

How the sensorimotor approach to consciousness bridges both relative and absolute explanatory gaps, and some refinements of the theory.

J. Kevin O'Regan
Laboratoire Psychologie de la Perception - CNRS
Université Paris Descartes, Paris, France

8 January 2016

Abstract

The problem of understanding how physical processes in the brain could give rise to consciousness has been divided into the problem of the "Relative Explanatory Gap" of explaining why different experiences differ the way they do, and the "Absolute Explanatory Gap" of explaining why anything can be conscious at all.

The main innovation of the sensorimotor theory is that it provides a very appealing way of bridging the Relative Gap by postulating that the quality of experiences corresponds to objective sensorimotor laws of interaction with the environment. Here I expound in greater detail how the approach deals with the Absolute Gap. I refine my previous efforts of understanding what we mean by "being conscious of something" by abandoning my previous hierarchical approach; by suggesting the concept of "mental manipulation"; and by relying more on a notion of self than I have done previously. I end up with a more variegated notion of consciousness than in previous work, that includes the ideas of a complex interwoven patchwork of consciousness, and is applicable to a variety of species. The approach suggests we need to review the links between consciousness and ethics.

Life, we know today, is not generated by a vital spirit. Indeed, it is a "category mistake" (Ryle, 1949) to think that life is the kind of thing that could be generated by anything, let alone by a vital spirit. Instead, today we take life to be simply a collection of capacities to act or interact that certain biological systems possess.

The sensorimotor approach takes a similar stance with regard to consciousness. It is a category mistake to think that consciousness is the kind of thing that could be generated by anything, let alone by the brain. Instead, consciousness should be taken to be a collection of capacities to act or interact that certain systems possess¹.

As a consequence of taking this view the sensorimotor approach claims that before even thinking about how to "explain" consciousness, the first task must be to define (or make precise) how the notion of consciousness can be cashed out in terms of capacities to act or interact, and what different capacities are involved. The enterprise of explaining consciousness must start as a definitional enterprise.

¹ (Dennett, 1991), (Pettit, 2003) and (Humphrey, 2006) among others, have also made the comparison with life, and/or favoured a dispositional or skill-based account of consciousness. However these authors seem not to have realized how such an approach facilitates an account of qualia, and in particular how it provides an obvious explanation of differences in experienced perceptual qualities within and between sensory modalities.

In a number of publications culminating in the most detailed account in (O'Regan, 2011), my collaborators and I had used this strategy to address the "explanatory gap" (Levine, 1983) or "hard problem" of consciousness (Chalmers, 1997).

As suggested by (Chalmers, 1997; Hurley & Noë, 2003), the explanatory gap can be divided into two parts: The Relative Gap: explaining why conscious experiences differ the way they do; and The Absolute Gap: explaining why "there's something it's like" at all to have a conscious experience.

The Relative Gap: Why conscious experiences differ the way they do

To explain why conscious experiences differ the way they do, the sensorimotor approach temporarily leaves aside the question of what makes an experience conscious, and concentrates only on what we mean by the quality of an experience. Furthermore, in talking about experiences, the sensorimotor approach avoids the temptation to talk about brain states or representations, since such language brings the danger of falling again into the category mistake of thinking that these states or representations might somehow generate the experiences. Instead the sensorimotor approach starts by espousing the counterintuitive stance that experiencing is not something that happens to us, but is an activity (or potential activity) that we are engaged in.

The stance is counterintuitive because it requires us to construe seemingly passive experiences such as the redness of a light or the smell of an onion in terms of (potential) activities. But it turns out that once we fully adopt this view, then we have available the apparatus to solve the problem of why experienced qualities differ the way they do: We can then identify experienced qualities with the sensorimotor laws (that is, the laws that link possible actions to ensuing sensory changes) that describe what we do when we are having the experience.

For example, the quality of softness of a sponge is constituted by -- it is nothing more or less than -- all the things that one can do with the sponge when one is checking for softness, and all the ways the sponge reacts when one handles it in this softness-checking way. Anything one can potentially say about the quality of softness of the sponge must ultimately boil down to something about how one potentially can interact with the sponge when one is checking whether it is soft.

Taking this stance for the quality of other sensory experiences, like hearing, smelling, tasting, and seeing may at first seem strained. Nevertheless the mental acrobatics necessary to adopt this view are rewarded by the advantage they give for bridging the Relative Gap of explaining why conscious experiences differ the way they do. Furthermore, the approach opens up scientifically productive research paradigms.

For example, the approach explains why different sensory modalities are associated with completely different experiences that cannot be easily compared, while within modalities experiences can be compared and contrasted (O'Regan, 2011; O'Regan, Myin, & Noë, 2004, 2005). This is something that would be impossible to explain if we thought experiences were generated by neural mechanisms: what could it possibly be about any particular architecture, connectivity, mode of synchrony or quantum phenomenon that generates the particular feel of seeing, as compared to the feel of hearing? And why are these feels incomparable, whereas within a sensory modality different feels can be compared?

The idea that sensory modalities are constituted by sensorimotor laws also opens up the possibility of new scientific paradigms. One example is sensory substitution, a domain which is today being actively pursued ((Bach-y-Rita & Kercel, 2003; Deroy & Auvray, 2012). The Rubber Hand Illusion (Botvinick &

Cohen, 1998) is another phenomenon that is an obvious prediction of the idea that what we mean by felt location on the skin can only be a consequence of possible ways of interacting at that skin location.

The approach can also explain why certain experiences -- we call them sensory experiences -- seem more perceptually "present" to us than mental experiences on the one hand, or than bodily goings-on like digestion or maintenance of visceral functions on the other hand. For example, really seeing a patch of red gives an experience of sensory presence, while just thinking about, remembering or imagining red does not. Likewise, really seeing a patch of red has sensory presence, whereas the sensory control systems controlling say oxygen in the blood or liver functioning have none. The approach explains this "sensory presence" by appealing to the notions of "bodiliness", "grabbiness" and "insubordinateness", which are objective, physically measurable aspects of our interactions with the world (O'Regan, 2011; O'Regan et al., 2004, 2005).

