



Neural Plasticity and Consciousness¹

SUSAN HURLEY

PAIS

University of Warwick

Coventry CV4 7AL

U.K.

E-mail: s.l.hurley@warwick.ac.uk

ALVA NOË

Department of Philosophy

University of California

Santa Cruz

U.S.A.

E-mail: anoe@cats.ucsc.edu

Abstract. We introduce a distinction between *cortical dominance* and *cortical deference*, and apply it to various examples of neural plasticity in which input is rerouted intermodally or intramodally to nonstandard cortical targets. In some cases but not others, cortical activity ‘defers’ to the nonstandard sources of input. We ask why, consider some possible explanations, and propose a dynamic sensorimotor hypothesis. We believe that this distinction is important and worthy of further study, both philosophical and empirical, whether or not our hypothesis turns out to be correct. In particular, the question of how the distinction should be explained is linked to explanatory gap issues for consciousness. Comparative and absolute explanatory gaps should be distinguished: why does neural activity in a particular area of cortex have *this* qualitative expression rather than *that*, and why does it have *any* qualitative expression at all? We use the dominance/deference distinction to address the comparative gaps, both intermodal and intramodal (not the absolute gap). We do so not by inward scrutiny but rather by expanding our gaze to include relations between brain, body and environment.

Key words: consciousness, cortical deference, cortical dominance, dynamic sensorimotor contingencies, explanatory gaps, neural plasticity, perceptual adaptation, phantom limbs, synaesthesia

Introduction

Why does neural activity in a particular area of cortex give rise to experience of red, say, rather than green, or to visual experience rather than auditory? Why, for that matter, does it have any conscious qualitative expression at all? These familiar questions point to the explanatory gap between neural activity and qualities of conscious experience.

In fact, these questions indicate that there are three different types of explanatory gap for consciousness.² There's the *absolute gap*: Why should neural processes be 'accompanied' by any conscious experience at all? And there are two comparative gaps. First, there's the *intermodal comparative gap*: Why does certain neural activity give rise to visual rather than auditory experience, say? Second, there's the *intramodal comparative gap*: Why does certain neural activity give rise to experience as of red, say, rather than experience as of green?

It seems natural to adopt an inward focus in response to such explanatory gap questions, to assume that they must be answered in terms of the intrinsic properties of the neural correlates of consciousness. But there are well-known grounds for skepticism about this strategy of response. Neural properties are qualitatively inscrutable.³ If you were to land in the visual system as a microscopic alien, you couldn't tell, by looking around at the local fireworks, whether experience was happening, or whether, if it was, it was visual experience, or whether, given that it was visual, it was visual experience as of something red. Mueller's theory (1838) of "specific nerve energies" recognized this. On his view, it is not the intrinsic character of the neural activity that makes it visual. Rather, it is the fact that the neural activity is set up by stimulation of the retina, and not, say, the cochlea. But this view still leaves an explanatory gap unbridged: why should differences in the peripheral sources of input, leading to differences in the cortical locations of the neural activity, make for the difference between what it is like to see and what it is like to hear?

We suggest that an inward focus in response to explanatory gap worries can be misleading. To find explanations of the qualitative character of experience, our gaze should be extended outward, to the dynamic relations between brain, body, and world. In this paper we apply this general strategy to the comparative explanatory gaps, both intermodal and intramodal. We set aside the absolute gap, dividing in hopes of conquering. We believe we can make progress by concentrating on the comparative gaps; whether our approach will help to bridge the absolute gap is a further question.

We take our start from consideration of neural plasticity. This phenomenon deserves serious attention from philosophers concerned with explanatory gaps, since it reveals that neural activity in a given area can change not just its function but also its qualitative expression. We introduce a distinction between *cortical dominance* and *cortical deference*,⁴ and apply it to various examples in which input is rerouted either intermodally or intramodally to nonstandard cortical targets. In some cases of rerouting but not others, cortical activity 'defers' to the nonstandard sources of input and takes on the qualitative expression typical of the new source.

This distinction is puzzling, and raises closely related empirical and philosophical issues. What explains why qualitative character defers to nonstandard inputs in some cases but not others? How does explanation of this difference address the comparative explanatory gaps?⁵ After laying out the dominance/deference distinction, with both intermodal and intramodal illustrations, we consider and criticize some possible explanations of it. We then put forward a dynamic sensorimotor account of the distinction. This promising hypothesis has the potential, if correct, to bridge the comparative explanatory gaps.

Whether or not our dynamic sensorimotor proposal turns out to be correct, our main claim here is that the dominance/deference distinction is important and worthy of further study, both empirical and philosophical. Explaining this distinction will help us to understand how qualities of consciousness are related to the rest of the natural world.

The distinction introduced: cortical dominance vs. cortical deference

What happens when areas of cortex receive input from sensory sources that would not normally project to those areas? When an area of cortex is activated by a new source, what is it like for the subject? Is the qualitative character of the subject's experience determined by the area of cortex that is active, or by the source of input to it?

Empirical work on neural plasticity shows that it can go either way. In cases of *cortical dominance*, cortical activation from a new peripheral input source gives rise to experience with a qualitative character normally or previously associated with cortical activity in that area. In such cases, we can say that cortical activity in a particular region *dominates*, that is, it retains its 'natural sign' or normal qualitative expression. In cases of *cortical deference*, in contrast, cortical activity in a given area appears to take its qualitative expression from the character of its nonstandard or new input source. In these cases, the qualitative expression of cortical activity in that area changes, *deferring* to the new input source.

Cortical dominance is illustrated by phantom limb cases in which there appears to be no change in the normal qualitative expression of activation of a given area of cortex, despite change in the source of input. Normally, tactile inputs from arm and face map onto adjacent cortical areas. After amputation of part of an arm, tactile inputs from the face appear to invade deafferented cortex whose normal qualitative expression is a feeling as of an arm being touched. When this area of cortex is activated from its new source, the face, it retains its normal qualitative expression, the touch-to-arm feeling. Thus, stroking the face is felt as the stroking of a phantom arm, as well as a stroking

of the face (Ramachandran and Hirstein 1998; see also Ramachandran and Blakeslee 1998: 28, 38).

