Minds & Machines (2012) 22:117–136
DOI 10.1007/s11023-012-9279-x

How to Build a Robot that is Conscious and Feels

J. Kevin O'Regan

Received: 31 January 2012 / Accepted: 18 May 2012 / Published online: 7 June 2012
© Springer Science+Business Media B.V. 2012

Abstract Following arguments put forward in my book (Why red doesn't sound like a bell: understanding the feel of consciousness. Oxford University Press, New York, USA, 2011), this article takes a pragmatic, scientist's point of view about the concepts of consciousness and "feel", pinning down what people generally mean when they talk about these concepts, and then investigating to what extent these capacities could be implemented in non-biological machines. Although the question of "feel", or "phenomenal consciousness" as it is called by some philosophers, is generally considered to be the "hard" problem of consciousness, the article shows that by taking a "sensorimotor" approach, the difficulties can be overcome. What remains to account for are the notions of so-called "access consciousness" and the self. I claim that though they are undoubtedly very difficult, these are not logically impossible to implement in robots.

Keywords Robot · Consciousness · Feel · Self · Sensorimotor theory

Can machines be conscious? More, can they really feel things? How do humans manage to do these things?

These are the questions that I try to answer in a recent book (O'Regan 2011). Here I shall present some of the arguments that I put forward in the book, coming to the conclusion that there is no *logical* obstacle to building conscious machines that really feel. I claim that the problem is easier than usually thought, but it still is difficult.

This paper is based on a talk given at PT-AI, Thessaloniki, 2011. Portions of the paper are taken from O'Regan (2011).

J. K. O'Regan (✉)
Laboratoire Psychologie de la Perception, CNRS UMR 8158, Université Paris Descartes, Paris, France
e-mail: jkevin.oregan@gmail.com
URL: http://www.kevin-oregan.net

Much of the problem in consciousness research over the last decades lies in a curious contradiction. On the one hand, everybody in (at least western) civilization seems to agree that we as humans intimately experience a thing we call consciousness—and that when we have sensory experiences (e.g. seeing, hearing, smelling, touching, tasting, suffering a pain), we really feel something. On the other hand, no one seems to be able to explain, in a way sufficiently precise to allow scientific investigation let alone robotic implementation, what they mean by such statements.

A way forward is to suggest definitions which are amenable to science, but which at the same time provide a reasonable degree of overlap with everyday use of the terms "consciousness" and "feel". I shall start by discussing what is generally thought to be the easier aspect of consciousness, namely what is called "access consciousness". This part of the paper is a necessary condition before coming to the more original contribution of this paper, namely the second part. This part addresses what (Chalmers 1997) calls the "hard" problem of "phenomenal consciousness" or "feel".

Being Conscious of Something

The philosopher Ned Block defines "Access Consciousness" as a first approach to understanding the more cognitive or reflective component of consciousness, which might also be called "awareness" (Block 1996). He suggests that what we mean by being (access) conscious of X is: being "poised" (meaning "ready" or "on the verge" or "prepared") to make use of X (i.e. to "access" X) in our rational actions, decisions, planning, or in our linguistic or communicative behavior.

Though used by some philosophers, I think this definition is not quite satisfactory. Surely what we would like to mean by the term conscious access requires more than just being poised to make use of something in rational actions and communicative behavior. To illustrate why, imagine a chess-playing automaton. As the machine plays chess with you, it clearly is "poised to make use of your chess moves in its current (admittedly limited!) actions, decisions, planning, and in its rational, (admittedly limited!) linguistic or communicative behavior". But we don't want to say that the machine is *conscious* of your moves. The reason is not because of the limitations of the machine's rational and communicative powers. The reason is that all the machine is doing is making a calculation and responding automatically to your moves. Furthermore the machine has no idea that it is doing what it is doing, nor that what it is doing could be put into a wider context of different other things it might be doing. Thus Block's criterion of being "poised to make use of something in current actions, decisions, planning and rational, linguistic or communicative behavior" does not seem to constitute a satisfactory characterization of conscious access.

On the other hand since the calculations necessary to play good chess are cognitively quite complex, it does seem reasonable to admit that the machine has what we might call "cognitive access" to your moves. Thus let me retain Block's definition of "being poised to make use of X in rational behavior", but let me use it to refer to what I shall call *cognitive access*.

So how would we build a Mark II version of the chess playing machine so that it had full-blown *conscious* access? First, the machine would have to not only have *cognitive* access to your move, but it would additionally have to be poised to reason or communicate about *the fact* that it had this cognitive access to your move. A second condition is that the machine would have to be aware of the wider context within which it was doing the cognitive accessing. For example you would assume a chess playing machine had knowledge of this wider context if the machine sometimes bantered with you about the move you made, or if it mentioned that it was deliberating whether or not to invoke a classical defense, or even if it ventured a word about the weather during a boring phase of the game. In other words, the machine would have to not only have cognitive access to your move, but it would additionally have to be poised to reason or communicate about the fact that it had this cognitive access to your move.

These conditions for conscious access can be expressed succinctly by using the notion of cognitive access in a hierarchical way. We can say a condition for an agent to have conscious access to X is:

The agent must have cognitive access to (*the fact that it has cognitive access to X*).

The idea that what we mean by having conscious access to something consists in a two-level hierarchy of access is not new in philosophy. David Rosenthal, referring back to Aristotle and Locke, defends what he calls the "Higher Order Thought" (HOT) theory of conscious awareness, in which a mental state is said to be conscious when a creature is aware it is in that state (Rosenthal 1997). The theory now exists in different variants, but has generally been well accepted by many philosophers (Rosenthal 2011; Lau and Rosenthal 2011; Gennaro 2004). The main difference with what I am suggesting here is that instead of speaking of "conscious mental states" as do HOT theorists, I prefer to use the notion of cognitive access, which seems easier to understand from the point of view of implementation in artificial cognitive systems.

The wider context of possibilities for action implicit in conscious access falls out of the definition of cognitive access: this is because the notion of cognitive access contains the notion of selection among a range of possibilities. This is already true in a limited way for the simpler version of the chess machine: its first order cognitive access to your move means that it is poised to make use of your move in its rational behavior. But what does that mean? It means that there are a variety of possible rational behaviors it could undertake (like moving this piece or that), and furthermore these are related to a variety of possible moves you might have made. Having cognitive access implies that the machine selects among these possibilities.

