

Why we are *not* conscious

The Problem

The Hard Problem of consciousness is the question of subjective experience, of why do the things we experience *feel* different? Why does the colour red have a different *quality* to green? Why does a pin prick feel different to a caress? Why does red feel different to the sound of middle-C, or to the taste of a lemon sherbet dib-dab; considering it's all neural activity underneath?

The “easy” problems of consciousness, for the record, are those that are susceptible to explanation by neurophysiological models. Examples include questions of how are memories stored, how do we perceive things, what is the causal role of pain, and what is behaviour? In principle, these can all be resolved by the uncovering of some neurophysiological mechanism. If so, we may get to the point where a neuroscientist can point to the brain and say, “here is the neural activity of someone watching a sunset” and “here is the neural activity of someone eating chocolate”. A great point to reach but it still doesn't explain why the neural activity of seeing red *feels* different to the neural activity of tasting chocolate; the Hard Problem.

Of course, many problems that seem hard often transform into easy problems once a better understanding of the issues has come to light, or a fresh perspective has been unearthed. I think we can get to the point where the mystery of consciousness is transformed into an understandable, if not easy, problem. Future generations may look back at us and mock us for not realising the obvious, just as we mock those who thought the Earth was the centre of the universe. And also, perhaps, for our narcissistic arrogance for assuming that something that, we think, is special and distinguishes us from other animals is especially *hard* rather than a simple and mundane feature of evolving nervous systems.

Why we are *not* conscious

It is a mysterious, interesting and important question to ask why we are conscious in this way. But there is an even more interesting question to ask, which I think can help us tackle the problem, and that is, why are we *not* conscious? If consciousness is a product of the brain and its neural activity why is there not consciousness associated with every instance of neural firing?

There are a number of ways in which we are not conscious:

- In fact, it seems, that most of what we do is done unconsciously. We carry out many tasks without much impinging on our consciousness, such as walking or maintaining

balance while standing. The behaviour of the iris system is something of which we are completely unaware. It just goes ahead by itself, unconsciously opening and closing to maintain a consistent measure of light falling on the retina, even though there is associated neural activity within the nervous system.

- There are some things of which we were once aware but are no longer. When learning a new task or skill, such as playing a guitar and even walking, perhaps, we focus our attention on the minutiae of what is required to pick up the skill, but once, after practice, we have learned the skill we are no longer aware of the details. When learning to play a basic guitar note we may, initially, concentrate on, and be conscious of, the pressure of a finger we are applying to a string to get a pure sound, void of buzzing. An accomplished guitarist is not conscious of such details; until a bum note is played, that is. The actions involved in the behaviour of walking are something of which we were once aware; we would consciously move and place a leg in such a way to achieve a particular goal. Now, as a grown up, it is all second nature and we are able to walk without being conscious of the details. This is in contrast to the behaviour within the iris system to which we, presumably, have never had conscious access, and, crucially, are not able to change how it operates; it is automatic, fixed and immune to processes of learning.
- Conscious awareness can also shift in an instant. What was not conscious can come into awareness, replacing a previous focus. One moment we may be happily focusing on a conversation with a friend in a restaurant and in the next our awareness is hijacked by the clatter of a waiter dropping plates. Awareness of a particularly vivid nature are those experiences of pain that capture our attention emphatically, such as a pin prick or shutting a finger in a door. There is also a sense in which awareness can shift level, from the general to the detail; from a symphony to that bum note, from walking through a forest to tripping over a root.
- We also seem to be able shift awareness at will, such as when following the steps of a recipe, asking a friend a series of questions or focusing on guitar chord finger positioning. At one moment we consciously focus on the brownness of onions frying in a saucepan then move on to adding tomatoes and concentrate on the mixture boiling; the colour of the onions no longer in focus. Although the nervous system is highly parallel, with many areas active all at once, it appears that our conscious experience is serial in nature. Generally conscious experience flows coherently from one focus to another; though for some people it may be quite erratic and disjointed.
- With the intriguing **Blindsight** condition people seem to be able to see, but have

no conscious experience. For example, a subject may claim that they cannot see an object in their field of view, but when asked are able to identify the object. Likewise, they may catch an object thrown to them though they claim not to be able to see it. As Blindsight results from damage to the brain it would seem to confirm that consciousness is a consequence of the structure of the nervous system, and that it can be lost if part of that structure is absent.

