Chapter 7 The cycle of thought

The brain is a spectacularly powerful biological computing machine, able to interpret the sensory world, move our bodies, communicate and navigate a vastly complex physical and social world far exceeds the power of anything yet created by artificial intelligence. But it is a very unconventional computer.

The power of the digital computers---including our familiar PCs, laptops and tablets---comes, largely, from carrying out many very simple calculation steps one after the other at phenomenal speeds: many billions of operations per second. By comparison, the brain is dreadfully slow. Neurons---the basic computational unit of the brain---calculate by sending electrical pulses to each other across hugely complex electro-chemical networks. The very highest rate of neural "firing" is about 1000 pulses per second; neurons, even when directly recruited to the task in hand, fire mostly at far more leisurely rates, from five and fifty times per second.[1] So neurons are painfully slow compared with the astonishing processing speed of silicon chips. But while neurons may be slow, they are certainly numerous. While a PC has one or at most a few processing chips, processing at a phenomenal rate, the human brain has roughly one hundred billion sluggish neurons, linked by roughly one hundred trillion connections.

So the spectacular cleverness of the human mind must come not from the frenzied sequences of simple calculations that underpin silicon computation. Instead, brain-style computation must result from *cooperation* across the highly interconnected, but slow, neural processing units, leading to coordinated patterns of neural activity across whole networks or perhaps entire regions of the brain.[2]

But it is hard to see how a vast population of interconnected neurons can coordinate on more than one thing at a time, without suffering terrible confusion and interference. Each time a neuron fires, it sends an electrical pulse to all the other neurons it is linked to (typically up to 1000). This is a good mechanism for helping neurons cooperate, as long as they are all focussing on different aspects of the same problem (e.g., building up different parts of a possible meaningful organization of a face, word, pattern or object). Then, by linking together, cross-checking, correcting and validating different parts of an organization (the parts of a face, the letters making up a word), it is possible gradually to build up a unified whole. But if interconnected neurons are working on entirely different problems, then the signals they pass between them will be hopelessly at cross-purposes—and neither task will be solved successfully: each neuron has no idea which signals it receives are relevant to the problem it is working on, or are irrelevant junk.

---

[1] A classic discussion is Feldman, J. A., & Ballard, D. H. (1982). Connectionist models and their properties. Cognitive science, 6(3), 205-254.

[2] This connectionist or "neural network" models of computation have been a rival to conventional "digital" computers since the 1940s. McCulloch, W. S., & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. The bulletin of mathematical biophysics, 5(4), 115-133, and exploded into psychology and cognitive science with books including Hinton, G. E., & Anderson, J. A. (1981). Parallel models of associative memory. Hillsdale, NJ: Erlbaum and McClelland, J. L., Rumelhart, D. E., & the PDP Research Group. (1986). Parallel distributed processing (2 Volumes). Cambridge, MA: MIT press. State-of-the-art machine learning now extensively uses neural networks---although, ironically, implemented in conventional digital computers for reasons of practical conventions. Building brain-like hardware is current just too difficult and inflexible.

So we have a general principle. If the brain solves problems through the cooperation computation of vast networks of individually sluggish neurons, then any specific network of neurons can work on just one solution to one problem at a time. To a rough approximation, the brain seems to be pretty close to one giant, highly interconnected network (although the connections between different regions of the brain are not equally dense). So we should expect, then, that a network of neurons in the brain should be able to cooperate on just one problem at a time.

This gives us the beginning of an explanation of the slow, step-by-step, nature of perception and thought that we met repeatedly in Part I: our ability to process one word, face, or even colour, at a time. Attending to a set of information to be explained is *setting the problem* that the brain is to solve. The problems can be very diverse: find the "meaning" in this pattern of black and white shapes; figure out what is being said in this stream of speech; visualize a cube balanced on one point; recall the last time you went to the cinema, and so on. We can think of this as involving specifying the values of some subset of the population neurons. Each step in a sequence of thoughts then involves cooperative computation to find the most meaningful organization of everything else we know to find an answer that best fits that "question." A single step may take many hundredths of a second---compared with the many billions of operations per second of a conventional computer, this might seem painfully slow. But a single computational step for a networks of billions of neurons may represent the integration of a vast amount of knowledge---so while each neuron may be sluggish (in comparison with the frenzied pace inside a silicon chip), the computational power achieved by drawing on the knowledge and processing power across a network of billions of neurons can be enormous.