This need for a selection to be made applies also to the higher level of cognitive access in the Mark II machine. The machine is poised to make use of the fact that it is poised to make use of your moves. In the same way as for the first level, the second level of cognitive access carries the implicit assumption that there are a variety of possible things the machine could do about the fact that it is poised to make use of your moves (it could carry on playing, but could also do other things, like talk about your move, or ignore your move and talk about the weather). Furthermore the things it might do are related to the possible different facts it might be poised to make use of (e.g. being poised to make use of your move must not be the only thing it could be poised to make use of, it must also be possible for it to be poised to make use of other things, like the expression on your face for example, or the fact that it is playing chess and not dominoes).

Thus we see that using the simple notion of cognitive access in a hierarchical fashion allows us to capture two conditions that seem necessary for conscious access: first that the machine must currently be cognitively accessing the fact that it is making the chess move; and second that the machine's chess-playing capacities should be embedded in a wider field of possible activities than just responding to moves on a blow by blow basis. Part of what we generally mean by having real conscious access to chess playing involves having richer cognitive capacities than just playing chess.

Does this requirement pose a problem for machines? Perhaps today it does, because having wider knowledge and richer cognitive capacities than required in a well-defined domain such as chess-playing, flight reservation, or medical diagnosis, say, is something where current AI programs fail. Common sense knowledge, like the problems of language comprehension, induction, metaphor and analogy, are areas where present-day AI has not yet made a decisive breakthrough (Dreyfus 1972, 1992; Kassan 2006). The use of very large databases and systems that can do sophisticated reasoning have helped us advance in these problems, but something still seems to be missing. A possibility that some AI researchers are putting their hopes in is embodied cognition (Pfeiffer and Bongard 2006; Mahon and Caramazza 2008; Wilson 2008) and epigenetic or developmental robotics (Lungarella et al. 2003). The idea is that having a body, and gradually learning how to use it in the way a human child learns over many years of development, may allow a way of grounding knowledge in the world.

There is therefore hope that with such approaches we will find the path to human-like intelligence. But what is important is that even if this is not the right path, it seems apparent that to the extent that human brains are computational devices, nothing magical can be going on in the human mind. Barring dualism, it should still *logically* be possible for machines to display the same type of intelligence, even if it is as yet difficult to achieve.

The Self

I have said that to be conscious of something, an agent must have cognitive access to the fact that *it* has cognitive access to that thing. Who or what is the "it"? Clearly the notion of conscious access requires not only certain cognitive capacities, but additionally it requires that the agent involved have at least a primitive version of what we call a *self*.

Is this a problem for science or robotics? Philosophers and psychologists have looked at the notion of self in many ways (Gallagher 2000; Gallagher and Shear 1999). From a scientific point of view we need to divide the notion into parts that can be instantiated in a mechanistic way. Following (Vierkant 2003), a first useful division distinguishes the *cognitive self* that can be described in terms of a single agent's cognitive capacities, from the *social self* that involves social interactions that the agent can have with other agents. Then, within the cognitive self we can further distinguish a hierarchy of cognitive capacities.

At the simplest level within the cognitive self is "self-distinguishing", that is the ability for a system or organism to distinguish its body from the outside world and from the bodies of other systems or organisms. There is clearly no problem in programming a robot to avoid bumping into other objects, or to avoid its arms hitting its own body. So self-distinguishing is clearly within the realm of possibilities for a robot.

The next level is "self-knowledge". Self knowledge in the very limited sense I mean here is something a bird or mouse displays as it goes about its daily activities. The animal exhibits cognitive capacities like purposive behavior, planning, and even a degree of reasoning. On the other hand the bird or mouse as an individual presumably has no concept of the fact that it is doing these things, nor that it even exists as an individual. Again, as a cognitive capacity self knowledge is within the realm of robotics.

At the top level of the classification we have "knowledge of self-knowledge". This is most typically a human capability, though some primates and possibly dogs, dolphins and elephants may have it to some extent. It may be most developed when the agent possesses language, and underlies what philosopher Daniel Dennett calls the "intentional stance" that humans adopt in their interactions with other humans (Dennett 1991). The individual can have a "Theory of Mind", that is, it can empathize with others, and interpret other individuals' acts in terms of beliefs, desires and motivations. This gives rise to finely graded social interactions ranging from selfishness to cooperation and involving notions like shame, embarrassment, pride, and contempt.

I have called all these forms of the cognitive self "cognitive" because they involve computations that seem to be within the realm of symbol manipulation and concept formation. There seems to be no logical difficulty involved in building these capacities into a robot, even if today's AI programs have difficulties, in particular with abstracting concepts.

The Perceived Reality of the Self

On the other hand there does still seem to be something missing. We as humans have the strong impression that there is someone, namely ourselves, "behind the commands". We are not just automata milling around doing intelligent things: there is a pilot in the system, so to speak, and that pilot is "I". It is I doing the thinking, acting, deciding and feeling. How can the self seem so real to us, and who or what is the "I" that has this impression? Here I want to appeal to current research in social and developmental psychology. Scientists in these fields agree that although we have the intimate conviction that we are individuals with a single unified self, the self is actually a construction with different, more or less compatible facets that each of us gradually builds as we grow up. The idea is that the self is a useful abstraction that our brains use to describe, first to others and then later to ourselves, the mental states that "we" as individual entities in a social context have. It is what Dennett has called a "narrative fiction" (Dennett 1991).

But then how can the self seem to us to be so real? The reason is that seeming real is part of the narration that has been constructed. The cognitive construction our brains have developed is a self-validating construction whose primal characteristic is precisely that we should be individually and socially convinced that it is real. It's a bit like money: money is only pieces of metal or paper. It seems real to us because we are all convinced that it should be real. By virtue of that self-validating fact, money actually becomes very real: indeed, society in its current form would fall apart without it. As pointed out by Vierkant (2003), the self is actually even more real than money because it has the additional property that it is self-referring: like some contemporary novels, the "I" in the story is a fiction the "I" is creating about itself.

In which case, shouldn't we be able to change the story in mid course? If our selves are really just "narrative fictions" then we would expect them to be fairly easy to change, and by ourselves furthermore. But actually this does not work. It is necessarily part of the very construction of the social notion of self, that we must be convinced that it is very difficult to change our selves. After all, society would fall apart if people could change their personalities from moment to moment.

