TOWARDS AN ANALYTIC PHENOMENOLOGY: THE CONCEPTS OF "BODILINESS" AND "GRABBINESS"[1]

J. KEVIN O'REGAN
Laboratoire de Psychologie Expérimentale, CNRS, Univ. René Descartes,
Boulogne-Billancourt, France

ERIK MYIN
Centre for Logic and Philosophy of Science, Vrije Universiteit Brussel, Belgium

ALVA NOË
Dept. of Philosophy, University of California, Santa Cruz, USA

1. Phenomenal Consciousness

In this paper, we present an account of phenomenal consciousness. Phenomenal consciousness is experience, and the *problem* of phenomenal consciousness is to explain how physical processes—behavioral, neural, computational—can produce experience. Numerous thinkers have argued that phenomenal consciousness cannot be explained in functional, neural or information-processing terms (e.g. Block 1990, 1994; Chalmers 1996). Different arguments have been put forward. For example, it has been argued that two individuals could be exactly alike in functional/computational/behavioral measures, but differ in the character of their experience. Though such persons would behave in the same way, they would differ in how things felt to them (for example, red things might give rise to the experience in one that green things give rise to in the other). Similarly, it has been held that two individuals could be functionally/computationally/behaviorally alike although one of them, but not the other, is a mere *zombie*, that is, a robot-like creature who acts *as if* it has experience but is in fact phenomenally unconsciousness. For any being, it has been suggested, the question whether it has experience (is phenomenally conscious) cannot be answered by determining that it is an information-processor of this or that sort. The question is properly equivalent to the question whether *there is anything it is like to be* that individual (Nagel 1974). Attempts to explain consciousness in physical or information-processing terms sputter: we cannot get any explanatory purchase on experience when we try to explain it in terms of neural or computational processes. Once a particular process has been proposed as an explanation, we can then always reasonably wonder, it seems, what it is about *that particular* process that make it give rise to experience. Physical and computational mechanisms, it seems, require some further ingredient if they are to explain experience. This explanatory shortfall has aptly been referred to as "the explanatory gap" (Levine 1983).

We suggest that the explanatory gap is a product of a way of thinking about consciousness which sets up three obstacles to an explanation, that is, three reasons for holding that the explanatory gap is unbridgeable. In this paper we propose ways of surmounting these obstacles, and in this way try to lay the foundations for a science of phenomenal consciousness.

What is it exactly about phenomenal consciousness which makes it seem inaccessible to normal scientific inquiry? What is so special about "feel"?

2. The first obstacle: the continuousness of experience

A first remarkable aspect about experience is that it seems 'continuous'. Experiences seem to be "present" to us, and to have an "ongoing" or "occurring" quality which we might picturesquely describe as like the buzzing, whirring, or humming of a machine.

Many scientists believe that to explain the ongoingness of experience we must uncover some kind of neural process or activity that *generates* this ongoingness. But this is a mistake (Dennett 1991; Pessoa, Thompson and Noë 1998; Hurley 1998). To see why, consider an analogy. Most

[1] This paper offers a theoretical overview of ideas developed in an a series of recent papers—O'Regan and Noë 2001a, b; c; Myin and O'Regan 2002; Noë and O'Regan 2000; 2002; Noë 2002; O'Regan 1992—and also in work in progress by the authors.

people would agree that there is something it is like to drive a car, and different cars have different "feels". You have the Porsche driving feel when you know that if you press the accelerator, the car will whoosh forwards, whereas nothing comparable happens in other cars. In a Porsche, if you just lightly touch the steering wheel, the car swerves around, whereas most other cars react more sluggishly. In general: the feel of driving a car, truck, tank, tractor or golf-cart corresponds to the specific way it behaves as you handle it.

Now as you drive the Porsche, you are having the ongoing Porsche driving feel. But notice that as you drive you can momentarily close your eyes, take your hands off the steering wheel and your foot off the accelerator, yet you are still having the Porsche driving feel even though you are getting virtually no Porsche-related sensory input. This is because the Porsche driving feel does not reside in any particular momentary sensory input, but rather in the fact that you are currently engaged in exercising the Porsche driving skill.

If the feel of Porsche driving is constituted by exercising a skill, perhaps the feel of red, the sound of a bell, the smell of a rose also correspond to skills being exercised. Taking this view about what feel is would have a tremendous advantage: we would have crossed the first hurdle over the explanatory gap, because now we no longer need a magical neural mechanism to generate ongoing feel out of nerve activities. Feel is now not "generated" by a neural mechanism at all, rather, it is exercising what the neural mechanism *allows the organism to do*. It is exercising a skill that the organism has mastery of.