The computational abilities of the brain are, then, both severely limited and remarkably powerful. The problem of interference implies that the cycle of thought is limited to proceeding one step at a time, and working on just one problem at a time. Yet by cooperatively drawing on a vast population of interconnected neurons, each contributing only a little to the overall solution to the problem in hand, each step in the cycle of thought may answer questions of enormous difficulty: decoding a facial expression, predicting what will happen next in a complex physical and social situation, integrating fast flowing input of speech or text, planning and initiating the spectacularly complex sequence of actions to return a tennis serve thundering down more than 100 mph. Each of these processes would, to the extent that can ever successfully be simulated on a conventional computer at all, correspond to millions or even billions of tiny steps, implemented with almost unimaginable speed, one after the other. But the brain takes a different tack: its slow neural units split up the problem into a myriad tiny fragments and share their tentative solutions in parallel across the entire, densely interconnected, network.

What is important is the very fact that the brain uses cooperative computation across vast networks of neuron implies that the brain networks---and indeed calculates one giant, coordinated, step at time, rather than, as in a conventional computer, through a myriad of almost infinitesimally tiny information processing steps. I shall call this sequence of giant, cooperative steps, running at an irregular pulse of several "beats" per second, *the cycle of thought.*[3]

---

[3] While the brain is interconnected into something close to a single network, this isn't quite the whole story. As with a PC, the brain seems to have some specialized hardware for particular

The comparison between conventional computers and the brain can, therefore, be highly misleading. We can write a document, or watch a movie, on our PC while, in the background, it is searching for large prime numbers, downloading music, crunching away on astronomical calculations, or any number of tasks, in the background. So as our conscious mind is focussed on making breakfast or reading a novel, all manner of profound and occult thoughts might be running along below conscious awareness, seems natural enough. But brains are very different from conventional computers---rather than being able to time-share a super-fast central processor, the brain works by cooperative computation across most or all of its neurons---and cooperative computation can only lock on to, and solve, one problem at a time.

We can sum up two important implications of the cooperative style of brain computation in slogan form (1) *one problem for one network* and (2) *no background processing in the brain.* Let's consider these in turn.

One problem for one network. Each network of cooperating neurons in the brain can only solve one problem at a time. This leaves open the possibility that the brain may naturally be divided into many independent networks; that it may function, instead, largely as a single interconnected system; and that it may be possible for the brain proactively to reconfigure itself into distinct networks to address the challenges in hand as effectively as possible. But whatever the nature of the division of the brain into cooperating networks of nerve cells (and we shall briefly touch on the question of how flexibly this can be done later), the key point is that each cooperative network of neurons can address precisely one problem at a time. And, in practice, the brain appears to be highly interconnected—so multitasking will be the exception rather than the rule.

No background processing in the brain. The brain is highly interconnected – whatever we are consciously attending appears to engage large areas of the brain. This implies that there will typically be severe interference between any two tasks or problems that engage conscious attention. So this means not only that we can only consciously attend to one problem at a time; but also that if we are consciously thinking about one problem, we cannot *even unconsciously*, be thinking about another—because the brain networks involved would be likely to overlap. Tasks and problems needing conscious attention engage swathes of our neural machinery; and each part of that machinery can only do one thing at a time. In particular, the unconscious cannot be working away on, say, tricky intellectual or creative challenge while we are consciously attending to some other task---because the brain circuits that would be needed for such sophisticated unconscious thoughts are "blocked" by the conscious brain processes of the moment. We shall return to the far-reaching implications of 'no background processing' below, leading us to revise our intuitions about "unconscious thoughts" and "hidden motives," and eliminate any suspicion that our behaviour is the product of a battle between multiple selves (e.g., Freud's id, ego, and superego).

---

problem, such as the "low-level" processing of images and sounds, and other sensory inputs, and planning basic movement control. And perhaps there are more somewhat independent networks specialized for other tasks too (e.g., processing faces, words, speech sounds). The questions of which "special-purpose" machinery the brain develops, whether such machinery is built in or learned, and, crucially, the degree to which such networks are "sealed off" from interference from the rest of the brain, are all of great importance.

So the brain is a set of highly interconnected, and cooperative, networks; but the *structure* of these networks turns out to particularly revealing. There is, it turns out, a relatively narrow neural bottleneck through which sensory information flows. This bottleneck restricts the possibility of doing many things at once---but, as we shall see, it also gives some fascinating clues about the nature of conscious experience.

Prodding the conscious brain

The eminent neurosurgeon, Wilder Penfield, pioneered brain investigations and brain surgery on people who were wide-awake.[4] From the patient's point of view, a little local anaesthetic to deal with any pain of the incision through the skull is all that is required. Though the brain detects pains of many and varied types throughout the body (pokes, abrasions, twists, excesses of heat and cold), the brain has no mechanism for detecting damage to itself. So Penfield's brain operations were entirely painless.