Another scientific implication of the sensorimotor approach is the phenomenon of Change Blindness (O'Regan, Rensink, & Clark, 1999; Rensink, O'Regan, & Clark, 1997; Simons & Rensink, 2005). Change blindness is produced by reducing the perceived "presence" of the visual field by interfering with the normal grabbiness of changes registered by visual transient detectors in the low-level visual system. It is predicted by the sensorimotor approach, with the idea that the experienced sensory presence of the whole visual field actually derives not from the whole detailed field of view being somehow passively represented in the brain, but from it being immediately available through action: the slightest flick of the eye or of attention (O'Regan, 1992, 2001).

While certainly much more remains to be done to fully exploit the sensorimotor-based account of sensory qualities -- and in a section below I will suggest that an additional quality linked to "mental manipulation" should be added -- it is clear that the account provides not just a philosophical advantage, but a richer outlook for science than current alternative accounts of sensory qualia. As an illustration of the latter point, take how the sensorimotor account provides a scientific basis for experienced sensory qualities of colour. (Philipona & O'Regan, 2006; Vazquez-Corral, O'Regan, Vanrell, & Finlayson, 2012; Witzel, Cinotti, & O'Regan, 2015) have shown how the special nature of the experiences of red, yellow, green and blue can be explained by referring to sensorimotor laws. In contrast, no neurophysiological account has succeeded in providing such an explanation (Chichilnisky & Wandell, 1999; Knoblauch & Shevell, 2001; Kuehni, 2004). Note that similarities and differences in neurophysiological codes or mechanisms cannot be used as an explanation of colour spaces, since neural codes are just codes: a principled reason for the particular link between each code and each quale must be provided by a theory of qualia. Psychophysical approaches like Austen Clark's (Clark, 1993), or Palmer's (Palmer, 1999), are also not satisfactory, since they simply observe the existence of a subjective colour quality space, without providing any principled reason for the observed similarities and differences. Finally, an account in terms of some kind of sensory primitives is also less parsimonious, since it needs to postulate a different ineffable sensory primitive for every different colour hue.

The Absolute Gap: Why there is "something it's like" to have a conscious experience

The preceding section has recalled very briefly how the sensorimotor approach deals with the *quality* of experience by identifying experienced quality with sensorimotor laws (for more detail cf O'Regan 2011). As said above, it is this strategy that gives the sensorimotor approach its advantage over other theories of phenomenal consciousness, and provides a scientific way of accounting for phenomenology, or the premise for what might be called an "analytic phenomenology" (O'Regan, Myin & Noe 2004). This is the

major innovation of the sensorimotor approach to consciousness, and it is this that I have emphasized in previous work.

However the notion of sensorimotor law applies to any arbitrary system that is interacting with its environment, e.g. a missile guidance system, or a bulldozer. I need additionally to answer the question that I had temporarily left aside in discussing qualities, namely the "Absolute Gap" question: Why is there something it's like to have a conscious experience at all?

It is hard to answer a question where the terms used are not clearly defined, and this is precisely such a question. In (O'Regan, 2010, 2011) I therefore tried to "operationalize" the question of "something it's like" by proposing that people will generally agree that real sensory perception like the sensation of a red patch of colour, have "something it's like", while just thinking about red, on the one hand, or having visceral body states like the amount of oxygen in the blood, on the other hand, has "nothing it's like". If this is accepted, then the term "something it's like" would refer to the same aspect of experience as the term "sensory presence". Sensory presence is one of the qualities of sensory experiences not shared by mental and visceral experiences. As explained in the section on the Relative Gap above, the degree to which experiences have sensory presence is well accounted for by the objective properties of bodiliness, grabbiness and insubordinateness possessed by sensorimotor interactions with the real physical world.

On the other hand, it could be argued that there is another aspect of experience that the phrase "something it's like" could characterise, namely whether the experience is "conscious". It is this second meaning of "something it's like" that I am interested in in the present paper: What makes an experience conscious?

In fact the question "What makes an experience conscious?" is formulated dangerously. Hidden in this formulation is the implicit assumption that there really is some *thing*, like a spirit, fluid, or mechanism, that actually "makes" the experience conscious. However, following the sensorimotor approach, *nothing* "makes" an experience conscious. Instead, there are ways of acting, or of interacting with the world which people tend to *designate* as being conscious ways of acting or interacting. The task of the sensorimotor approach is to find out what is particular about those ways of acting or interacting that people designate as "conscious" ways. The task is a task of being precise about definitions: what capacities to act do we want to *call* cases of acting consciously? What kinds of modes of acting or interacting do they *consist in*, precisely? And we shall see that once we are precise about definitions, the task of "explaining" consciousness melts away.

Thus a better formulation of the question "What makes an experience conscious?" would be: "What do we *mean* when we say that a person is experiencing something *consciously*?" and "What do we *designate* as conscious ways of behaving?" In (O'Regan 2011) I had already sketched a way of undertaking this definitional enterprise. In the following paragraphs I wish to explain in a different way the steps that I had taken there, and add some qualifications to my previous account.

To determine what is prototypically meant when we say that an adult human is interacting with something in a conscious way, let us proceed by considering a succession of cases where we would be increasingly willing to apply the word "conscious". And let us take the concrete example of squishing a sponge and being conscious of its softness.

The sponge testing automaton

Consider first a sponge-testing automaton with a hand that squishes sponges and sorts them according to whether they are soft or hard. Say the automaton has squished the sponge and is about to sort it into

the pile of soft sponges. Clearly we do not want to say it is conscious of the softness of the sponge. Why? First, the automaton is simply executing a pre-programmed stimulus-response reaction. To be conscious, one of the most obvious and basic requirements is that the machine should have the possibility of linking inputs to outputs in different, more flexible ways, and these different ways should be determined by something like goals that could be attributed to the system. Furthermore these goals themselves should depend on the current context and the past history of the system.

Second, being "conscious" of something is not a matter of what one does *immediately*. On the contrary, there is the notion of "taking note" of something, without reacting now, but being ready to make use of that thing at a later time. The "making use" can be a physical action or a mental action.

(Block, 1996) has a notion of "Access Consciousness" that captures these two points, although it is phrased in words that I do not want to adopt. Block says a mental state is conscious if a representation of its content is poised for use as premise in reasoning, for rational control of action, and possibly for rational control of speech. In this formulation Block is talking about "mental states" and using the terms "representation" and "content" which are all ways of thinking that the sensorimotor approach wishes to avoid, since we are attempting to qualify not (alleged) representational states, but conscious interaction or engagement with the environment. But there is an idea in his definition which we want to retain, and that is the idea of being "poised". In fact Block's definition captures the fact that when one is conscious of something, one may actually not be doing anything physical at all. On the contrary, one is ready or "poised" to do something, and furthermore, that "doing" can be, not physical doing, but "mental" doing. So we can partially adopt Block's language and say that part of what is generally meant when someone is conscious of something, is that the person is "poised" to use that thing in their rational control of action, or as a premise in their reasoning.