An analogy can be made with "life": life is not something which is generated by some special organ in biological systems. Life is a *capacity* that living systems possess. An organism is alive when it is *has the potential* to do certain things, like replicate, move, metabolize, etc. But it need not be doing any of them right now, and still it is alive.

It may seem very peculiar to conceive of say, the feel of red, as a skill being exercised, but we shall see the possibility of this position, as well as its advantages, in the next sections. The idea and its implications has been developed in our previous papers (O'Regan & Noë 2001a; O'Regan & Noë 2001b; O'Regan & Noë 2001c; Myin & O'Regan 2002; cf. also Clark 2000; Grush 1998; Järvilheto 2001; Myin 2001, Pettit 2003a,b for similar recent views).

A consequence of the "skill" idea: change blindness

When we look out upon the world, we have the impression of seeing a rich, continuously present visual panorama spread out before us. Under the idea that seeing involves exercising a skill however, the richness and continuity of this sensation are not due to the activation in our brains of a neural representation of the outside world. On the contrary, the ongoingness and richness of the sensation derive from the knowledge we have of the many different things we can do (but need not do) with our eyes, and the sensory effects that result from doing them (O'Regan 1992). Having the impression of a whole scene before us comes, not from every bit of the scene being present in our minds, but from every bit of the scene being immediately available for "handling" by the slightest flick of the eye.

But now a curious prediction can be made. Only part of the scene can be being "handled" at any one moment. The rest of the scene, although perceived as present, is actually not being handled. If such currently un-handled scene areas were to be surreptitiously replaced, the change should go unnoticed.

Under normal circumstances any change made in a scene will provoke an eye movement to the locus of the change. This is because there are hard-wired detectors in the visual system that react to any sudden change in local luminance and cause attention to focus on the change. (We will come back to this important property of the visual system under the heading of "grabbiness" in Section 3.)

But by inserting a blank screen or "flicker" (Rensink, O'Regan & Clark 2000), or else an eye movement, a blink, "mudsplashes" (O'Regan, Rensink & Clark 1999), or a film cut between successive images in a sequence of images or movie sequence (for a review see Simons 2000);, the sudden local luminance changes that would normally grab attention and cause perceptual handling of a changing scene aspect are drowned out by the mass of other luminance changes occurring in the scene. There will no longer be a single place that the observers' attention will be attracted to, and so we would expect that the likelihood of "handling" and therefore perceiving the location where the scene change occurs would be low.

And indeed that is what is found: surprisingly large changes, occupying areas as large as a fifth of the total picture area, can be missed. This is the phenomenon of "change blindness" (demonstrations can be found on http://nivea.psycho.univ-paris5.fr and http://viscog.beckman.uiuc.edu/change/).

3. The second obstacle: the qualitative character of experience

In the previous section we showed that by taking the view that experiences depend on the exercise of skills, we can forego the search for neural processes that are, like the experiences themselves, ongoing. We no longer need to postulate a magical neural process that "generates" phenomenal consciousness, because, we claim, phenomenal consciousness is not generated: rather it is a skill people exercise.

We now come to the second difficulty in explaining experience.

Suppose you are thinking about your grandmother. You can cast your attention on the color of her eyes, the sound of her voice, the smell of her perfume. Nevertheless, thinking about your grandmother is nothing like actually seeing her: thinking has no perceptual phenomenal quality. Why is this? Why is there something it is like to have a perceptual experience (Nagel 1974)? This question forms the second obstacle that would seem to bar our path towards understanding phenomenal consciousness.

The key, we propose, has to do with distinct properties of the kinds of skills that we exercise when we undergo conscious experience and that make these skills different from other skills (practical skills such as the ability to drive, cognitive skills, etc). These aspects are bodiliness and grabbiness.

Bodiliness

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 if you are merely thinking about your grandmother.

Bodiliness 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 from the two ears, etc.

Bodiliness is one aspect of *sensory* stimulation which makes it different from other forms of stimulation, and contributes to giving it its peculiar quality. Because of bodiliness, sensory information has an "intimate" quality: it's almost as though it were part of your own body.