But couldn't we by force of will just mentally overcome this taboo? If the self is really just a story, changing the self should surely in fact be very easy. It turns out that we can under some circumstances break the taboo and flip into altered states where we become different, or even someone else. Such states can be obtained voluntarily through a variety of "culturally bound" techniques like possession trances, ecstasies, channeling provoked in religious cults, oracles, witchcraft, shamanism, or other mystical experiences, and states like latah, amok, koro and hypnosis (Simons and Hughes 1985; Somera 2006). Sometimes such states can also be provoked involuntarily under strong psychological stress caused by physical abuse, brainwashing by sects, in religious cults and in war, giving rise to conditions such as post traumatic stress disorder, and dissociative identity disorder (i.e. multiple personality) (Lifton 1986, 1989, 1991). There is also a suggestion that social networking and gaming on the internet may to a certain extent allow people to create alternative selves (Floridi 2011; Ward 2011).

Hypnosis is interesting because it is so easy to induce, confirming the idea that the self is a story we can easily control if we could only decide to break the taboo. Basic texts on hypnosis generally provide an induction technique that can be used by a complete novice to hypnotize someone else. This suggests that submitting to hypnosis is simply a matter of choosing to play out a role that society has familiarized us with, namely "the role of being hypnotised". It is a culturally accepted loophole in the taboo, a loophole which allows people to explore a different story of "I". Another indication that it is truly cultural is that hypnosis only works in societies where the notion is known. You can't hypnotise people unless they've heard of hypnosis.

This is not to say that the hypnotic state is a pretense. On the contrary, it is a convincing story to the hypnotized subject, just as convincing as the normal story of "I". So convincing, in fact, that clinicians are using it more and more in their practices, for example in complementing or replacing anesthesia in painful surgical operations (Patterson and Jensen 2003; Faymonville et al. 1995; Hardcastle 1999).

There is also the fascinating case of Dissociative Identity Disorder (formerly called Multiple Personality Disorder). A person with Dissociative Identity Disorder may hear the voices of different "alters", and may flip from "being" one or other of these people at any moment. The different alters may or may not know of each others' existence. The surprising rise in incidence of Dissociative Identity Disorder over the past decades signals that it is a cultural phenomenon (Hacking 1986; Hartocollis 1998). Under the view I am taking here, Dissociative Identity Disorder is a case where an individual resorts to a culturally accepted ploy of splitting their identity in order to cope with extreme psychological stress. Each of these identities is as real as the other and as real as a normal person's identity—since all are stories.

In summary, the troubling idea that the sense of self is a social construction seems to be the mainstream view of the self in social psychology. If this view is correct, then we can confirm that there really is logically no obstacle to us understanding the emergence of the self in brains and thus to recreate it in machines. Like the cognitive aspect of the self, the sense of "I" is a kind of abstraction that we can envisage would emerge once an agent, biological or non-biological, has sufficient cognitive capacities and is immersed in a society where such a notion would be useful.

Finally, if we have the possibility of providing a non-biological agent with a self, and if it has sufficient cognitive capacities for it to be able to cognitively access the fact that it is cognitively accessing something, then we have all the ingredients necessary for the agent to be (access) conscious of that thing.

The Problem of Feel

The above discussion has concerned the notion of access consciousness. The claim was that there is no logical obstacle to making machines that display this capacity, even though there remain formidable difficulties before this can actually be done. To the extent that what is involved in having access consciousness is a cognitive capacity that can be described in computational terms, access consciousness should be possible in non-biological agents. It may however be necessary for such agents to have a real body developing over time in a physical environment, and possibly also to have a social environment where the agent interacts with other agents.

So let us imagine we had constructed such an agent. It would be like the highly evolved robot played by Arnold Schwarzenegger in the film *Terminator,* or like the replicants in Ridley Scott's film "Blade Runner". Now the question arises: When these androids get hurt, do they have real pain? Do they really feel things as humans do? Clearly it would be conceivable to build robots that manifest all the external reactions of pain: they could display avoidance reactions and even cry "ouch": but would they actually be feeling anything? Would they not merely be play-acting? What is it about the human brain that has the consequence that some processes are accompanied by a real feel. And when there is a feel, why is it the way it is?

Differences Between Feels

Take the case of sound. Why is a sound perceived like a sound rather than providing some other experience, like say a visual experience, or no experience at all? The natural thing is to suppose that there must be something about the mechanisms in auditory cortex that gives rise to *auditory* rather than visual experiences. But suppose that we had found there was something special about the connectivity, firing rates, or some other characteristic of these mechanisms. Say it was a special neurotransmitter that was involved. Then we could always ask: what is it about the neurotransmitter that gives the auditory feel, rather than the visual feel? And now suppose upon further investigation we discovered that the neurotransmitter for auditory feel had an extra nitrogen atom as compared to the one involved in the visual feel. Then we could ask what it was about the nitrogen atom that provided the particularly auditory quality. And so on indefinitely, we could keep asking: and what is it about this neural aspect that provides the auditory rather than the visual feeling?

The same kind of problem would be encountered in trying to explain the nature of any sensory experience. For example, why, when one is touched on the arm, does the sensation have an *arm* feeling rather than, say, a *leg* feeling? We know there is a somatosensory representation of the skin in parietal cortex, where arm and leg are represented at different locations. But appealing to this fact does not explain why activation in the arm area of the map gives arm-like sensations.

And the same problem exists for the question of why red looks red rather than green, for example. If we search for an explanation in terms of neural mechanisms, we are faced with a logical problem: why do the neurons sensitive to red light give the red feeling rather than the green or any other feeling?

In all these cases there seems to be an infinite regress of questions, leading to what Joseph Levine has called an "explanatory gap" (Levine 1983). It is this which has led philosophers to agree that feel, or, as they often call it, phenomenal consciousness, is the "hard" problem of consciousness (Chalmers 1997). Explaining why sensations differ the way they do is a problem that logically seems to have no possible solution in neurophysiological terms. A fortiori therefore it seems difficult to imagine a robotic implementation of feel.

Why there is a Feel at All

But there is an even more fundamental question: why do sensory experiences like seeing, hearing, touching, tasting and smelling have a feel at all? Consider the fact that there are mechanisms in the body that control autonomic processes like the glucose or oxygen content in the blood, one's digestion, or one's balance, to name but a few. These systems involve sensory receptors just as vision or touch do, and yet there is no feel associated with them. What is it about the classic five sensory modalities of vision, hearing, touch, smell and taste which gives them a *sensory presence*, whereas the majority of brain processes, even involving sensory receptors, involve no feel? Thinking and imagining are another example. Thinking of red or imagining red or a pain are in themselves neither red nor painful. Like autonomic processes in the nervous system, thoughts and imaginings do not possess sensory presence, or at least not the same kind of presence as real sensory stimulation. Why not?