The futuristic dishwashing robot

Suppose therefore that instead of a sponge testing automaton, the system is a futuristic household dishwashing robot. It has hands and can grasp the sponge to wash the dishes. When it fishes the sponge out of the soapy water it has to make sure that it has not picked up the soap or the scouring pad by mistake. Let's say at this moment it is squishing the sponge to make sure it is soft. Can we now say that the robot is conscious of the softness?

We are getting closer to the normal use of the word "conscious". The robot is "poised" to make use of the fact that its sponge squishing activity obeys the laws of softness in its future behaviour: Because the sponge is soft, it is going to identify that it has the sponge and not the scouring pad, and use it to clean the plate and not to scour the pan. The robot has a variety of action-plans among which it is choosing, so it is not just a pure stimulus-response system: It can do different things, like squirt dishwashing liquid into the water, sort through different things it needs to wash, etc. At this moment one might reasonably attribute to the robot the "goal" of cleaning the plate with the sponge (rather than scouring the pan with the pad). Would we now be willing to say that the system is conscious of the softness of the sponge? Again, we would not.

Perhaps part of the problem is that albeit sophisticated, the system has still been programmed to do all these things. Although it is "poised" to make use of the softness of the sponge, for us to admit it was conscious of the softness, the system would have to somehow escape from any pre-programmed behaviour, be it ever so complex. Our intuition about what we mean when we use the phrase "conscious of something" is, I think, that the machine would have to have meta-knowledge about what it is doing, so that we could say of it that it was "knowingly" doing what it was doing, and was not simply selecting among and executing pre-programmed plans of action.

What does this mean? Presumably to "knowingly" do something the system would have to know that it has a number of possible things it might do, and that this particular thing is what it has chosen to do now. My intuition is that there must be something like a "selection" or "choice", allowing for certain things *not* to be chosen. Furthermore the system should also know that "it" exists as a system, so that we could say of it that it has a cognitive, "knowing", "observing" self, and that this self is not only making the choices, but also observing that it is making these choices. It is self-aware. (I will use the word "self" here even though there exist more basic, non-cognitive, non-narrative notions of self like the "minimal self" (Gallagher, 2006; Zahavi, 2014) which are not relevant to the prototypical sense of "being conscious of something" applied to adult humans that I am trying to characterise here.

The self-aware self

Note that a requirement for a notion of a self-aware self in the definition of being conscious of something is not generally emphasized in current theories of consciousness, be they philosophical, psychological or neuroscientific (with some exceptions e.g. (Baars, Ramsøy, & Laureys, 2003; James, 1890)). Most current theories, while not denying the the existence of the self as a cognitive and social construct, consider it to be largely orthogonal to the question of consciousness. The reason is perhaps that most current theories take consciousness to be a special phenomenon that is somehow generated by brains. From that point of view it is natural to look for the special brain properties or functioning (coherence, synchrony, recurrent networks, availability of information..) that give rise to consciousness, and these would not have any necessary link to the notion of self. However, once we take consciousness no longer to be generated by the brain, but a capacity that humans with selves possess, then understanding the involvement of the self becomes essential.

But what is a "self"? Are we in danger of being accused of circularity by introducing a self that is self-aware, or of creating a concept that is not amenable to science in our definition of what it is to be conscious of something? Not, I believe, if we can specify in behavioural terms *what is meant* by having a self-aware self.

Many thinkers have addressed the question of a cognitive, knowing, observing self, both from a philosophical, (social-)psychological and neuroscientific perspective (Baars et al., 2003; Carruthers, 2000; Dennett, 1992, 2003; Graziano & Webb, 2015; Hacking, 1986; James, 1890; Metzinger, 2004; Prinz, 2012; Rochat, 2003, 2015; Vierkant, 2003). While the approaches differ in important ways, many concur in suggesting an account of the self as a "story" or "model" that is progressively constructed by accumulating information from ongoing activities, knowledge, attitudes, beliefs, desires and motivations. Not everything that happens to the system gets integrated into the self. Part of what we mean by being conscious of something is that that particular thing among possible other things is integrated into the record that is being constructed, and can thus be made use of at a later time by the system.

In addition to theoretical efforts to characterise the self, there is also a rich empirical literature in psychology which studies the Theory of Mind capacities that underlie the self. Tests have been devised like the mirror task, and false belief tasks like the Sally-Anne task (Baron-Cohen, Leslie, & Frith, 1985; Leslie & Frith, 1988) to discern to what extent an animal or human has a self, and psychologists have studied how these notions develop through childhood (Rochat, 2003). Social psychologists have studied how the self is constructed through interaction of a person in their social and cultural setting (cf. review by (Prinz, 2012)). Common to all the approaches is the idea that the self is not a single thing "generated" by the brain, but a collection of capacities that allow modes of interaction with the world that we globally designate as being part of having a self. We need postulate no special "phase-transition-like" phenomenon in the brain that somehow engenders the self out of complexity, recurrence or self-

reference: Instead, each of the various cognitive and social capacities underlying behaviours corresponding to what we designate as having a self can be enabled by a variety of brain mechanisms working together (Baars, 1996; Baars et al., 2003). While a detailed mechanistic account has not been given, there is little doubt that science has the apparatus needed to account for each capacity. In sum, there is no circularity involved in postulating that the self is required as part of an account of consciousness. And there is no need to appeal to arcane as-yet-not-understood mechanisms to account for the fact that we use a self-referring notion of self to describe our individual and social behaviours.

What might seem puzzling about an account of the self as a continuously updated "story" is the question of how such a story could have causal effects on itself, and how it could influence the very physical system which it is instantiated in, that is, the human brain or body. But in fact many computational architectures in artificial intelligence do both of these things routinely (e.g. "virtual machines", cf. (Sloman, 2001, 2010)): Programs running on a computer can change their own code, and of course a program can affect the hardware that it is running on -- the simplest example is by turning off the computer.