The purpose of Penfield's surgery was to relieve severe epilepsy by attempting to isolate, and remove, portions of the brain from which the seizures originated. In an epileptic attack, the cells across large areas cease any complex cooperative computation to solve the problem of the moment and start instead to 'fire' in slow synchronized waves and hence, being entrained by each other, become disengaged from their normal information processing function. A slightly fanciful parallel is to imagine the population of a busy city suddenly dropping their varied and highly interconnect activities (buying, selling, chatting, building, making) to join in a single, continuous, coordinated, but entirely involuntary, Mexican wave---wherever the wave spreads, work will come to a complete halt. In severe epileptic attacks, the whole, or large areas of, the cerebral cortex become entrained and hence entirely non-functional , until the brain is somehow able to reset itself; and severe epileptics can suffer such debilitating attacks many times each day. The entrainment in epilepsy typically starts in a specific region of the cortex. It is as if the inhabitants of one district are particular prone to spontaneously launching into Mexican waves---and the nearby neighbourhoods are then inexorably drawn in, and the wave spreads inexorably across the city. Penfield's logic was that, if only the troublesome district could somehow be isolated from the rest of the city, the Mexican wave would be unable to spread---and normal life would resume unhindered. In practice, Penfield found that most effective treatment often turned out to be rather extreme: rather than a few subtle surgical cuts to key regions of the cortex, the equivalent of closing a few bridges or main road, Penfield was often driven to remove large regions of the cortex, analogous to completely flattening often huge areas of a city.

[4] Penfield, W. & Jasper, H. H. (1954). *Epilepsy and the functional anatomy of the human brain*. Little, Brown.

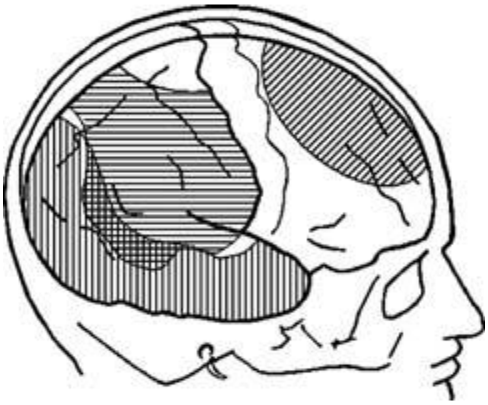

Figure 7-1. The areas of cortex removed in three of Wilder Penfield's epileptic patients, shown on a single diagram. The area removed to the front of the brain is here shown on the right hand side for visibility, but was actually on the left, in mirror-image.[5]

I can scarcely imagine anything more alarming than having areas of one's brain removed with nothing more than local anaesthetic. But the use of local anaesthetic turned out to be crucial: electrically stimulating different parts of the cortex, the awake patient was able to communicate a good deal of information about the function of that area; some areas of the cortex are much more crucial than others, and Penfield was therefore able to operate, as far as possible, without inadvertently causing paralysis, or loss of language.

Penfield was, therefore, in the uncanny position of being able to converse with patients throughout the period that substantial areas of their brains were being removed. Perhaps rather disconcertingly, the patients not only maintained consciousness throughout, but typically maintained a fluid conversation, neither exhibiting, nor reporting, any disruption to their conscious experience.

One might conclude that perhaps these areas of cortex, and perhaps the cortex more generally, is simply not relevant to conscious thought. But a wide range of other considerations ruled out this possibility. For example, recall that the visual neglect patients we described in Chapter 2 had no conscious experience of, frequently, one half of the visual field--- and the area of absence of conscious processing maps neatly onto the area of visual cortex that has been damaged. Similarly, patients who have, often through a localised stroke, suffered damage to cortical regions known to process colour, motion, taste, and so on, show the consequences of such damage in their conscious experience: they no longer perceive colour normally, may perceive the world as intermittently frozen and moving jerkily, or report they have lost their sense of taste. In short, the cortical processing machinery we possess seems to map directly into the conscious phenomenology we are able to experience. And Penfield himself was able to collect new and very direct evidence of the close connection between cortex and consciousness, through his electrically stimulation of difference parts of cortical surface. Such stimulation often did intrude on conscious experience, and in very striking ways. Depending on the area and the stimulation, patients would often report visual experiences, sounds, dreamlike strands, or even what appeared to be entire flashes of memory

[5] Reprinted by permission from Merker, B. (2007). Consciousness without a cerebral cortex: A challenge for neuroscience and medicine. *Behavioral and Brain Sciences*, *30*, 63–134. [get permission]

(most famously, one patient volunteered the strangely specific "I can smell burnt toast!," in response to a particular electrical probing of the brain).

So how could it be that stimulating a brain area leads to conscious phenomenology, but removing that very same area brain appears to leave conscious experience entirely affected? Penfield's answer was that consciousness isn't located in the cortical surface of the brain, but in the deeper brain areas into which the cortex projects---in particular, to the collection of evolutionarily old "sub-cortical" structures that lie at the core of the brain and which the cortex enfolds.