Grabbiness

Suppose that minor brain damage destroys your knowledge about your grandmother's eyeglasses. Are you immediately aware that this has happened? No, the loss of the memory of your grandmother's glasses causes no whistle to blow in your mind to warn you. 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 would trigger local motion detectors in your low-level visual system, and an eye saccade would immediately be peremptorily programmed towards the location of the change. Your attentional resources would be mobilized and you would orient towards the change. This "grabbiness" of sensory stimulation, that is, its capacity to cause automatic orienting responses, is a second aspect which distinguishes it from other types of neural activity in the brain. Grabbiness is the fact that sensory stimulation can grab your attention away from what you were previously doing.

Towards an analytic phenomenology

Our claim is that bodiliness and grabbiness are jointly responsible for giving the particular qualitative character to the exercise of sensorimotor skills which people have in mind when they talk of the "feel" of sensation or experience. Because of bodiliness, you are in a way "connected" to sensory stimulation: it changes with your minutest body motion. Because of grabbiness, you somehow can't get away from sensory stimulation: it has the capacity to monopolize your attention and keep you in contact with it. Bodiliness and grabbiness ensure that, unlike thoughts and memories, sensory stimulation has a "clinging" quality. Unlike thoughts and memories, experiences 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 bodiliness and grabbiness may be the reason why there is something it's like to have a sensation.

Note an important point about the concepts of bodiliness and grabbiness: they are physically measurable quantities. A scientist should be able to come in and measure how much bodiliness and how much grabbiness there is in different types of sensory stimulation. The amount of bodiliness is determined by the degree to which sensory input depends on body motions. The amount of grabbiness is determined by the extent to which an organism's orienting responses and processing resources are liable to be grabbed by the input.

If bodiliness and grabbiness are objectively measurable quantities, and if we are right in saying that they determine whether a sensory input has "feel", then we should be able to predict how much "feel" different mental phenomena have.

We have already seen that memory phenomena, like the memory of your grandmother, or thoughts or knowledge, have little or no bodiliness and no grabbiness. They have little feel, therefore. This seems to correspond with what people say about memory, thoughts and knowledge.

We have also seen that experiences, like the experience of seeing the color of your grandmother's eyes, have bodiliness and grabbiness, and should be perceived as possessing "feel".

Now it is interesting to consider whether there exist intermediate cases. If we are right about the relation between bodiliness, grabbiness and feel, then cases of a little bit of bodiliness and grabbiness 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's 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 can't 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 in the perceptual system involved causes drastic sensory changes (high bodiliness). Moreover, if anything in your visual field suddenly turns red, or if suddenly a bell starts ringing near you, you will be immediately alerted (high grabbiness).

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, because the latter have bodiliness and grabbiness to a much higher degree.

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 lots of money; it is knowing you can go on an expensive trip and that you needn't worry about the price of dinner.

Thus being rich has a certain degree of bodiliness, 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 not nearly as direct as in true sensory stimulation like seeing, when 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 bodiliness.

Similarly, being rich also has no grabbiness. 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 purposely check the news.

Further interesting cases concern obsessive thoughts and experiences like worry and anxiety, as well as embarrassment, fear, love, happiness, sadness, loneliness and homesickness. These are more grabby than normal thinking, because you cannot but help thinking about them. Some of these phenomena also have a degree of bodiliness, because there are things you can do to change them: for homesickness you can go home, for happiness you can remove the things that make you happy. Clearly there is "something it's like" to experience these mental phenomena, but the quality they have is not of a sensory nature[2].

It is interesting to consider also the case of proprioception: this is the neural input that signals mechanical displacements of the muscles and joints. Motor commands which give rise to movements thus necessarily produce proprioceptive input, so proprioception has a high degree of bodiliness. On the other hand, proprioception has no grabbiness: body position changes do not peremptorily cause you to attend to them. Thus, as expected from the classification we are putting forward, while we generally know where our limbs are, this position sense does not have a sensory nature.

The vestibular system detects the position and motion of the head, and so vestibular inputs have bodiliness. They also have some grabbiness, since sudden extreme changes in body orientation immediately result in re-adjusting reactions and grab your attention, sometimes provoking dizziness or nausea. In this sense then, the vestibular sense has a limited degree of sensory feel.

