Phenomenology and the Cognitive Sciences (2005) 4: 369–387
DOI: 10.1007/s11097-005-9000-0

Sensory consciousness explained (better) in terms of 'corporality' and 'alerting capacity'

J. KEVIN O'REGAN[1], ERIK MYIN[2] and ALVA NOË[3]
[1]*Laboratoire Psychologie de la Perception, Centre National de Recherche Scientifique, Université Paris 5 René Descartes, Paris, France (E-mail: oregan@ext.jussieu.fr)*
[2]*Centre for Philosophical Psychology, Department of Philosophy, University of Antwerp, Belgium and Centre for Logic and Philosophy of Science, Department of Philosophy, VUB Brussel, Belgium (E-mails: Erik.Myin@ua.ac.be; emyin@vub.ac.be)*
[3]*Department of Philosophy, University of California, Berkeley, CA (E-mail: noe@socrates.berkeley.edu)*

Abstract. How could neural processes be associated with phenomenal consciousness? We present a way to answer this question by taking the counterintuitive stance that the sensory feel of an experience is not a thing that happens to us, but a thing we do: a skill we exercise. By additionally noting that sensory systems possess two important, objectively measurable properties, corporality and alerting capacity, we are able to explain why sensory experience possesses a sensory feel, but thinking and other mental processes do not. We are additionally able to explain why different sensory feels differ in the way they do.

Key words: qualia, consciousness, sensorimotor, skill, sensation, action

Introduction

When you look at a visual scene, the light from the scene is brought together in your eyes; the information is processed by your visual system; an internal representation of the scene is formed in your brain; and then something magical happens: you become *conscious* of the scene. It is not simply registered in your brain, it does not just affect the things you can do and say at this moment: there is an extra "feel," and this extra has been called "phenomenal consciousness."

How can the emergence of "feel" in humans be explained in neural terms? Various hypotheses have been proposed, involving, among others, concepts such as synchronous neural activity across widely separated cortical areas (Crick and Koch 1990; Singer 1993), cortico-thalamic reverberations (Edelman 1989) or quantum processes in neuron microtubules (Penrose 1994; Hameroff 1994), to name but a few.

But surely even if it had been demonstrated that it was, say, quantum processes in microtubules which provided the "feel," one would still be left with the problem of exactly what the special thing is about quantum processes in microtubules which brings about the magic.

There seems to be what Levine (1983) has called an "explanatory gap" between the explanations offered by traditional scientific methods, and the "raw feel" associated with sensation: a kind of barrier that prohibits explanation of phenomenal consciousness in terms of brain mechanisms.

In this paper we will present a way of thinking about phenomenal consciousness which overcomes this problem. It is not a completely new approach, since it combines ideas that have previously been suggested by certain philosophers and scientists.[1] But the parts have been assembled in a new way which is based on empirical investigations, and can be subjected to experimental test.

Our paper is based on the "sensorimotor" theory put forward by O'Regan and Noë (2001b), but it introduces two new ideas (corporality and alerting capacity) which were not developed in that paper, though discussed in O'Regan and Noë (2001a) and, in greater detail by Myin and O'Regan (2002).

A counterintuitive proposal

A neural theory of phenomenal consciousness has to account for why and how particular neural mechanisms provide the particular feels they are supposed to provide. A first aspect of this task concerns the problem of what we call "phenomenality", that is, of accounting for the fact that certain types of neural activity are associated with a sensory feel while others are not. Thus, why is it the case that the brain activity associated with seeing red possesses a *sensory feel,* whereas remembering, thinking, or talking about red does not? A second aspect is the problem of phenomenal quality: Given that certain neural activity does have phenomenality, that is, given that it produces a sensory feel, what makes the feel the same or different from other feels? What exact characteristic of the neurons in *visual* brain areas is it that produces *visual* rather than, say, auditory sensations? What makes certain neurons produce the feeling of red whereas others produce the feeling of green?

Such questions are particularly tricky because it is actually not even clear whether they could ever be answered. Suppose I had observed that oscillations at 100 Hz reliably produced the sensation of red, and 200 Hz produced the sensation of green. Then the question would still remain unanswered of exactly why it is that way. Some reason would have to be provided why particular frequencies gave the particular colors that they give. Surely neural activation is simply like a code that allows information to be transmitted. Seeking the explanation for sensation in particular types of neural activation is like seeking for the meaning of words in the particular shapes of the letters that compose it. Obviously words require letters to be written and differentiated, but the meaning of the words is not caused by the letters. Similarly, it is difficult to see how patterns of neural activation could themselves cause the qualitatively experienced aspects of sensations.

To overcome these difficulties we propose a counterintuitive move. We propose to take the stance that neural activity plays a different role than is usually supposed. Rather than directly producing phenomenal experience, neural activity is involved in phenomenal experience because of what it allows organisms to do. To illustrate this consider the case of pain.

You have cut your finger and it hurts. The hurting seems to you to be going on in a continuous fashion like the buzzing, whirring, or humming of a machine. This ongoingness is a salient phenomenal aspect that calls out for an explanation. Scientists seeking neural mechanisms to explain pain therefore would naturally seek a mechanism which also has this ongoing nature. But consider the case of the light in the refrigerator: you open the refrigerator door – the light is on. You close the door. You quickly sneak open the door again: yes, the light is still on. Every time you check, the light is on. So you are led to believe that the light is on continuously. But in fact it is only on when you check.[2]

Perhaps the situation with pain and sensations in general could be similar: perhaps the apparent "ongoingness" of sensations is actually created by the fact that whenever you attend to them, information about the sensory stimulus is indeed available. If that were so, then, in order to account for the ongoingness of a sensation we would not have to postulate an ongoing brain process. We would merely have to postulate that whenever persons asked themselves about a sensation, they could immediately have access to information about the sensation.