- The “other minds” problem concerns whether or not other people have minds and are conscious as we ourselves are, considering that we can only see their behaviour and do not have direct access to their mental states. It *is* a problem, but it is generally assumed that other people do have minds and are conscious, as minds are a product of the brain and other people do have brains. This would seem sensible in general, but when we look at the detail the picture is not so clear. We find that some people are not conscious of things that other people are, and also that people think about and experience some things in very different ways.

Aside from the case, already mentioned, of awareness related to skill acquisition, people can acquire very different ways of experiencing the world. This seems particularly striking in attitudes towards our fellow humans, related to division and prejudice. For example, in India a person of a high caste may have visceral negative feelings to a low caste person that they do not have to someone of their own caste. These sort of subjective prejudicial conscious experiences can be seen in many societies throughout the world and have been seen throughout history.

Furthermore, one person may be aware of the implications of Einstein’s theory of relativity that another simply does not appreciate. Some people may experience great anger if they see an image of Muhammad, the founder of Islam, whereas others are indifferent. The content of experience is largely irrelevant in the current context, but demonstrates that some minds are conscious in a way that other minds or not. Others do, more likely than not, have minds but not all minds are equally conscious. The “other minds” issue is of a different nature than the other points in this list, and less relevant in the current context, as the lack of conscious experience is a consequence of the lack of development of associated neural structures.

A New Hope

Why are we not consciously aware of the operation of the iris and similar systems? Why do we cease to be aware once something has been learned? Why does the focus of attention shift to an unconscious state? How is it that Blindsight subjects can respond to what is in

the world without being consciously aware? In essence, why is there so much of our own activity of which we are unaware; why are we partial zombies?

Psychology, Cognitive Science, Neurobiology and Philosophy of Mind have produced a vast body of work over the decades accumulating much knowledge and many theories, techniques and methodologies related to behaviour and the functions and operation of the nervous system and the mind. Surely these are of use to give some insight into the nature of consciousness, and non-consciousness?

Well, I wouldn't, as they say, start from there. To tackle any problem it is essential to have the correct foundational understanding of the conceptual model of the problem. In order to properly understand, for example, the motions of the planets it is necessary that the problem is framed in the heliocentric rather than geocentric model. If the wrong model is used then the wrong phenomenon will be addressed or steps will be overlooked which may have led to a resolution.

I contend that the current, conventional input-output conceptual model of the operation of living systems is not valid. The correct conceptual model to use [of course] is Perceptual Control Theory (PCT). The perceptual control conceptualisation can lead to different, previously unknown, ways of thinking about behaviour and the operation of living systems, and, so, to the nature and role of consciousness, and suggest potential resolutions to the Hard Problem.

According to PCT the neural architecture of living systems comprises a hierarchy of sets of simple feedback control systems whereby action is varied in order to control what is perceived. For example, in a shower we turn (vary) the heat mixer tap until we feel (perceive) the desired water temperature. Action is driven by the difference (known as error) between the desired perception (comfortable water temperature) and what we currently perceive (actual water temperature). When there is no error (the two match) then no further action is taken. The same process occurs at all levels of behaviour; e.g. control of our perceived position of the car between the white lines on the road while driving or our self-image as a kind person or the political system we want to see governing our lives. The higher-level, more complex systems are not, of course, present at birth but develop through a process of learning; reorganisation of the structure of the nervous system.

A Flexible Mind

Here we come onto an intriguing link between consciousness and learning as hinted to earlier with respect to what is no longer conscious.

