Explaining what people say about sensory qualia

J. Kevin O'Regan

Laboratoire Psychologie de la Perception

Université Paris Descartes, CNRS and Ecole Normale Supérieure

May 31st, 2009

To appear in: N. Gangopadhyay, M. Madary, F. Spicer (Eds) Perception, Action, and Consciousness: Sensorimotor Dynamics and Two Visual Systems. Oxford: OUP.

Introduction

There is an argument promulgated by certain philosophers (notably Dennett 1988, 1991), claiming that from a logical and philosophical point of view the notion of "qualia" makes no sense. On the other hand, other philosophers (e.g. Nagel 1974, Peacocke 1983, and Block 1990) say that qualia must exist since otherwise there would be "nothing it's like" to have sensations: we humans would merely be empty vessels making movements and interacting with our environments, but there would be no inside "feel" to anything.

Independently of this debate there are things normal people say about their sensory experiences that relate to the notion of qualia. People say that they cannot completely describe the “raw”, basic, ultimate aspects of their sensations (e.g. the redness of red) to others (this is usually termed "ineffability"). They say that even if they cannot describe these aspects, they can be compared and contrasted (I shall say they have "structure"). And people say that there is "something it's like" to have these raw sensory experiences (they have "sensory presence").

Whether or not qualia should be taken to exist from a philosophers' point of view, these three things that people say about their sensory experiences need to be explained.

In this essay I show how, under the “sensorimotor” view of phenomenal conciousness (O'Regan & Noë (2001) and further developed by O'Regan, Myin & Noë (2005, 2006) we can understand what people might mean when they say these things, independently of whether qualia actually exist. This angle on how the sensorimotor view can be applied was not stressed in my original papers.

In addition to providing an explanation for what people say, the sensorimotor view has the additional advantage that it makes new empirical predictions and opens new theoretical horizons.

Because I want to remain neutral about whether qualia exist, I shall be defining the term “raw feel”[1]. I shall be trying to find a definition which is such that if you believe qualia exist then you could take raw feel to be equal to qualia. If you do not think qualia exist, then you should take raw feel to be what people say about the most basic aspects of their sensory experiences. Raw feel may be part of what Block (1995) calls “phenomenal consciousness”.

After defining raw feel below, I shall detail the three aspects of raw feel which are problematic for philosophy or science, and go on to show how the sensorimotor approach provides an explanation for them. I shall then look at some empirical consequences of conceiving of raw feel in sensorimotor terms. At the end of the paper I shall consider what it means to consciously feel.

Feel, and “raw” feel

Suppose I consciously see a red patch of colour: I have the feel of red. What exactly is this feel? What do I *experience* when I feel the feel of red? Let us peel away the components of this feel of red until we get to the core, “raw” component.

One aspect of the feel of red is the cognitive state that red puts me into. This state might involve mental associations with roses, ketchup, blood, red traffic and brake lights, and might include knowledge about how red is related to other colours: for example that it's similar to pink and quite different from green. There are also the thoughts and linguistic utterances that seeing redness might provoke, and the changes it might cause in my knowledge, my plans, opinions or desires.

Another component of the feel of red is the learnt bodily reactions that redness may engender, such as habits that I have established and that are associated with red: for example pressing on the brake when I see red brake lights.

Yet another aspect of the feel of red may be the physiological states or tendencies it creates. Certainly such states are found in the case of jealousy, love, and hate, which involve certain, often ill-defined, urges to do things or to modify the present situation. Similarly, emotions like fear, anger, and shame, and states like hunger and thirst may involve reactions of the autonomic nervous system, and may be accompanied by drives to engage in certain particular activities. Thus it may be that even colour sensations may involve physiological states or tendencies. For example, red may be a colour that excites you whereas blue may calm you down.

But many people will say that all these components of red do not constitute the "raw" feel of red itself. They are extra components, add-ons, or products of some more basic thing that most people think exists, namely the primal, raw feel of red, which is at the core of what happens when I look at a red patch of colour. In fact most people would say that the raw feel is actually the cause of the other components. It is the raw feel of the redness of the traffic light that engenders my recognition of the redness and my urge to press on the brake.

Three problematic aspects of “raw feel”

In the following sections I discuss the three things people say about “raw feel”, namely their ineffability, structure, and “what it’s like”, and discuss how they are problematic from a philosophic or neuroscientific point of view.

Ineffability

One of the first things people say about raw feel is that their raw feels are private and ultimately impossible to communicate to anyone else. For example: how can I ever know whether red looks the same to me as it does to you?

This "ineffability" of raw feel has led philosophers to conclude that special theoretical apparatus might be needed to understand it. Ineffability is a first critical aspect of raw feel that we would like to account for.

Structure, and why it is problematic for neurophysiology

A second thing people will say about raw feels is that they differ among each other. There is red, green, pink, black. There is the sound of a tractor, of a violin, of middle C, of the wind in the willows, there is the smell of a rose and the scratch of an itch.

Furthermore, there may be structure in the differences.

One aspect of the structure arises through the fact that certain feels can be compared and others not. For example, red and pink can be compared. We say they belong to the same sensory modality. But red and the sound of a bell cannot reasonably be compared. We say they are in different sensory modalities.

Another aspect of the structure is the fact that when (within a sensory modality) comparisons are possible, the comparisons can sometimes be described in terms of dimensions. An example of a linear dimension is intensity of sound, that goes from no sound to very strong sound. An example of a circular dimension is colour hue, which goes around from red to orange to yellow to green to blue to purple and back to red again. In some cases, such as smell, taste, or musical timbre, the dimensions may be difficult to pin down, or there may be a large number of them (in smell, up to 30 dimensions may be necessary -- cf. Madany Mamlouk and Martinetz 2004).

Structure poses problems for neurophysiological accounts. Here I will discuss the problems encountered by an account in terms of "neural correlates", and by another neurophysiological account based on the idea of "isomorphism".

Take the example of colour. At the very first level of processing in the retina, information about colour is coded in three "opponent" channels: the Blue-Yellow and Red-Green channels measure the equilibrium between blue and yellow and between red and green. There is also a light-dark channel measuring the equilibrium between black and white.

Information in the opponent channels gets transmitted via the optic nerve into the brain, through the lateral geniculate nucleus up into area V1, where it is then relayed to a multitude of other brain areas, and in particular to area V4 which is sometimes thought to be particularly concerned with colour sensation (see Gegenfurtner and Kipers 2003 for a review and a critique of this claim). What is it in these pathways or in the higher areas that explains the nature of the raw feel of red and green and the difference between them?

It is not hard to imagine how differences in the activity in the Blue-Yellow and Red-Green channels could create differences related to what I have called the "extra" aspects of the feel of red and green. Thus we can easily imagine that the neuronal circuits responsible for the feel of red activate cognitive states like those responsible for the memory of strawberries and red traffic lights. We can imagine that red sensitive channels are linked to learned habits linking red sensory input to muscular output, like pressing the brake. We can imagine physiological states that are caused by red and green stimulation, and that such states affect our tendencies to act in different ways.

But what about the "core" aspect of the sensations of red and green: the "raw feel"? How could brain mechanisms such as activation in the opponent channels generate these?

There can be no way of answering this question satisfactorily. It constitutes an instance of what Chalmers (1995) called the "hard problem": For imagine we had actually hit upon what we thought was an explanation. Suppose that it was, say, the fact that the red-producing group of neurons generated one particular frequency of oscillations in wide-ranging cortical areas, and that the green-producing neurons generated a different frequency.

Then we could ask, why are things that way? Why does this particular frequency of oscillation produce the raw feel of red, and that frequency the raw feel of green? Couldn't it be the other way round? We could look further to see what was special about these frequencies. Perhaps we would find that the red producing frequencies favoured metabolism of a certain neurotransmitter, whereas the green producing frequencies favoured metabolism of a different neurotransmitter. Is that sufficient? No, for we would still have to explain why these particular neurotransmitters generated the particular raw feels that they do. Clearly no matter how far we go in the search for mechanisms that generate raw feel, as noted by Dennett (1991), we will always end up being forced to ask an additional question about why things are this way rather than that way. There is what Levine (1983) calls an "explanatory gap". We would not have any difficulty if we were simply interested in understanding how the neurotransmitters determined the extra components associated with red and green by influencing brain states and potential behaviours and bodily states. But we would not be able to explain the way they determined the raw feel itself. Any account in terms of "neural correlates" will come up against this problem.