Let us take another example. Most people would agree that there is something it is like to drive a car, and different cars have different "feels". What then is the distinctive feel of driving a Porsche? Does it come from the wind blowing in your hair or the smell of the leather seats, or the purr of the motor as you speed along the Autobahn? Although those are indeed aspects of Porsche driving, people will agree that the critical thing about Porsche driving is how the car *handles*. You have the Porsche driving feel when, while driving, you know that if you press the accelerator, the car will whoosh forwards, whereas if you are driving a tractor more or less nothing happens. In a Porsche, if you just lightly touch the steering wheel, the car swerves around, whereas in the tractor nothing much happens.

But notice something very important: you are speeding along the Autobahn in your Porsche. You briefly close your eyes, take your hands off the steering wheel and your foot off the accelerator. You now are getting very little sensory input. Yet you still are having the Porsche driving experience.

This shows that the Porsche driving feel does not derive from immediate sensory input, or for that matter, from activation of brain mechanisms. It derives – no, it is *constituted by* – the fact that you are engaged in the process of driving the Porsche; it derives from – no, it *consists in* – the fact of relying on your implicit knowledge of what *would* happen now if you did certain things, like press the accelerator or touch the steering wheel (but you need actually do nothing at all).

Let us take another example. There is a party game in which someone puts a household object like a comb, a cork, a potato, or a harmonica, into an opaque bag. You put your hand into the bag and attempt to guess what the object is. When you play this game, you at first have the impression of feeling roughnesses or smoothness, corners, edges, textures on the tips of your fingers. Then suddenly the veil falls, so to speak, and you no longer have the impression of textures localized on your fingertips: instead you feel a *whole object*. And this is despite the fact that your fingers are in contact with only a few small areas of the object.

It is worth playing this game to experience the surprising sudden transition from localized textures on your fingertips to the definite sensation of having a whole object in your hand.[3]

Whence comes the feeling of having a *whole* object, when in fact you are in contact with only a minute portion of it? The answer is, it comes from being *at home* with the things you *might* do (but you need not do them) and with the things that would happen to the sensory input when you do those things.

The object-in-the-bag example shows that the feeling of spatial wholeness derives, not from activation of sensory input coming simultaneously from all parts of the object, but rather from the *implicit knowledge* of what things you might now do in exploring the object, and of what changes would come upon your senses if you did so.

To summarize: the Porsche driving example shows that the phenomenal ongoingness of sensation might derive, not from an ongoing neural mechanism, but from the implicit knowledge of the way sensory input will change when certain actions are undertaken. Similarly, the object-in-the-bag example shows that the spatial extendedness or richness of perception also does not require concomitant richness of sensory input: on the contrary, again, the implicit knowledge of how sensory input will change when you move around, without having to actually move, could be what *constitutes* the feeling of sensing a whole object.

From the analogies discussed above it becomes apparent that we could make a case that neural activity plays a different role than generally supposed with respect to phenomenal consciousness. Instead of *generating* phenomenal consciousness, neural activity *enables the exploratory activity* which experiencing sensory feel *consists in*. By taking this counterintuitive stance we immediately avoid the difficulty of having to explain why one type of neural activity causes phenomenal consciousness while another type of neural activity does not: this is because under the new approach, the differences *reside in the different things you do* rather than in something special about the different neural activities themselves. The neural activities no longer play a direct causal role in phenomenal consciousness. They are necessary in that they enable the exploratory activity, but they are themselves not the direct cause of the associated feels. Instead, the feels are not being directly caused at all, they

are simply *constituted* by the fact that particular skills are being exercised. Take the analogy with meaning again: the letters of a word are necessary for a word to have meaning, but the letters are not causing the meanings. The meanings are constituted by the way the particularly spelled word is used.

It should at this stage be mentioned that the refrigerator light and object in the bag examples are in fact somewhat misleading. For example, in the refrigerator case it was suggested that the sensation of ongoing stimulation could be obtained by sensory input being immediately available whenever it is sought. This can be interpreted as meaning that when the sensory input is found, there is indeed a momentary sensation. But we wish to avoid this way of thinking, since it would leave us with the job of having to explain how each of these momentary sensory stimulations itself provides a feel. We would be back to the explanatory gap.

To avoid this, what we want to say is that sensation is not created by availability of a mass of localized, momentary sensations, but instead that having a sensation *consists in* the fact of being-engaged-in-the-process of manipulating the sensory input. The Porsche driving example is therefore a better analogy than both the refrigerator light and the object in the bag, because it illustrates how a feel is constituted by the exercising of the skill involved in driving, by the give-and-take the driver exercises when he engages in the sensorimotor skill of driving.[4]

Phenomenality

So let us now see how this counterintuitive proposal can be applied to deal with some of the problems of sensory consciousness. The first problem we will discuss is what we call "sensory phenomenality," or "phenomenality" for short.[5]

As an example, consider the difference between actually seeing a red patch and merely thinking about a red patch. The seeing experience clearly has a phenomenal sensory quality that the thought lacks. We suggest that this "phenomenality" of sensory experience can be accounted for by realizing that in animals, neural activity in sensory channels possesses two very special properties which differentiate it from other neural activity in the animal's nervous system and which correspond to specific exploratory relations with the environment: "corporality" and "alerting capacity". These properties account for the fact that sensations have a special quality which is not shared by other mental phenomena like thoughts and knowledge.