Cortical deference is illustrated when congenitally blind persons read Braille. During Braille reading, visual cortex is active. Moreover, stimulation of visual cortex via transcranial magnetic stimulation (TMS), which in normal subjects distorts visual but not tactile perceptions, in such blind subjects distorts tactile perceptions (Sadato et al. 1996, 1998; Buchel et al. 1998; Cohen et al. 1997a). In these subjects, visual cortex seems not only to perform a tactile perceptual function, but also to have tactile qualitative expression. Visual cortex defers qualitatively to its nonstandard tactile inputs. We describe this and further illustrations of cortical deference below.

It may be natural to expect cortical dominance to be the norm. Perceptual scientists may assume that for every type of experience, there is a locus in the brain such that experience of that type supervenes on neural activity at that locus. Activity at such a neural locus may be held to be necessary and/or sufficient to produce experience of the relevant type *however* that activity is produced, whether by normal perceptual processes, by direct stimulation, or by stimulation from a nonstandard source. Such a locus is sometimes called a *neural correlate of consciousness (NCC)*, or a *bridge locus*.⁶ It may seem that if there is such a bridge locus for a given type of experience, then cortex should dominate in the event of rerouting.

There are two points to note about this assumption of dominance as the norm. First, an empirical point: cortex does not in fact always dominate, as we have noted. It is important to recognize that cortical deference occurs as well. Second, a philosophical point: the supervenience of types of experience on neural properties at bridge loci does not entail dominance, but is equally compatible with deference, since neural activity at a given locus can have different neural properties. We argue this point elsewhere, where we claim that our account of the dominance/deference distinction, though compatible with neural supervenience, addresses explanatory gaps in a way that neural supervenience does not. To avoid distraction for purposes of this article, it is helpful to keep in mind that we do not regard cortical deference, or our account of the dominance/deference distinction, as threatening to the neural supervenience of experience.

Why does cortex defer in some cases but dominate in others? How could an explanation of this difference contribute to bridging the comparative explanatory gaps? This article makes an initial approach to these questions. We lay the groundwork in the next two sections by giving a general schematization of the distinction and applying it to various examples.

The distinction schematized

We schematize the distinction in terms of the relations between two mappings, one from peripheral sources of input to cortical target areas, the other from cortical areas to qualitative expressions.

Suppose there are two different peripheral (i.e., proximal) sources of afference or input, A and B. A and B can be specified broadly to give an intermodal comparison, as in visual stimulation (a pattern of light hitting the eye) vs. tactile stimulation (a pattern of touch to the skin). Or they can be specified more narrowly to give an intramodal comparison, such as tactile stimulation to the face vs. tactile stimulation to the arm. (The intermodal/intramodal distinction need not always be sharp.) Suppose also there are two different target areas of cortex to which afference of kinds A and B normally project, respectively. A normally projects to area 1, and B to area 2. These areas are identified anatomically. The normal qualitative expression associated with area 1 is the A-feeling and that associated with area 2 is the B-feeling.⁷

Now we hypothesize a rerouting of afference. First, suppose that area 2, to which B normally projects, is deafferented so that the projection from B to 2 is eliminated. It may be removed through surgical intervention (severance of the afferent channel, or by removal of the limb or organ which is the peripheral source of the afference), or accident, or it may be congenitally missing. Second, suppose that afference from source A now somehow comes to project to area 2. A general question then arises.

Will activation of area 2 by afference with Source A give rise to experiences or feelings of the same type as activation of area 2 by afference with source B? That is, will area 2 retain its normal qualitative expression, the B-feeling, when it is activated by the new, nonstandard source? If so, then *cortex dominates* input source in the determination of 'what it is like' for the subject.

Or, will activation of area 2 by afference with source A give rise instead to sensation of the same type as activation of area 1 by afference with source A? That is, will the qualitative expression of area 2 change from the B-feeling to the A-feeling, reflecting the new source of afference? If so, then the qualitative character of the subject's experience depends in some way on the character of the source of input, rather than just on whether cortical area 1 or 2 is active: *cortex defers*, apparently to the source of input, in the determination of 'what it is like' for the subject (see Figure 1).

Notice that cortical dominance thus schematized involves illusory experience, while cortical deference involves veridical experience. If A and B are stimulations of different sensory organs, such as tactile and visual stimulations, then we can speak of *intermodal* dominance or deference. If A and B




Figure 1.

are stimulations within one modality, such as touch to the face and to the arm, we can speak of *intramodal* dominance or deference.

The distinction at work

We now explain how the distinction applies to the examples mentioned above and to further examples.

Phantom limb cases provide examples of intramodal dominance. Suppose A and B are touches to face and arm, respectively. Area 1 of somatosensory cortex normally receives tactile afference from the face, while the adjacent area 2 of somatosensory cortex normally receives tactile afference from the arm. The normal qualitative expression of area 1 is the feeling of the face being touched, while that of area 2 is the feeling of the arm being touched.

The arm is then lost, so there is no longer any afference from the arm reaching area 2. Instead, area 2 appears to be invaded by afference from the face, which now projects to area 2 as well as the adjacent area 1. Touches to

the face now activate both cortical areas 1 and 2. The question arises: will touches to the face feel only like touches to the face, or will they also feel like touches to the missing ‘phantom’ arm? The answer is: in many cases they feel also like touches to the missing arm (Ramachandran and Hirstein 1998; Ramachandran and Blakeslee 1998: 29, 33, 45).⁸ This suggests that activation of area 2 retains its normal qualitative expression, touch-to-arm feeling. Cortex dominates the new source in determining what it is like for the subject.

Synaesthesia may provide examples of *intermodal* dominance. Color-graphemic synaesthetes experience vivid sensations of color when reading or hearing words, letters, or digits. Particular colors can be associated with particular letters or digits. There is evidence that synaesthetic experience is automatic and truly perceptual, rather than merely a matter of metaphorical association. For example, in a variant on the usual Ishihara tests for color-blindness, synaesthetes who see numerals as colored were presented with a collection of 2s and 5s such that the 5s were mirror images of the 2s. The numerals were arranged so that the 5s made a pattern. Normal subjects could not see the pattern. But it simply ‘popped out’ for the synaesthetes, since they saw the 2s and the 5s in different colors (Ramachandran and Hubbard 2001).