A possibility to partially escape this situation has been proposed by some scientists (e.g. Palmer 1999) and philosophers (Hayek 1953; Hardin 1988; Austen Clark 1993) The idea is to give up trying to explain why raw feels are the way they are, and instead just concentrate on explaining why the differences and similarities between the raw feels are the way they are. For example there might be a set of brain states corresponding to colour sensations, such that the structure of the similarities and differences between the different states can be mapped on to the similarities and differences between colour sensations. Palmer (1999) has suggested that such an "isomorphism" between the structure of brain states and the structure of colour sensations would provide at least part of an explanation of colour sensations.

But this is not the case. As a concrete example, suppose some brain state produced the raw feel of red and furthermore, that brain states near that state produced feels that perceptually are very near to red. Suppose this happened in a lawful way that corresponded to people's judgments about the experienced proximity of red to other colours.

The trouble is, what is meant by saying a brain state is "near" another brain state? Brain states are activities of ensembles of neurons and there is no single way of saying this brain state is "closer to" or "further from" another brain state. In particular, if there is some way of ordering the brain states so that their similarities correspond to perceptual judgments about similarities between colours, then the question remains of why it is this way of ordering the brain states, rather than that, which predicts sensory judgments. A justification still needs to be given for the choice of metric used to compare brain states. This point even applies in the simplest of all cases, namely the case of intensity. Suppose the perceived intensity of a sound correlated perfectly with the logarithm of spiking frequency in a particular brain mechanism. Then we can always ask: Why exactly that particular law? Why the logarithm instead of a power law or any other law? Spiking frequency is just a code used by the brain, and to satisfactorily explain the feel of intensity, we need an explanation for why the link between the code and the phenomenology is what it is (c.f. Teller's (1984) notion of "linking proposition").

So to summarize more generally about neural circuits and the structure of the qualities of raw feel: finding neural correlates that are necessary and sufficient to generate a particular raw feel would be very interesting, but it would not adequately explain what we would really like to know, namely how these brain structures make the feel the way it is, and how they make it similar or different to other feels.

Even discovering an isomorphism between perceptual judgments and certain associated brain states, though interesting (and in fact inevitable in any scientific explanation of how the brain determines feel), is not sufficient as an explanation of raw feel.

The underlying problem is that there is no way of making a link between feel and physical mechanisms. Neural firing rates, or any physically definable phenomenon in the brain are incommensurate with raw feels.

So the fact that people say that the qualities of raw feels have structure is a second critical aspect of feels we need to explain.

Sensory presence, and why it is problematic for neurophysiology

After ineffability and structure as things people say about raw sensory feels, we now come to what philosophers consider to be perhaps the most mysterious thing, namely that, using the term of Nagel (1974), raw sensory feels "feel like something", rather than feeling like nothing.

Another term that has been used to qualify the particular nature of sensory feels is "presence". The term was used in the phenomenological tradition of Bergson, Heidegger, Husserl and Merleau-Ponty (see Natsoulas 1997,1999 for a recent approach; see also Matthen's contribution to the present volume) and has recently become a key concept in virtual reality (see for example Ijsselsteijn 2002). There are similarities and differences in the usages of the different authors, but this notion of "presence" may be quite close to the notion of "feeling like something".

The trouble is, what do these terms really mean? Certainly they are evocative: for example, we all believe that there is "nothing it's like" for a mere machine to capture a video image of red, whereas we as people really experience redness. The redness is "present" to us, whereas for the machine it is just information that can be used for actions.

In order to make progress we need an operational definition of what it means for there to be “something it’s like” to have a sensory experience, or for an experience to have “presence”. Since we can never be sure that everybody means the same thing, it will be useful to invent a new term and try to define it clearly. I shall use the term “sensory presence” which seems to be close to what we want in referring to the “what it’s like” or presence of sensory experience. The strategy I will use will be to contrast mental states which we can agree have sensory presence with experiences which do not possess it, or which possess it to a lesser degree. I shall consider autonomic physiological states and thoughts, and consider again how neurophysiological explanations must fail to account for the difference in sensory presence of these states as compared to raw sensory feels.

Autonomic functions

Consider the fact that your brain is continually monitoring the level of oxygen, carbon dioxide and sugar in your blood. It is keeping your heartbeat steady and controlling other bodily functions like your liver and kidneys. All these activities involve biological sensors that register the way different systems are functioning in your body. These sensors signal their measurements via neural circuits and are processed by the brain. And yet this autonomic neural processing has a very different status than the redness of the light: Essentially whereas you feel the redness, you do not feel any of the goings-on that determine internal functions like the oxygen level in your blood. The redness of the light is perceptually present to you, whereas states measured by probably the majority of sensors in your body also cause brain activity but generate no such sensory presence.

Why should brain processes involved in processing input from certain sensors (namely the eyes, the ears, etc.), give rise to a felt sensation, whereas other brain processes, deriving from other sensors (namely those measuring blood oxygen levels etc.) do not give rise to a felt sensation?

The answer that comes to mind is that autonomic systems like those that detect and control blood oxygen are simply not connected to the areas that govern conscious sensation.

At first this seems to make sense, but it is not a satisfactory explanation, because it merely pushes the question one step deeper: why do those areas that govern conscious sensation produce the conscious sensation? Supposing we had produced an explanation: say, that there is something special about the wiring, the neurons or the interactions this area has with other brain areas. Then just as was the case to explain the structure of sensory experiences, we can always ask, when this special thing is activated, why does conscious sensation ensue? Any argument based on brain functions is again going to encounter the infinite regress in which we can always ask the question: "And then what happens?" (Dennett 1991; cf. also Chalmers' (1995) "hard problem").

We need a different way of looking at the question of why some neural processes are accompanied by a sensory "presence" and others are not.

Thoughts

Another case we can examine as regards the degree to which we want to attribute to it the notion of "sensory presence" concerns mental activities like thinking, imagining, remembering, and deciding -- I will call all these "thoughts", for short. As in the situation for sensory inputs, you are aware of your thoughts, in the sense that you know that you are thinking about or imagining something, and you can, to a large degree, control your thoughts. But thoughts don't have the same, specifically sensory presence as sensory experiences. Indeed I suggest that as concerns the sensory aspect of what they feel like, thoughts are more like blood oxygen levels than like sensory experiences.

Of course thoughts are *about* things and so come with mental associations: the thought of red might be associated with blood and red traffic lights and red cough drops. But the thought of red does not *itself* have a red quality or indeed any sensory quality at all. Thoughts may also come with physical manifestations: the thought of an upcoming examination might make me feel nervous. But any such nervousness is a consequence of the thought, not a quality of the thought itself. People sometimes say they have painful or pleasureful thoughts, but what they mean is that the content of the thoughts are painful or pleasureful. The thoughts themselves have no sensory quality.

There has been some debate about whether thoughts should be considered as having phenomenal quality -- for example G. Strawson (1994; also Horgan and Tienson, 2002) consider that they do. Ultimately this is a matter of what we mean by phenomenal quality. But I think no one can dispute that thoughts do not have the same kind of "sensory" quality that sensory experiences have. When I say that there is nothing it's like to have thoughts, or that there is a sense in which they have no presence, what I mean is that they do not possess the particular "sensory" presence of sensory experiences.

So now we can ask: Why? What is it about the brain mechanisms that cause thoughts which makes them fail to have the sensory presence that sensory experiences have? And again, as was the case for autonomic functions (and for the structure of sensory experiences), any explanation in terms of particular properties of brain mechanisms will lead to an infinite regress.

To conclude: an operational definition of what we mean by raw feels having sensory presence is to note that this statement is being made in contrast with brain processes that govern autonomic bodily functions, and in contrast with thoughts or imaginings: neither of these impose themselves on us with the same sensory presence as sensory feels; for neither of these is there, in a sensory way, "something it's like" to have them.