Another puzzlement about the self is our intimate feeling that the self feels real, and not like a "story". However many cognitive and social "stories" are very real to humans: mental phenomena like love, embarrassment, or patriotism, perhaps even the hurt of pain are felt to be very real even though they are clearly social constructs. It could be part of the story of our selves for us to believe that our story is real (O'Regan, 2011; Vierkant, 2003).

In conclusion on the self: appealing to the self in an account of what it means to say that an agent is conscious of something does not constitute an appeal to magical mechanisms, and it does not introduce circularity into the theory of consciousness. If information processing systems can be constructed that refer to their own actions and are able to affect their own hardware, then it makes sense that in the "social" context of other systems like themselves, they should be able to use their notion of self in communicating with other systems, and ultimately communicate even with themselves and be self-aware.

The human dishwasher

So let us now envisage a more sophisticated dishwashing system, that has a self. Let us actually take a human being, who has just grasped the sponge in the soapy water and squishes it to make sure it is not the scouring pad. All this with the purpose of cleaning the dirty dish. Is the human now conscious of the softness?

Not necessarily. The human could be unconsciously squishing the sponge, while he or she is conscious of something else, like say the temperature of the water, or a discussion with a friend. A person can be involved in doing many things simultaneously -- e.g. keeping their balance, holding the dish to be wiped in the other hand, listening to the radio. Each of those things may in one way or another simultaneously be poised to determine present or future (rational, physical or mental) behaviour.

What we generally mean by saying a person is conscious of something is that that thing is being incorporated into the self-story that the person is elaborating at that particular moment. In that way the person would say that the thing is playing a role in what they as a person are doing, in the sense that their self is poised to be able to make use of the thing in the planning, decisions, rational and possibly language behaviour that this person is undertaking at that moment.

Note also that the thing we are conscious of is perceived within the more global context of our current activity, which we are partially aware of. Furthermore we are simultaneously marginally aware of our own awareness of the thing and its context. Thus one should not just conceive of being conscious of a thing in terms of a single object of awareness, but rather as a whole vista in which the thing we are attending to is in the foreground, and additional intricate contextual, meta-knowledge and self-referring aspects form a patchwork of inter-relations in the background or "fringe" (Bailey, 1999; James, 1890; Mangan, 2009).

Bridging the Relative and Absolute Gaps

Let us now summarize all the requirements we have assembled that suffice to say of a human that the person has a conscious experience with a particular quality (or one might say: that the human is conscious of the quality of an object), and then show how the account explains the two problems of the Relative Gap: why experienced qualities differ the way they do; and Absolute Gap: why there "is something it's like" to experience the quality. Note that we are looking at the prototypical case of the use of the phrase "Conscious of something", as generally used in everyday parlance for adult humans. My proposal is to adopt this prototypical use of the notion, as a "target" definition. Later I will discuss how some of the conditions I am putting forward here must be relaxed.

Let us take the "signature" case of an adult human consciously experiencing the softness of the sponge.

We are first presupposing some necessary background conditions:

- the human is a body (including a brain) immersed in the physical world
- the human (considered simply as a physical system), has sensors and effectors that give information about its current interaction with the world and its own internal bodily states
- the interaction of the human with the world (including its internal body states) can be described by many simultaneously applicable sensorimotor laws. As examples: some of the laws are constitutive of the softness of the sponge (the softness laws), or of the weight of the dish held in one hand; others are constitutive of oxygen level in the blood, vestibular interactions ensuring balance, limb movements...

The following conditions guarantee that the human is making some kind of selection in what it is applying itself to, and also in what it is doing or going to do as a consequence. The idea of selection is related to the idea of deliberation, and guarantees that the human is not just a stimulus-response automaton executing pre-programmed behaviours:

- softness of the sponge is just one of many qualities that is playing a role in the human's activities.
- there are many possible roles that softness can be put to in the human's future (mental or physical) behaviour. One possible role is to determine whether the item being grasped is a sponge or a scouring pad.
- the setting or context (specifically: dishwashing) in which the human is doing what it is doing is one of many possible contexts.

The requirement for a self:

- the human has sufficient cognitive capacities to have a self, which is a record being continually updated that contains knowledge about the events that affect the human, and includes dispositions like beliefs, desires and motivations. Part of the knowledge possessed by the human is also the fact that it exists as a human, and that it has a record of its own properties. As a consequence, in this prototypical case, we can call the human, and the human calls itself: a person.

Then the person (and an outside observer) will say they are conscious of the experience of softness (i.e. the person will ascribe the experience to himself/herself) if:

- the person has selected to verify the applicability of one particular set of sensorimotor laws among many possible, namely the softness-of-the-sponge laws, and the person is poised to make use of the fact that they are relevant to the person's present or future decisions and physical or mental behaviour.

And now with respect to the two explanatory gap questions we can ask:

Relative Gap: Why is the softness experience the way it is? The answer is that the human is interacting in a certain way with the sponge, and the laws that determine that interaction *constitute* what we mean by the quality of the experience. Everything that can be said about the feeling of softness is contained in the things that the human is or might be doing with respect to the soft sponge. In particular the laws describing softness are different from the laws of hardness. They are also very different from the laws that characterise auditory or visual sensations, for example. Finally the laws have the properties of bodiliness, grabbiness and insubordinateness which guarantee that the experience of softness has sensory presence. This is in contrast to the laws that underlie mental processes, which have no bodiliness, grabbiness and insubordinateness and so no sensory presence. And it is in contrast to the laws that determine oxygen level and vestibular interactions ensuring balance, which lack some or all of the properties of bodiliness, grabbiness and insubordinateness, and are therefore not experienced as having sensory presence.

Absolute Gap: Why is the softness experience conscious -- or better: why is the person conscious of the softness? One could say that *unconsciously* the human as a global physical system is experiencing all sorts of things, for example the weight of the dish held in one hand, the oxygen level in the blood, vestibular interactions ensuring balance, or limb movements. Why, among all these things, is the person experiencing the softness *consciously*? The answer is that what we *mean* by being conscious of something, is that the self of the person is poised in a particular way with regard to the softness. Since this is the case here -- with the person not simply being a stimulus-response system, and actively making use of the softness to choose the sponge and wipe the plate, and with the self incorporating this use of the softness into the record that constitutes the self -- our definition is satisfied and the person is conscious of the softness.