Corporality

Suppose at this instant you are in the process of remembering your grandmother. You can cast your attention on the color of her eyes, the sound of her

voice, the smell of her perfume. Any bit of information that you have in memory about your grandmother is immediately available to you whenever you think about her. In this sense then, the process of remembering your grandmother is similar to the object-in-the-bag example discussed above. When you remember your grandmother, though at any one moment you are only recalling one particular aspect of your grandmother, you *know* that any other aspect is available to you. Consequently, when you remember your grandmother you have the same sort of feeling of presence of your grandmother as you did when you suddenly conceived you were holding a whole object in the bag. Nevertheless, your memory of your grandmother somehow lacks perceptual reality. We do not perceive memory as having a truly perceptual quality. Why is this?

Part of the answer is corporality. If you really are looking at your grandmother and you turn your eyes, blink, or move your body, there will be an immediate and drastic change in the incoming sensory information about your grandmother. On the other hand, nothing at all will happen to your remembrances of your grandmother.

Corporality is the fact that when you move your body, incoming sensory information immediately changes. The slightest twitch of an eye muscle displaces the retinal image and produces a large change in the signal coming along the optic nerve. Blinking, moving your head or body will also immediately affect the incoming signal. As concerns auditory information, turning your head immediately affects the phase and amplitude difference between signals coming into the two ears; approaching a sound source increases the amplitude, etc.

Corporality is an aspect of sensory stimulation which makes it different from other forms of neural activity in the brain, and contributes to giving it its peculiar "sensory" quality. Because of corporality, sensory information has a certain "intimate" quality: because sensory information is so exquisitely sensitive to body motions, it is almost as though it were part of you.

Alerting capacity

Suppose that through some minor brain event in the neural areas that are implicated in the memory of your grandmother, you lose the knowledge about what kind of eyeglasses she has. What happens? Are you immediately aware that this has happened? No, you are not. The loss of the memory of your grandmother's glasses causes no whistle to blow or bell to ring in your mind to warn you that this has happened. Only when you cast your mind upon the memory of your grandmother do you actually realize that you no longer know what her glasses were like.

But consider what happens if instead of thinking about your grandmother, you are actually looking at her. Even if you are not paying attention to her

glasses in particular, if they should suddenly disappear, this would inevitably grab your attention: The sudden change in local contours would create a visual transient which would trigger local motion detectors in your low-level visual system, and an eye saccade would immediately be incontrovertibly programmed towards the location of the change. Your attentional resources would be mobilized and orient towards the change.

Animals' sensory systems are wired up so that sudden changes in incoming signals cause the organisms' processing resources to be directed to the sudden event. The mechanisms that allow this are genetically built into sensory systems at the lowest level. Transient detectors in the low level visual system cause your eyes to saccade involuntarily to any sudden flash or motion. In the auditory system there are head-orienting reflexes which cause you to turn to sudden noises, and the same is true for sudden tactile stimulation.

This "grabbiness" of sensory stimulation, that is, its capacity to cause automatic orienting responses, is therefore another aspect of sensory stimulation which distinguishes it from other types of neural activity in the brain. Alerting capacity is the fact that sensory stimulation can grab your attention away from what you were previously doing.

Degrees of feel

Our claim is that corporality and alerting capacity together contribute to giving sensory stimulation a peculiar "intimate" quality which corresponds to what people call "feel." We think that what people mean when they say a mental phenomenon has "feel" is precisely that the mental phenomenon has corporality and alerting capacity. Because of corporality, you are in a way "connected" to sensory stimulation: it changes with your minutest body motion. Because of alerting capacity, you somehow can't get away from sensory stimulation: it has the capacity to monopolize your attention and keep you in mental contact with it. Corporality and alerting capacity ensure that, unlike thoughts and memories which you can mentally summon or discharge like discrete servants, sensory stimulation has a "clinging" quality. Unlike thoughts and memories, sensations stay with you like your shadow. They follow you around like a faithful dog. Furthermore, like the dog, they force themselves upon you by grabbing your attention whenever anything unexpected happens in the world.

We suggest that corporality and alerting capacity may be the reason why sensory experience has phenomenality.

Note an important point about the concepts of corporality and alerting capacity: they are objectively measurable. A scientist with an oscilloscope, say, or a voltmeter should be able to measure how much corporality and how much alerting capacity there is in different types of sensory stimulation. The amount of corporality is determined by the amount of change that occurs in

sensory input when the body moves. The amount of alerting capacity is, for example, determined by the amount of orienting response that can be provoked by sudden changes in sensory input.

If corporality and alerting capacity are objectively measurable quantities, and if we are right in saying that they determine whether a sensory input is perceived as possessing "feel," then we should be able to predict how much "feel" different mental phenomena will be perceived as having.