And we also note that appealing to brain mechanisms to explain these differences provides no help. Thus “sensory presence” is mystery number three concerning raw feel.

The sensorimotor approach to feel. Step one: the quality of feel

I have pinpointed three things that most people will say about raw feel; ineffability, structure and sensory presence. The problem is that these things seem to admit of no explanation in terms of brain mechanisms. I shall show now how the sensorimotor approach provides a way of thinking that accounts naturally for these things.

The sensorimotor approach involves two steps (in earlier papers on the sensorimotor approach, I had not amply stressed this distinction between the two steps). The first step involves characterising the quality of sensory feels. The claim will be that the quality of a feel is the *quality of the sensorimotor interaction involved*. The second step concerns what is required for an agent to be conscious of this quality -- conscious in the sense of "access conscious": We shall see at the end of this chapter what this involves exactly.

Most of the work of the sensorimotor approach, and the part that elucidates the three problems of ineffability, structure and presence, depends on the first step, and it is this that I shall be detailing in the next sections. To understand the idea, I want to consider a case which is not a prototypical case of sensory feel, but which, when extended to more typical sensory feels in general, provides the key to a solution. It is the case of the tactile sensation of softness[2].

Imagine a person squeezing a sponge and experiencing the feel of softness. What brain mechanism might generate the softness feel?

Clearly the question is ill-posed: The word "generate" is inapplicable. Asking what generates the softness feel is the wrong kind of question to ask about softness, because softness is not the kind of thing that can be generated. The reason is that softness is not a thing, it is a quality: namely the quality of the interaction you have with a soft object like a sponge. The experienced quality of softness consists in a particular fact about the sensorimotor interaction you are having, namely the fact that when you press on the sponge, it squishes.

And now we can consider how by taking this view of what softness is and applying it to sensory feels in general, we can solve the three mysteries of ineffability, structure of qualities and sensory presence. It will become clear that the reason the new view provides a solution is that we are no longer looking in the brain for something that generates feel. Instead we are identifying qualities of feel with qualities of modes of interaction with the world.

How the sensorimotor approach explains ineffability

Obviously when you squish a sponge there are all sorts of muscles that you use, different parts of your fingers are involved, and there are all sorts of precise ways that the sponge squishes under your pressure. It's inconceivable for you to have cognitive access to all these details. It's like executing a skiing manoeuver, or whistling: you don't know what you do with your various muscles, you *just do it*.

The precise laws of the sensorimotor interaction are thus ineffable, they are not cognitively available to you, nor can you describe them to other people.

Applying this idea to feels in general, we can understand that the ineffability of feels is a natural consequence of thinking about feels in terms of ways of interacting with the environment. Feels are qualities of the sensorimotor interactions which we are currently engaged in. We do not have cognitive access to each and every aspect of these interactions. As an example, the particular muscle bundles we use to move our eyes or sniff or move our head are involved in feeling the feels of seeing, smelling and hearing, but usually we cannot know what they are.

The sponge analogy thus accounts naturally for the ineffability of feels. I now come to the structure and then to the "what it's like" or sensory presence.

How the sensorimotor approach explains structure

Consider now how the sponge analogy deals with the second mystery of feel, namely the fact that feels are sometimes comparable and sometimes not, and that when they are comparable, they can sometimes be compared along different kinds of dimensions.

Take sponge squishing as compared to, say, whistling. There is little objectively in common between the modes of interaction constituted by sponge squishing and by whistling. They are not comparable.

On the other hand the feels can be compared within the gamut of variations of sponge-squishing and within the gamut of ways of whistling. In the case of sponge squishing for example, some things are easy to squish, and other things are hard to squish. There is a continuous dimension of softness. Furthermore, it is part of the laws that define what we mean by softness, that “hardness” is the opposite of what we mean by softness. It is thus simply a matter of fact, deriving from the very definitions of hard and soft, that there should be a continuous dimension going from very soft to very hard.

So here we have examples that are reminiscent of what we noticed about raw feels. Sometimes comparisons are nonsensical: just as between sponge squishing and whistling, nothing very much can be said about the comparison between the feel of seeing and that of hearing or of touch. And sometimes comparisons can be made: just as within touch there is a clear relation between hard and soft, within seeing, comparisons can be made, for example between bright and dim, or between red and pink.

Thus, if we take the view that the qualities of feel are qualities of sensorimotor interactions, the existence of a complex structure defining how these qualities can be compared and contrasted is a natural consequence.

Furthermore this structure derives from objective physical facts concerning the modes of interaction involved. For example the reason softness and hardness can be compared, but softness and whistling cannot be compared is that objectively, the modes of interaction we engage in have precisely this structure. In the cases of sensory experiences like seeing, hearing, touching, tasting and smelling, the claim is that their structure is also an objective consequence of -- indeed it is constituted by the structure of -- the modes of interaction that we engage in when we have the associated sensory experiences. To illustrate this claim I shall be showing later how it can work to explain the structure of colour.

Note that if we considered that the differences in structure in feels were caused by neural mechanisms, we would have to explain why the neural mechanisms generate the particular structure of feels that they do. Even if we managed to find some brain mechanisms that were isomorphic to the observed structure in feels, then we would have to explain the particular choice of classification scheme or metric that we used in order to make the isomorphism apparent.

But in the sensorimotor view we escape from this problem, since we do not consider that brain mechanisms are generating the feel. They are merely enabling an interaction, and it is the occurring of this interaction that constitutes the experience of feeling. It is the quality of the interaction, that is, the laws that describe its structure, that constitute the quality of the feel.

Thus, in the sensorimotor approach, the similarities and differences between feels are no longer described in terms of neural similarities and differences for which we have no natural way of choosing a classification or metric. We no longer need to ask whether we should take this or that similarity measure among the many possible ways of comparing neural states -- whether we should, for example, take neural firing rate or its logarithm or inverse. Instead, similarities and differences in feels are described in terms of the metrics or classifications that humans already use every day to describe the way they interact with the world (e.g. an object is softer when it cedes more under your pressure).

Brain mechanisms do of course play a role in the account of feel in terms of an interaction with the environment. The brain mechanisms enable the particular interactions that take place when we experience something. Interestingly, once we have found objective descriptions of the sensorimotor interactions that constitute a sensory experience, then perforce we will find brain mechanisms that enable them. To the extent that there are objective differences and similarities in the sensorimotor interactions, there will necessarily be corresponding differences and similarities in the brain mechanisms. Thus there will necessarily be an isomorphism between the brain mechanisms and the quality of the feel. And there is the risk that having discovered such an isomorphism or correlation, one might naïvely be misled into thinking that somehow it was the brain mechanism that "generated" or was the cause of the associated feel. But this would be wrong. And taking that stance would immediately lead to the the classic infinite regress underlying the explanatory gap, since one would have to explain why the brain mechanism did what it did.

But the problem is obviated when we take the sensorimotor view according to which the quality of a feel is not generated by anything, but rather is constituted by the laws that describe our sensorimotor interaction.

How the sensorimotor view explains sensory presence

Now I come to the question of sensory presence, or why people say that there's "something it's like" to have a feel. I had suggested that an operational way to understand what is meant by this statement is to make the contrast with autonomic processes in the nervous system, and thoughts, since these are brain processes that lack the "sensory presence" of a real sensory experience. Even if we wish to argue that there is "something it's like" to think, or to digest, or even simply to exist, whatever "it's like" is not the same as the "what-it's-like" of having a real sensory presence. The sensorimotor approach provides a natural way of accounting for such differences by appealing to three concepts: bodiliness, insubordinateness and grabbiness.

Let us take again the example of sponge squishing. What is it about sponge squishing that gives that activity a real sensory presence? Obviously sponge squishing involves actually engaging in an interaction with the world. But why should actually engaging in an interaction with the world give this impression of realness and presence?

I suggest the answer has to do with the fact that the degree of voluntary control we exercise when we have a sensory feel is only partial. When we have what we call a sensory feel, we have a certain degree of voluntary control over what we are doing. We can exercise this control by modifying sensory input through movements of our body (I call this "bodiliness"). However the control is partial, because sensory input derives from the outside world and the the outside world escapes our control to a certain extent (I call this "insubordinateness"). Finally a third factor is "grabbiness", namely the fact that the outside world has the capacity to grab our cognitive ressources. Let us look at these three aspects in more detail.

