye is similar to ours in basic design, but the receptors are distributed evenly across the retina, and most of the ganglion cells receive input from a great many receptors, so that there's no area of special acuity. Furthermore the majority of the ganglion cells are geared to respond to moving contrasts, with preferences as to the direction of movement.

One type of ganglion is excited by moving stimuli with small, dark, tightly curved outlines. A stimulus pattern of this sort following a zigzag path nearly always means a flying insect, which equates with a potential meal. When such a stimulus is registered the frog unrolls its long tongue at ballistic speed, and with luck catches a fly on the sticky end of it and rolls the prize back into its mouth. Another variety of ganglion is sensitive to a larger darkness crossing the visual field, or a sudden dimming of the overall light level. Activation of these cells prompts the frog to make a jump, and avoid being trodden on by an animal which was oblivious to its presence. (Maybe you too have been startled, in fading light, when a bit of the ground you were about to tread on suddenly leapt away.)

What we have here, then, is a pair of stimulus patterns with specifications that are hardwired in the retina. Both of them depend on movement, which is a strength in one way, since plenty of things might be mistaken for a stationary fly, but not much moves like one. But a fly that stays still is safe from the frog's tongue, and if you notice a frog before its threat-sensitive ganglions register you it's fairly easy to creep up on it. Identifying movement in the outside world is pretty straightforward for the frog, incidentally, because no allowance has to be made for eye movements. The eye doesn't swivel in its socket, except to adjust for movement imposed on the animal by external forces, as when the lily pad it's sitting on is lifted by a ripple. Another sort of ganglion cell provides the information about global movement that allows the view to be stabilised.

The axons from the ganglion cells that respond to flying insects and to danger run not to the thalamus but to the optic tectum (a structure known as the superior colliculus in mammals). The target neurons here conserve the retinal relationships and form a map of the retinal field, in the usual fashion, so the trajectory of the stimulus is reflected by the neurons it excites. And from there the

onward route to motor neurons which trigger the fly-catching action or the escape jump is short. Consequently the response to the appearance of a fly can be very fast - which is just as well, since the fly has even shorter neuronal pathways from eye to muscle, and can manoeuvre very fast to escape. The rapidity with which the tongue-flip is initiated no doubt helps to explain why frogs are so much better at catching flies in midflight than we are. Visual consciousness, in this context, could only make for disastrous delay.

Frogs don't seem to be so good at navigating their way around. On several occasions I've seen one jump straight into a solid fence or wall. After bouncing off once the policy seemed to be to try again. When that failed they made a right-angled turn and set off in a new direction. However, frogs do possess a type of ganglion cell that responds to stripey textures, which provides some guidance for navigation. One species jumps towards horizontal lines, away from vertical ones, while another, which lives in grassland, prefers to jump towards vertical texture. Again, the axons run to the midbrain, and again, it's only moving stimuli which can engage the mechanism, which must be a little inconvenient for a species which gets around by jumping rather than running. However, frogs are light, tough, and bendable, so they don't damage easily.

There's yet one more type of ganglion in the frog's retina, not very numerous, which records the general pattern of light and darkness in terms of the blue wavelength. This seems similar to the mammal's blue-responsive ganglion, also uninterested in contrasts. It must be what enables the frog to head for shade when it needs to cool down, for sunshine and blue sky when it needs to warm up. This ganglion sends its message to the thalamus, and the hypothesis, of course, is that applying its information may perhaps involve a modicum of consciousness.

More recent research has focussed on toads, which have basically similar visual systems, but with variations. While frogs sit around waiting for their food to come to them, toads are hunters, seeking out worms, caterpillars and beetles, with the aid of retinal ganglion cells which are sensitive either to small blobs or to long thin horizontal objects. It takes movement to activate either sort of cell, but toads can seize a prey-sized object after it has stopped moving. They can also identify a motionless object revealed by their own movement. And although they have an escape reaction to larger moving objects they can learn in time to associate the moving object which is the experimenter's hand with the delivery of food - which must make them more convenient animals to maintain in the laboratory.

There are four or five different types of ganglion in the toad's retina, with different-sized receptive fields and different preferences as to the speed of movement, and the connections from retina to thalamus and optic tectum are much the same as in the frog. In addition, an area has been identified which contributes to distinguishing prey or predator from background, and movements in the external world from the effect of the animal's own movement. There are circuits for managing orientation towards the prey, and approach towards it, and a loop which seems to function to suppress the prey-seizing action when appropriate.