We have already seen that memory phenomena, like the memory of your grandmother, or thoughts or knowledge, have no corporality and no alerting capacity. They should consequently be perceived as having no sensory feel: this seems to correspond with what people say about memory, thoughts and knowledge.[6]

We have also seen that sensations, like the color of your grandmother's eyes, have corporality and alerting capacity, and should be perceived as possessing "feel". And indeed most people will agree that it feels like something sensory to see the color of your grandmother's eyes, whereas it does not feel like anything sensory (or at least much less so) to be thinking about her eyes. This difference in phenomenality, we suggest, seems well captured by the difference in corporality and alerting capacity between seeing on the one hand, and thinking on the other: seeing has corporality and alerting capacity, thinking has none.

Now it is interesting to consider whether there exist intermediate cases. If we are right about the relation between corporality and alerting capacity and feel, then cases of a little bit of corporality and alerting capacity should correspond to a little bit of feel.

Indeed a case in point is Porsche driving. In Porsche driving, some of your body movements produce immediate changes in sensory input – pressing the accelerator, touching the wheel, etc. But most of your body movements do not change sensory input related to the Porsche driving experience. Turning your head changes visual input, but the change is not specific to the Porsche driving feel – rather it constitutes the feel characteristic of vision. Sniffing your nose gives you the smell of leather, but that is specific to the sense of smell. Those very particular sensorimotor contingencies which determine the feel of Porsche driving are restricted to a very particular set of behaviors which are specific to *driving*, namely those to do with how touching the wheel or pressing the accelerator affects what the car does. You cannot get the feel of a car by just waving your hands around in the air. You have to actually be exercising the car-driving skill.

The situation is quite different from the feel of seeing red or hearing a bell, say, where almost any small body twitch or muscle movement cause drastic sensory changes.

We thus expect – and this corresponds well with what people say about the feel of driving – that it makes sense to say that Porsche driving has a feel, but

the feel is less intimate, less direct, less "present" than the sensation associated with seeing red or hearing a bell.

Another interesting intermediate case is the feeling of being rich. What is being rich? It is knowing that if you go to your bank you can take out money; it is knowing you can go on an expensive trip and that you need not worry about the price of dinner.

Thus being rich has a certain degree of corporality, because there exist things you can do with your body which have predictable sensory consequences (e.g. you can make the appropriate maneuvers at the cash dispenser and the money comes out). But clearly, again, the link with body motions is nothing like as direct as in true sensory stimulation like seeing, in which the slightest motion of virtually any body part creates immediate changes in sensory input. So being rich can hardly be said to have very much corporality.

Similarly, being rich also has no alerting capacity. If your bank makes a mistake and suddenly transfers all your assets to charity, no alarm-bell rings in your mind to tell you. No internal mind-siren attracts your attention when the stock market suddenly goes bust: you only find out when you purposefully check the news.

Thus being rich has neither corporality nor alerting capacity. It therefore makes sense that when people say they feel rich, they will agree that the "feel" they experience lacks the presence, or the raw sensory quality that they experience with true sensory experiences like sights, sounds and smells.

Dreams, hallucinations and mental imagery

Hallucinations and dreams are characterised by the fact that when people are hallucinating or dreaming they are convinced that they are having something that shares at least some of the phenomenal properties with full-blown perceptual experiences that they have in real life. Clearly however neither hallucinations nor dreams can involve corporality or alerting capacity in the normal fashion, since there is no sensory input at all (in dreaming), or input not normally related to the experienced content (in hallucination).

On the other hand it is also clear that it is precisely corporality which ultimately allows people to realise that they are actually hallucinating or dreaming – the classic way of knowing that you are dreaming is to try to switch on the light: this kind of "reality-checking" is nothing more than testing for corporality – checking that your actions produce the normal sensory changes expected when you are having real sensory experiences.[7]

It is important to note however that what counts in giving the particular "sensory" feel of sensation is not the actual sensory input itself, but the reliance on *implicit knowledge* that the sensory input possesses corporality and alerting capacity. This means, as in the example of Porsche driving or the object-in-the-bag, that an observer can have a sensation even though he is, at a given

moment, doing *nothing at all*, and even though he is receiving *no sensory input at all*. It suffices for this that he be in the same condition which he would usually be in when he is putting to active use his implicit knowledge that the sensorimotor contingencies associated with a sensation are currently applicable.

It is as though, in the Porsche driving example, someone surreptitiously removed everything but the seat of the Porsche out from under you precisely at the moment when you had closed your eyes, lifted your foot from the accelerator and your hands from the steering wheel. At that moment you would still be enjoying the Porsche driving experience, even though you were getting no sensory input at all, and though there was no Porsche at all.

We can therefore understand how it might easily come about that a person would have experience as of reality without any sensory input, and therefore no corporality and alerting capacity. The person merely has to rely on the implicit assumption (in point of fact incorrectly) that if he were to move, then those changes would occur that normally occur when he moves. He just has to implicitly expect (incorrectly) that were there to be a sudden event, his attention would be incontrovertibly attracted to it.