It is unclear whether synaesthesia results from nonstandard neural projections. However, recent imaging work (Nunn et al. 2002) has found that when synaesthetes with colored hearing listen to spoken words, there is clear activation in an area of visual cortex that has been identified as a color-experiencing area (Hadjikhani et al. 1998; cf. Zeki et al. 1998; Tootell and Hadjikhani 1998). Activation of this area under the same conditions is not found in normal subjects. This suggests that language inputs get routed in synaesthetes not just to their normal destinations but also to this area of visual cortex, where they elicit experiences as of color. Cortical activation dominates over the source of stimulation.

To spell out this suggestion about synaesthesia in terms of our schematized distinction: Suppose that A is stimulation of auditory channels generated by the spoken word “Wednesday” and B is a pattern of light entering the eye from a yellow visual stimulus. Input from A activates area 1, whose normal qualitative expression is ‘sounds like ‘Wednesday’’. Input from B activates area 2, whose normal qualitative expression is ‘looks yellow’. Here, there is no disconnection of area 2 from input B. Rather, there are additional nonstandard neural projections: input from A also activates area 2, perhaps via area 1, again eliciting area 2’s normal qualitative expression ‘looks yellow’. Since area 2 retains its normal qualitative expression ‘looks yellow’ even when activated by an input from a different modality, synaesthesia thus interpreted would count as a case of intermodal dominance.

Visual to auditory rerouting in ferrets provides an example of intermodal deference. In newborn ferrets, nerves from the retina that would normally project to visual thalamus and visual cortex have been surgically rerouted to project instead to auditory thalamus and auditory cortex. Auditory areas are thus deprived of their normal auditory inputs, and provided instead with inputs the source of which is visual stimulation. Here, A and B are retinal and auditory inputs, and areas 1 and 2 are visual and auditory cortex, respectively.

As a result of this rerouting, 2-dimensional retinotopic maps (similar to those normally found in visual area V1) form in auditory cortex (Roe et al. 1990, 1992). Some single cells in auditory cortex develop orientation and direction selectivity normally found in cells in visual cortex. Groups of cells in auditory cortex form orientation modules and acquire some visual field properties (Roe et al. 1990, 1992; Pallas and Sur 1993; Sur et al. 1999; Sharma et al. 2000; Elman et al. 1996: 273ff. See also and compare Rauschecker on auditory to visual rewiring in the cat). However, auditory cortex with visual input did not make ectopic connections with visual cortex, but maintained its connections with other auditory cortical areas (Pallas and Sur 1993).

Moreover, the visual information thus carried in rewired auditory cortex can be made to mediate visual behavior. Unilaterally rewired ferrets are trained to respond differently to light stimuli and sound stimuli presented to the non-rewired hemisphere. Then, when light stimuli are presented to the rewired hemisphere, in the portion of the visual field that is ‘seen’ only by this induced projection to auditory cortex, the rewired ferrets respond as though they perceive the stimuli to be visual rather than auditory. The researchers suggest that the functional specification and perceptual modality of a given cortical area can be instructed to a significant extent by its extrinsic inputs; as a result, “. . . The animals ‘see’ with what was their auditory cortex” (Mezernich 2000: 821; cf. Carman et al. 1992; von Melcher et al. 2000). It is argued that the different characteristics of input activity from specific sources (visual vs. auditory) generate not just representational structure specific to that source but also source-specific sensory and perceptual qualities. To put the point in our terminology, this recognition of cortical deference is seen as a striking departure from the traditional and widely held assumption cortical dominance (Mezernich 2000).

However, this work on ferrets, striking though it is, still leaves room for skepticism about whether there has really been an alteration in the qualitative expression, as opposed to representational and functional roles, of a given area of cortex (expressed to us, for example, by Ned Block). Other work, on human subjects, leaves no such room for skeptical manoeuvre.

Early blind readers of Braille provide examples of intermodal deference in human subjects. Brain imaging work on congenitally and early blind subjects reveals activation in visual cortex during tactile tasks, including Braille reading, whereas normal controls show deactivation (measured by PET scans) of visual cortex during tactile tasks (Sadato et al. 1996, 1998; see also Buchel 1998; Buchel et al. 1998).⁹ The researchers suggest that the neuronal mechanisms of cross-modal plasticity include unmasking of normally silent inputs (here, projections from tactile input to visual cortex), stabilization of normally transient connections, and axonal sprouting. Referring back to our schematism, A and B are here peripheral tactile and visual stimulations, and areas 1 and 2 are somatosensory cortex and visual cortex, respectively.

The question arises how the early blind experience such activation of visual cortex: what is its qualitative expression? This question is directly addressed by work that uses TMS to produce transient interference with visual cortex activity during Braille reading. In the early blind subjects, TMS applied to visual cortex produced both errors in Braille reading and reports of tactile illusions (“missing dots,” “extra dots,” and “dots don’t make sense”) (Cohen et al. 1997a).¹⁰ By contrast, in normal subjects TMS to visual cortex had no effect on tactile tasks or sensations, whereas similar stimulation is known to disrupt the visual performance of normal subjects. In our terms, the qualitative expression of area 1 in normal subjects is visual experience, while in these early blind subjects it is tactile experience: the qualitative expression of activation of visual cortex here defers, apparently to the source of input. The researchers view their work as supporting “the idea that perceptions are dynamically determined by the characteristics of the sensory inputs rather than only by the brain region that receives these inputs” (Cohen et al. 1997a: 182; cf. Maudlin 1989: 408).¹¹

How can the distinction be explained?

The fact that both dominance and deference occur needs explanation. Why do some cases of neural rerouting result in dominance while others result in deference? What explains whether qualitative expression goes one way or the other in particular cases? What explains why activity in a certain cortical area is experienced as *like this* rather than *like that*? To take one of our intramodal examples, why is cortical activity in a certain area expressed as a touch-to-arm feeling rather than merely as a touch-to-face feeling? And in the intermodal examples, why is cortical activity in a certain area expressed as tactile rather than visual feeling, or as visual rather than auditory? These questions express comparative explanatory gap issues, but they are open to empirical answer.

An initial hypothesis might be that we find dominance in cases that involve *intramodal* plasticity and deference in *intermodal* cases. For example, the experience of touch to the face as touch to the phantom limb is a case of dominance, and involves only the sense of touch. By contrast, the experience of tactile distortion as a result of TMS applied to visual cortex is a case of deference. This is a case of cross-modal plasticity in which tactile inputs find a nonstandard target in visual cortex.