Again, as was the case for the differences between sensory experiences, we are tempted to appeal to neural mechanisms in order to account for why only certain brain processes are accompanied by sensory presence. But again any such effort will ultimately confront the explanatory gap. Any neural explanation for how sensory presence is generated is open to the infinite regress question of: Why does this particular neural mechanism provide presence, whereas others do not?

Although over the past decades consciousness researchers have come up with a myriad of hypotheses to explain consciousness [see e.g. Chalmers (1997) for a long list], all must ultimately confront both the logical impossibility of invoking neural mechanisms to explain the differences between sensory experiences, and also the fact that only certain brain processes provide sensory presence. Hypotheses like 40 Hz oscillations, quantum coherence in microtubules or re-entrant loops, among many others, do no more than add further mystery to what is already mysterious.

The problem of the explanatory gap between physical processes in the brain and the particular feel associated with sensory experiences is reminiscent of the problem of explaining the origin of life, as it was posed at the beginning of the 20th Century. Vitalists proposed the existence of an élan vital or vital spirit to explain why biological organisms were imbued with life. But today we know that life is not the kind of thing that is generated by anything. It is a *word* that describes the way certain systems interact with their environment. Thinking that life was the kind of thing that could be generated was to make what philosopher Gilbert Ryle (1949) called a category error. It was to commit the error of reification, that is to think that life was a thing, when in fact it is a way of interacting with the environment.

Perhaps we are making the same error with respect to feel.

The Sensorimotor View of Feel

The "sensorimotor" view of feel takes the stance that it is an error to think of feels as being the kind of thing that is generated by some physical mechanism, and that it is an error to look in the brain for something that might be generating feel. Instead, the sensorimotor view suggests that we should think of feel in a new way, namely as a *way of interacting* with the world (O'Regan 2011; O'Regan and Noë 2001; O'Regan et al. 2004). To illustrate, take the example of softness.

When I press a sponge, where is the feeling of softness generated? Obviously this is somehow the "wrong" question. We know that softness is not really generated anywhere. Indeed, softness is not the kind of thing that can be generated. The softness of the sponge *lies in the way the sponge squishes under the pressure of my fingers*. I feel the softness when a particular sensorimotor law is currently in effect, namely the law that describes how when I press on it, the sponge squishes. The sensorimotor approach to feel takes the stance that more generally, the quality of all feels, in all sensory modalities, is constituted by such laws of sensorimotor interaction with the environment.

In the next section I shall be showing how the sensorimotor approach provides a natural way of overcoming the first logical problem we confronted above, namely of explaining why the differences between different feels are the way they are. This will lead us to several interesting applications of the sensorimotor approach. In the subsequent section I will then treat the second logical problem, namely the question of why feels have sensory presence. In all what follows I shall only be discussing the *quality* of feels, and shall not adjudicate on whether the feels are *conscious*. This will be the subject of the last section in this part of the paper.

The Difference in Quality Between Feels

Let us first consider how taking the sensorimotor approach answers the question of why we *hear* sounds and *see* sights rather than the other way around. Under the sensorimotor approach, the role of the brain is not to *generate* sensations, but rather to *enable* the different modes of interaction with the world that constitute the particular sensorimotor interactions that are characteristic of different feels. To explain the different qualities of hearing and seeing, we must not appeal to neural processes, but to different sensorimotor laws governing auditory versus visual interactions. So what are the sensorimotor laws involved in seeing and hearing? Examples might be the fact that when you blink and you are hearing, there is no change, but when you blink and you are seeing, there is a very big change. When you are hearing, moving forward creates a particular pattern of increased intensity of sensory input, whereas the pattern when you are seeing is an expanding flowfield. These are just two examples of a multitude of other laws.

But now note that this very simple philosophical idea predicts something interesting. If it were possible to simulate the laws of seeing by using some other sense instead of vision, then we would predict that it should be possible to see through this other sense modality, provided the correct sensorimotor laws were maintained.

Sensory Substitution

This is the idea of Sensory Substitution. Paul Bach y Rita in the 1970s had already hooked up a video camera manipulated by a blind person, through some electronics to an array of vibrators that the blind person wore on their stomach or back (Bach-y-Rita 1972). He had found that immediately on moving with the device, observers were able to navigate around the room, and had the impression of an outside world, rather than feelings of vibration on the skin. With more practice they were able to identify simple objects in the room. There are reports of blind people referring to the experience as "seeing" (Guarniero 1974, 1977). With modern electronics, sensory substitution is becoming easier to arrange and a variety of devices are being experimented with (Bach-y-Rita and Kercel 2003; Auvray and Myin 2009). Bach-y-Rita and his collaborators have developed a tongue stimulation device which, though it has low resolution, has proven useful for substitution of vestibular information in people with vestibular impairments. There is work being done on the neural consequences of using a visual-to-auditory substitution device invented by Peter Meijer in Eindhoven, where information from a webcam is translated into a kind of "soundscape" that can be used to navigate and identify objects (Pascual-Leone et al. 2005; Proulx et al. 2008). A link to a video showing how an observer learns to use such a device is http://www.nstu.net/malika-auvray/demos/substitutionsensorielleST3.mpg.

Touch

Let us apply the sensorimotor approach to the question of why touch on the arm feels like it's on the arm rather than on the leg, say. Under the sensorimotor approach, to answer such questions, we have to consider what the sensorimotor laws are that underlie a touch on the arm as opposed to the leg. On reflection it is then evident that what we mean by a touch on one's arm is, among other things, the fact that if one moves one's arm, the sensory input is modified, whereas if one is being touched on the arm, and one moves one's *leg,* there is no change in the input. Similarly if one turns one's eyes towards one's *arm*, one will generally see something touching the arm whose action is correlated with the sensory input, whereas this is not so if one turns one's eyes towards one's *leg*.

A striking confirmation of these ideas which has received attention over recent years. In the "rubber hand illusion" (Botvinick and Cohen 1998; Makin et al. 2008), a person watches a rubber hand being stroked while at the same time their own, out-of-sight, real hand is simultaneously stroked in synchronous way. Most people after a few minutes get the peculiar impression that the rubber hand belongs to them.

A related phenomenon concerns the well known fact that when one drives a car, one has the impression that one extends the boundaries of one's body to the boundaries of the car. One can feel the tires against the curb when one parks, for example. And yet the information one gets about one's tires comes very indirectly from the contact one has with the seat where one is sitting, or one's fingers on the steering wheel. How can this come about? The sensorimotor theory provides an obvious explanation: the sensorimotor laws that govern the tactile input from the seat of the car and from one's fingers are best described in terms of tires against the curb.