Dreaming and hallucinating therefore pose no problem for the sensorimotor approach that we are proposing. Indeed the approach actually makes it easier to envisage brain mechanisms that engender convincing sensory experiences without any sensory input, since the sensation of richness and presence and ongoingness can be produced in the absence of sensory input merely by the dreamer implicitly 'supposing' (in point of fact incorrectly) that if the eyes were to move, say, they would encounter more detail. This state of 'supposing you can get more detail' would be a much easier state to generate than having to actually recreate all the detail somewhere in the brain. In dreaming, furthermore, the state would be particularly easy to maintain because what characterizes dreaming would seem to be a lack of attention to the absence of disconfirming evidence, which is quite unsurprising, since you are asleep. This lowering of epistemic standards implies that, while dreaming, you are easily led into thinking you are perceiving, while – if only you were to pay attention – it would be obvious that you are not. Thus you can remain convinced for the whole duration of your dream that you are experiencing reality. A whole series of different bizarre dream events may be taken at face value simply because nothing disconfirms them.

Similar remarks apply to mental imagery. As for dreams and hallucinations, mental imagery would correspond to a kind of perceptual action without an actual stimulus and without 'going through' the motions – it would involve having implicit expectancies[8] without these being actually fulfilled by worldly responses. Notice that both the unrealness of dreams and the fleetingness of mental imagery can be attributed precisely to the absence of the implicitly expected responses from the world. There is *just* implicit expectation, and

without support from the world outside, implicit expectation alone cannot sustain for long something which is truly like real perception. In the Porsche-driving example when everything but the seat is removed, the driver will very soon find out that he's no longer driving a Porsche, and the Porsche-driving feeling will vanish.

In recent papers, Thomas (1999) has developed, more fully than we can do here, a strongly related account of mental imagery and imagination. Situating himself within what he calls 'Perceptual Activity Theory', he shows how his 'Active Perception Approach' to imagery and imagination[9] has, besides considerable theoretical advantages, the potential to account for a wide range of empirical data.

Conclusion on phenomenality

We have seen that, when added to the idea that feels correspond to exploratory activity involving sensorimotor skills, the concepts of corporality and alerting capacity allow the fundamental difference to be captured between mental phenomena that have no sensory feel, like memory and knowledge, and mental phenomena that have sensory feel, like sensations. Corporality and alerting capacity furthermore allow us to understand why certain intermediate situations, like driving or being rich can also be qualified as possessing a certain, but lesser, degree of "feel." Corporality and alerting capacity are objectively measurable quantities that determine the extent to which there is a specific sensory quality to the experience of having a sensation.

Phenomenal quality: Intermodal differences

But if sensory stimulation has the specific sensory quality we call phenomenality, then the question immediately arises of what exactly that quality is, and why and how the quality can differ, in particular, depending on the sensory modality involved.

For example hearing involves a different quality as compared to seeing, which has a different quality as compared to tactile sensations.

To a neuroscientist, the obvious way of explaining these different phenomenal qualities would be in terms of different neural mechanisms: one might note that seeing involves activation in visual cortex, whereas hearing involves activation in auditory cortex. Johannes Müller in the 19th Century had proposed that sensory nerves possess a "specific nerve energy" which lies at the origin of the different sensory qualities – the modern idea that different cortical mechanisms are involved is today's version of the same idea.

But this will not do. After all, neural activation, be it here or there in the cortex, is simply neural activation. Something more is needed to explain why

a particular neural activation purveys a particular hearing-like quality, and another activation purveys a particular seeing-like quality. How could different neural activations possibly give rise to different feelings? Neural activation is simply a way of coding information in the brain. How can differences in the code ever give rise to differences in feel?

Clearly the point is that it is not neural activation which is generating sensation, rather it is what that neural activation *allows the organism to do*.

Said in another way, the specific nature of a feel resides in the specific things you do, or can do, when you have that feel. The differences between different feels are accounted for by the fact that they involve *doing different things*.

Consider our example of driving a Porsche again, and think about how its feel differs from the feel of riding a bicycle. Does the difference come from the different brain regions that are involved in the two activities? No. The difference comes from the fact that you do different things when you drive a Porsche as compared to riding a bicycle. Among other things, Porsche driving is knowing that you can get the whoosh by pressing on the accelerator; bicycle riding is knowing that you must peddle to stay up.

We can now understand the origin of the difference between the feel of seeing and the feel of hearing. The difference comes from the different ways your sensory input depends on the things you can do. You are seeing if, when you blink, there is a large change in sensory input; if when you turn your head, the optic flow on your retina shifts sideways; if when you move forward there is an expanding flow field; if when you make an eye movement there is a drastic perturbation and shift of the retinal image. On the other hand, you are hearing if, when you blink, nothing happens; if when you turn your head the input along the auditory pathways from each ear changes intensity and phase in a certain lawful way; if when you make an eye movement nothing much happens in the auditory pathways; if when you move forward the amplitude of the input along the auditory pathways increases, etc.

Just as the feel of Porsche driving is constituted by the things you might (but need not) do when you drive a Porsche, the feel of hearing and seeing are constituted by the different things you might (but need not) do when you explore the world following auditory or visual sensorimotor skills.

Of course the notion of sensorimotor law can be conceived at many levels, and the laws just illustrated are situated at a high level, common across individuals. On the other hand there are aspects of sensorimotor laws that are particular to particular individuals. Thus, in a given individual, the particular input/output relationships that hold between his eye muscle commands and the resulting sensory changes in his optic nerve will be very idiosyncratic, depending on the particular way that individual's visual apparatus and muscle command pathways are wired up. Part of the phenomenal quality of vision will be constituted by this particular idiosyncratic set of input/output

relationships. This will be part of the feel of vision that the person will never be able to describe verbally or share with other people – it is a possible specification of what philosophers call the "ineffable" part of visual sensation.