Bodiliness. When you are interacting with the world, body motions generally cause systematic changes in sensory input. In the visual modality this is clearly the case: moving your eyes causes dramatic changes of the retinal image. In the case of touch, we note that touch is an exploratory sense, since only by moving our hand, for example, can we accurately recognize an object. But even for passive touch, when someone touches you, if you then voluntarily remove or shift the body part being touched, there will be a change in sensory input.

For hearing, moving the body modifies auditory input by changing the amplitude of the sound impinging on the ears, and in the case of rotation of the head, by changing the relative delay between signals coming into the two ears. The complicated shape of the earlobes creates micro-reflections that change with head orientation and are also important factors in sound localization and identification. Thus, moving the body is a way of testing whether a sound comes from the outside world. When you have a ringing in your ears, you identify the sensation as not coming from the outside world because turning your head doesn't change the sound.

Recent research on smell has also shown that humans, like other animals, can use body and head movements to monitor the delays and differences between the smells coming into the two nostrils in order to follow scents (Schneider and Schmidt 1967; Porter et al 2007; Sobel et al. 1998). But more fundamentally, we know that we are really smelling and not just imagining it, when we can confirm that sniffing, and moving our body (in particular our head) changes the signal coming from our olfactory receptors.

All these examples show that susceptibility of sensory input to voluntary body motion is an essential feature of what it is like to be experiencing external-world stimulation. It is perhaps actually a logical consequence of the fact that what we mean by the outside world is what is delimited by our body. In all cases, moving the body and the eyes is a way of checking whether a sensation comes from the world. If no change occurs when we move, then what we are experiencing probably does not originate from the outside world.

We can now understand why it is that information provided by sensors involved in autonomic functioning, on the one hand, and neural activity involved in thinking, remembering, imagining and other reflective activities, on the other hand, are not accompanied by a sensation of sensory presence: it is (partly) because they have no bodiliness. Voluntarily moving your body only has an indirect effect on autonomic body parameters like blood oxygen, carbon dioxide and sugar levels, and on functioning of internal organs. Moving your body has no direct effect on thought processes[3].

Bodiliness is not complete: Insubordinateness. Whereas dependence on voluntary bodily motion is an essential feature of sensory input originating in the outside world, not all changes in sensory input from the outside world are caused by our voluntary body motions. The outside world has a life of its own: objects move, sounds change, smells appear and disappear, with the consequence that we are not complete masters of the changes that the world causes in our sensory input. This insubordinateness of the world is another factor that characterizes "real" sensory inputs and distinguishes them from our own thoughts and imaginings. Thoughts and imaginings are entirely the property of our own minds. They are under our voluntary control, they are predictable and completely controlled by us and so we do not perceive them as corresponding to real-world events[4].

Grabbiness. Grabbiness consists in the fact that sensory input signals have the power, in certain circumstances, to deflect our cognitive ressources and grab them incontrovertibly so that it is difficult for us to voluntarily attend to anything else. One kind of signal that has this property is so-called "transients", that is, sudden changes in sensory input. For example a bright flash in peripheral vision, a sudden loud noise, a sudden poke, punch, or tickle, grab our attention, provoke an immediate orienting reflex, and cause our cognitive processing to be deviated toward the source of the change. In some sensory modalities certain kinds of non-transient stimulation can also grab our attention. For example in hearing, very loud continuous sounds can prevent one from thinking properly. In touch, stimulations signalling tissue damage like burns and aches can grab our attention. In smell, certain pungent or obnoxious odors have the property that we cannot avoid paying attention to them.

I suggest that grabbiness is a general property of sensory systems[5], and is not shared by other brain systems. In particular, consider the brain systems that deal with autonomic functions like keeping blood pressure stable, with holding blood sugar levels constant, with adjusting breathing, digesting, and with keeping a host of other body functions working properly. These systems do not have the faculty of interrupting cognitive processing. This I would claim is part of the reason we do not directly "feel" our internal vital functions as having a sensory quality: e.g. why blood sugar level does not appear to us as having the quality of a sensory experience like seeing or hearing.

Or consider thoughts: thoughts do not grab your attention like loud noises, pungent smells or intolerable pain. Except in pathological cases, you are not possessed by thoughts, you possess them.

Summary on sensory presence

To summarize why raw sensory feels feel like something rather than feeling like nothing, that is, why they have "sensory presence":

Experiencing a raw sensory feel involves engaging with the real world. Doing so involves having control, but not complete control of this engagement: control derives from *bodiliness*, that is, the fact that voluntary bodily changes provoke systematic variations in sensory input. But control is not complete because our sensory input is not exclusively determined by these bodily motions: The real world is *insubordinate* and has a life of its own that creates variations in our sensory input that we cannot cause through our voluntary body motion. A final way in which our engagement with the real world escapes our control derives from *grabbiness*: the fact that our sensory input systems are wired up to be able to grab our cognitive ressources incontrovertibly in certain circumstances, making us tributary to the outside world.

Why the sensorimotor approach works better than the neural correlate approach

Note the advantage of considering feels to be modes of interaction with the environment. If we thought feels were generated in the brain, we would have to go looking in the brain for something special about the neural mechanisms involved that generates the feels. We would have to postulate some kind of special neurons, special circuitry or chemical basis that provide the feeling of “presence” or "what it's like". And then, as explained in previous sections, we would be led into an infinite regress, because once we had found the special neurons, we could always then ask what exactly it was that made them special.

But if experiencing a sensory feel involves engaging in a particular mode of interaction with the environment, then since you are doing something, there will be laws that characterize the particular mode of interaction you are involved in. These laws constitute the quality of the feel. Thus by the very sensorimotor definition of feel, if you are having a feel, there must be something it is like for you to have it. What’s more, if the interaction you are having involves using your sensory apparatus to get information about the outside world, then the interaction will have the hallmark of sensory feels: Because of the bodiliness and grabbiness of sensory channels, and because of the external world's inherent insubordinateness to your will, the control the brain exercises over the interaction will only be partial. This partial control corresponds to the quality of “sensory presence” possessed by sensory stimulation. Such sensory presence is not possessed by mental activities like thoughts or imaginings (here control is complete), nor by autonomic control systems that keep the body functioning normally (here we have no voluntary control).

The role of action

The role of action in the sensorimotor approach is often misunderstood.

If we take feel to be defined as the quality of an interaction with the environment, then because the notion of interaction necessarily implies action, feel also necessarily implies action. However this statement should not be taken to mean that experiencing a feel at a given moment necessarily requires an action to be occurring at that moment. Saying that feel implies action should be understood to mean that in order to have a feel, action must potentially play a role (Noë 2004 has also commented on the role of action in perception; c.f. also Noë, this volume).

As an analogy, consider the dancer poised instantaneously in a choreography, the acrobat poised at the top a jump, the mountain explorer resting at the base camp.

What we mean by being engaged in an activity does not require that one should be continuously acting, only that one's current situation should be part of and should fit correctly into a potential activity.

Take as an example the feel of touch on my arm. Feeling touch on my arm is: being poised to verify that my current situation is part of what happens normally when I am being touched on my arm. In particular, if I move, there will be a change in tactile input. More precisely such a change will occur if I move my arm, but not if I move my foot. It is contingencies such as this that are constituent of a stimulation in the tactile modality, and more particularly, with a stimulation on the arm rather than on the foot.

So when I receive tactile stimulation on my arm, even without me moving at all, comparison mechanisms in the brain will register that such input has been received before, and that when this had previously been the case, such input was also associated with systematic changes that occurred when I moved, and more precisely, when I moved my arm and not other parts of my body. For this reason, I will perceive the sensation as a tactile sensation, and more precisely, as coming from my arm, and not from other parts of my body. And this will be the case even without me moving.

An interesting prediction is made from this approach. It is that sensory input that has never in the past been observed to be systematically modifiable by voluntary body motions should not be experienced as being of a sensory nature. This may provide part of the explanation of why we do not experience feels in our visceral organs: we generally cannot voluntarily make movements that modify the neural signals that they provide.