The examples given here are simply a first attempt to use the notions of bodiliness and grabbiness to make a classification of phenomenal processes (but see also O'Regan & Noë 2001a). Further work is needed in this direction. Additionally it may be useful to consider the possibility that there are other objective dimensions that may be useful in creating what could be called an "analytic phenomenology' based on objectively measurable quantities like bodiliness and grabbiness. In particular, to deal adequately with pain and emotions we may additionally need the concept of "automaticity", which measures the degree to which a stimulation provokes an automatic behavior on the part of the organism.

Summary

We have seen that, when added to the idea that feels correspond to having mastery of skills, the concepts of bodiliness and grabbiness allow the fundamental difference to be captured between mental phenomena that have no feel, like memory and knowledge, and mental phenomena that have feel, like sensations. Bodiliness and grabbiness furthermore allow us to understand why some intermediate situations, like driving or being rich can also be qualified as possessing a certain, but lesser, degree of "feel". Bodiliness and grabbiness are objectively measurable quantities that determine the extent to which there is something it's like to have a sensation. We suggest that bodiliness and grabbiness therefore allow us to pass the second obstacle to overcoming the explanatory gap. They explain why there is something it is like to feel.

4. Third obstacle: modality and sensory quality

To explain the nature of experience it is necessary not only to explain why there is something it is like to have an experience, one must also explain why it is like this, rather than like that (Hurley and Noë, submitted; Chalmers 1995).

For example hearing involves a different quality as compared to seeing, which has a different quality as compared to tactile sensation. Furthermore, within a given sensory modality there are differences as well: for example, red has a different quality from green. This is the third major obstacle to an account of phenomenal consciousness.

Explaining these differences in neural terms will not work: Neural activation is simply a way of coding information in the brain. As of now, we have no clue how differences in the code could ever give rise to differences in feel.

[2] But note that the grabbiness involved in these phenomena is a "mental" or "psychological" rather than sensory: it is not automatic orienting of sensory systems, but rather uncontrollable, obsessive mental orienting.

But if we consider experiences as skills, then we can immediately see where their differences in phenomenal quality come from: they come from the nature of the different skills you exercise. Just as Porsche driving is a different skill from tractor driving, the difference between hearing and seeing amounts to the fact that among other things, you are seeing if, when you blink, there is a large change in sensory input; you are hearing if nothing happens when you blink, but, there is a left/right difference when you turn your head; the amplitude of the incoming auditory signal varies in a certain lawful way when you approach a sound source, etc. We call these relations between possible actions and resultant sensory effects: sensorimotor contingencies (O'Regan & Noë 2001b).

Sensory substitution

From this follows a curious prediction. We claim that the quality of a sensory modality does not derive from the particular sensory input channel or neural circuitry involved in that modality, but from the laws of sensorimotor contingency that are involved. It should therefore 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, and a number of new devices are being tested with the purpose of substituting different senses: visual-to-tactile (Sampaio, Maris, & Bach-y-Rita 2001); visual-to-auditory (Veraart, Cremieux, & Wanet-Defalque 1992); auditory to visual (e.g. Meijer 1992); auditory-to-tactile (cf. for review Richardson & 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 having sensations in the modality being substituted.

Related phenomena which also support the idea that the feel of a sensory modality is not wired into the neural hardware, but is rather a question of sensorimotor contingencies comes from the amusing experiment of Botvinick & Cohen Botvinick & Cohen (1998), where the "feel" of being touched can be transferred from your own body to a rubber replica lying on the table in front of you (see also other interesting work on the body image in tool use (Yamamoto & Kitazawa 2001, Iriki, Tanaka, & Iwamura 1996). The finding of Roe, Pallas, Hahm, & Sur (1990) according to which embryonically "rewired" ferrets can see with their auditory cortex can also be interpreted within the context of our theory.

Intramodal sensory 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 are necessary for feel, but cannot by themselves generate the feel of hard and soft. On the contrary, we claim the difference between hard and soft comes from the different skills that you implicitly put to work when you touch hard and soft surfaces: the fact that 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 constituted by the things you implicitly know about the way the surface will react to your ongoing 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?

But 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 one aspect of the feel of red

is: knowing the laws that govern the changes in the light reflected off the paper as you turn it (cf. Broackes (1992)).

Another aspect of the skill involved in the feel of red concerns retinal sampling. Retinal sampling of a centrally fixated red patch is done by a densely packed matrix of short, medium and long-wavelength sensitive cones. There is also a yellowish macular pigment which covers the central retina. When an eye movement brings the patch into peripheral vision, the cone matrix that samples the patch is interspersed with rods, the distribution is slightly different, and there is no macular pigment. The resultant change in quality of the incoming sensory stimulation is another aspect of what it is like to be looking at a red patch.