Another example is the roughness or smoothness of the paper at the tip of one's pen as one writes. The information from the pen comes in a distributed way across the palm of one's hand and on one's fingers. How does the brain project this very complex information to the tip of the pen, and how does it do this independently of the way one holds the pen? Under the sensorimotor approach this is easy to understand. The job of the brain is to abstract sensorimotor laws from input and output, and the best way of characterizing the input–output relation that holds when one writes with one's pen is in terms of the pressure at the tip of the pen.

Color

An exciting application of the sensorimotor approach is to the question of the sensory qualities of color. Color is the philosopher's prototype for a pure phenomenal experience. For that reason, a good test of the sensorimotor approach is to try to account for why red looks red rather than like another color. Appealing to neural channels can only come up against an infinite series of questions like: "Well then what is it about the red channel that gives that red sensation?" The sensorimotor account on the other hand must seek some sensorimotor law that characterizes individual color sensations and distinguishes them from other colors.

Usually one considers color sensations not to involve any kind of action, and so taking this stance at first seems counterintuitive. However by considering color not as a property of light per se, but as a property of reflecting surfaces, it is possible to conceive of color in a way that depends implicitly on action. The sensorimotor approach considers that the sensation of color is not a thing generated by the brain, but an abstract law: the law that describes the way surfaces change incoming light [the philosopher (Broackes 1992) has also made this suggestion for color]. When one moves a surface about, tilting it this way and that so it reflects light from different illuminants, the reflected light changes in a systematic way. The law that governs such changes is what characterizes a color. With a collaborator I have applied this idea to two classical problems in color science. The first is the problem of why certain colors tend to be considered prototypical or "basic" or "focal" throughout all cultures in the world (in particular red, yellow, green and blue are the four colors most frequently considered basic). The second problem is the problem of exactly why certain very precise hues of red, yellow, green and blue are perceived as being "pure" or "unique", that is, not containing other colors. Without any parameter adjustments and purely from first principles, our results (Philipona and O'Regan 2006) provide more accurate predictions for these two problems than classic neurophysiology-based approaches, thus giving credence to the sensorimotor approach. Of course more work needs to be done for further sensorimotor specification of color experiences, but this is an encouraging first step. Other work (Bompas and O'Regan 2006a, b; Richters and Eskew 2009) using an adaptation paradigm adds further evidence of the implicit role of action in color sensation by showing that perceived color quality can be modified by modifying the usual dependence of color on eye movements.

Sensory Presence

The preceding sections have considered examples of how the sensorimotor approach provides an explanation of why sensory feels are the way they are. Why hearing is like hearing and not like vision; why touch is felt where it is felt; why colors look the way they do. The other question that appeared logically impossible to explain from a neurophysiological point of view was the question of why sensory feels feel like something, whereas other brain processes, like autonomic functions and thoughts and imaginings have no sensory presence.

Under the sensorimotor view we must not look in the brain for a way to explain this difference. We must look for differences in the laws that govern our interactions with the environment when we engage in a sensory experience, as compared to engaging in autonomic functions or thoughts.

To approach the problem, let us again take the example of softness. Why is there "something it's like" to feel softness? The reason seems obvious: it is because when one presses something soft, one is really doing something, not just thinking about it or letting one's brain deal with it automatically. But then what is it about a real interaction with the world that allows us to know that it is real? How does one know, when one is squishing a sponge, that one really is squishing it, and not just thinking about squishing it, hallucinating or dreaming about it? Why do seeing, hearing, touching, tasting and smelling provide us with real sensory feelings, whereas autonomic processes in the nervous system produce no feeling? I suggest the answer lies in four aspects of real-world interactions: richness, bodiliness, insubordinateness and grabbiness.

First of all, the world is *rich*. There is so much information in the world that one cannot possibly imagine it. If one is just thinking about squishing a sponge, one cannot imagine all the different possible things that might happen when one presses here or there. If one really is seeing a visual scene, wherever one looks, the world provides infinite detail. Richness is a first characteristic of real-world interactions that distinguishes them from imagining or thinking about them.

Bodiliness is the fact that voluntary motions of the body systematically affect sensory input. This is an aspect of sensory interactions which distinguishes them from autonomic processes in the nervous system and from thoughts. Sensory input deriving from visceral autonomic pathways is not generally affected by voluntary actions. Digestion, heartbeat, glucose in the blood, although they do depend somewhat on movements, are not as intimately linked to them as is sensory input from visual, auditory and tactile senses. If one looks at a red patch, and moves one's eyes or body, then the sensory input changes dramatically. If one is listening to a sound, any small movement of the head immediately changes the sensory input to one's ears in a systematic and lawful way. If one is thinking about a red patch of color or about listening to a sound, then moving the eyes, the head, the body, does not alter the thought.

But note that in real world interactions, bodiliness is not actually complete. This is because what characterises sensations coming from the world is the fact that precisely they are not completely determined by body motions. The world has a life of its own, and things may happen: mice may move, bells may ring, without us doing anything to cause this. I call this *insubordinateness*. The effect of the world on our sensors partially escapes our control, because not all sensory changes are caused by our movements. Some come from the world.

And then there is *grabbiness*. This is the fact that sensory systems in humans and animals are hard-wired in such a way as to peremptorily interfere with cognitive processing. When there is a sudden flash or loud noise, we react automatically by orienting our attention towards the source of interruption. This fact is an objective fact about the way some of our sensors—namely precisely those of the five classic sense modalities, are wired up. Visual, auditory, tactile, olfactory and gustatory systems are able to interrupt ongoing cognitive activities and cause an automatic orienting response. On the other hand a sudden change in blood sugar or in other autonomic pathways like a sudden vestibular or proprioceptive change, will not cause exogenous orienting. Of course such changes may make one fall over, or become weak, for example, but they do not directly prevent my cognitive processing from going on more or less as normal—although there may be indirect effects of course through the fact that I fall over or become weak.

The idea is that what we call our real sense modalities are precisely those that are genetically hard wired so as to be able, in cases of sudden change, to interrupt our normal cognitive functioning and cause us to cognitively orient towards the change. Those other, visceral, autonomic sensing pathways, are not wired up this way.