The proposal that visceral organs should not be perceived as being the location of sensations is compatible with the phenomenon of "referred pain", where damage to internal organs is perceived as originating in superficial body locations that share nerve pathways with them. Best known among types of referred pain is the pain associated with myocardial ischemia (heart attack) which can be felt in the upper chest, arm, or even hand or jaw, and appendicitis which begins as a pain near the navel. In neither case is the pain felt at the actual visceral location of the malfunction.

We do however sometimes have aches inside our body, such as stomach aches, headaches and toothaches. How can this be accounted for? Under the approach I am suggesting, this can only occur in cases when body motion does actually modify the sensory signals. It is true that shaking the head when one has a headache modifies the pain, and that the location of the pain may perhaps be assimilated to similar sensations on the surface of the scalp that occur when one receives stimulation there. Similarly, pressing on the stomach or tooth can modify stomach and toothache, thereby providing information as to the location of the sensation. Interestingly, and compatible with the suggestions being made here, a toothache in a tooth in the lower jaw is sometimes incorrectly attributed to the tooth in the upper jaw that generally is in contact with it, and vice versa.

Empirical evidence

Experiencing a raw sensory feel involves engaging with the real world in a sensorimotor interaction. The laws that govern such an interaction constitute the quality of the associated raw feel. Raw feels are ineffable because they are skills that we engage in, and like all skills they are not completely accessible to cognitive analysis. Raw feels have structure because the sensorimotor interactions that constitute them are governed by the complex constraints of our sensory and motor apparatus, as well as by the constraints of the real world. In contradistinction to thoughts and autonomic processes in our central nervous systems, raw feels have sensory presence or "what it's like-ness" by virtue of the fact that sensorimotor engagement with the world involves having control, but not complete control of this engagement: control derives from *bodiliness*, that is, the fact that voluntary bodily changes provoke systematic variations in sensory input. But control is not complete because our sensory input is not exclusively determined by these bodily motions. The real world is *insubordinate* and has a life of its own: it creates variations in our sensory input that we cannot cause through our voluntary body motion. A final way in which our engagement with the real world escapes our control derives from *grabbiness*: the fact that our sensory input systems are wired up to be able to grab our cognitive ressources incontrovertibly in certain circumstances, making us tributary to the outside world.

The sensorimotor approach has an advantage over neural correlate aproaches to the nature of raw feel. The view in terms of neural correlates must always appeal to some linking hypothesis that connects feel to neural mechanisms, and justification for the choice of such linking hypotheses is not forthcoming. The sensorimotor view on the other hand provides a natural way of describing and explaining sensory feels in terms of the objective laws that govern our sensorimotor interactions with the world.

In addition to this logical or philosophical advantage, the sensorimotor approach provides a way of thinking about feel that generates interesting avenues of empirical research. In particular several empirical results provide converging evidence in favor of the way the sensorimotor approach accounts for what people say about the quality of sensory feels.

Sensory Substitution

A first result concerns the perceived quality of sensory feels in different sensory modalities: why are they the way they are, and why are the similarities and differences structured in a certain way? The sensorimotor approach provides a natural answer in terms of the different sensory interactions you engage in. For each sensory modality there exist laws of sensorimotor interaction that characterise that modality. When you see, for example, and you close your eyes, there is a large change in sensory input, whereas closing your eyes has no effect on sensory input when you hear. When you see, moving your eyes produces a particular kind of flow field on your retinas, whereas when you hear, moving your eyes has no effect.

A very interesting prediction ensues from these ideas. If what determines the sensory quality of a sensory modality is not some characteristic of the neural channel involved, but the laws that sensory input obeys when you move, then it should be possible to get the impression of seeing, for example, through channels other than the visual channel, provided that the laws being obeyed are visual-type laws.

This is exactly what happens in *sensory substitution*. The classic example is the work of Bach y Rita (1972), who equipped blind (or blindfolded) persons with an array of tactile vibrators that the persons wore on their abdomen, and which created a vibratory "image" of information registered by a video camera held by the person. Despite the device's very poor resolution (20 by 20 vibrators) compared to the 150 million receptors in normal human vision, users reported that they had an impression of "seeing", provided they were allowed to actively wield the camera.

Such devices have proven cumbersome in the past, but with today's technical advances there is renewed interest in developing them (for a review, see Bach y Rita and Kercel 2003). Particularly successful devices have been constructed to compensate for vestibular deficits and substitution of vision with hearing. Nevertheless, the intrinsic limitations of the skin or the ear as an input channel prevent obtaining the same resolution as with the eyes. It is therefore understandable that the feel experienced by users of such devices should be far from what a normal sighted person experiences (see Auvray et al. 2007a,b for examples of studies to investigate these issues).

Colour

As already explained, most neuroscientists consider that colour experience is generated by neural activation in colour opponent channels in the visual system. But this idea, though appealing, has the usual "explanatory gap" difficulty, since it immediately raises the question of exactly what in the channels causes the accompanying feels.

The idea is also questioned by empirical results (Jameson and D'Andrade 1997). If it were true, then one would expect that what people experience as "pure" colours -- what colour scientists call "unique hues" -- should correspond to maximal activations in the corresponding channels. For example, the maximum activation of the opponent red-green channel in the red direction should provoke a sensation of pure red, and maximum activation in the green direction should provoke a sensation of pure green. Similar statements should be true for maximum activation in the blue-yellow channel.

However these predictions are not bourne out: it is found that to experience a sensation of pure blue, activation of the blue-yellow channel in the blue direction is not sufficient: some activation of the red-green channel in the green direction is needed. Similarly, to get a sensation of pure yellow, activation of the blue-yellow channel in the yellow direction is not sufficient: some activation of the RG channel in the red direction is needed (Chichilnisky and Wandell 1999; Valberg 2001; Knoblauch and Shevell 2001)

Another problem with the classic view comes from Berlin and Kay's (1969) and Kay and Regier's (2003) classic anthropological data on naming coloured surfaces (also Regier, Kay, and Cook 2005). In a sample of informants from 110 unrelated cultures throughout the world, these authors observed that there is a regular structure to the way their informants named the colours. Certain particular surfaces stood out from all the others in that they were systematically given names across all the different cultures. The first four of these special or "focal" surfaces were precisely "red", "yellow", "green" and "blue". Unfortunately however, although the link with the red-green and blue-yellow opponent system seems tempting, the precise structure of the data and the precise colours hues of red, yellow, green and blue that judged to be focal have not been explained by appeal to neurophysiological opponency.

On the other hand by taking the sensorimotor approach to feel, both the unique hues and the naming data can be satisfactorily explained.

The idea is to say that the "feel" of colour is constituted by a quality of our interaction with coloured surfaces. Under this view, suggested by Broackes (1992), colour is the law describing how a surface changes incoming light.

Such laws have been described by physicists, who use the notion of "reflectance function" to characterize how each incoming wavelength of light is absorbed or reflected by the surface. But this characterisation does not correspond to the biological reality of our sensory systems, which cannot register light energy in individual wavelengths, but only in the wide bands of wavelength absorbed by our particular photoreceptors. For that reason Philipona and O'Regan (2006) defined a "biological reflectance function" which corresponds to the reflectance function as registered by human photoreceptors. When this function is examined for different surfaces, it is found that those surfaces that are frequently named across multiple cultures are precisely those surfaces that have a "singular" biological reflectance function. By "singular" is meant the fact that, mathematically, the function has a simpler behaviour than is usually the case. To be more precise, a biological reflectance function usually maps the three-dimensional space of possible photoreceptor values corresponding to incoming light, into another three-dimensional space of photoreceptor values corresponding to reflected light. But when the function is singular the reflected light is projected into a smaller space, namely either a one-dimensional space (the case for red, blue and green) or a two-dimensional space (the case for yellow).

Philipona and O'Regan found that when they looked at the singularity of the different surfaces used in the Kay and Regier world colour survey, there were four strongly singular surfaces, and they were almost exactly those surfaces that were most often given a name in the Kay and Regier data. The precision of the agreement observed is very striking, and strongly supports the idea that the sensorimotor approach to surface colour is on the right track (see Figure).