Being conscious of things that are not sensory qualities: mental manipulation

Above I have attempted to characterise what *people prototypically mean* when they say that they are conscious of the quality of a sensory experience. I have given the example of the quality of the softness of a sponge, but I intend the account to apply to any sensory quality, for example the particular red hue of the light or the smell of truffles in the omelette. As said in the section on the Relative Gap, it may require some mental acrobatics to accept that seemingly passive sensations like a hue of red or the smell of truffles can be treated in the same active way as the softness of a sponge, but this is the wager of sensorimotor theory.

In addition to being conscious of experiences, one can additionally be conscious of other types of things. Indeed being conscious of an experience itself is probably a minor part of what it is usually to be conscious of something. When I say I'm conscious of the softness of the sponge, usually I do not actually mean that I am savouring the feeling of the sponge's softness: rather I am mainly conscious of the *fact* that the sponge is soft rather than hard. The sensorimotor law, the procedure of pressing and testing that the sponge squishes softly, is not normally the main object of my consciousness. Similarly when I'm driving and stop at the red light, it is generally not the experiencing of that hue itself that I am conscious of. It is the *fact* that the light is red (rather than orange or green), and that this implies I must stop the car.

When it comes to experiences then, when we say we are "having a conscious experience of softness", the phrase is ambiguous. It might mean that I am consciously attending to my softly-squishing of the sponge; or it might mean that I am conscious of the fact that the sponge is soft, or to some mixture of the two.

Furthermore, the quality of the accompanying experiences will be consequently different in each case. When I am consciously attending to my softly-squishing activity, the experience will be constituted by the sensorimotor laws of softly-squishing. But when I am consciously attending to the fact that the sponge is soft, while there is the sensorimotor interaction of softly-squishing going on in the background, the main component of my experience presumably involves a *mental* kind of manipulation. I am mentally implicitly probing the question of whether the sponge is hard or soft, and coming to the conclusion that the sponge is soft.

This mental manipulation could probably be said to have a quality itself. Just as the quality of a sensory experience is constituted by the sensorimotor laws involved, we could conceive of the quality of mental manipulation to be constituted by the "laws" of mental manipulation. It is not entirely clear how to make the analogy with sensorimotor laws, since for mental manipulation there seems to be no obvious correlate of the notions of actions and sensory consequences that we have for sensorimotor laws. Perhaps one could invoke mental "actions" as consisting of querying one's own knowledge and memory, and the "consequences" being the information that one retrieves from such queries. From this there might be a way of understanding what could be considered qualities of mental activity such as: "certitude", "doubt", "puzzlement", "surprise", "complexity", "humour", "boredom", "simplicity", "clarity", "elegance", "analogy", "metaphor", etc.

Furthermore it should be noted that mental "manipulations" very often give rise to bodily effects. Presumably emotions are a prototypical case where mental evaluation of one's situation, conditions, motivations, beliefs, desires, can give rise to "states" (actually modes of potential interaction) such as "happiness", "sadness", "fear", where both our bodily propensities to act are modified, and our visceral functioning may change. Thus there may be bodily and visceral products of mental states which through their sensorimotor effects spill over into our sensory experience.

In conclusion, the quality of a conscious experience may be constituted partly by the ongoing sensorimotor laws, and partly by the quality of mental manipulation that is involved. In most cases there will be a mixture of the two. When we are conscious of an experience itself, without a significant mental classificatory component, then the quality is entirely constituted by the sensorimotor laws. This may be the exception rather than the rule (and possibly accessible only to expert meditators!). When we are conscious of a purely mental thing, like a thought, an idea, a fact, then what we experience is mainly the quality of the accompanying mental manipulation. This mental activity may in turn produce changes in propensity to act and in visceral functioning that modify sensorimotor laws and so sensory experience.

Usually we are conscious of things for which there is both a sensory and a mental component, and the associated experience is partly sensory and partly mental. Furthermore part of the mental component of the experience will derive from the self-referential knowledge that it is we ourselves who are currently conscious, and that we are engaging in a particular task in a particular situation. Consciousness of something should thus be conceived as involving a complex patchwork or vista of inter-related foreground and background aspects of that thing, including self-referential knowledge about our selves being conscious of it.

Note that the complexity or self-referent aspects must not be taken in any sense to be "generating" consciousness. It is simply what we usually mean by being conscious, and it can involve different

components and degrees of complexity. Deciding which of such variants we want to properly ascribe the term "conscious" to is a matter of definition, and not a matter of fact.

And we should probably resist the temptation to think that at some level of complexity or self-reference or recurrence, some special phenomenon comes into play that suddenly gives rise to the possibility of consciousness. The self, whose gradual appearance during development is certainly facilitated by the onset of language and a rich social context, is a progressively maturing and complex collection of capacities, and there seems to be no reason to think that there is some "critical point" where it somehow coalesces into full existence like a phase transition in physics. Even if it is true that there may be plateaux of selfhood as measured by developmental psychologists, ascribing consciousness to one or other level of the plateau is a matter of definition and not a matter of fact.

Consciousness without a body

The notion of "being conscious of an experience" or "conscious of something" that I have dissected above plausibly corresponds to one prototypical use of the word "conscious" for adult humans.

But the way the characterisation is formulated, it assumes that the person who is conscious has a body in addition to their brain. What then shall we say of people who no longer have access to their bodies, like locked-in patients, or people born with inoperative bodies? Further, what is the role of the body's internal functioning and metabolism in this characterisation? Do we need to attribute some special role to the fact that human bodies are equipped with all sorts of homeostatic, autopoietic or even thermodynamic mechanisms that ensure survival (Bickhard, 2009; Deacon, 2011; Maturana, 2014)? And what about possible future artificial agents with no physical, acting bodies, yet capable of symbolic thought?

The need for a body in the prototypical definition of consciousness I gave above derived from the intuition that what we mean by sensory experience necessarily requires bodily interaction with the world. This was the "trick" which allowed us to provide an objective account of the quality of sensory experience in terms of the sensorimotor laws that constitute them. But while the body was necessary to account for *sensory* qualities, as suggested above, *mental* qualities might analogously be accounted for in terms of laws of "mental manipulation".

From these points we can therefore argue that the conscious experiences of agents deprived of bodies should differ from what normal adult humans experience, since such agents would not be simultaneously interacting with the world and experiencing the accompanying sensorimotor laws. They would only be experiencing the quality of the mental manipulation. Furthermore, in normal humans, mental states often give rise to bodily and visceral changes, which then participate in changing the sensory experience. A person without a body would not experience these accompanying sensorimotor effects of mental manipulations².