Note that grabbiness allows us also to understand why thoughts are not perceived as real sensations. If one is seeing or hearing something, any change in the environment immediately creates a signal in the transient detectors and alerts one that something has happened. But suppose for example that overnight some neurons die in one's brain that code the third person of the latin verb "amo". Nothing wakes one up to tell one this has happened. To know it, one has to actually think about whether one still remembers the third person of amo. In general, except in the case of obsessions, thoughts and memory do not by themselves interrupt cognitive processing in the way that loud noises and sudden flashes or pungent smells cause automatic orienting.

An example of the application of the notions of richness, bodiliness, insubordinateness and grabbiness is Change Blindness.

Change Blindness

When we see a visual scene, we have the impression of seeing everything, simultaneously, continuously, and in all its splendor. Why do we have this impression of reality? Richness, bodiliness, insubordinateness and grabbiness provide an explanation.

The visual world is very rich, much more detailed than any imaginable scene. It has bodiliness because whenever we move our eyes or body, the input to our eyes changes drastically. And it is insubordinate because our own movements are not the only thing that can cause changes in input: all sorts of external changes can also happen. And there is also grabbiness.

Usually, if something suddenly changes in the visual scene, because hard-wired transient detectors in the visual system automatically register it and orient one's attention to it, one sees the change. But if the change is made so slow that the transient detectors are not triggered, then an enormous change can happen in a scene without one's attention being drawn to it (Simons et al. 2000; Auvray and O'Regan 2003). An example is shown in this video: http://www.kevin-oregan.net/sol_Mil_cinepack.avi, where almost a third of the picture changes without most people noticing it.

This 'progressive change' demonstration is part of a whole literature on "change blindness". Change blindness can also occur if the transient that normally triggers attention to move to the location of a change is drowned out by a large flicker or localized "mudsplashes" on the image, or by eye saccades, blinks, or film cuts, or even by real life interruptions (Simons and Ambinder 2005; Simons and Rensink 2005). Change blindness illustrates how—along with richness, bodiliness, and insubordinateness—grabbiness is important in determining the feeling of presence we have of our visual world. When we artificially remove grabbiness by using different change blindness paradigms, we no longer have the impression of seeing things.

The Phenomenality Plot

Change blindness is an example of how the notion of grabbiness helps explain why visual sensations have sensory presence, whereas just thinking or imagining does not. Ultimately my claim is that what people call the "what it's like" or the "presence" of any real sensory experience can be related to the degree to which objective characteristics such as bodiliness, grabbiness, insubordinateness and richness, and possibly others, are involved in the sensorimotor interaction. My claim is that if we probe carefully into what people actually mean when they say an experience has "something it's like", we will find that each and every aspect of people's statements will correspond to such objective characteristics of sensory interactions.

The fact that the degrees of bodiliness, grabbiness, insubordinateness and richness can be deployed upon a continuum further suggests that sensory presence is not all-or-none, and that it may be present to different degrees in different experiences. In fact it is possible to classify experiences, even those that are not part of the classic list of five senses, for example emotions, pain, hunger, and even states such as embarrassment or loneliness, on what I call a "phenomenality plot". Plotting the amount of bodiliness, grabbiness, insubordinateness and richness involved in an interaction predicts the degree to which people will agree that such experiences possess "something it's like" or "presence" or "phenomenality". Such a phenomenality plot allows us to understand why and when different experiences will be taken to have the "presence" of sensory experience. Why for example, sometimes even thoughts (like alarming or obsessive thoughts) may have a high degree of "presence".

A particularly interesting case is pain. As for all sensations, it makes no sense to search for brain structures that might generate the qualitative "hurt" associated with pain. It is no use saying that the "hurt" is generated by special nociceptive pathways or pain circuits. To hope to find an explanation for "hurt" in these pathways and circuits themselves would lead to an infinite regress of questions, as I have already emphasized. Moreover, the argument some make that our species has survived because the hurt of pain has prevented us from continually damaging ourselves only tells us that the hurting must cause avoidance. It fails to answer the question of why the hurting has a feel: after all, we could imagine having automatic avoidance reactions, with negative cognitive valence, without them having any feel.

As an alternative, the idea of the sensorimotor approach is to examine what the hurt of pain really consists of, and in particular to examine how pain fares with respect to its sensory presence, as indicated in the phenomenality plot. Clearly, compared to other sensations, pain has only a moderate amount of bodiliness. Moving one's body can somewhat modify the pain (one can remove one's finger from the fire, rub the aching limb, and thereby change the sensory input), but some pains change only moderately (one cannot remove them completely) when one moves the body part itself or when one touches the affected body part. I suggest that this reduced ability to modulate the sensory input of some pains by body motions, that is, its reduced bodiliness, explains the unlocalized quality of some pains, contrary to the clearly localized quality of touch, vision, hearing, and taste (smell is also a case that depends even less on body motions, and as a consequence has a spatially unlocalized quality i.e. a more "inner" nature).

If pains have only moderate bodiliness, then why are they still experienced as having a sensory quality? Why are they not more like thoughts, which have very little bodiliness and so are not perceived as spatially localized? Because unlike thoughts, pains have the other characteristics of sensory experiences: They have richness and partial insubordinateness, and, most important, they have grabbiness to an extraordinary degree. It is virtually impossible to prevent oneself from being distracted by a noxious stimulation. One's cognitive processing is interrupted and one cannot function properly as long as there is a pain.

So we see that pain is different from the basic senses in having less bodiliness, but much more grabbiness. Is this sufficient to account for the very strong specificity of pain: for the fact that pain actually "hurts"? Intuitively, extreme grabbiness does not seem sufficient. Hurting is more than just having one's attention very strongly grabbed by a sensory stimulation whose localization may not be strongly defined. As evidence, consider the sound of chalk or fingernails scraping on a blackboard. This stimulation is extraordinarily grabby for some reason and makes many people cringe, though it is clearly not painful. We need a better characterization than just extreme grabbiness to account for the "hurt" of pain. A possibility suggested in O'Regan (2011) is that there are two components to pain: a purely sensory one which in some sense does not "hurt": it simply causes avoidance reactions and autonomic stress reactions in the body such as tachycardia, hypertension, sweating and midriasis. And a "hurt" or "suffering" component which is added in over and above the sensory component and provides a strong negative cognitive and social evaluation. A similar idea of a dual nature to pain has been suggested by Grahek (2007), using evidence from hypnosis, morphine, cingulectomy and lobotomy, where patients' autonomic reactions are normal: they say they feel the pain, but they say that the pain "does not bother them"; and also from patients with congenital insensitivity to pain and pain asymbolia, who have the converse affection: the normal withdrawal reactions are absent [see also Aydede (2010)].