Philipona and O'Regan extended the approach to coloured lights by assuming that lights are perceived as though they are coloured surfaces illuminated by standard white light. Predictions for "unique hues" made in this way account for existing data better than neurophysiological approaches, accurately predicting the curious facts mentioned above about how to obtain the sensations of pure blue and pure yellow.

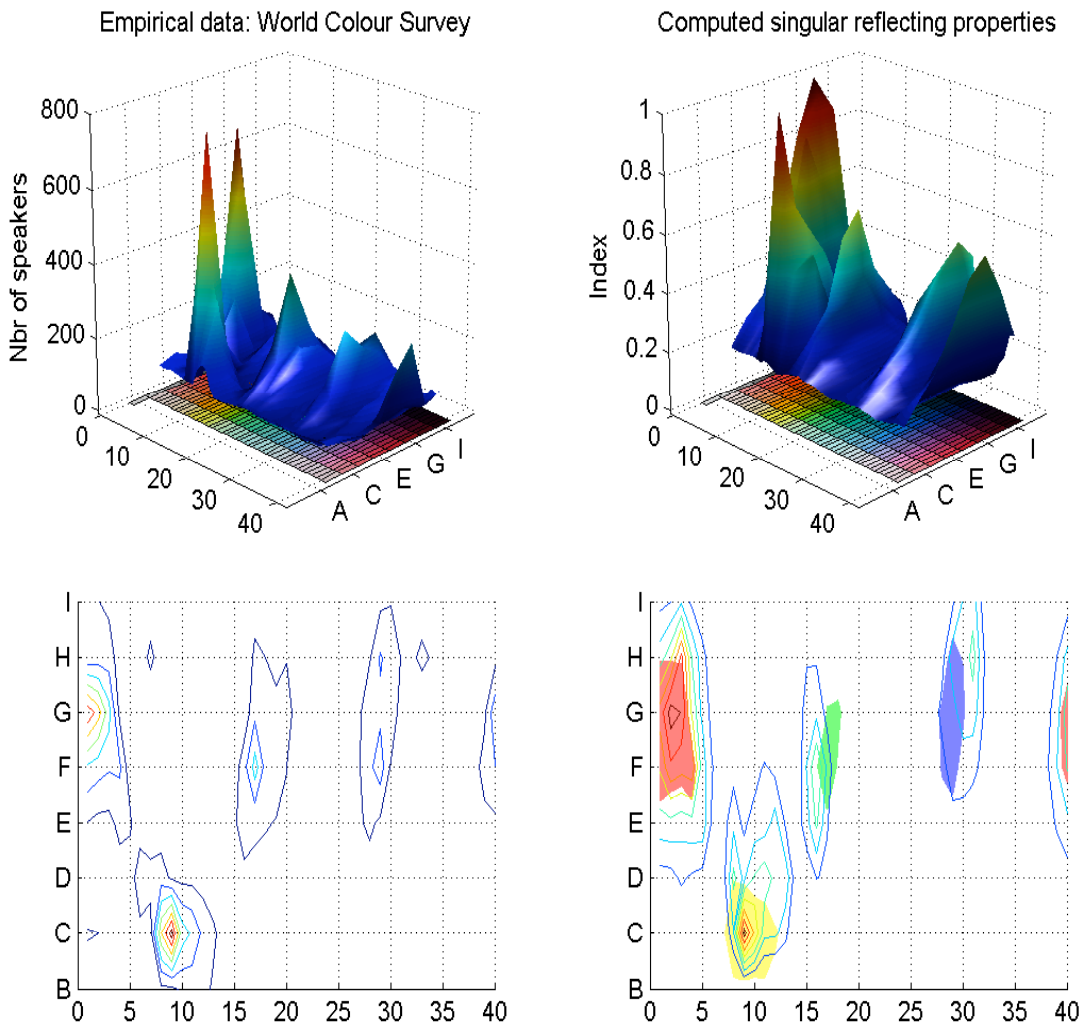

Top left and bottom left*: histogram and contour plot showing the number of speakers in Kay and Regier's (2003) 110-culture survey of colour naming, who had a name for the colours in the selection of Munsell chips labelled A-I and 0-40 shown in the ground plane of the top graph. Strong peaks are visible at G1 and C9 corresponding to red and yellow. Weaker peaks at F17 and F29 or H29 correspond to green and blue. The iso-contours for the top 10% of these data are re-plotted as red, yellow, green and blue colored surfaces in the bottom right graph.* ***Top right and bottom right****: the degree of singularity (singularity index) of the 3 x 3 matrices representing the linear mapping between incoming and outgoing human cone photoreceptor absorptions for the same Munsell chips, as calculated by Philipona and O'Regan from known physical data for Munsell reflectances and cone absorption spectra. The peaks of singularity shown in the contour plot fall very close to the colored surfaces corresponding to the top 10% of the Kay and Regier data. This shows that the colors that tend to be given names across many cultures are very close to those that change incoming light in a "simpler" fashion (see text). Adapted from Philipona and O'Regan (2006).*

Spatial and temporal presence and continuity of the visual world: change blindness

When we look at a visual scene, we have the impression of seeing everything in it, simultaneously and in an occurrent, ongoing fashion. According to the traditional view of seeing, this "present" quality of visual experience arises because an internal representation of the scene has been created and is "active" in the brain.

Under the sensorimotor view, the idea is that we internally store very little information about the scene. Instead, we access the information in the world when we need it: as suggested by O'Regan (1992), the world serves as a sort of "outside memory". The impression of seeing "everything" and of doing so in a continual fashion derives from the immediate accessibility of the information at the slightest flick of attention or the eye, and not from the existence of the information in an internal brain representation.

In this way of obtaining the sensation of continuity, the grabbiness of the visual sensory modality plays a reinforcing role. This is because if anything should change in the scene, the visual transients produced in the visual channels create an automatic, incontrovertible alerting reaction which orients our attention to the change location. When there are no such transients, our conviction that everything is still there is upheld.

These ideas make the curious prediction that if we could make changes in the visual scene without the change provoking a visual transient, observers would still go on seeing what was originally there, and not see the change: unless of course they were explicitly attending to the item that changed.

Exactly this is the manipulation achieved in the experimental paradigm of change blindness. In one variant of this paradigm, a large change is made in a picture, but it is made so slowly that no visual transients occur that can attract the observer's attention to the changed location. In the better known "flicker" and "mudsplash" paradigms, a change is made in a scene, but the change is accompanied by other, large and sudden changes which also create visual transients. These act as decoys which prevent attention going to the location of the to-be-sought change[6].

Results in a large literature using these paradigms (see reviews by Simons and Levin 1997 and Simons and Rensink 2005[7]) very clearly demonstrate that observers tend not to see the changes, confirming the prediction and providing converging evidence for the sensorimotor approach.

The sensorimotor approach to feel. Step Two: Consciously experiencing

Up until now I have been discussing the first step of the sensorimotor approach, which was concerned with the quality of sensory experience. The approach claims that experienced sensory quality is constituted by the (physically objective) sensorimotor laws that underlie the interaction. But whereas this characterises the quality of the experience, we still need to define what makes a sensory experience conscious.

For this we require the second step in the sensorimotor approach. The claim here is simply that a sensory experience becomes conscious when a person has a certain form of access (namely "conscious access") to the fact that the person is engaging in a sensorimotor interaction.

To understand this, note first that when we say we are conscious of something, there is an implicit presupposition that there is indeed a "we" to be conscious. But who or what is this "we"? This question is essentially the question of the self.

Discussion of the nature and origin of the notion of the self is an active subject today. Philosophers are trying to decide what precise components the notion of self boils down to. Developmental psychologists and psycholinguists are trying to ascertain how the notion of self develops in the maturing child; cognitive anthropologists look at whether the notion is different in different human societies, cognitive ethologists study which species possess the notion, and social psychologists investigate how the self is determined by an individual's social environment[8].

Unlike the question of raw feel, for which there seemed to be an explanatory gap, the problem of the self -- which is certainly not one problem but many problems, with organismic, cognitive and social aspects -- is today being approached using the idea that brains are computational devices with the capacity to abstract, to generalize, and in the case of humans, to use language. The social aspect of the notion of "I", with its self-validating and self-referring properties, could emerge from social interactions of agents with brains having such computational faculties.