We can think of evolved living systems in terms of flexibility related to learning.

- Comparatively simple organisms, and early in evolutionary terms, may have highly

efficient control systems but are inflexible in that they are unable to learn during their lifetime. A fly, for example, has a very successful avoidance control system, but it is not flexible as the fly is unable to change its nature through learning. Any learning that does take place, let's say it involves the *gain* of a control system, occurs only through change over successive generations. Those, within the fly population, who have a good gain value are more likely to survive, and procreate, than those with a poor gain value (acts more slowly). In this way the structure of the fly's control system reorganises over the generations converging on an optimal gain value for its environment (or at least a value that is good enough for inheritance).

- A more flexible arrangement may allow the change of gain value during a single generation. Plasticity within the neural control systems may allow short-term change of a gain value. The change could be driven by persistent error within a control system. So, if a control system does not control well initially, then there will be persistent system error. This error drives the reorganisation of the gain value until the persistent error is low, therefore suspending further change.

This sort of system would enable some lifetime flexibility where an organism would adapt to the circumstances it encounters after birth. It would also enable adaptation to changing circumstances, as renewed persistent error would lead to the reorganisation to a different gain value. These self-organizing systems would add flexibility to an organism, but could make them more vulnerable (poor control) during the periods of reorganisation. However, an organism could comprise a combination of both the fly-like inflexible systems and the more flexible adaptive systems. This would seem quite plausible from an evolutionary point of view where new systems get added on to the more primitive systems.

- If there were many of the flexible control systems, at different levels, all reorganising at the same time things may get a little messy with conflict between systems and the reorganisation in one affecting control within another. In this case it would be beneficial if only certain systems were reorganised at any one time rather than all at once. It may be the case that not all control systems are "active" at the same time, and that reorganisation only occurs while a system is active. The switch between active systems may be triggered by what happens in the environment. For example, an animal may be drinking at a water hole, but then sees a crocodile, activating a flight control system. If this system then has the ability to reorganise then it can improve the crocodile-avoidance performance. With this type of system the shift in activity, and reorganisation, would appear to be driven by circumstances external to

the organism.

- We can envisage a similar, but much more useful form of reorganisation; self-directed reorganisation. If an organism could direct, at will, the control systems to be reorganised then learning could take place in a far more sophisticated way, and in terms that are meaningful to that organism, rather than triggered by external events.

The first and last of these do seem to correspond to actual living systems. It is unclear whether the other two exist in actuality, but are included as possibilities for the sake of completeness. Humans may embody all these forms of flexibility.

Quality Control

What is becoming apparent is a significant link between what we are and are not conscious of and learning, expressed as a process of reorganising the structure of the associated neural systems.

Some control systems within the nervous system operate successfully, though may have momentary error driving action, but are not associated with conscious experience. Examples are the iris control system within the human eye, and the avoidance control system of the fly. These systems are inflexible with no processes of lifetime learning involved in the formation of their structure.

Some, other, systems are flexible in that the parameters of control are variable enabling adaptation to changing circumstances or needs. For these the conscious awareness that springs forth appears to be associated with error within control systems. The clatter of dropped plates signifies error with an aural localisation system. Perception of a bum note represents error within a melody control system. The experience of pain is error within an avoidance system. Learning a new task involves the focusing of conscious awareness on the poor control, the error, of a system.

This connection between error, awareness and learning suggests a role of consciousness as a *quality* control system. By being able to perceive the quality of a standard control system the quality control system is able to initiate reorganisation, if control is inadequate. If the perceived quality of a guitar chord sounds poor, reorganisation (which can be random) of the control parameters takes place until the quality improves.

It is not difficult to envisage the form of such quality control systems, which is coherent within the context of evolutionary development. Rather than the inputs to the systems forming perceptions derived from aspects of the environment the inputs form perceptions of the quality of the normal perceptual control systems. In simple terms, the inputs could derive from the error signals of the normal systems instead of the perceptual signals. Also,

rather than the outputs contributing to the references of other systems they act upon the *parameters* of those systems. Although this functionality represents a major evolutionary leap forward in terms of the abilities it provides to the lifetime flexibility of an organism it is similar enough to the existing structures to be plausible as a simple evolutionary progression.