But the *intermodal deference/intramodal dominance hypothesis* is not satisfactory, for at least two reasons.

First, it does not accommodate all the cases we've considered, even so far. We have seen that some intermodal cases are plausibly regarded as examples of dominance, such as synaesthesia. Moreover, there is evidence that the referral of sensations to phantom limbs may be highly unstable over time, so even here they may be departures from strict dominance. However, the hypothesis could be reformulated, so that intermodal rerouting is necessary but not sufficient for deference.

But second, even if this reformulated hypothesis is correct, we'll still want to know *why*. Indeed, even if it were to turn out, on further reflection, that the intermodal/intramodal distinction does coincide with the deference/dominance distinction, we would still want to know *why*. The intermodal deference/intramodal dominance hypothesis, if correct, would still not be *explanatory*, would be too close to mere redescription of data (though it might provide clues to a more explanatory account).

A quite different suggestion turns on whether damage or rerouting has occurred early or late. The hypothesis is that deference to nonstandard sources of input tends to result from early rerouting of inputs to nonstandard targets, while dominance results from late rerouting. The intuition behind this *early deference/late dominance hypothesis* is that dominance is the norm for a mature brain with established qualitative expressions, while deference results from early rerouting, before the brain has settled into a quality space.

However, consider the fact that patients born without arms may nevertheless have phantom arms (Ramachandran and Blakeslee 1998: 40–42). The early deference/late dominance hypothesis would predict that such patients should not experience the kind of referred sensation experienced by amputees with phantoms, such as in the case of dominance we described above in which a touch to the face is felt also as a touch to the phantom arm. We do not know if referred sensation is found in congenital phantoms as well as in late-acquired phantoms. Again however, if it were, the hypothesis could be reformulated, so that early rerouting is necessary but not sufficient for deference. This reformulation would also be prompted by synaesthesia, an

apparent example of dominance that starts very early in development (from as far back as synaesthetes can remember).¹²

The reformulated prediction would then be that late rerouting should give rise to cortical dominance. Here, the evidence at present seems less than decisive. Sadato et al. (1998) studied 8 blind subjects, 4 of whom were blind at birth and 4 of whom became blind later (on average, at 8.5 years). “. . . [T]he critical point is that the primary visual cortices of both early and later blind groups are activated during Braille reading . . .,” irrespective of the onset of time of blindness (Sadato et al. 1998: 1215; see and compare Buchel 1998; Buchel et al. 1998).

However, this study did not directly address the question of how the later blind subjects experienced this activation by applying TMS to visual cortex, as did Cohen et al. (1997a). It would be interesting to know whether Sadato’s later blind subjects would experience tactile distortions from TMS to visual cortex. Even if visual illusions also resulted, tactile distortions from TMS to visual cortex in these subjects would show cortical deference resulting from relatively late rerouting, contrary to the present prediction.¹³

Cohen et al. (1997b, 1999) studied blind subjects who lost their sight later in life still (14–15 years) than Sadato’s subjects. But in these subjects, activation was not found in visual cortex during Braille reading. Moreover, TMS to visual cortex did not disturb Braille reading. However, since there is no imaging evidence of tactile to visual rerouting in these subjects in the first place, the issue of dominance vs. deference is not raised by Cohen’s late blind studies.

However, if the prediction that late rerouting leads to cortical dominance holds up, further explanation would still be needed. We’d still want to know why: what is it about early but not late rerouting that permits deference? We’d want an explanation at a deeper level, one that sheds light on *why* qualitative expression can come to reflect the source of input in early but not late rerouting.¹⁴

Thus, both the intermodal deference/intramodal dominance and the early deference/late dominance hypotheses are explanatorily shallow, even if they turn out to contain elements of truth. What could give us a deeper level of explanation?

Intermodal plasticity without neural rerouting: adaptation to TVSS

In order to move toward a deeper explanation, let’s consider some examples of plasticity that do not involve neural rerouting. These examples are not captured by the dominance/deference distinction as we have schematized it so far, because they involve external rather than neural rerouting. What is altered

in these cases is the *external* relation between the objects of perception – the distal sources of input – and the perceiver’s sensory organs – the peripheral source of inputs –, rather than the *internal* relation between peripheral sources of input and cortical targets. Even so, these examples illustrate a distinctive feature of cases of cortical deference, namely, changes in qualitative expression as a result of rerouting. We’ll consider both intermodal and intramodal examples of such external rerouting leading to deference. These examples lead us to extend the dominance/deference distinction and motivate a dynamic sensorimotor account of the distinction.

Consider first perceptual adaptation to a tactile-visual substitution system, which involves an intermodal external rerouting. In a well-known series of studies by Bach-y-Rita (1972, 1984, 1996; Sampaio et al. 2001), blind patients are outfitted with a tactile-vision substitution system (TVSS).¹⁵ Vibrators or electrodes on the back, thigh or tongue receive inputs from a camera fitted on the subject’s head or shoulder. Visual input to the camera produces tactile stimulation of the skin, which in turn gives rise to activity in somatosensory cortex, and to tactile experience.

After a period of adaptation (as short as a few minutes), subjects report perceptual experiences that are distinctively non-tactile and quasi-visual. For example, objects are reported to be perceived as arrayed at a distance from the body in space and as standing in perceptible spatial relations such as “in front of” or “partially blocking the view of,” etc. However, Bach-y-Rita emphasizes that the transition to quasi-visual perception depends on the subject’s exercising active control of the camera (1984: 149). If the camera is stationary, or if someone else controls it while the subject passively receives tactile inputs from the camera, subjects report only tactile sensation.

In our schematism, A is here peripheral tactile input (patterns of vibration and pressure on the skin by mechanical fingers) and B is peripheral visual input (patterns of light falling on the eye). Call cortical target area 1 *somatosensory cortex* and cortical target area 2 *visual cortex*. Note that TVSS involves no rerouting from peripheral source of input to cortical areas, either before or after perceptual adaptation. Peripheral tactile input continues to stimulate cortical activity in somatosensory cortex throughout.