I shall take it on faith that a conclusion from this literature is that while the problem of the self is very complex, it is nevertheless amenable to a scientific approach. Let us then consider that we have a scientific way of dealing with the notion of "we" in the phrase "we are conscious of X". The next thing to do is to look at what is meant by the word "conscious".

Again here the philosophers have raised tremendous debates, but there seems to be agreement that one main use of the notion of "being conscious of X" is amenable to a scientific account, namely the notion of "access consciousness" (see Carruthers 2009 for a review). An agent is access conscious of X when it is ready or "poised" to make use of X in its decision-making processes (Kirk 1994; Dretske 1995; Tye 1995), and (some philosophers require) also in its rational behavior (Block 1995). The fact of being poised in this way to make use of X might additionally have to be the object of some kind of higher-order representation (Armstrong 1984; Rosenthal 2005; Carruthers 2005; Lycan 1996).

From the psychologists' point of view, the notion of being access conscious of X has been taken to correspond to the idea of attending to, or devoting one's processing ressources to X in such a way that X becomes globally available to a wide variety of behavioural capacities. The idea is an old one, suggested by Baars (1988) in his "global workspace" model, and taken up more recently in neuroanatomical terms by, for example, Dehaene and Naccache (2001).

The issues involved in defining access consciousness constitute an immense literature, but as was the case for the notion of self, the consensus is that access consciousness is a notion that can be functionally defined. There is therefore general agreement that a reasonable, scientific account can be given for it.

The second step of the sensorimotor account thus involves saying that "we are conscious of a sensory experience" when "we" (as defined in cognitive and social terms) become "access conscious" (as defined in one of the functionally defined forms referred to above) of the fact that we are engaging in the sensorimotor interaction which constitutes that experience.

Consciously experiencing a sensory feel

In sum then, we are now in a position to functionally characterise what is meant by consciously experiencing a sensory feel.

The essential idea or what I have been calling step one of the sensorimotor approach, is the claim that experiencing a sensory feel involves engaging in a sensorimotor interaction with the environment. The experienced quality of the feel is constituted by the laws that physically, objectively, describe this interaction. By virtue of the fact that the interaction is essentially a sensorimotor skill, the laws that describe it will escape complete description. This yields the ineffable quality of sensory feels.

The fact that nevertheless the laws that characterise the interaction will be constrained by the physics and biology of the situation will result in objective similarities and differences between the interaction and other interactions. This yields the fact that the qualities of different sensory feels will fall into what may be a complex structure of objective similarities and differences.

Finally, if the laws of interaction are characterised objectively by bodiliness, insubordinateness and grabbiness, then the experienced quality will have the property which is the hallmark of sensory feels, namely sensory presence. Said another way, there will seem to be "something it's like" to have the feel, in contrast to autonomic body functions and thought processes, for which this is not true -- at least in the sense I am using the term "sensory presence" or "what it's like".

Step two of the approach is then to say that if an agent has sufficient cognitive capacities and is sufficiently integrated into a social environment to have a notion of self, and if this agent is suitably poised to cognitively access the fact that it is engaging in the sensorimotor interaction, then the agent is consciously experiencing the associated feel.

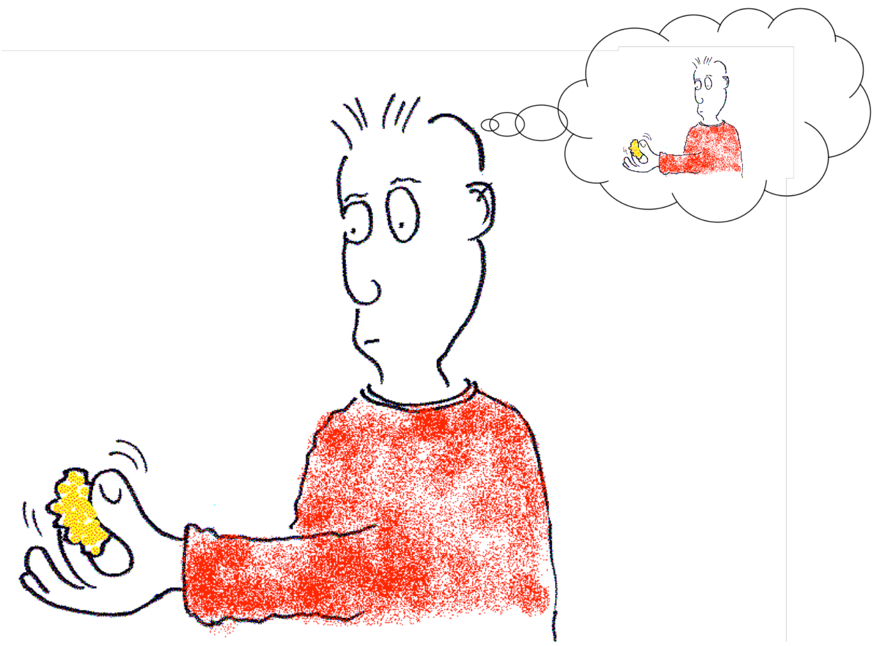

Have I explained qualia?

I have expressly avoided the word "qualia" in this treatment, on the grounds that qualia are such a contentious topic. Instead I have used the notion of raw feel, giving it what I hoped was a definition as close as possible to qualia but corresponding to a plausible version of what people say everyday about the most basic aspects of their sensations. There is little doubt that the man in the street believes that at the basis of every feel there is a "raw" component that is producing its behavioural effects: for example the raw feel of pain is what produces avoidance reactions, and not the opposite: the man in the street rejects the

Jamesian or behaviouristic idea that being in pain is just the sum total of pain-like behavioural manifestations.

Thus whether or not raw feels exist, people talk about them as though they did. Furthermore people claim that they have certain properties, and among these the problematic ones (for philosophers) have been those of ineffability, structure in their qualities, and, most important, "presence" or "what it's like". The sensorimotor approach provides an account of why people say their raw feels have these properties, even if certain philosophers might want to claim that qualia do not exist.

What is the relation between this approach and Dennett's (1991) eliminativist view of qualia?

The difference between the two views is that, as explained in the introduction, the sensorimotor view remains neutral[9] on whether or not qualia exist, and simply tries to explain what people will tend to say about them (or at least, about "raw feels", if qualia do not exist). Because it provides a natural explanation (in terms of physical and biological laws) of the qualities of these raw sensory feels (their ineffability, their structure and their "presence" or "what it's like"), and because it opens new empirical research programs and generates testable hypotheses such as those summarized here about sensory substitution, change blindness, and colour, the sensorimotor view is more scientifically productive than the eliminatist view.

Acknowledgements

I warmly acknowledge the help of Erik Myin in commenting on an earlier version of this chapter, the help of the participants at the December 2008 CONTACT workshop in Bristol in helping clarify some of the ideas here, and of the editors of the book in making final suggestions for improvement.

References

Armstrong, D. (1984) Consciousness and causality. In D. Armstrong and N. Malcolm (eds), Consciousness and Causality, Oxford: Blackwell.

Auvray, M., and Myin, E. (in press). Perception with compensatory devices: From sensory substitution to sensorimotor extension. *Cognitive Science*.

Auvray, M., Hanneton, S., and O'Regan, J.K. (2007a). Learning to perceive with a visuo-auditory substitution system: Localisation and object recognition with 'The vOICe'. *Perception*, 36,3,416-430.

Auvray, M., Philipona, D., O'Regan, J.K., and Spence, C. (2007b). The perception of space and form recognition in a simulated environment: the case of minimalist sensory-substitution devices. *Perception*, 36,12,1736-1751.

Baars, B. (1988) A Cognitive Theory of Consciousness. Cambridge: Cambridge University Press.

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

Bach-y-Rita, P. and Kercel, S.W. (2003), Sensory substitution and the human-machine interface, *Trends in Cognitive Sciences*, 7, 12, 541-546.