The anatomy of the brain gives, perhaps, an initial clue. Sub-cortical structures (such as the thalamus and the superior colliculus) have rich neural projections that fan out into the cortex that enfolds it; and these connections allow information to pass both ways. Intriguingly, most information from our senses passes through the brainstem before projecting on to the cortex; and information passes through to deep subcortical structures to drive our actions. So what we may loosely call the 'deep' brain (the relevant anatomy is too complex to outline here; and the functions of different components are both intricate and contested) is serves as a relay station between the sensory world and the cortex; and from the cortex back to the world of action. Here, perhaps, somewhere in these deep brain structures is the crucial bottleneck of attention; and whatever passes through the bottleneck is consciously experienced.

Penfield's viewpoint was more recently elaborated and extended by Swedish neuroscientist Bjorn Merker.[6] Merker highlights a number of further observations that fit with Penfield's suspicion that conscious experience requires linkage between diverse cortical areas and a narrow processing bottleneck deep in the brain.

If conscious experience is controlled by the deep, sub-cortical, brain structures, for example, then one might expect those structures to control the very presence of consciousness: in particular, to act as a switch between wakefulness and sleep. Indeed, such a switch does appear to exist--- or, at least, electrically stimulating highly localised deep brain structures in animals (specifically parts of the midbrain and the pons) can lead to sharp reduction in activity across the whole of the cortex. The animal lapses into a quiescent state.[7] Moreover, if this brain area is surgically removed, the animal is comatose---as if unable to wake up. By contrast, wakefulness is unaffected in either animals or humans by removal of large areas of cortex.

Might temporary disruption deep in the human brain cause the switch momentarily to be thrown---so that consciousness might cease abruptly perhaps just for a matter of seconds or minutes? Following Penfield, Merker points out that so-called 'petit mal' or absence epilepsy seems to have just this character. A petit mal episode involves a person, during everyday activity, suddenly adopting a vacant stare and become completely unresponsive to their surroundings. If the patient is walking, she will slow and freeze in position, while remaining upright; if she is speaking, the speech may continue briefly, although typically slowing, and ceasing entirely; if eating, a forkful of food may hover in space, between the plate and the mouth. Attempts to rouse the patient during a period of "absence" are usually ineffective, although they can, on occasion, cause the patient to "awake"

[6] Merker, B. (2007). Consciousness without a cerebral cortex: A challenge for neuroscience and medicine. *Behavioral and Brain Sciences*, *30*, 63–134.
[7] Moruzzi, G. & Magoun, H. W. (1949) Brain stem reticular formation and activation of the EEG. Electroencephalography and Clinical Neurophysiology 1: 455–73.

suddenly. Usually, though, consciousness returns spontaneously often within seconds. The sufferer typically has no immediate knowledge of having suffered an epileptic episode---conscious experience appears, bizarrely, to pick up uninterrupted from where it left off, as if time, from the point of view of the patient, has stood still. Patients have, in particular, complete amnesia during the period of "absence."

Recordings of the electrical activity in the cortex during periods of "absence" show the typical slow-wave pattern--- but this pattern appears to synchronise simultaneously across the cortical surface from the very beginning of the episode, rather than, as with many other forms of epileptic seizure, propagating from one area to the next. It is if a Mexican wave were to begin across the entire city, rather than rippling from one district the next--- and this would suggest some external communication signal, perhaps a radio broadcast, simultaneously instructing the population to act in unison. Penfield's conclusion is that the deep subcortical brain structures, with their rich fan of nervous connections to the cortical surface that enfolds them, play just this role--- perhaps, indeed, the very structures that, when stimulated electrically, seem to shut down the cortex in animals. Indeed, one of Penfield's closest collaborators, Herbert Jaspars, showed, in the 1940s, that electrical stimulation to a deep brain structure, the mid-line thalamus, in the cat produced an analogue of absence epilepsy 'on demand.'[8] So the "switch" of consciousness, which can shut the cortex down, or bring it back to life, seems to lie in the deep brain regions to which the cortex is densely connected, rather than being part of the cortex itself.

A further clue from Penfield's neurosurgical investigations comes from the results of Penfield probing with electrical stimulation as he operated. Such electrical stimulation would frequently trigger epileptic episodes of many and varied kinds. Patients were, after all, suffering epileptic attacks of the most extreme and frequent form--- or they would never have been referred to Penfield for such radical surgery---and so it is not surprising that their brains were readily triggered into an epileptic state. Yet Penfield reports that the one form of epilepsy that was never induced, whichever region of the cortex might stimulated, was "absence" or petit mal epilepsy--- no electrical stimulation of the cortex itself could trigger the crisp, immediate, total shutdown of the cortical system, because, he suggested, the "switch" of consciousness lies not in the cortical surface, but deep within the brain.