What has been rerouted is the external relationship between distal sources of visual input, objects in space, and peripheral sources of tactile input. So, we can add to our schematism a new lowest level of distal sources of inputs, A’ and B’. Let A’ be a distal source of tactile input, an object that is touching the skin, and B’ be a distal source of visual input, namely, an array of objects in space. TVSS effects a new external intermodal mapping from distal sources of visual input to peripheral tactile inputs and on to somatosensory cortex. As a result, the qualitative expression of somatosensory cortex after adaptation

appears to change intermodally, to take on the visual character of normal qualitative expressions of visual cortex. Such a change in qualitative expression involves no neural rerouting and so does not fit our original characterization of cortical deference. That is, in contrast to the cases of deference we considered earlier, here there is no apparent deference to a nonstandard source of peripheral input, for there is no change in peripheral sources of input. However, because there is still a change in qualitative expression of activity in given area of cortex, this case prompts us to extend our characterization of deference to include cases of external rerouting (see Figure 2).

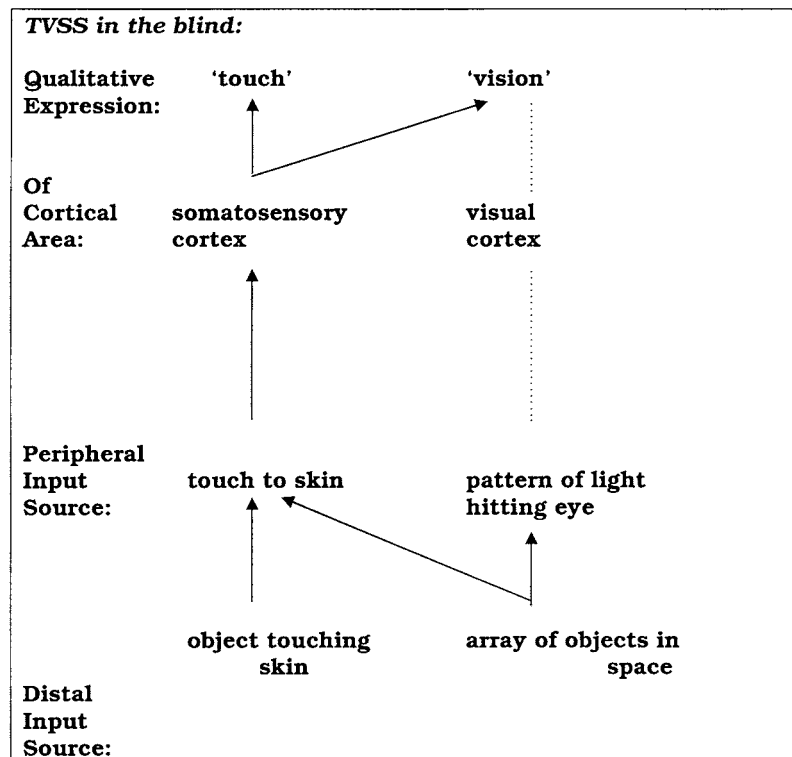


Figure 2.

In structural respects, TVSS-perception after adaptation is more like vision than it is like touch. In vision, and in TVSS-vision, we make perceptual contact with objects arrayed out before us at a distance in space. Neither vision nor TVSS-vision requires immediate physical contact with perceptual objects. In contrast, touch is a perceptual mode that proceeds by bringing a touched object into direct contact with the surfaces of the body. When a person is first outfitted with TVSS, she feels tactile sensations on her back

(say). When adaptation is complete, she may still continue to feel sensations on the back (at least if she were to attend to them), but she now also “feels” the presence of objects in space around her.

Someone might concede that the qualitative expression of somatosensory cortex changes character in adaptation to TVSS, but be skeptical that there is really an intermodal change, from tactile to visual. Of course, there are various ways which the qualitative upshot of adaptation to TVSS isn't like vision. For example, color is missing. Nevertheless, TVSS-perception is like vision in ways that are not captured merely by its spatial significance; after all, audition provides spatial information too, but in a quite different way. TVSS-perception provides spatial information in a distinctively visual way (Noë 2002; O'Regan and Noë 2001a, b, c). In saying this we cannot appeal to the proximal source of stimulation, or to the fact that *visual* areas of the brain are activated. For TVSS-perception is *like vision* despite the fact that eyes and visual cortex are not directly activated.

Consider further similarities of TVSS-perception to vision. TVSS subjects report experiencing characteristically visual effects and illusions, such as parallax, perspective, looming, zooming and depth effects, and the waterfall illusion (Bach-y-Rita et al. 1969; Bach-y-Rita 1995). When you *see* an object, you make perceptual contact only with the facing side or aspect of the object. You only see what is in view. This fact has important dynamic sensorimotor implications. To bring more of an object into view, it often suffices to move in relation to the object. This pattern of dependence of what you perceive on what you do holds for TVSS-perception in much the same way that it holds for vision. In a similar vein, both vision and TVSS-perception are governed by laws of occlusion for which there are no good analogues in touch or audition (even though audition provides spatial information). You see, or TVSS-perceive, objects around you only if they are not block “from view” by other opaque objects.

The distinctively visual character of TVSS-perception stems from the way perceivers can acquire and use practical knowledge of the common laws of sensorimotor contingency that vision and TVSS-perception share (O'Regan and Noë 2001a). For example, as you move around an object, hidden portions of its surface come into tactile-visual view, just as they would if you were seeing them. As you move closer to an object, its apparent tactile-visual size increases, just as it would if you were seeing it. As you turn to the left, objects in “view” swing to the right in your tactile-visual field, just as they would if you were seeing them. What it is like to see is similar to what it is like to perceive by TVSS because seeing and TVSS-perception are similar *ways of exploring the environment*: they are governed by similar sensorimotor constraints, draw on similar sensorimotor skills, and are directed toward

similar visual properties, including perspectively available occlusion properties such as apparent size and shape. These similarities go beyond just providing spatial information; they extend to the distinctively visual way in which dynamic sensorimotor interactions with the environment provide information to the TVSS-perceiver. We claim that the changed qualitative expression induced by TVSS is qualitatively unlike tactile experience and like visual experience in broad and significant respects, which our account illuminates.