So, such a system can tell when other systems are working poorly and change the structure of those systems so that they perform better. Instead of perceiving aspects related to the environment those systems perceive the quality of control of internal systems; they are meta-perceptual control systems.

In The Zone

For this quality control system to be viable it would have to be selective, moving around to particular subsystems rather than operating on all systems at the same time, otherwise it is likely that there would be significant interference between systems. This is certainly consistent with the experience of the focus of conscious attention shifting around. The actual mechanism is not yet clear but may involve a neural structure that extends across vast areas of the brain encompassing the basic control systems in different domains, and at different levels. In this way the focus could shift to specific systems according to the current goals of the hierarchy.

Conscious attention may be directed to a specific task to be learned, such as juggling. Initially the focus would be on throwing and catching a single ball. The aspiring juggler would be aware of the quality (error) of the consequences of their throwing and catching behaviour, and directed reorganisation would automatically adjust the parameters until the quality improves. Later, once the quality of control of juggling with one ball, then two balls, has enabled progression to all three, then the focus switches from the details of each ball to the perception of quality of perceptions associated with the juggling as a whole. The focus, for example, may be on the perceived rhythm of the juggling behaviour which is crucial for maintaining consistency and endurance. Focus has shifted to a higher level of perception, once the lower level control systems have been learned. The quality of control is good (error is low) and, so, consciousness ceases.

This seems to tie in with the psychological phenomenon known as *Flow*, or being *in the zone*. It is a mental state, often joyful, of being completely immersed in an activity, whereby everything is operating smoothly without having to think (be conscious of) the details (lower-level control systems) of the activity. This state of mind might be experienced in a variety of tasks such as juggling, basketball or dancing. I have experienced this myself with Latin American partner dancing. While learning there is a great deal on which to focus conscious attention; the placement of the feet, how to hold your partner, where to place

the arms, how to position the fingers for a firm and decisive lead, how and when to turn your partner and how to synchronise with both your partner and the timing of the music. It would be onerous and impractical if you had to think about all these things every time you danced. Fortunately, after a few years of learning the details become second nature and slip out of consciousness. Once accomplished you get the point of dancing when you are concentrating only on the moment, the generality of the dance and of the joy of everything working without needing think about it. That's being *in the zone*.

It can seem almost mystical in a way in that this complex and highly-skilled behaviour is being produced from within yourself of which you have no conscious awareness. It is almost as if it is another person producing that behaviour. Yet is it is merely the *zombie* within, the consequence of a process of focused awareness and reorganisation culminating in post conscious, high quality control systems. Here is highlighted a characteristic of control systems whether they be zombie flies, or internal zombies like the iris system or the accomplished systems for dancing or playing a piano, in that successful, performant systems have no need for consciousness.

Can zombies juggle?

The beauty of this current hypothesis of consciousness is that it can, in principle, be empirically tested and verified. If the proposed quality control systems exist in the brain then it should be possible to identify the associated neural structures. A further consequence of the hypothesis is that there would be a correspondence between the subject of experience of conscious awareness and the activity of the neural quality control systems and the target perceptual control systems under reorganisation.

These may be very difficult to verify given the state of current knowledge and technology regarding the examination and identification of neural structures. So, there may be a quicker way to verify, or falsify, the idea. A central concept in the theory is that conscious awareness is necessary for the reorganisation that takes place during learning. Naturally, it follows that someone who lacks awareness could not learn a new skill. Therefore, it should be possible to take a Blindsight zombie (no insult intended, some of my best friends are zombies) and try to teach them to juggle. In this case the suggestion is that they are able to unconsciously perceive the balls, but would not be aware of the qualitative consequences of their actions, so reorganisation would not take place and the skill would remain unlearned. This would provide support for the hypothesis that the role of consciousness is as a quality control system. If, on the other hand, the zombie could learn how to juggle then the hypothesis would be wrong.