There is presumably also a distinction to be made between different cases where an agent lacks a body. One is the case of a person who has previously had a body and is now deprived access to it, perhaps because they are recently paralyzed, or because they are being prevented from moving. A second case is the case of a person who has never had a body -- say born without a body or tetraplegic -- or a future disincarnated artificial intelligence.

When stimulated by some sensory input, the brain of a person who has previously had a body will be able to classify that input as belonging to this or that sensorimotor law that the person has previously

² However a brain in a vat with a simulated body, would, if the simulation was properly done, have the sensorimotor experience.

experienced. Part of the experience of such a person will therefore not be so different from a person who actually still has a physical body. On the other hand, as said above, mental functioning often gives rise to bodily and visceral changes, which in turn produce a sensory experience. People without bodies would not experience such accompanying sensory effects.

The case of an agent who had never had a body would be quite different, I suggest, because their brains would never have learnt the sensorimotor laws associated with exteroceptive bodily interactions, and there could be no visceral repercussions of their mental states. An example might be the case of a hypothetical artificial cognitive agent interacting with other agents on the internet. It is conceivable that such an agent might have symbolic interactions with the "world" of communication that it inhabits, and could possess its own identity, its own internal processes and states. The capacities of the agent might be so rich, and its cognitive capacities so great that it would make sense to say that it had a self. In that case, the notion of being conscious of something that I have elaborated above would also apply, except that what we called sensorimotor interactions for humans, should be replaced by the input-output relationships that describe the interactions of the agent with the symbolic world "outside" of it. The agent's "experience" would then correspond to the laws that describe these input-output relationships.

Consciousness without a self

Another requisite for being "conscious of something" in the characterisation I have proposed is the need to have a self. But then what of babies whose (narrative) selves are presumably not well developed till they are 3-4 years old (Rochat, 2003)? And what of animals like dogs and cats which have very limited notions of self as compared to humans? Shall we say they are not conscious at all, or shall we say they are "partially" conscious? There are also neuropsychological patients whose memory (either long- or short-term) is so limited that they can hardly construct a "self", understood as constituting a record of prior and accumulating knowledge about the self. Or people in a vegetative state whose mental activity may be even further reduced. Are they conscious?

If we want to remain consistent in our definition of "being conscious of something", understood in the way the phrase is normally applied to adult humans, then we have to allow that agents with restricted selves must have correspondingly restricted abilities to be conscious of things. If being conscious of something is understood as a self "contemplating" a patchwork of foreground and background items, among these items being the realization that its own self is part of the patchwork, then we can envisage that agents with less well established selves have a correspondingly less intricate or complexly structured vista. But because of the possible interrelations of reference and self-reference within this patchwork, degrees and variants of consciousness probably should not be considered on a linear scale from "not conscious" to "completely conscious", but rather as being able to manifest themselves in a more or less richly variegated way.

Then what of the question of whether agents with restricted selves can have conscious experiences?

Under the usage I am recommending of the notion of "being conscious of something", agents with restricted selves will only be able to have experiences precisely to the limited degree that they are capable of being conscious of something.

While this way of thinking has the advantage of providing a consistent use of the word "conscious", it has the disadvantage that it suggests that we should consider that very young babies, adults who have very severe mental pathologies, and most animals, do not have conscious experiences in the normal sense of the phrase as applied to adult humans. Given that in our society consciousness is often taken as a criterion for being treated with special respect, this labelling incurs ethical dangers.

On humanist grounds we may thus be relieved that authors such as (Gallagher, 2006; Zahavi, 2014) propose that even organisms like babies and animals can have "minimal selves", in which case perhaps we can still apply the notion of consciousness to them. My conviction is that this fuels confusion, because the normal sense of the phrase "conscious of something" requires a "self" that is sufficiently cognitively developed to be able to reason, cause actions, create memories, store knowledge. Without such kinds of mentality, the notion of consciousness is a completely different breed from that used when we say we are "conscious of something" in everyday parlance.

Instead of corrupting the notion of consciousness, and thereby promoting continuance of confused, reified notions of consciousness that precipitate scientists into a quest for spooky mechanisms that "generate" consciousness, I suggest we modify our ethical reasoning, and remove the issue of consciousness from the criteria that underlies ethical judgments.

Comparison with O'Regan (2011)

In O'Regan (2011) I had made a similar attempt to characterise the everyday use of the phrase "being conscious of something" by imagining a chess playing automaton, and building into it progressively more capacities. I had introduced the notion of what I called "cognitive access". Cognitive access was understood to incorporate the ideas I have sketched here concerning the necessity for the system to be "selecting" both among possible inputs to consider, and among possible (potential) actions to perform. I had then employed the notion of cognitive access in a hierarchical way, as in: "the system has cognitive access to the fact that it has cognitive access to something" as a way of capturing what people *mean* when they say that the system is conscious of that thing. I had then further added the condition that the system concerned required a notion of "self" in order for it to make sense to say that the system was conscious of something.

The hierarchical use of cognitive access seemed intuitively appealing in that it provided a way of incorporating the notion of meta-knowledge that seems to be present when one is conscious of something. It resembled the notion of higher order thought well-known in the philosophy of consciousness (Carruthers, 2011; Rosenthal, 2004), except that there the hierarchy refers to mental states and contents, which are notions that cannot be used in the sensorimotor theory.

It now seems to me that the need for there to be meta-knowledge and self-knowledge involved in the characterisation of being conscious of something might actually already be satisfied by postulating the presence of an (observing) self. In that case postulating a hierarchy of cognitive access would be superfluous. Further work should investigate the comparative advantages of the different formulations, but it would seem that the new one I am suggesting here is more economical.

The new formulation has the further advantage that it more easily accommodates William James's well known idea of "fringe" of consciousness (Bailey, 1999; James, 1890; Mangan, 2009), which is that a person is not generally conscious of a single thing. Rather they have what could even more accurately be called a vista of consciousness, where one or a few things are in the foreground, and a variety of other things are present more in the background, but in a patchwork of complex inter-relations, including self-referring aspects.

A third new idea that I have presented here as compared to O'Regan (2011) is the idea that mental manipulation might be considered to have a quality. Furthermore I have suggested that such mental manipulation very often impacts on a person's bodily actions and states, indirectly creating properly

sensory experiences through the accompanying sensorimotor laws. Such an approach might be further developed in an account of emotions, which I have hardly mentioned here.