Berlin, B., & Kay, P. (1969). Basic Colour Terms: Their Universality and Evolution (Berkeley. *University of California Press. Bornstein, MH (1975). Qualities of colour vision in infancy. Journal of Experimental Child Psychology,* 19, 401-419.

Block, N. (1990) "Inverted Earth," Philosophical Perspectives, 4, J. Tomberlin, ed., Northridge: Ridgeview Publishing Company.

Block, N. (1995). On a confusion about a function of consciousness. *Behavioural and Brain Sciences,* 18(2), 227-247.

Block, N. (2008). Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioural and Brain Sciences,* 30(5-6), 481-499.

Broackes, J. (1992) The autonomy of color, in *Reduction, Explanation and Realism*, K. Lennon and D. Charles. (Eds.) Oxford: Oxford University Press, 421-465.

Carruthers, P. (2005) Consciousness: essays from a higher-order perspective. Oxford: Oxford University Press.

Carruthers, P. (2009) Higher-Order Theories of Consciousness. In E. N. Zalta (Éd.), The Stanford Encyclopedia of Philosophy. http://plato.stanford.edu/archives/spr2009/entries/consciousness-higher/.

Chalmers, D. J. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies*, 2(3), 200-219.

Chichilisky, E. and Wandell, B.A. (1999) Trichromatic opponent colour classification. *Vision Research* **39**, 3444-3458.

Clark, A. (1993) *Sensory Qualities*, Clarendon Press, Oxford.

Clark, A. and Chalmers, D. (1998) *The Extended Mind*. Analysis, 58(1), 7-19.

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

Dehaene, S. and Naccache, L. (2001) Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition,* 79, 1–37.

Dretske, F. (1995) Naturalizing the Mind. Cambridge, MA: MIT Press.

Gegenfurtner, K.R. and Kipers, D.C. (2003) Colour Vision, *Annual Review of Neuroscience*. 26,181–206.

Hardin, C.L. (1988) *Colour for philosophers: Unweaving the rainbow*, Hackett Publishing Company, Indianapolis.

Hayek, F.A. (1953) *The Sensory Order*, University of Chicago Press, Chicago.

Horgan, T. and Tienson, (2002) The Intentionality of Phenomenology and the Phenomenology of Intentionality, in *Philosophy of Mind: Classical and Contemporary Readings*, D. Chalmers, ed., 520-33.

Ijsselsteijn, W.A. (2002) Elements of a multi-level theory of presence: Phenomenology, mental processing and neural correlates. *Proceedings of PRESENCE 2002*, Universidade Fernando Pessoa, Porto, Portugal, Oct. 9-11, pp. 245-259.

Jameson, K., & D'Andrade, R. G. (1997). It's not really red, green, yellow, blue: An inquiry into perceptual colour space. In C. L. Hardin, L. Maffi (Eds.) *Color categories in thought and language*. Cambridge University Press. 295–319.

Kay, P., and Regier, T. (2003). Resolving the question of colour naming universals. *Proceedings of the National Academy of Sciences,* 100(15), 9085-9089.

Kirk, R. (1994) Raw Feeling. Oxford: Oxford University Press.

Knoblauch, K. and Shevell, S.K. (2003) Relating cone signals to colour appearance: failure of monotonicity in yellow/blue. *Visual Neuroscience* 18, 901-906.

Levine, J. (1983). Materialism and Qualia: The Explanatory Gap. *Pacific Philosophical Quarterly* 64, 354-361.

Lycan, W. (1996) Consciousness and Experience. Cambridge, MA: MIT Press.

Madany Mamlouk, A. and Martinetz, T. (2004) On the dimensions of the olfactory perception Space, *Neurocomputing* 58–60, 1019 – 1025

Metzinger, T. (2003). *Being No One: The Self-Model Theory of Subjectivity*. Cambridge, Mass.: Bradford Books, MIT Press.

Nagel, T. (1974) What is it like to be a Bat? Philosophical Review, 83, 435-456.

Natsoulas, T. (1997). The Presence of Environmental Objects to Perceptual Consciousness: A Difference It Makes for Psychological Functioning, *American Journal of Psychology*, 110, 4, 507-526.

Natsoulas, T. (1999) A rediscovery of presence. *Journal of Mind and Behaviour*. 20, 1,17-42.

Noë, A. (2004). Action in perception. 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. and Noë, A. (2001) A sensorimotor account of vision and visual consciousness. Behavioural and Brain Sciences, 24(5), 883-917.

O'Regan, J. K., Myin, E., and Noë, A. (2005) Phenomenal consciousness explained (better) in terms of bodiliness and grabbiness. *Phenomenology and the Cognitive Sciences,* 4,4, 369-387.

O'Regan, J. K., Myin, E., and Noë, A. (2006) Skill, corporality and alerting capacity in an account of sensory consciousness. *Progress in Brain Research*, 150, 55-68.

Palmer, S.E. (1999) Colour, consciousness and the isomorphism constraint. *Behavioural and Brain Science*, 22, 6, 923-43; discussion 944-89.

Peacocke, C. (1983) Sense and Content, Oxford: Oxford University Press.

Porter, J., Craven, B., Khan, R. M., Chang, S., Kang, I., Judkewitz, B., et al. (2007). Mechanisms of scent-tracking in humans. *Nature Neuroscience*. 10(1), 27-29.

Regier, T., Kay, P., & Cook, R. S. (2005). Focal colours are universal after all. *Proceedings of the National Academy of Sciences,* 102(23), 8386-8391.

Rosenthal, D. (2005) Consciousness and Mind. Oxford: Oxford University Press.

Schneider, R. A., and Schmidt, C. E. (1967). Dependency of olfactory localization on non-olfactory cues. *Physiology & Behaviour*, 2(3), 305-309.

Simons, D. J., and Levin, D. T. (1997). Change Blindness. *Trends in Cognitive Sciences*, 1(7), 261-267.

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

Sobel, N., Prabhakaran, V., Desmond, J. E., Glover, G. H., Goode, R. L., Sullivan, E. V., et al. (1998). Sniffing and smelling: separate subsystems in the human olfactory cortex. *Nature*, 392(6673), 282-286.

Strawson, G. (1994) *Mental Reality*, Cambridge, Mass: the MIT Press, Bradford Books.

Teller, D. Y. (1984). Linking propositions. Vision Research, 24(10), 1233-1246.

Tye, M. (1995) Ten Problems of Consciousness. Cambridge, MA: MIT Press.

Valberg, A. (2001) Unique hues: an old problem for a new generation. *Vision Research,* 41, 1645-1657.

Vierkant, T. (2003). *Is the Self Real?* Münster: LIT Verlag.

---

[1] Herbert Feigl has also used the term raw feel in a technical sense, and Kirk (1994) uses "raw feeling". I'm not sure of the relation to my use.

[2] I thank Erik Myin for coming up with this excellent example.

[3] The case of "epistemic actions" (e.g. the Tetris player mentioned by Clark and Chalmers 1998) do not contradict this, since the body motions modify thoughts only indirectly through the way they change our perceptual perspective.

[4] A borderline case is hallucinations and dreams. Though they are also not real, they often do seem real to us. Perhaps part of the reason for this is precisely that they are not completely under our voluntary control.

[5] It would be interesting to find a special class of nerve projections from sensory areas in the brain to the frontal areas of the cortex where the higher cognitive processing is done. These special circuits could provide a kind of "interrupt" command that causes normal cognitive functioning to stop and orient towards the source of the interrupting signal. Such circuits would only be present for sensory channels, and not for systems in the brain that control our autonomic functions.

[6] Demonstrations can be seen on http://nivea.psycho.univ-paris5.fr

[7] Note that this last article has been incorrectly interpreted by Block (2008) as rejecting the view that our internal representations are sparse.

8 For a bibliography on different approaches to the self see the web site maintained by Shaun Gallagher: http://www.philosophy.ucf.edu/pi/. See also a special issue of Annals of the New York Academy of Sciences in Vol 1001, 2003. See also (Vierkant, 2003) for a review. There is also a rich literature on the self within psychoanalysis.

[9] The point that the sensorimotor view is neutral on whether qualia exist, and only attempts to account for what people say about qualia, was not stressed in previous presentations of the sensorimotor approach.