Nevertheless, the *common* laws of vision, such as those mentioned above related to eye movements, body movements and the resultant changes in optic flow, will generally be shared by people with visual apparatus similar to ours, and will be constitutive of the phenomenal quality of vision, as opposed to audition or the tactile sense, in organisms such as humans.

Sensory substitution and other tests

There is an important consequence of these ideas. The hypothesis is that, barring the "ineffable" aspects of a particular sensory modality, the shared aspects do not derive from the particular sensory input channel or neural circuitry involved in that modality, but from the laws of sensorimotor contingency that are obeyed by that particular modality. A counterintuitive prediction follows from this: It should be possible to obtain a visual feel from auditory or tactile input, for example, provided the sensorimotor laws that are being obeyed are the laws of vision (and provided the brain has the computing resources to extract those laws).

Such "sensory substitution" has been experimented with since Bach-y-Rita (1967) constructed the first device to allow blind people to see via tactile stimulation provided by a matrix of vibrators connected to a video camera. Today there is renewed interest in this field,[10] and a number of new devices are being tested with the purpose of substituting different senses: visual-to-tactile (Sampaio et al. 2001); visual-to-auditory (Veraart et al. 1992); auditory to visual (e.g. Meijer 1992); auditory-to-tactile (cf. review by Richardson and Frost 1977).

Such devices are still in their infancy. In particular, no systematic effort has been undertaken up to now to analyze the laws of sensorimotor contingency that they provide. In our opinion it will be the similarity in the sensorimotor laws that such devices recreate which determines the degree to which users will really feel they are receiving stimulation in the modality being substituted.

A very interesting experiment by Botvinick and Cohen (1998) also argues in favor of the view that sensation is determined by sensorimotor contingencies. In this experiment, an observer sits with a rubber replica of an arm in front of him on a table. The observer puts his own arm on the table next to the rubber arm, but hidden from his view by an opaque screen. With a paintbrush, for example, the experimenter simultaneously strokes the observer's arm and the rubber arm, in a repetitive fashion, taking care to make the same movements for the real arm and the rubber arm. After about ten minutes, the observer has the distinct feel of being touched *on the rubber arm*.

This result shows that the localization of sensation on the body is determined by the sensorimotor contingencies between tactile and visual input. Our approach would predict that the phenomenon would be even stronger if the observer's own actions were involved, but this remains to be tested.

Hurley and Noë (2003) detail how a basically sensorimotor approach allows to predict and explain which type of experience a neural zone will be involved in after experimental or natural rewiring.

Phenomenal quality: Intramodal differences

We have seen that the feel of different sensory modalities can be accounted for by the different things you do when you use these modalities. But what about the differences *within* a given sensory modality: can we use the same arguments?

Within the tactile modality, this idea seems quite plausible. Consider the feel of a hard surface and the feel of a soft surface. Does this difference come from different kinds of tactile receptors being activated, or from the receptors being activated in different ways? No, we argue, since receptor activations are only codes that convey information – they cannot by themselves be the origin of the feel of hard and soft. On the contrary, we claim the difference between hard and soft comes from the fact when you push on a hard surface it resists your pressure; when you push on a soft surface, it gives way. The feel of hard and soft are constitued by the different skills that you implicitly put to work when you touch hard and soft surfaces; the feel of a surface is constituted by the things you implicitly know about the way the surface will react to your exploration.

Now while this makes sense for tactile exploration, it might seem difficult to apply the same approach to other sensory modalities: what has the difference between red and green for example, got to do with sensorimotor contingencies? How can the feel of red consist in *doing something*, and the feel of green consist in *doing something else*?

Consider what happens when you look at a red piece of paper. Depending on which way you turn the paper, it can reflect more of bluish sky light or more of yellowish sunlight from your window, or more of reddish lamplight from your desk. We suggest that the feel of red is: knowing the laws that govern the changes in the light reflected off the paper as you turn it.

Another point: when you look directly at a patch of red, the light is sampled by the central part of your retina, which is densely populated with photoreceptors sensitive to long, medium and short-wavelength light. But when you move your eyes off the patch, it falls on the peripheral retina, where photoreceptors are less densely packed, and where the distribution of photoreceptor

types is slightly different. Seeing red is: knowing that laws typical of red are currently governing the changes in sampled light as you move your eyes.

If we are correct about the fact that the phenomenal quality of color is determined, not by the particular neural pathways which are activated by the color, but by the laws of sensorimotor contingency which are potentially obeyed when the color is explored, then we can make a curious prediction.

We could set up an experiment in which observers are hooked up to apparatus that measures their eye movements as they look at a patch of color on a computer screen. We could arrange things so that when the observers look directly at the color, the patch is red, but as soon as the observers' eyes move off the patch, its color goes green. By this means we have set up a new law of sensorimotor contingency. We predict that after a while, observers should adapt to this new sensorimotor contingency, and should come to see the central red patch as being *the same color* as the peripheral green patch. This and other predictions about the phenomenology of color are currently being tested with some degree of success (Bompas, Clark and O'Regan 2002, Bompas and O'Regan, in press; Philipona and O'Regan, in press).