Recall that if the camera is stationary, or if someone else controls it, TVSS subjects do not adapt to achieve vision-like experience of objects in space, but continue to report only tactile sensation. Adaptation to TVSS does not occur unless the subject actively controls the camera. This is what a dynamic sensorimotor approach would predict: active movement is required in order for the subject to acquire practical knowledge of the change from sensorimotor contingencies characteristic of touch to those characteristic of vision and the ability to exploit this change skillfully.¹⁶

In TVSS, somatosensory cortex defers: but to what? Our original cases of deference make it clear that qualitative expression cannot be explained just in terms of the area of cortex activated. Extended deference in TVSS shows that it is not enough to appeal in addition to the character of peripheral sources of input. In TVSS, there is an intermodal change in the qualitative expression of somatosensory cortex, yet there is no rerouting of peripheral inputs to somatosensory cortex. Rather, external rerouting between distal and peripheral input sources induces the change in qualitative expression.

However, we do not suggest that this external rerouting in itself explains the change in qualitative expression. Rather, the external rerouting effects a change in the dynamic pattern of sensorimotor contingencies in which peripheral tactile inputs and their cortical targets participate, a change from the pattern characteristic of touch to the pattern characteristic of vision. This change makes distinctively visual know-how and skills newly available to the active subject. In TVSS, somatosensory cortex defers to distinctively visual qualities of distal objects, but this deference is mediated by the perceiver's new sensorimotor skills. It is the perceiver's practical knowledge of distinctively visual patterns of sensorimotor contingency that give TVSS *visual* objects.

A dynamic sensorimotor hypothesis

These insights about TVSS can be generalized. According to our dynamic sensorimotor hypothesis, changes in qualitative expression are to be explained not just in terms of the properties of sensory inputs and of the

brain region that receives them, but in terms of dynamic patterns of interdependence between sensory stimulation and embodied activity. What drives changes in qualitative expression of a given area of cortex, and hence what explains the difference between dominance and deference, is not simply a remapping from the sources of input, whether internal or external, to that area of cortex, but rather higher-order changes, in relations between mappings from various different sources of input to different areas of cortex and from cortex back out to effects on those sources of input, which are in turn fed back to various areas of cortex. Note that there is an essential and inextricable motor element to our account; intramodal, intermodal, and sensorimotor relations are all potentially relevant. Qualitative adaptation depends on a process of sensorimotor integration.

A general account of intermodal differences in qualitative expression is thus suggested by a dynamic sensorimotor approach.¹⁷ Different sensory modalities are governed by different, rich and systematic patterns of dynamic interdependence between sensory stimulation and active movement. For example, to *see* something is to interact with it in a way governed by the dynamic patterns of sensorimotor contingency characteristic of vision, while to *hear* something is to interact with it in a different way, governed by the different patterns of sensorimotor contingency characteristic of audition.¹⁸ Your visual impressions are affected by eye movements and blinks in specific, lawlike ways, while eye movements and blinks are irrelevant to the character of your auditory impressions. Again, as you approach an object, visual field flow expands, while as you withdraw visual field flow contracts. By contrast, as you approach the source of a sound slowly, the amplitude of the auditory stimulus increases, while as you withdraw the amplitude decreases; there are also Doppler effects. Perceivers are familiar with these distinctively different dynamic patterns of sensorimotor contingency and know how to exploit them to explore and negotiate their environments. According to the dynamic sensorimotor view, perceptual experience is a skillful activity, in part constituted by such practical know-how.

If the dynamic sensorimotor view is to provide a bridge across comparative explanatory gaps, it should not be explanatorily shallow. It must be explanatorily satisfying. And it is.¹⁹ When it is brought to our attention that certain sensorimotor contingencies are characteristic of vision, others of hearing, others of touch, there is an ‘aha!’ response. What we have learned does not have the character of a brute fact. Rather, it is intelligible why it is like seeing rather than hearing to perceive in a way governed by the sensorimotor contingencies characteristic of vision rather than those characteristic of audition. It is *not* intuitively tempting to respond: “Yes, that correlation of sensorimotor contingencies with vision may well hold, but *why* does it

hold? Why do *those* patterns of interdependence go with what it is like to *see*, rather than to hear or to touch?" By contrast, if it is brought to our attention that activity in a certain brain area is correlated with vision, we do indeed still want to ask: "But *why* does brain activity *there* go with what it is like to *see*, rather than to hear or touch?" Sensorimotor contingencies are more promising in respect of intermodal comparative explanatory gaps than neural correlates of consciousness.

Readers should keep in mind, when consulting their own intuitions for the 'aha!' response, that we are here addressing comparative rather than absolute explanatory gaps. For present purposes the relevant question is *not*: when we describe certain sensorimotor patterns, do we make it compelling that there is something it is like for the agent at all? It is rather: assuming that there is something it is like for the agent, when we describe the sensorimotor patterns characteristic of, say, vision by contrast with those characteristic of audition, does our description make it compelling that it is like seeing rather than like hearing for the agent?

Because TVSS effects a change from patterns of sensorimotor contingencies characteristic of touch to patterns characteristic of vision, a dynamic sensorimotor view predicts deference and an intermodal change in qualitative expression of somatosensory cortex in this case. But can a dynamic sensorimotor approach be extended to intramodal differences of qualitative expression?

It might be suggested that the dynamic sensorimotor hypothesis is on stronger ground in predicting intermodal deference than intramodal deference, since intermodal rerouting results in larger-scale, more global changes in sensorimotor contingencies than does intramodal rerouting. For example, changes in the sensorimotor contingencies between touching the face and touching the arm, or between looking at something red and looking at something green, are relatively minor, restricted and subtle, compared with those between looking at something and listening to it, or between looking at something and touching it. If this suggestion were to hold up, then perhaps the dynamic sensorimotor hypothesis could provide a deeper level of explanation for the intermodal deference/intramodal dominance hypothesis considered above, which we said was explanatorily shallow.