When we think of the human brain, and the astonishing intelligence it supports, we conjure up an image of the tightly folded walnut-like surface of the cortex, lying just beneath our skull. Indeed, in humans, the cortex is of primary importance: while, in many mammals, such as rats, the cortex is of rather modest size in comparison with other brain regions, in primates, such as chimpanzees and gorillas it dominates the brain, and is expanded particularly spectacularly in humans. But the cortex receives its input, and sends its output, through deeper, subcortical brain structures---and these subcortical structures may determine the contents of the 'flow' of consciousness, and, indeed, whether we are conscious at all.

To see how this might work, the link between perception with action may be especially illuminating. Suppose that we are picking fruit from a tree. Our brain needs to alight on the next fruit to be picked, determine whether it sufficiently ripe and not yet rotten, perhaps identifying it through a veil

[8] Jasper, H. H. & Droogleever-Fortuyn, J. (1947) Experimental studies on the functional anatomy of petit mal epilepsy. Proceedings of the Association for Research in Nervous and Mental Disease 26:272–98

of foliage; and needs to plan a sequence of movements which will successfully grasp the fruit and twist it from its stem. Or, in the case of a person, the action might consist merely of suggesting that someone else pick the fruit, or describing it, perhaps inwardly. But, in any case, it is crucial that the action is connected to the visual input concerned with this particular piece of fruit, and that the different pieces of the visual input (the different fragments of the fruit visible through foliage, perhaps) are integrated into a whole. If we reach out to the fruit and touch it, then information about the location of our arm, and the sensations as we brush through the leaves and then feel the surface of the fruit we grasp it must be linked up with our visual input (or so that we know that we are reaching and grasping the very fruit that we are looking at). And all of this information must be linked, in turn, to memory: of our earlier decision to pick apples (and, perhaps, only particularly ripe apples), our memory of past visual experience, which allows us to identify apples, leaves and branches; and we may, in turn, be reminded of apple picking incidents in our childhood, or any other myriad of agricultural or biological facts about apples, and so on. And the action reaching and grasping an apple is, itself, potentially complex, requiring not just coordination of a single arm and hand, but of course, potentially stretching actions, standing on tiptoe is if necessary, and making a suite of postural compensations to maintain balance.

Now actions occur, roughly, one at a time; but each action required integrating potentially a large amounts of information from the senses, from memory, and from our motor system. So perhaps the role of one or more structures deepen our brains is to serve as the focal point for such integration; to probe the diverse areas of the surrounding cortex, devoted to the processing of sensory information, the storage and retrieval of memory, or the control of movement, and to bring them to bear on the same problem. But the sequential nature action would imply sequential nature of our focus; and this might, perhaps, underpin the limited, unitary, flow of conscious experience.

If this is right, then electrical stimulation to an area of cortex might very well be expected to intrude conscious experience: connections between cortex and deep brain areas are two-way. A sudden surge of activity in a particular area of the cortex may lead to signals into deep brain areas which disturbs, and even overrides, their current activity, generating, for example, strange sensory experiences and fragments of memory. But the complete disappearance of an area of cortex, unless it happens to be engaged directly in some current mental activity, will pass entirely unnoticed, without even the merest ripple in conscious experience.

This perspective explains, too, why patients with visual neglect, where cortex corresponding to a large area of the visual field may be damaged or entirely inoperative, can nonetheless be entirely unaware of their deficit. We are consciously aware, perhaps, only of the specific perception-action task on which we are currently focused. So, if engaged in fruit picking, a person with visual neglect will focus their attention only on visual information in parts of the visual cortex which are intact, and link with memory and action systems through the coordinating power of structures deepen the brain, just as a person with normal visual processing. There conscious experience may, therefore, be entirely normal. They will not, of course, pick or describe, fruit whose visual positions projects into the ‘blind’ area of visual cortex---so their visual phenomenology, while entirely normal moment by moment, will be restricted to, say, the left half visual field.

So our brain is fully engaged with making sense of the information it is confronted *at each moment*. Consciousness, and indeed, the entire activity of thought, appears to be guided, sequentially, through the narrow bottleneck---deep, subcortical structures, search for, and create, patterns in sensory input, memory, and motor output, one at a time. The brain's task is, moment by moment, to link together different pieces of information, and to integrate and act on them right away. Our brain will, of course, lay down fresh memories as this processing proceeds; and draw on the richness of memories of past processing.

So our *no background processing* slogan is reinforced. Or, at least, if there are brain processes which are scurrying about behind the scenes, contemplating, evaluating, and reasoning about matters that we appear not to be thinking about at all, then neuroscience has found no trace of them. The brain appears, instead, primarily to be concentrating on making sense of immediate experience, and generating sequences of actions, including language (whether spoken aloud or inner speech), through the narrow bottleneck of conscious thought. The deep subcortical "crucible" of thought can, we may assume, integrate and transform the entire gamut of information represented by the cortex; but, to return to our first slogan, it can only integrate and transform to solve *one problem at a time*.