Conclusion on phenomenal quality

One of the barriers preventing a scientific explanation of phenomenal consciousness was that it was difficult to understand how different types of neural activation, or activation of different cortical areas could give rise to different qualities of sensation – neural activations are just arbitrary codes for information, and information in itself has no feel.

We have circumvented this problem by construing feel, not as neural activation, but as what that neural activation *allows*: that is, a kind of *skill*, like driving or bicycle riding. The differences between the feels of the different sense modalities are accounted for naturally by realizing that sensory exploration in different sense modalities corresponds to different skills, and so *constitutes* different feels. The idea can also be made to work *within* a given sense modality, explaining the what-it-is-like of red versus that of green in terms of the different things you do when you are exploring red and green.

Conclusion

The fact that sensations appear to happen to us in an ongoing, occurrent way has led many scientists, in their quest for the neural correlate of consciousness, to seek brain mechanisms which are themselves also ongoing. In particular, scientists are searching for brain mechanisms whose *activity* gives rise to phenomenal consciousness. But we claim that any such quest is doomed, since the question will always ultimately remain of how activity of a physical system,

no matter how complex or abstruse, can give rise to something which seems inherently nonphysical, namely "feel." Invoking phase transitions, qualitative changes of functioning, or other principles linked to exceeding a critical degree of complexity or connectedness only pushes the mystery back into a deeper hiding place.

Our solution is to show that it is an error to think that the sensory experience requires a concomitant ongoing brain mechanism. On the contrary, we show how the experience can be obtained by the organism having implicit knowledge about what would happen to the organism's sensory inputs if the organism were to move. We made this at first highly counterintuitive idea more palatable by giving various analogies: that of the refrigerator light, the Porsche driving experience, and the object-in-the-bag game. We suggested that "feel" should be construed, not as a brain activity, but as a skill that the organism excercises. The skill we refer to involved having mastery of what we called the sensorimotor contingencies – that is, the laws of co-variation between actions of the organism and resulting changes in sensory input.

Taking this counterintuitive stance we can adress the question of sensory feel. We showed that two objectively measurable aspects of animals' sensory systems accounted for the peculiar quality of sensations: corporality and alerting capacity. Corporality is the fact that the minutest body movements immediately alter neural activation deriving from sensory input, but such movements do not affect knowledge, beliefs, or other non-sensory mental phenomena.

Alerting capacity is the fact that in animals, sensory systems are wired up so that sudden changes in sensory input cause the animal to immediately orient its processing resources towards the change. Only sensations possess alerting capacity: knowledge and beliefs possess no alerting capacity.

Because corporality and alerting capacity come in varying degrees, it is possible to make what might be called an "analytic phenomenology" of mental phenomena, going from those that have a high degree of corporality and alerting capacity, to those that have none. We showed that as expected, the amount of corporality and alerting capacity that a mental phenomenon possesses correlates well with the amount of "feel" that is generally associated with the phenomenon. So the color red, which has a lot of corporality and alerting capacity is perceived as being a proper sensation, with a lot of feel. The "feel" of driving a car and the "feel" of being rich possess intermediate amounts of corporality and alerting capacity, and as expected the "feels" involved are of an intermediate nature, being perceived as progressively less sensory, as having less "phenomenality." Finally mental phenomena like knowledge, memories and beliefs possess no corporality or alerting capacity at all, and are not perceived as having any phenomenal quality.

Another hurdle to understanding phenomenal consciousness was phenomenal quality. Just as brain activity could not explain the origin of the "what

it is like" of sensory experience, brain activity in itself can also not explain the differences in the "what it is like" of different sensory modalities. The sensorimotor contingency theory however very naturally accounts for this. The reason is that under this theory, the feel of a sensation is *constituted* by the skill that is involved in experiencing it. Thus the different sensory modalities, as for example hearing, seeing, and touch, to the extent that they involve different skills, must perforce be associated with different sensory qualities.

Of course, whenever a creature is in a perceptual situation, countless sensorimotor contingencies apply, and some of these are actualised by associated exploratory movements. But, either by being alerted by something significant that happens out there, such as a sudden movement, or a loud noise, or by the perceiver's own decision, one of the ongoing patterns of exploration is allowed to play a prominent role in the perceiver's mind. The perceiver bundles all the capacities he disposes of and reorganises them with respect to what thereby becomes the prominent theme of the ongoing perception (note that it might be prominent only fleetingly). Then, the perceiver is perceptually aware.

Acknowledgments

Erik Myin wishes to thank, for support, the Fund for Scientific Research-Flanders (Belgium) (FWO project G.0175.01). Alva Noë gratefully acknowledges the support of faculty research funds granted by the University of California, Santa Cruz as well as that of a UC President's Research Fellowship in the Humanities.