The accessibility to science comes from framing consciousness as exercise of skills

The account of what it is to consciously experience something that I have given could certainly be developed much further. However the important point to note is that the account completely bridges both the relative and the absolute explanatory gap questions, and is completely accessible to science.

The reason for this success lies in the fact that we have been precise about what we *mean* when we talk about the quality of experience, and by what we *mean* when we say of a person that they are conscious of something. We describe these things in terms of (potential) actions or interactions, so that we can easily conceive of computational and physical mechanisms that enable these behaviours. There is no special fluid or essence that we need to postulate, no quantum gravity effects or widespread synchrony of brain oscillations or reverberations, no global brain communication, no highly integrated information. Such mechanisms may actually exist in the brain, and they may contribute to enabling the capacities that we designate as corresponding to having a conscious experience. But if they do, then it is not the special, complex, arcane, or yet-to-be-understood properties of these mechanisms which somehow "generate" consciousness, or make it emerge from complexity. It is simply because the mechanisms contribute to enabling those particular potential behaviours that we designate as corresponding to conscious experiences.

As one illustration of the last point, consider the Global Workspace Theory of Baars and its more modern neuroscientific development by Naccache et al. (Dehaene & Naccache, 2001). Or take Tononi's information integration theory of consciousness (Tononi, 2004). These approaches propose the idea that consciousness ensues when information in the brain is "widely available" or "integrated". From the point of view of the sensorimotor approach, while possibly not false, making such a claim is misleading. It is misleading because consciousness does not "ensue". It is not generated. It is not a product or emanation of the brain. Consciousness is a word that applies to people who have certain capacities with respect to certain things they are attending to. These capacities, in order to be present, probably require information to be widely available or integrated in the brain. But it is not this wide availability or integration that "generates" consciousness. Now it may be that the proponents of these theories would actually agree with my statements. But then, where lies the interest in these theories, if they are merely saying that consciousness involves the simultaneous participation or integration of information from many brain systems? My suspicion is that the fascination for such theories comes from an unavowed slip into the fateful category error: There is the idea that some kind of new phenomenon, analogous to a phase transition in physics, takes place when wide-scale brain activity or information integration "ignites", passes some threshold, and "generates" consciousness.

Furthermore, such theories, while they may be pinning down the neural mechanisms that underlie what allows people to be conscious of something, have nothing to offer in an explanation of why conscious experiences have the particular sensory qualities that they do (e.g. why the sponge feels soft rather than hard, or why red has a visual rather than tactile quality, and why thought and visceral processes have no properly sensory qualities). On the other hand, the sensorimotor account provides an obvious explanation for such differences in terms of sensorimotor laws.

One final point concerns ethical issues. If we accept the sensorimotor approach's suggestion that consciousness should, like life, be conceived of as the exercise of skills, including mental skills, and not as an emanation of brains, and if we accept that the relevant skills are variegated and inter-related in complex, sometimes self-referring ways, then the issue of what we call properly conscious becomes the

issue of selecting somewhat arbitrary criteria in what may not even be a linear continuum between conscious and "non-conscious". We no longer have a factual basis for saying that this or that creature is conscious. Many ethical theories are based on the notion of sentience, if not consciousness, and these will be impacted by this realization. It will no longer be so helpful to base ethical decisions about pain, anesthetics, locked-in patients, capital punishment, animals we sacrifice, on whether these people or creatures are conscious, since consciousness is not a single thing, but a patchwork of variegated capacities, within which it would be difficult to define one or other as being primordial.

Acknowledgements

I gratefully acknowledge support from ERC Advanced grant "FEEL", No. 323674. I am extremely grateful to Peter Boltuc and Jan Degenaar for extensive comments.

Bibliography

- Baars, B. J. (1996). Understanding subjectivity: Global workspace theory and the resurrection of the observing self. *Journal of Consciousness Studies*, 3(3), 211–216.
- Baars, B. J., Ramsøy, T. Z., & Laureys, S. (2003). Brain, conscious experience and the observing self. *TRENDS in Neurosciences*, 26(12), 671–675.
- 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
- Bailey, A. R. (1999). Beyond the fringe: William James on the transitional parts of the stream of consciousness. *Journal of Consciousness Studies*, 6(2-3), 141–153.
- Baron-Cohen, S., Leslie, A. M., & Frith, U. (1985). Does the autistic child have a "theory of mind" ? *Cognition*, 21(1), 37–46. doi:10.1016/0010-0277(85)90022-8
- Bickhard, M. H. (2009). The biological foundations of cognitive science. *New Ideas in Psychology*, 27(1), 75–84. doi:10.1016/j.newideapsych.2008.04.001
- Block, N. (1996). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18(02), 227–247.
- Botvinick, M., & Cohen, J. (1998). Rubber hands "feel" touch that eyes see. *Nature*, 391(6669), 756. doi:10.1038/35784
- Carruthers, P. (2000). *Phenomenal Consciousness: a naturalistic theory*. Cambridge University Press.

- Carruthers, P. (2011). Higher-Order Theories of Consciousness. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*, (Fall 2011 Edition). Retrieved from <http://plato.stanford.edu/archives/fall2011/entries/consciousness-higher/>.
- Chalmers, D. J. (1997). *The conscious mind: in search of a fundamental theory*. Oxford University Press US.
- Chichilnisky, E. J., & Wandell, B. A. (1999). Trichromatic opponent color classification. *Vision Research*, 39(20), 3444–3458. doi:10.1016/S0042-6989(99)00033-4
- Clark, A. (1993). *Sensory Qualities*. Clarendon.
- Deacon, T. W. (2011). *Incomplete Nature: How Mind Emerged from Matter*. W. W. Norton & Company.
- Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition, The Cognitive Neuroscience of Consciousness*, 79(1–2), 1–37. doi:10.1016/S0010-0277(00)00123-2
- Dennett, D. C. (1991). *Consciousness Explained*. Boston: Little, Brown & Co.
- Dennett, D. C. (1992). The self as a center of narrative gravity. In F. Kessel, P. Cole, & Johnson (Eds.), *Self and consciousness: Multiple perspectives* (pp. 103–115). Hillsdale NJ: Erlbaum. Retrieved from <http://cogprints.org/266/>
- Dennett, D. C. (2003). The Self as a Responding—and Responsible—Artifact. *Annals of the New York Academy of Sciences*, 1001(1), 39–50. doi:10.1196/annals.1279.003
- Deroy, O., & Auvray, M. (2012). Reading the world through the skin and ears: a new perspective on sensory substitution. *Frontiers in psychology*, 3. Retrieved from <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC3491585/>
- Gallagher, S. (2006). *How the Body Shapes the Mind*. Oxford University Press, USA.
- Graziano, M. S. A., & Webb, T. W. (2015). The attention schema theory: a mechanistic account of subjective awareness. *Frontiers in Psychology*, 6. doi:10.3389/fpsyg.2015.00500
- Hacking, I. (1986). Making up people. In T. C. Heller, M. Sosna, & D. E. Wellbery (Eds.), *Reconstructing Individualism* (pp. 222–36.). Stanford: Stanford University Press.
- Humphrey, N. (2006). *Seeing red: a study in consciousness*. Belknap Press.
- Hurley, S., & Noë, A. (2003). Neural plasticity and consciousness. *Biology and Philosophy*, 18(1), 131–168.