So we have some tentative answers to how the cooperative style of the brain computation operates. In the Penfield/Merker vision of the brain, both the questions that the brain faces and the answers that it provides are represented in subcortical structures, including the thalamus and superior colliculus, which serve as relay stations between cerebral cortex and the senses and our motor systems---essentially as a gateway the hemispheres of the brain and the outside world. And, we might suspect, both the questions and their answers are concerned primarily with the organisation of sensation and movement; and the rich interconnections between these structures and the cortex provide the networks of cooperative computation that are able to solve the problems posed by these subcortical structures. Yet while the cortex is crucial in processing visual information, planning movement, and drawing on memories, we are, it appears, aware of the results of the vast cooperative enterprise across the brain only in so far as the results of such compiutations reach the subcortical "gateway" structures---these, and not the cortex itself, are the locus of conscious experience.

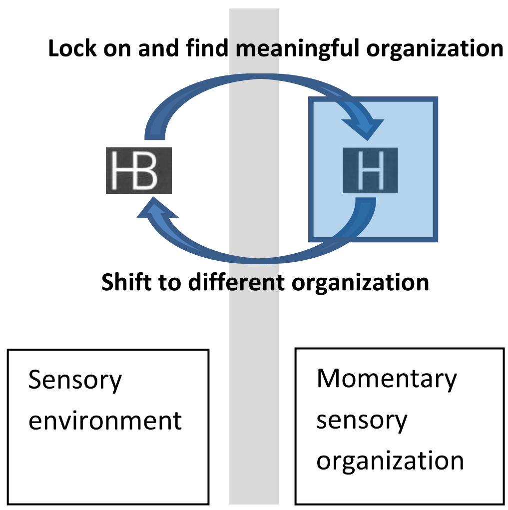

Figure 7-2. *The cycle of thought*. Upper arrow: The brain locks on to and organizes a fragment of the visual stimulus; we are conscious of, and can report, this organization. Lower arrow: But the brain---and the eye---is restlessly struggling to break free of its current organization and lock on to a different fragment of the image. This cycle is so rapid and fluent that we can have the sense of awareness of a complex object---or even an entire scene, in full detail and colour. Our stream of consciousness is of the successive sensory organizations---consciousness is entirely confined to the blue shaded box: we have no conscious access to the information being picked up by our senses (left hand side of the figure) or how that information is interpreted (blue arrows); or how our brain shifts to lock on to different information, e.g., by shifting our attention (with or without moving our eyes).

Four principles of the cycle of thought:

So here is an outline of how the mind works, in the form of a picture (Figure 7-2) and four proposals:

1. ***Attention as interpretation***. At each moment, we "lock on" to (or, to use the psychologist's vocabulary, attend to) a set of information, which the brain then attempts to organize and interpret. The target might be aspects of sensory experience, a fragment of language, or a memory.[9] Figure 7-2 illustrates the case where the brain has momentarily locked onto the "H" in complex stimulus. Moments later, it might alight on the "B." Crucially, our brain locks on to one target at a time---the organization of just one target can pass through our mental bottleneck at any moment. This implies that the line shared between the "H" and "B" in the stimulus is, at any given moment, interpreted as belonging to one or the other---but not both. Following Penfield and Merker's conjectures, we might suspect that such information is represented in subcortical structures buried deep in the brain (e.g., the thalamus and superior colliculus), with connections across the entire cortex, so that the full range of past experience and knowledge can be brought to bear on finding meaning in the current target.

Remember that we can lock on to, and integrate, information of all kinds. We can use all and any pieces of information and great steps of ingenuity and imagination to find meaning in the world; but we can only create one pattern at a time.

2. ***Our only conscious experience is our interpretation of sensory information***. The *result* of the brain's "interpretations" of sensory input is conscious---we are aware of the brain's interpretation of the world. But the 'raw materials' from which this interpretation is constructed, and the process of construction itself, are not consciously accessible. (Conscious experience is the organization represented in deep brain structures, with inputs from across the cortex—we are not directly conscious of cortical activity itself). So, in Figure

[9] Psychologists and neuroscientists will recognize these ideas as drawing on a range of prior ideas, from the emphasis on organization in Gestalt psychology, and Bartlett's "effort after meaning" in human memory, to Ulric Neisser's perceptual cycle, the vast range of experiments on the limits of attention, to O'Regan and Noe's theory of consciousness, to the astonishing results from Wilder Penfield's early experiments in brain surgery and Bjorn Merker's theorizing about the central role of "deep" (sub-cortical) brain structures in conscious experience. My own attempt, to lock on to, and organize, these and other findings and ideas into a cohesive pattern probably doesn't correspond precisely to any previous theory, though it has strong resemblances to many.

7-2, we perceive the “H” or the “B,” but have absolutely no awareness of the process through which these were constructed.