Notes

1. E.g. Bergson, Ryle, Merleau-Ponty, G. Evans, D.M. MacKay, J.J. Gibson.
2. This very nice example was suggested by Thomas (1999) in an approach quite similar to ours.
3. The impression is similar to what happens when you look at some kinds of hidden figure. Suddenly you see what the figure represents, and you see it as a whole, whereas before you were looking at the same figure, and all you saw was disconnected bits.
4. Our take on this is thus different from Block's position in his discussion of the 'refrigerator light' metaphor (see Block 2001, pp. 208–209). Block there opposes the idea that consciousness of more than the momentary local input (the whole object in our bag example, a whole scene in his visual example) is dispositional, because, according to him, we feel that we are phenomenally conscious of the whole object (the whole scene in his example), not that we could become phenomenally conscious if we would make this or that movement. But this doesn't represent our position. The claim is not that we know that we could be conscious, but that we are conscious because we know we can obtain such and such effects with such and such movements. There is no residual mysterious notion of phenomenal

consciousness nested within the disposition; the disposition IS the consciousness. Our position seems to have been anticipated by Merleau-Ponty (1945/1976), p. 81: "Voir un objet, c'est ou bien l'avoir en marge du champ visuel et pouvoir le fixer, ou bien répondre effectivement à cette sollicitation en le fixant. Quand je le fixe, je m'ancre en lui, mais cet 'arrêt' du regard n'est qu'une modalité de son mouvement: je continue à l'intérieur d'un object l'exploration qui, tout à l'heure, les survolait tous," Translation: "To see an object, is either to have it in the fringe of the visual field and be able to fixate it, or is to actually respond to this solicitation and fixate it. When I fixate it, I anchor myself in it, but this standstill of the gaze is but a modality of its movement: I continue, inside the object, the exploration which before encompassed all of them (objects)."

5. "Sensory phenomenality" refers to the fact that perceptual experience has a particular conscious 'feel' of being perceptual. We will, for reasons of readability, further drop the "sensory", but still mean to refer to "sensory phenomenality"
6. Perhaps there are cases, like anguish or obsessive thoughts which might be considered to have 'feel', to some extent. This can be understood when one notes that what characterizes such thoughts is that they wholly occupy your mental resources: they have a form of alerting capacity, even though they have no corporality.
7. Interestingly, alerting capacity, on the other hand, seems more difficult to check when you are dreaming or hallucinating, since this would require you to mentally generate an unexpected event, something which is impossible if it is you yourself who decides to generate it. Although of course you can dream that you check!
8. We use the word expectancy but qualify it with the adjective "implicit" so as to emphasize that we do not mean that the observer is consciously expecting. What we mean is that the observer is tuned to the sensorimotor contingencies.
9. In a further paper, Thomas (2001), the approach sketched out in Thomas (1999) is applied to color perception.
10. There is currently a European Community initiative financing such work.

References

Bach-y-Rita, P. 1967. Sensory plasticity. Applications to a vision substitution system. *Acta Neurol Scand.* 43(4): 417–426.

Bompas, A. and O'Regan, J. K. in press. Evidence for a role of action in color vision. *Perception*.

Bompas, A., Clark, J. and O'Regan, J. K. 2002. Color perception in the sensorimotor contingency theory. *Perception* 31: 90–90.

Block, N. 2001. Paradox and cross purposes in recent work on consciousness. *Cognition* 79: 197–219.

Botvinick, M. and Cohen, J. 1998. Rubber hands 'feel' touch that eyes see [letter]. *Nature* 391: 6669: 756.

Crick, F. and Koch, C. 1990. Toward a neurobiological theory of consciousness. *Seminars in the Neurosciences* 2: 263–275.

Edelman, G. M. 1989. *The Remembered Present.* New York: Basic Books.

Hameroff, S. R. 1994. Quantum coherence in microtubules: A neural basis for emergent consciousness? *Journal of Consciousness Studies* 1: 91–118.

Hurley, S. L and Noë, A. 2003. Neural plasticity and consciousness. *Biology and Philosophy* 18: 131–168.

Levine, J. 1983. Materialism and qualia: The explanatory gap. *Pacific Philosophical Quarterly* 64: 354–361.

Meijer, P. B. L. 1992. An experimental system for auditory image representations. *IEEE Trans. Biomed. Eng.* 39(2): 112–121.

Merleau-Ponty, M. 1945/1976. *Phénoménologie de la perception*. Paris: Gallimard.

Myin, E. and O'Regan, J. K. 2002. Perceptual consciousness, access to modality, and skill theories. *Journal of Consciousness Studies* 9(1): 27–45.

O'Regan, J. K. and Noë, A. 2001a. Acting out our sensory experience. *Behavioral and Brain Sciences* 2(5): 955–975.

O'Regan, J. K. and Noë, A. 2001b. A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences* 24(5): 883–917.

Philipona, D. and O'Regan, J. K. in press. Color naming, unique hues and hue cancellation predicted from singularities in reflection properties. *Visual Neuroscience*.

Penrose, R. 1994. *Shadows of the Mind.* Oxford: Oxford University Press.

B. J. Richardson, B. L. and Frost, B. J. 1977. Sensory substitution and the design of an artificial ear. *J. Psychol.* 96 (2nd Half): 259–285.

Sampaio, E. Maris, S. and Bach-y-Rita, P. 2001. Brain plasticity: 'visual' acuity of blind persons via the tongue. *Brain Res.* 908(2): 204–207.

Singer, W. 1993. Synchronization of cortical activity and its putative role in information processing and learning. *Annu. Rev. Physiol.* 55: 349–374.

Thomas, N. 1999. Are theories of imagery theories of imagination? An active perception approach to conscious mental content. *Cognitive Science* 23: 207–245.

Thomas, N. 2001. Color realism: Toward a solution to the 'hard problem'. *Consciousness and Cognition* 10: 140–145.

Veraart, C., Cremieux, J. and Wanet-Defalque. 1992. Use of an ultrasonic echolocation prosthesis by early visually deprived cats. *Behav Neurosci* 106(1): 203–216.