5. Summary: how we have crossed the gap

We have presented arguments showing how three obstacles to understanding experience can be circumvented.

The first obstacle was the fact that experiences appear to be ongoing, occurrent processes inside us. This has led scientists to seek for brain mechanisms which are themselves also ongoing, and whose activity gives rise to feel. 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 "feel".

Our solution is to show that feel is not directly generated by a brain mechanism, but consists in the active exercising of a skill, like driving or bicycle riding. The ongoingness of feel is not "produced" or "secreted" by brain activity, but resides in the active doing, the give-and-take that is involved in exercising a particular skill.

The second barrier to explaining feel is the question of there being something it is like to have the experience, that is, of the experience having a qualitative character. We showed how the concepts of bodiliness and grabbiness allow the fundamental difference to be captured between mental phenomena that have no feel, like memory and knowledge, and mental phenomena that have feel, like experiences or sensations. Bodiliness and grabbiness are objectively measurable quantities that determine the extent to which there is something it's like to have a sensation. Bodiliness and grabbiness allow us to pass the second obstacle to overcoming the explanatory gap. They explain why there is something it's like to feel.

The third obstacle preventing a scientific explanation of the experience was that it was difficult to understand how different types of neural activation could give rise to different types of experience, e.g. experiential differences within and between sensory modalities -- neural activations are just arbitrary codes for information, and information in itself has no feel.

A natural solution comes from the idea that differences in the feel of different sense modalities correspond to the different skills that are involved in exercising each modality. This idea can also be made to work within a given sense modality, explaining the what-it-is-like of red versus green in terms of the different things you do when you are exploring red and green.

How to make a robot feel

With these tools in hand, can we build a robot that feels?

We provide the robot with mastery of the laws that govern the way its actions affect its sensory input. We wire up its sensory receptors so that they provide bodiliness and we ensure grabbiness by arranging things so that sudden sensory changes peremptorily mobilize the robot's processing resources. Will the robot now have "feel"?

No, one more thing is necessary: the robot must have access to the fact that it has mastery of the skills associated with its sensory exploration. That is, it must be able to make use of these sensory skills in its thoughts, planning, judgment and (if it talks) in its language behavior.

Reasoning, thought, judgment and language are aspects of mind where AI and robotics have not yet reached human levels. But there is no a priori, logical argument that prevents this from being possible in the future. This is because there is no barrier in principle that prevents reasoning, thought, judgment, and language from being described in functional terms. They are therefore in principle amenable to the scientific method and can theoretically be implemented by an information-processing device. Of course, because human reasoning is intricately linked with human culture and social interaction, it may not be possible to satisfactorily replicate human

reasoning without also replicating the social and developmental process through which each human goes.

But when we manage to do this, then if we make a robot whose sensory systems possess bodiliness and grabbiness, then the robot will feel. Indeed, it will feel for the same reasons that we do, namely because we have access to our mastery of sensory skills, and because of the bodiliness and grabbiness of sensory inputs.

Acknowledgments

E. M. wishes to thank, for support, the Fund for Scientific Research-Flanders (Belgium) (FWO project G.0175.01). A.N. 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.

References

BACH-Y-RITA, P. (1967). "Sensory plasticity. Applications to a vision substitution system", *Acta Neurologica Scandinavica,* 43, 4, pp. 417-426.

BLOCK, N. (1995). "On a confusion about a function of consciousness", *Behavioral and Brain Sciences,* 18, 2, pp. 227-247.

BOTVINICK, M., & COHEN, J. (1998). "Rubber hands 'feel' touch that eyes see [letter]", *Nature,* 391, 6669, pp. 756.

BROACKES, J. (1992). "The autonomy of colour". In K. Lennon & D. Charles (Eds.), *Reduction, Explanation, and Realism*, Oxford University Press, pp. 421-465.

CHALMERS, D. J. (1996). *The conscious mind: In search of a fundamental theory*, Oxford University Press.

CLARK, A. (2000). "Phenomenal immediacy and the doors of sensation", *Journal of Consciousness Studies,* 7, 4, pp. 21-24.

DENNETT, D. C. (1991). *Consciousness Explained*, Little Brown.

GRUSH, R. (1998). "Skill and spatial content", *The Electronic Journal of Analytic Philosophy* (Vol. 6), http://www.phil.indiana.edu/ejap.