However, it does not hold up. In principle, the dynamic sensorimotor account applies to intramodal as well as intermodal differences in qualitative character. Perhaps sensorimotor contingencies change in subtler, less global ways within modalities than between modalities than within modalities. But at the qualitative level also, there is a subtler difference between seeing red and seeing green than there is between seeing and touching, or between seeing and hearing. There are nevertheless significant differences

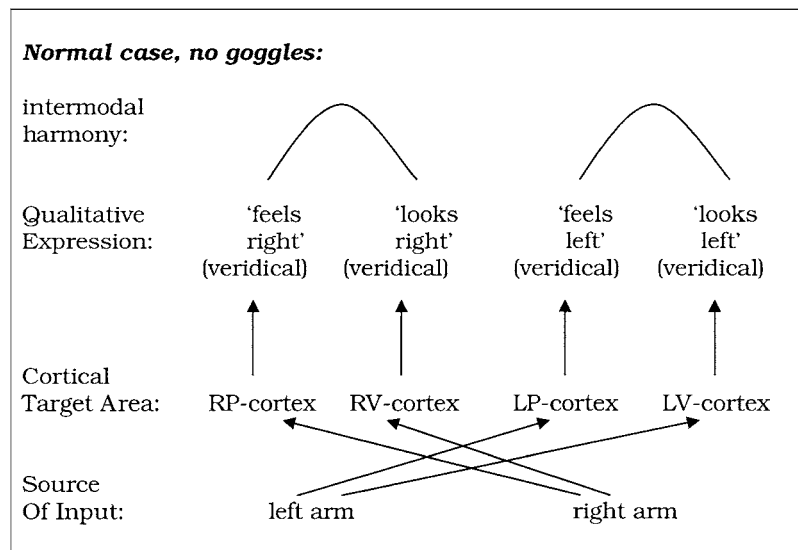


Figure 3.

in sensorimotor contingencies between qualities within one modality.²⁰ So the dynamic sensorimotor hypothesis does not necessarily predict intramodal dominance. And this is all to the good, since we find striking intramodal deference when we consider adaptation to goggles, to which we now turn.

Intramodal plasticity without neural rerouting: reversing goggles

Consider the results of experiments on the long-term effects of wearing left-right reversing goggles (Taylor 1962; see also Stratton 1897; Kohler 1951). Here again there is an external rather than a neural rerouting. And here it results in intramodal deference.

The initial effect of left-right reversing goggles is to produce a left-right reversal in perceptual content. The goggle wearer's right hand looks as if it is on the left and vice versa. The explanation is straightforward. Normally, without goggles, a rightward distal object would produce certain peripheral visual inputs that would in turn project to a certain area of visual cortex, which we'll call right visual cortex, or RV-cortex.²¹ The normal qualitative expression of RV-cortex is 'looks rightward'. Similarly, a leftward distal object would produce different peripheral visual inputs that would project to left visual cortex, or LV-cortex, the normal qualitative expression of which is 'looks leftward'. The goggles effect an external intramodal rerouting: now a rightward distal object produces peripheral visual inputs that project to LV-

cortex, and a leftward distal object produces peripheral visual inputs that project to RV-cortex. So the goggles initially make the right hand look as if it is on the left and vice versa. Notice that again there is no internal neural rerouting of the projections from peripheral inputs to cortical targets; these are unchanged (see Figure 3).

However, putting on the goggles initially disrupts vision dramatically. Movements of eyes and head and body give rise to surprising, unanticipated, confusing sensory effects. For example, when you rotate your head, the world, dizzily, seems to move around you. It used to be that you had to move your head leftward to bring leftward objects more into view, but that no longer works. When you try to move your right hand rightward, it *feels* as if your right hand is moving rightward, but it *looks* as if your left hand is moving leftward. Moving your right hand still activates right proprioceptive cortex, the qualitative expression of which remains 'feels right', even though it looks leftward. So there is a conflict between vision and proprioception when they are co-stimulated by the same right hand movement. The external intramodal visual rerouting effected by the goggles results in an intermodal conflict between vision and proprioception, such that proprioception is veridical and vision is not (see Figure 4).

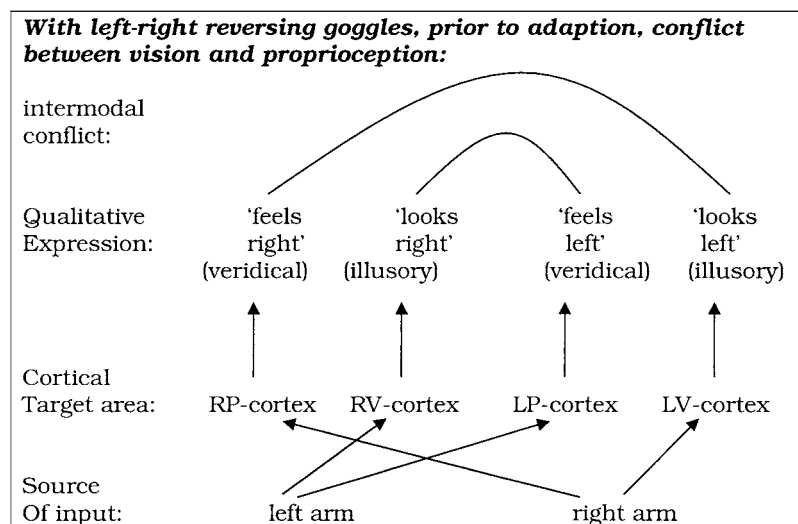


Figure 4.

Thus, movements of your eyes, head, limbs, and whole body quickly demonstrate that the old patterns of sensorimotor contingency no longer apply. While the mapping from distal objects, such as your moving hand, to visual cortex has altered, the mapping from your moving hand to proprio-

ceptive cortex has not altered. Nor has the mapping from motor cortex to your moving hand. Since the mapping to visual cortex has changed, but not the mappings to proprioceptive cortex or from motor cortex, the higher-order relations between these mappings, or sensorimotor contingencies, have changed.

According to our view, for you to experience something as visually leftward is for it to present itself to you as occupying a certain position in a space of sensorimotor possibilities, through which you have the skills to navigate. For example, when something looks to you as though it is on the left, you know how to move your eyes, or turn your head, to bring the thing more into view, you know how to raise your arm and hand in order to block it from view, or to rearrange things that block your view of it so that they no longer do, and so on. When you put this goggles on, you initially lose this know-how.

If perceptual experience depends on the perceiver's sensorimotor know-how, then one way to interfere with perceptual experience would be to alter the dynamic patterns of sensorimotor contingency his know-how exploits. As we've seen, this is just what putting on left-right reversing goggles does. However, as the goggle-wearer learns how to navigate through the new space of sensorimotor contingencies, our view predicts that his perceptual experience should adapt accordingly.

This prediction is born out by Taylor's (1962) description of the experience of a left-right reversing goggle subject, who moved freely through and interacted with his environment over a long period. At first his vision was disrupted as described above. But over time, in addition to seeing a leftward object on the right, he began to see a ghostly version of the same object on the left in its true position. According to Taylor, his visual experience eventually adapted, so that leftward objects came once more to look as if they were on the left and not on the right at all. After adaptation, vision was again veridical, and did not conflict with proprioception (see Figure 5).