- James, W. (1890). *The principles of psychology*. New York: Dover Publications.
- Knoblauch, K., & Shevell, S. K. (2001). Relating cone signals to color appearance: Failure of monotonicity in yellow/blue. *Visual Neuroscience*, 18(06), 901–906.
doi:10.1017/S0952523801186062
- Kuehni, R. G. (2004). Variability in unique hue selection: A surprising phenomenon. *Color Research and Application*, 29(2), 158–162.
- Leslie, A. M., & Frith, U. (1988). Autistic children's understanding of seeing, knowing and believing. *British Journal of Developmental Psychology*, 6(4), 315–324. doi:10.1111/j.2044-835X.1988.tb01104.x
- Levine, J. (1983). Materialism and qualia: The explanatory gap. *Pacific Philosophical Quarterly*, 64(354-361).
- Mangan, B. (2009). Cognition, fringe consciousness, and the legacy of William James. *The blackwell companion to consciousness*, 671–685.
- Maturana, H. R. (2014). Understanding Social Systems? *Constructivist Foundations*, 9(2), 187–188.
- Metzinger, T. (2004). *Being no one: The self-model theory of subjectivity*. The MIT Press.
- O'Regan, J. K. (1992). Solving the “real” mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology*, 46(3), 461–488.
- O'Regan, J. K. (2001). Thoughts on change blindness. In L. R. Harris & M. Jenkin (Eds.), *Vision and Attention* (pp. 281–302). New York: Springer.
- O'Regan, J. K. (2010). Explaining what people say about sensory qualia. In N. Gangopadhyay, M. Madary, & F. Spicer (Eds.), *Perception, Action, and Consciousness: Sensorimotor dynamics and two visual systems* (pp. 31–50). Oxford University Press.
- O'Regan, J. K. (2011). *Why red doesn't sound like a bell: Understanding the feel of consciousness*. New York: Oxford University Press, USA.
- 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). Kluwer.

- O'Regan, J. K., Myin, E., & Noë, A. (2005). Skill, corporality and alerting capacity in an account of sensory consciousness. *Boundaries of Consciousness : Neurobiology and Neuropathology. Progress in brain research, 150*, 55–68.
- O'Regan, J. K., Rensink, R. A., & Clark, J. J. (1999). Change-blindness as a result of “mudsplashes.” *Nature, 398*, 34.
- Palmer, S. E. (1999). Color, consciousness, and the isomorphism constraint. *Behavioral and Brain Sciences, 22*(06), 923–943.
- Pettit, P. (2003). Looks as Powers. *Philosophical Issues, 13*(1), 221–252. doi:10.1111/1533-6077.00013
- 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.
- Prinz, W. (2012). *Open Minds: The Social Making of Agency and Intentionality*. MIT Press.
- Rensink, R. A., O'Regan, J. K., & Clark, J. (1997). To see or not to see: the need for attention to perceive changes in scenes. *Psychological Science, 8*(5), 368–373.
- Rochat, P. (2003). Five levels of self-awareness as they unfold early in life. *Consciousness and Cognition, 12*(4), 717–731. doi:10.1016/S1053-8100(03)00081-3
- Rochat, P. (2015). Layers of awareness in development. *Developmental Review, Theories of development*, 38, 122–145. doi:10.1016/j.dr.2015.07.009
- Rosenthal, D. M. (2004). Varieties of higher-order theory. *Higher-Order Theories of Consciousness: An Anthology* (Johns Benjamins Publishing Company, 17-44., pp. 17–44). Amsterdam: R. J. Gennaro.
- Ryle, G. (1949). *The concept of mind*. University of Chicago Press.
- Simons, D. J., & Rensink, R. A. (2005). Change blindness: Past, present, and future. *Trends in Cognitive Sciences, 9*(1), 16–20.
- Sloman, A. (2001). Beyond shallow models of emotion. *Cognitive Processing: International Quarterly of Cognitive Science, 2*(1), 177–198.
- Sloman, A. (2010). Phenomenal and access consciousness and the “hard” problem: A view from the designer stance. *International Journal of Machine Consciousness, 02*(01), 117. doi:10.1142/S1793843010000424

- Tononi, G. (2004). An information integration theory of consciousness. *BMC Neuroscience*, 5, 42.
doi:10.1186/1471-2202-5-42
- Vazquez-Corral, J., O'Regan, J. K., Vanrell, M., & Finlayson, G. D. (2012). A new spectrally sharpened sensor basis to predict color naming, unique hues, and hue cancellation. *Journal of Vision*, 12(6).
doi:10.1167/12.6.7
- Vierkant, T. (2003). *Is the Self Real? An Investigation into the Philosophical Concept of "self."* Münster: LIT Verlag.
- Witzel, C., Cinotti, F., & O'Regan, J. K. (2015). What determines the relationship between color naming, unique hues, and sensory singularities: Illuminations, surfaces, or photoreceptors? *Journal of Vision*, 15(8), 19. doi:10.1167/15.8.19
- Zahavi, D. (2014). *Self and other: Exploring subjectivity, empathy, and shame*. OUP Oxford. Retrieved from
https://books.google.fr/books?hl=en&lr=&id=UExVBQAAQBAJ&oi=fnd&pg=PP1&dq=zahavi+self+and+other+exploring&ots=2cQI5DwMW3&sig=OUqKaLmFhU_TdIG33anC1tS_MT4