Siliconophobia

This discussion suggests that the problem of consciousness is not as mysterious, or hard, as has been thought. If we can locate the neural structures of perceptual control systems and of quality control systems and identify associated reorganisation connections, and also identify the correspondence between conscious awareness and the neural activity of the quality control and the transition to a post-conscious state of good quality systems then would there be any mystery remaining?

If we were to endow a robot with perceptual control systems as well as the overlying quality control systems such that it could direct reorganisation to particular perceptual control systems enabling it to learn new abilities and adapt to dynamic circumstances, as well as reporting its subjective experience of its shifting focus could we say that it was not conscious? My first instinct would be that it was not, but also recognise that I would have little justification for such a view, and that I am just being robot-ist (as opposed to a roboticist) and prejudiced against silicon-based lifeforms, and must resist the tendency to assume that human consciousness has some special status that cannot be replicated in other, non-carbon, non-biological forms.

Evolution has shown that the immense diversity of life has emerged from a simple process, of natural selection. PCT shows that the complexity of observed behaviour emerges from another simple process, of perceptual control. There is no reason to suspect that consciousness is also not a product of a similarly simple process, in this case of quality control.

The Pursuit of Quality

This perspective of the role of consciousness adds an extra dimension to the nature of control systems, as a way of directing the improvement of quality. Rather than systems just controlling perceptions in a quantitative way (e.g. signal intensity) there also exists systems that control in qualitative way. That is, a distinction is perceived between different systems even though the degree (quantity) of neural activity within the different systems may be equal. The associated conscious experience may be the phenomenon know as *qualia*, and an emergent consequence of the quality control systems.

Consciousness is generally experienced as a continuous and contiguous (apart from sleep) flow of awareness. This may be something of an illusion and an unintended consequence of the quality control system process. We could imagine a creature with some nascent flexible learning systems acquired as part of its evolutionary development. While one of these quality control systems is reorganising there may be some associated primitive conscious awareness.

When the reorganisation ceases so does the awareness. There may be a gap in time before another system reorganises and there is another bout of conscious experience. The creature is likely to entirely unaware of the period of non-awareness. Human consciousness may work something like this, though we may realise that we have no memory of certain periods, such as when driving.

When we consider the nature of the content of consciousness it could be seen as a constant striving for quality within our lives. Learning a skill concerns the improvement of abilities. Seeking and maintaining employment is about job security. Cultivating social relationships is about the quality of interactions with others and not feeling lonely. Arranging activities such as skiing or going to a comedy club are about experiencing a higher degree of quality associated with such experiences.

When conscious awareness is not focused on a specific task or activity it may be focused in a more general and high-level way on the existential nature of the Self. This could also be seen as a perception of the quality of our control systems in general. If it is the case that the control systems are performing well (achieving their goals with minimal error) then we experience that as happiness or contentment. If, on the other hand, many of our systems are in a state of persistent error, without reorganisation having any effect, then we call that experience stress, or depression.

In summary, what we identify as conscious experience is the process of the perception of the quality of our standard perceptual control systems and the reorganisation of the structure of those systems. Conscious awareness ceases when the reorganisation has subsided to the extent that the quality control system no longer contains error, indicating that the operation of the perceptual control systems is of a high quality.

Conscious awareness shifts around during the lifetime of a sentient being constantly improving the quality of specific systems as necessary by transforming the structure of those systems. Once the quality has been achieved awareness moves on to another system that requires improvement. In this way conscious awareness can shift between the details of one system and the details of another as well as up and down levels to perceive the quality of our existence at different scales of perspective.

With humans this results in a complex and sophisticated, rich contextual dynamic experience, but may be something of an illusory, fortuitous accident rather than a coherent sense of self.