Many mammals also have some retinal ganglion cells with complex requirements. The rabbit retina, for instance, includes quite a lot of ganglion cells sensitive to the direction of movement. In some species even orientation is identified in the retina, and only a rather modest proportion of ganglions is responsive merely to small, unoriented contrasts, or to the presence or absence of light.

In many cases, moreover, the larger proportion of the axons from the retina goes not to the thalamus but to the superior colliculus. In rats many of the fibres divide, so there is a considerable overlap; over 90% of the ganglion axons project to the superior colliculus and about 30% to the thalamus. And the arrangement is much the same in the hamster, with the interesting twist that the upper part and the periphery of the visual field (the locations most important for spotting impending danger) are particularly well represented in the superior colliculus.

A likely deduction is that hardwired stimulus patterns are reported to the superior colliculus. Such patterns are pretty certain to include those that portend danger and require a rapid response. For some species they will also include the stimuli that represent food. And for many they seem likely

to include a particular combination of edges and perhaps colours, moving in a characteristic rhythm, all adding up to an adequate definition of a conspecific. The superior colliculus is organised in layers, and the number of layers varies considerably between species. It's possible that the number of layers is related to the number of visually directed behaviours that are hardwired.

Some mammals appear to be seriously handicapped by the removal of the superior colliculus, but get by quite well without a visual cortex - at least on the tests that experimenters have thought to try them with. It may be that the data from the ganglion cells with the less demanding requirements is reported to the lateral geniculate nucleus, and thence to the visual cortex, where it can be put together in flexible ways so that varied sorts of information can be extracted from it. That information can be used to guide the behaviours that can't easily be triggered by innately determined stimulus patterns, the most obvious of which is acquiring geographical knowledge of the environment. For many species the information that goes to the visual cortex is quite limited, and likely to be useful only for a very broad-scale mapping. There must be considerable support from senses such as touch and olfaction.

The rat's visual cortex, for instance, is reported to be largely concerned with overall patterns of light and dark, and much less concerned with featural detail than a monkey's or a cat's. There are orientation-sensitive neurons, but vertical and horizontal orientations are apparently more strongly represented than obliques. In the mouse the cortical neurons sensitive to the orientation of lines and edges are distributed randomly, but orientation-tuned cells are also found in the superior colliculus, where they are arranged in columns as in the primate visual cortex. One might guess, perhaps, that rodents are pretty good on visual textures, not so good on visual objects. They're probably more dependent on their whiskers for learning about shapes.

In us only a modest proportion of the optic fibres go to the superior colliculus - which is probably why it was at first thought to be responsible only for co-ordinating information relevant to eye movements, and for registering the sort of sudden movement that always attracts visual attention. But we have larger eyes than rats and rabbits, more densely packed with receptors and ganglion cells, so it's possible that the number of ganglion axons conveying information to the superior colliculus is no lower in us, merely overwhelmed by the vastly increased number communicating to the lateral geniculate.

Recapitulation

The process which creates the moment's conscious experience also establishes a brief, very accurate record of it, so that a succession of inputs is interpreted as a unit. The products of several fixations can be assembled, or the effects of a changing viewpoint, or the sensations that result from running one's fingers across an object.

This makes it possible to deal with the many sensory stimulus patterns that are inherently evanescent - with the flow of auditory input, and with movement, either seen or felt.

We can replay the sequence preserved in this short-term memory, even if the input only triggers attention after the full sequence has registered. It's consequently termed echoic memory.

Most sensory schemata are created out of such a succession of inputs. The echoic memory allows us to discover the relationships between the various components of a stationary or changing stimulus pattern. It can be termed a comparing memory.

The capacity of echoic memory is hard to measure. It no doubt differs somewhat in different sensory channels. It must vary considerably between species, and somewhat between individual members of a species, though it can be expanded by training. Its capacity must determine the maximum degree of complexity possible for a sensory schema.

Humans have evolved a particularly capacious auditory short-term memory as part of their language ability.

The fact that the conscious experience takes a good 50 milliseconds or so to create means that even the first syllable of a sentence can be interpreted in the context of what follows.

An important function of sensory consciousness is that it allows the percepts made in different sensory channels to be related to each other, regardless of whether there is any immediate profit in discovering the correlation.

A notable consequence is that distance senses can come to function as feedback senses.

Conscious visual feedback and a consciousness of muscle stretch, joint position and so on make it possible for some species to be born with a very rough-hewn repertoire of movements, and learn how to shape them into a wide variety of subtly differentiated ones.