HURLEY, S. L. (1998). *Consciousness in Action*, Harvard.

IRIKI, A., TANAKA, M., & IWAMURA, Y. (1996). "Coding of modified body schema during tool use by macaque postcentral neurones", *Neuroreport,* 7, 14, pp. 2325-2330.

JÄRVILHETO, T. (2001). "Feeling as knowing -- Part 2. Emotion, consciousness, and brain activity.", *Consciousness & Emotion*, in press.

LEVINE, J. (1983). "Materialism and qualia: The explanatory gap.", *Pacific Philosophical Quarterly,* 64, 354-361, pp. .

MEIJER, P. B. L. (1992). "An experimental system for audiory image representations", *IEEE Transactions on Biomedical Engineering,* 39, 2, pp. 112-121.

MYIN, E. (2001), "Color and the duplication assumption", *Synthese*, 129**:** 1. pp. 61-77.

MYIN, E., & O'REGAN, J. K. (2002). "Perceptual consciousness, access to modality and skill theories: A way to naturalize phenomenology?", *Journal of Consciousness Studies,* 9, 1, pp. 27-45.

NAGEL, T. (1974). "What is it like to be a bat?", *Philosophical Review,* 83 , pp. 435-456.

NOË, A. (2002). "On what we see.", *Pacific Philosophical Quarterly*, 83:1

NOË, A. & O'REGAN, J.K. (2000). "Perception, attention and the grand illusion", *Psyche*, 6 (15). URL: http://psyche.cs.monash.edu.au/v6/psyche-6-15-noe.html.

NOË, A. & O'REGAN, J.K. (2002). "On the brain-basis of visual consciousness". In A. Noë & E. Thompson: *Vision and Mind: Selected Readings in the Philosophy of 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**, pp. 461-488.

O'REGAN, J. K., RENSINK, R. A., & CLARK, J. J. (1999). 'Change-blindness as a result of 'mudsplashes'.", *Nature*, **398**, p. 34.

O'REGAN, J. K., & NOË, A. (2001a). "Acting out our sensory experience", *Behavioral and Brain Sciences,* 2, 5, pp. 955-975.

O'REGAN, J. K., & NOË, A. (2001b). "A sensorimotor account of vision and visual consciousness", *Behavioral and Brain Sciences,* 24, 5, pp. 883-917.

O'REGAN, J. K., & NOË, A. (2001c). "What it is like to see: A sensorimotor theory of perceptual experience", *Synthese,* 129, 1, pp. 79-103.

PESSOA, L.,THOMPSON, E, & NOË, A. (1998). "Finding Out About Filling In: A Guide to Perceptual Completion for Visual Science and the Philosophy of Perception", *Behavioral and Brain Sciences*, 21: 6, 1998.

PETTIT, P. (2003a). "Looks red", *Philosophical Issues,* 13, in press.

PETTIT, P. (2003b). "Motion blindness and the knowledge argument.". In P. Ludlow, Y. Nagasawa, & D. Stoljar (Eds.), *The Knowledge Argument* , MIT Press.

RENSINK, R. A., O'REGAN, J. K., & CLARK, J. J. (2000). 'On the failure to detect changes in scenes across brief interruptions. *Visual Cognition*, 7, 1-3, pp. 127-145**.**

RICHARDSON, B. L., & FROST, B. J. (1977). "Sensory substitution and the design of an artificial ear", *J Psychol,* 96, 2d Half, pp. 259-285.

ROE, A. W., PALLAS, S. L., HAHM, J. O., & SUR, M. (1990). "A map of visual space induced in primary auditory cortex", *Science,* 250, 4982, pp. 818-820.

SAMPAIO, E., MARIS, S., & BACH-Y-RITA, P. (2001). "Brain plasticity: 'visual' acuity of blind persons via the tongue", *Brain Research,* 908, 2, pp. 204-207.

SIMONS, D. (2000). Current approaches to change blindness. *Visual Cognition***, 7**, pp. 1-16.

VERAART, C., CREMIEUX, J., & WANET-DEFALQUE, M. C. (1992). "Use of an ultrasonic echolocation prosthesis by early visually deprived cats", *Behavioral Neuroscience,* 106, 1, pp. 203-216.

YAMAMOTO, S., & KITAZAWA, S. (2001). "Sensation at the tips of invisible tools", *Nature Neuroscience,* 4, 10, pp. 979-980.