Consciously Experiencing a Feel

The previous section provides a way of understanding the perceived qualities of feels: why they differ from one another in the way they do, and why people tend to say that feels have "sensory presence" or "something it's like". But in discussing the qualities of feel, I have not adjudicated on whether these feels are *conscious*. Indeed, the definition of feel I have used simply involves an agent being involved in a sensorimotor interaction with the environment. As such, everything I have said could apply to any kind of agent, even as simple as a thermostat, since even a thermostat senses the environment and acts in accordance with it, thereby defining a sensorimotor interaction. What then is needed for a feel to be conscious?

I suggest the answer is provided in the first part of the paper: to be consciously experiencing a feel, an agent simply has to have conscious access to the ongoing sensorimotor interaction. By “conscious access” I mean precisely what I described in the first part of the paper, namely cognitive access to the fact that the agent (with a self) has cognitive access to the feel. A thermostat not only does not have the concept of self, it does not have the cognitive capacities to cognitively access what it is doing, let alone cognitively access the fact that it has such access. So a thermostat is not conscious.

But to the extent that an agent has the necessary cognitive capacities, I would say that the agent is consciously experiencing a feel.

To illustrate, consider the case of an absent minded driver, discussing with a partner but nevertheless expertly maneuvering in the traffic and stopping at a red light. The driver has made use of the red light in his behaviour, planning and decision-making, and so by the definition of cognitive access, he had cognitive access to the light. But one would probably only want to say that the driver experienced the red light *unconsciously*. In order for the experience additionally to be conscious, the driver must have been poised to mentally stand back from his experiencing of the red light so as to be able to make further use of the redness of the light in a wider context of rational behaviour, planning or decisions: e.g. he must have been poised to do such things as note to himself that the shade of red was pleasant, or to say to his partner: “We’re going to be late because of that red light”.

Thus, following what was explained in the first part of this paper, normal use of the phrase “having a conscious experience”, apart from requiring an agent with the concept of self, involves the two-tiered hierarchy of cognitive access. A lower level where the agent is poised to make use (in its behaviour, planning or decisions) of information gleaned from what it is doing. And a higher level where the agent is poised to mentally stand back and make further use (in more rational behaviour, planning or decisions) of the information that he has gleaned.

Conclusion

Can robots be conscious and feel? The argument presented here holds that the answer is: potentially, yes. Curiously however the difficulty involved is not where mainstream opinion would place it. Indeed, in current philosophical thought, consciousness has been divided into an easy, or at least not-so-hard, problem, and a “hard” problem (Chalmers 1997). The “easy” problem is generally considered to be the problem of access consciousness. The “hard” problem is the problem of phenomenal consciousness, or “feel”. What I have shown here is that the reason why the “hard” problem has been considered hard may be that people have been thinking about feel the wrong way. They have made a category error, analogous to the error made by vitalists in thinking that life was a thing that could be generated by something like a vital spirit. If we correct this error and instead consider “feel” to be a way of interacting with the environment, the difficulty of explaining the properties of feel dissipates. We can understand why feels differ in the way they do, and why they have “something it’s like” or sensory presence. Consciously experiencing a particular feel then amounts to having conscious access to the interaction with the world which constitutes that feel.

The problem then remains of understanding conscious access, the supposedly "easy" problem. In fact we have seen here that a reasonable notion of conscious access involves a two-tiered, hierarchical mode of cognitive access, and it requires that the agent involved possess the notion of "self". Although these requirements are logically amenable to scientific investigation and to implementation in robots, practically we are quite far away from achieving such implementation. The reason is that the higher level of cognitive access would seem to involve the capacity to generate concepts and abstractions that we do not know how to implement today. It also involves a body of common sense knowledge that may be difficult to incorporate into a robot. It may be necessary for robots to have bodies, for them to grow up over time, and for them to have a social insertion in a way analogous to humans.

Acknowledgments This work was partially financed by French-Japanese ANR BINAAHR. I thank Tony Morse for useful comments.

References

Auvray, M., & Myin, E. (2009). Perception with compensatory devices: From sensory substitution to sensorimotor extension. *Cognitive Science, 33*(6), 1036–1058. doi:10.1111/j.1551-6709.2009.01040.x.

Auvray, M., & O'Regan, J. K. (2003). L'influence des facteurs sémantiques sur la cécité aux changements progressifs dans les scènes visuelles [The influence of semantic factors on blindness to progressive changes in visual scenes]. *L'Année Psychologique, 103*, 9–32.

Aydede, M. (2010). Pain. In E. N. Zalta (Ed.), *Stanford encyclopedia of philosophy (Spring 2010 Edition)*. Retrieved from http://plato.stanford.edu/archives/spr2010/entries/pain/.

Bach-y-Rita, P. (1972). *Brain mechanisms in sensory substitution*. New York: Academic Press.

Bach-y-Rita, P., & Kercel, S. W. (2003). Sensory substitution and the human-machine interface. *Trends in Cognitive Sciences, 7*(12), 541–546. doi:10.1016/j.tics.2003.10.013.

Block, N. (1996). On a confusion about a function of consciousness. *Behavioral and Brain Sciences, 18*(02), 227–247.

Bompas, A., & O'Regan, J. K. (2006a). Evidence for a role of action in colour perception. *Perception, 35*, 65–78.

Bompas, A., & O'Regan, J. K. (2006b). More evidence for sensorimotor adaptation in color perception. *Journal of Vision, 6*(2), 145–153.

Botvinick, M., & Cohen, J. (1998). Rubber hands 'feel' touch that eyes see. *Nature, 391*(6669), 756. doi:10.1038/35784.

Broackes, J. (1992). The autonomy of colour. In K. Lennon & D. Charles (Eds.), *Reduction, explanation, and realism* (pp. 421–465). Oxford: Oxford University Press.

Chalmers, D. J. (1997). *The conscious mind: In search of a fundamental theory*. Oxford: Oxford University Press US.

Dennett, D. C. (1991). *Consciousness explained*. Boston: Little, Brown & Co.

Dreyfus, H. L. (1972). *What computers can't do*. New York: Harper & Row.

Dreyfus, H. L. (1992). *What computers still can't do*. Cambridge: MIT Press.

Faymonville, M. E., Fissette, J., Mambourg, P. H., Roediger, L., Joris, J., & Lamy, M. (1995). Hypnosis as adjunct therapy in conscious sedation for plastic surgery. *Regional Anesthesia and Pain Medicine, 20*(2), 145.