Perception always works this way: we “see” objects, people, faces given a pattern of firing from light-sensitive cells triggered by light falling on our retina; we hear voices, musical instruments, and traffic noise based on picking up the complex patterns of firing of vibration-detecting cells in our inner ear. But we have no idea, from introspection alone, where such meaningful interpretations spring from---how our brain makes the leap from successive waves of cacophonous chatter from our nervous system to a stable and meaningful world around us. All we ‘experience’ is the stable and meaningful world: we experience the result, not the process.[10]

I claim that we have absolutely no introspective access to the information on which these meanings are based, or to the process by which our brain reaches these conclusions.

Of course, we can sometimes find little meaning in what we are looking at or thinking about---even so, our conscious experience reflects whatever partial and contradictory patterns we can find; we can never ‘confront’ the raw sensory input.

So far, the account has focussed on consciousness of meaning in sensory information. The third claim is that there is nothing else: *all conscious thought concerns the meaningful interpretation of sensory information.*

3. **Conscious experience is exclusively the meaningful organization of sensory information**. We have no conscious experience of non-sensory information, although we may be conscious of their sensory ‘consequences’ (i.e., I have no conscious experience of the abstract number 5, although I may conjure up a sensory representation of five dots, or the shape of symbol “5”). Deep brain regions are, after all, relay stations for conveying sensory information to the cortex---so if they are locations of conscious experience, then we should expect conscious experience to be sensory experience, and nothing more.

The claim that we are aware of nothing more than meaningful organization of sensory experience isn’t quite as restrictive as it sounds. Sensory information need not necessarily be *gathered* by our senses, but may be invented in our dreams or by active imagery. And much sensory information comes, of course, not from the external world, but from our own bodies---including many of our pains, pleasures, sensations of effort or boredom. We are conscious of the words we use to encode abstract ideas; or imagery which accompany them. But we are not conscious of the abstract ideas themselves, whatever that might mean. I can imagine (just about) three apples or the symbols “3,” “iii” or “three”; and I can imagine various rather indistinct triangles and the word “triangle.” But I surely can’t imagine, or in any sense be conscious of, the abstract number 3, whatever that might mean; or the abstract mathematical concept of ‘triangularity.’ I can hear myself say “Triangles have 3

[10] Indeed, precisely because we see only the stable, meaningful world, and have no awareness whatever of the vastly complex calculations our brain is engaged in to weave meaning out of nervous stimulation, new-comers to psychology and neuroscience are often surprised that the brain even needs to make such calculations. We can imagine that the world merely presents itself, fully interpreted, to the eye and ear. Yet the opposite is the case: about half of the brain is dedicated, full-time, to what is fairly uncontroversially agreed to be perceptual analysis. But, as we shall see, the reach of perception may be greater still.

straight sides" or "The internal angles of a triangle add up to 180 degrees"---I surely don't have any additional conscious experience of these abstract truths.

Similarly, as we have seen already, it is a mistake to think that we are conscious of any of our beliefs, desires, hopes or fears. I can say to myself "I'm terrified of water" or I can have visions of myself struggling desperately as rip-tide pulls me out to sea. But it is words and images that are the objects of consciousness---not the "abstract" belief. Just in case you doubt this viewpoint, reflect on what beliefs you are conscious of *right now*; how many are there, exactly; can you feel when one belief leaves consciousness or a fresh belief 'comes into mind?' I suspect not.[11]

So we can tie together our three proposals, into a fourth proposal about the stream of consciousness:

4. ***The stream of consciousness*** . A thought is the process of the creation is a meaningful organisation of sensory input. The stream of consciousness is of a succession of thoughts: an irregular cycle of experiences which are the results of meaningful organisation of sensory input---the shifting contents of the right-hand box in Figure 7-2. (In the Penfield/Merker story about the brain, sub-cortical structures deep in the brain form a 'crucible' onto which the resources of the whole cortex can be focussed to impose meaning on fragments of sensory information---but only one pattern can be placed in the crucible at a time).

Note, in particular, that the cycle of thought is sequential: we lock on to and impose meaning on one set of information at a time. Now of course your brain can control your breathing, heart-rate and balance independent of the cycle of thought, to some extent at least (we don't topple over when particularly engrossed in problem). But the brain's activities beyond the sequential cycle of thought are, we shall see, surprisingly limited---we can manage, roughly speaking, just one thought at a time.

From this point of view, many of the strange phenomena we saw in Part I fall into place:

- The brain is continually scrambling to link together scraps of sensory information to which we lock on (and the ability, of course, to gather more information, with a remarkably quick flick of the eye). We 'create' perception of an entire visual world from a succession of fragments, picked up one at a time (Chapters 2). Yet our conscious experience is merely the output of this remarkable process; we have little or no insight into the relevant sensory inputs or how they are combined.
- As soon as we query some aspects of the visual scene (or, equally of our memory), then the brain immediately locks onto relevant information and attempt to impose meaning upon it. The process of creating such meaning is so fluent that we imagine ourselves merely to be reading off pre-existing information, to which we already have access, just as, when scrolling

[11] The question of whether we have so-called imageless thoughts was hugely controversial, very early in the development of psychology. Otto Külpe (1862-1915) and his students at the University of Würzburg famously reported that they experienced ineffable and indescribable states of awareness when thinking about abstract concepts. These mysterious experiences, supposedly lacking any sensory qualities, were viewed of great theoretical significance by Külpe. Other early psychologists, including the British psychologist, Edward Titchener (1867-1927), who had studied in Germany and set up a laboratory at Cornell University in upstate New York, reported that they had no such experiences. Perhaps remarkably, the resulting transatlantic controversy shook the psychological world. I, for one, have no idea what it would be like if I did have an impalpable non-sensory experience, any more than I know what it would be like to see a square triangle.

down contents of word processor, or exploring a virtual reality game, we have the illusion that the entire document, or labyrinth or the monsters we may encounter, pre-exist in all their glorious pixel-by-pixel detail (somewhere 'off-screen'), rather than being created for us by the computer software at the very moment we needed (e.g., when we scroll down or 'run' headlong down a virtual passageway). In particular, as soon as we wonder about the colour, or detail, of an object, our eyes lock onto that object, and fill in the answer, given all the available scraps of information. This is the sleight of hand that underlies the Grand Illusion (Chapter 3).

- In perception, we focus on fragments of sensory information and impose what might be quite abstract meaning: the identity, posture, facial expression, intentions of another person, for example. But we can just as well reverse the process. We can focus on an abstract meaning, and create a corresponding sensory image: this is the basis of mental imagery. So just as we can recognize a tiger the slightest of glimpses, we can also *imagine* a tiger--- although, as we saw in Chapter 4, the sensory image we reconstruct is remarkably sketchy (e.g., our imaginary tiger won't have a specific number or even direction of stripes).
- Now feelings are just one more thing we can pay attention to. An emotion is, as we saw in Chapter 5, the *interpretation* of a bodily state; and the relevant aspects of our bodily feelings are very simple, varying on just two dimensions: energetic-relaxed and positive-negative. So experiencing an emotion requires attending to bodily state and relevant aspects of the outer world: the interpretation imposes a 'story' linking body and world together. Suppose Lestrade feels the physiological traces of negativity (perhaps he draws back, hunches his shoulders, his mouth turns down, he looks at the floor) as Holmes explains his latest triumph. The observant Watson attends, successively to Lestrade's demeanour and Holmes' words. searching for the meaning of these snippets, perhaps concluding: "Lestrade is jealous of Holmes' brilliance and success." But Lestrade's reading of his own emotions works in just the same way: he must attend to, and interpret, his own physiological state and to Holmes' words and conclude that he is jealous of Holmes' brilliance and success. Needless to say, Lestrade may be thinking nothing of the kind---he may be trying (with frustratingly little success) to find flaws in Holmes' explanation of the case. If so, while Watson may interpret Lestrade as being jealous, Lestrade is not experiencing jealousy (of Holmes' brilliance and success, or anything else)---because experiencing jealousy results from a process of interpretation, in which jealous thoughts are the 'meaning' generated; but Lestrade's mind is attending to other matters entirely: in particular, the details of the case.
- Finally, consider choices (Chapter 6). Recall how the left hemisphere of a split-brain patient fluently, though often completely spuriously, "explains" the mysterious activity of the left hand---even though that hand is actually governed by the brain's right hemisphere. This is the left, linguistic, brains attempt to impose meaning on the right hand movements: to create such meaningful (though, in the case of the split-brain patient, entirely illusory) explanation requires locking onto the activity of the left hand and to make sense of it. It does not, in particular, involve locking onto any hidden inner motives lurking within the right hemisphere (the real control of left hand) because of left and right hemispheres are, of course, completely disconnected. But notice that, even if the hemispheres were connected, the left hemisphere would not be able to attend to the right hemisphere's inner workings—because the brain can *only* attend to perceptual input (including the perception of one's own bodily state), but not to any aspect of its own inner workings.

We are, in short, relentless improvisers, powered by a mental engine which is perpetually creating meaning on sensory input, step by step. Yet we are only ever aware of the meaning created; the process by which it arises, and the sensory information and memory traces it draws upon, is utterly unavailable to introspection. Our step-by-step improvisation is so fluent that we have the illusion that the 'answers' to whatever 'questions' we ask ourselves were 'inside our minds all along.' But, in reality, when we decide what to say, what to choose, or how to act, we are, quite literally, *making up* our minds, one thought at a time.

So that is the cycle of thought, and its implications, in outline. The story needs, of course, a little unravelling, clarifying, explaining, and defending. This will be the task of the rest of this book.