Floridi, L. (2011). The informational nature of personal identity. *Minds and Machines, 21*, 549–566. doi:10.1007/s11023-011-9259-6.

Gallagher, S. (2000). Philosophical conceptions of the self: Implications for cognitive science. *Trends in Cognitive Sciences, 4*(1), 14–21.
Gallagher, S., & Shear, J. (1999). *Models of the self*. Exeter: Imprint Academic.
Gennaro, R. J. (2004). *Higher-order theories of consciousness*. Amsterdam, Philadelphia: John Benjamins.
Grahek, N. (2007). *Feeling pain and being in pain*. Cambridge: MIT Press.
Guarniero, G. (1974). Experience of tactile vision. *Perception, 3*(1), 101–104.
Guarniero, G. (1977). Tactile vision: A personal view. *Visual Impairment and Blindness, 71*(3), 125–130.
Hacking, I. (1986). Making up people. In T. C. Heller, M. Sosna, & D. E. Wellbery (Eds.), *Reconstructing individualism* (pp. 222–236). Stanford: Stanford University Press.
Hardcastle, V. (1999). *The myth of pain*. Cambridge, MA: MIT Press.
Hartocollis, L. (1998). The making of multiple personality disorder: A social constructionist view. *Clinical Social Work Journal, 26*(2), 159–176. doi:10.1023/A:1022819001682.
Kassan, P. (2006). AI gone awry: The futile quest for artificial intelligence. *Skeptic, 12*(2), 30–39.
Lau, H., & Rosenthal, D. (2011). Empirical support for higher-order theories of conscious awareness. *Trends in Cognitive Sciences, 15*, 365–373. doi:10.1016/j.tics.2011.05.009.
Levine, J. (1983). Materialism and qualia: The explanatory gap. *Pacific Philosophical Quarterly, 64*, 354–361.
Lifton, R. J. (1986). *The Nazi doctors: Medical killing and the psychology of genocide*. New York: Basic Books.
Lifton, R. J. (1989). *Thought reform and the psychology of totalism: A study of "brainwashing" in China*. Carolina: University of North Carolina Press.
Lifton, R. J. (1991). *Death in life: Survivors of Hiroshima*. Carolina: University of North Carolina Press.
Lungarella, M., Metta, G., Pfeifer, R., & Sandini, G. (2003). Developmental robotics: A survey. *Connection Science, 15*(4), 151–190. doi:10.1080/09540090310001655110.
Mahon, B. Z., & Caramazza, A. (2008). A critical look at the embodied cognition hypothesis and a new proposal for grounding conceptual content. *Journal of Physiology Paris, 102*(1–3), 59–70. doi:10.1016/j.jphysparis.2008.03.004.
Makin, T. R., Holmes, N. P., & Ehrsson, H. H. (2008). On the other hand: Dummy hands and peripersonal space. *Behavioural Brain Research, 191*(1), 1–10.
O'Regan, J. K. (2011). *Why red doesn't sound like a bell: Understanding the feel of consciousness*. New York, USA: Oxford University Press.
O'Regan, J. K., Myin, E., & Noë, A. (2004). Towards an analytic phenomenology: The concepts of "bodiliness" and "grabbiness." In A. Carsetti (Ed.), *Proceedings of the international colloquium: "Seeing and thinking. Reflections on Kanizsa's studies in visual cognition." University Tor Vergata, Rome, June 8–9, 2001* (pp. 103–114). Dordrecht: Kluwer.
O'Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences, 24*(5), 883–917.
Pascual-Leone, A., Amedi, A., Fregni, F., & Merabet, L. B. (2005). The plastic human brain cortex. *Annual Review of Neuroscience, 28*, 377–401.
Patterson, D. R., & Jensen, M. P. (2003). Hypnosis and clinical pain. *Psychological Bulletin, 129*(4), 495–521.
Pfeiffer, R., & Bongard, J. C. (2006). *How the body shapes the way we think*. Cambridge: MIT Press.
Philipona, D. L., & O'Regan, J. (2006). Color naming, unique hues, and hue cancellation predicted from singularities in reflection properties. *Visual Neuroscience, 23*(3–4), 331–339.
Proulx, M. J., Stoerig, P., Ludowig, E., & Knoll, I. (2008). Seeing "Where" through the ears: Effects of learning-by-doing and long-term sensory deprivation on localization based on image-to-sound substitution. *PLoS ONE, 3*(3), e1840. doi:10.1371/journal.pone.0001840.
Richters, D. P., & Eskew, R. T. (2009). Quantifying the effect of natural and arbitrary sensorimotor contingencies on chromatic judgments. *Journal of Vision, 9*(4), 1–11. doi:10.1167/9.4.27.
Rosenthal, D. (1997). A theory of consciousness. In N. Block, O. Flanagan, & G. Güzeldere (Eds.), *The nature of consciousness: Philosophical debates* (pp. 729–753). Cambridge, MA: MIT Press.
Rosenthal, D. (2011). Exaggerated reports: Reply to Block. *Analysis, 71*(3), 431–437.
Ryle, G. (1949). *The concept of mind*. Chicago: University of Chicago Press.
Simons, D. J., & Ambinder, M. S. (2005). Change blindness. *Current Directions in Psychological Science, 14*(1), 44–48. doi:10.1111/j.0963-7214.2005.00332.x.
Simons, D. J., Franconeri, S. L., & Reimer, R. L. (2000). Change blindness in the absence of a visual disruption. *Perception, 29*(10), 1143–1154.

Simons, R. C., & Hughes, C. C. (1985). *The culture-bound syndromes: Folk illnesses of psychiatric and anthropological interest*. New York: Springer.

Simons, D. J., & Rensink, R. A. (2005). Change blindness: Past, present, and future. *Trends in Cognitive Sciences, 9*(1), 16–20.

Somera, E. (2006). Culture-bound dissociation: A comparative analysis. *Psychiatric Clinics of North America, 29*(1), 213–226.

Vierkant, T. (2003). *Is the self real? An investigation into the philosophical concept of "self"*. Münster: LIT Verlag.

Ward, D. (2011). Personal identity, Agency and the multiplicity thesis. *Minds and Machines, 21*, 497–515. doi:10.1007/s11023-011-9256-9.

Wilson, M. (2008). How did we get from there to here? An evolutionary perspective on embodied cognition. *Handbook of cognitive science. An embodied approach* (pp. 375–393). San Diego: Elsevier Ltd.

Copyright of Minds & Machines is the property of Springer Science & Business Media B.V. and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.