Taylor emphasizes that this result was achieved as a result of a rigorous and extended training program during which the goggle wearer engaged in intensive sensorimotor interactions with his environment. Visual adaptation was modular, and reflected specific practice sessions. For example, after lots of bike riding with goggles, buildings on the left came to look leftward, while the writing on signs on those buildings was still left-right reversed. After subsequent practice at reading with goggles, the writing appeared normal as well. Just as intermodal adaptation to TVSS depends on the active control and sensorimotor interactions, so intramodal visual adaptation to left-right reversing goggles depends on the subject's activity.²²

Like TVSS, Taylor's goggle-adaptation result fits our extended characterization of deference: the qualitative expression of activity in RV-cortex

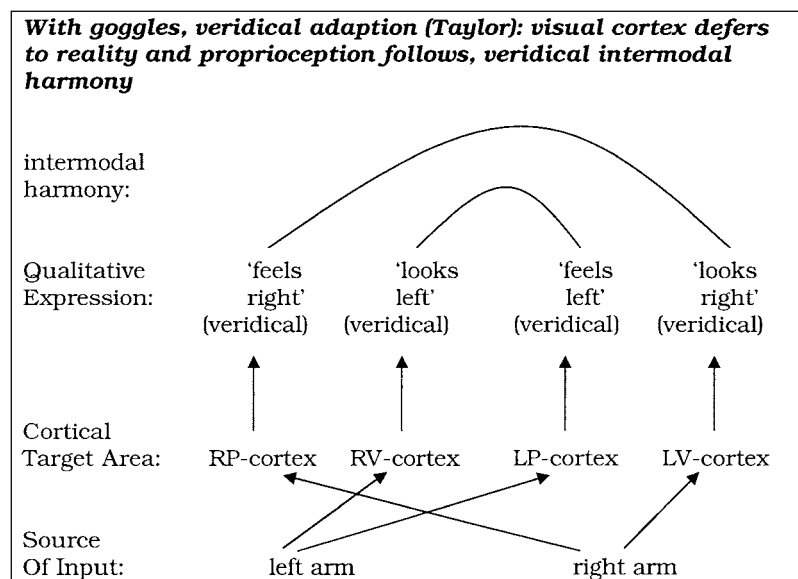


Figure 5.

changes as a result of external rerouting. In this case, by contrast with TVSS, the change is intramodal: from 'looking rightward' when the goggles are first put on to 'looking leftward' after adaptation.

Again, we ask: deference *to what*? The change in qualitative expression cannot be explained in terms of a change in the peripheral source of visual inputs, as there has been no rerouting between peripheral inputs and their cortical targets and so there is no change in peripheral input source. Rather, the change in qualitative expression is induced by external rerouting between distal and peripheral input sources.

Our response is again that the external rerouting effects a change in the dynamic pattern of sensorimotor contingencies in which peripheral inputs and their cortical targets participate. As a result, the subject temporarily loses the know-how on which his visual experience depends. But with practice he regains it, having learned a new way of navigating skillfully through a restructured space of sensorimotor contingencies with which he eventually has become familiar. Leftward objects once more present themselves to him as in terms of a certain characteristic trajectory in this multidimensional space, as related to changing proprioceptive and motor mappings in characteristic ways. Now, when he moves his head leftward, it does bring objects that look (and are) leftward more into view. When he tries to move his right hand rightward, it feels as if his right hand is moving rightward and it also looks as if his right hand is moving rightward. Vision here defers to the true position

of distal objects, but this deference is mediated by the perceiver's reconfigured dynamic patterns of sensorimotor skills in relation to such objects. The perceiver's practical knowledge of the dynamic patterns of sensorimotor contingencies characteristic of seeing leftward objects, within the larger space of sensorimotor contingencies that characterize vision generally, is what makes them look *leftward*.

A challenge

Charles Harris (1965, 1980) questions whether there is genuinely *visual* adaptation to reversing goggles of the kind claimed by Taylor. Some long-term goggle wearers may judge that their visual experience has veridicalized, that leftward objects eventually come to look leftward even while wearing goggles. But Harris suggests that any such judgements are mistaken. He explains two aspects of adaptation to reversing goggles.

He argues, first, that adaptation to the goggles is the result not of vision righting itself, but rather of the adaptation of proprioception to reversed vision. The right hand, which looks as if it is on the left, now comes to *feel* (proprioceptively) as if it is on the left too. Behavioral dispositions are also reversed and so brought into accord with reversed visual experience, thus eliminating the intermodal discord and confusion induced initially by putting on the goggles. Harris says that

. . . so many visual judgments and visually guided behaviors are affected [by processes of adaptation] that one *could* talk about a modification of visual perception, as long as one bears in mind that here too what is actually modified is the interpretation of nonvisual information about positions of body parts" (1980: 113).

In his view, if perceivers interpret the adaptive change in their experience as a change in *visual* experience, their judgments about their own experience are wrong; nonvisual experience rather than visual experience has really adapted.

Secondly, Harris argues that a process of familiarization can explain why even though proprioception, not vision, has really adapted, it nevertheless can seem to subjects that vision has reverted to normal. Just as with practice mirror-writing can come to seem "normal," so, with practice, reversed vision can come to seem normal and familiar. But Harris suggests that visual experience really remains left-right reversed even though it comes to seem normal. Harris' hypothesis is a kind of left-right positional version of an inverted spectrum hypothesis.

If Harris is right, then proprioception adapts intramodally, not vision. We don't get a change in the visual qualitative expression of RV-cortex, as

Taylor's work suggests. Rather, thanks to the power of vision to influence proprioception, we get a change in the proprioceptive qualitative expression of neural activity in what we can call *right-proprioceptive cortex* (RP-cortex). RP-cortex changes its qualitative expression from 'feels rightward' to 'feels leftward'. Note that here again there has been no rerouting between sources of proprioceptive inputs and their targets in proprioceptive cortex. The right hand projects to RP-cortex both before and after adaptation.

So, if Harris is right, we still have a kind of deference. RP-cortex defers; but *to what?* Here again it defers not to a new source of proprioceptive input (there is no new source), but to a new co-stimulation relationship with visual cortex. Actions are intended, felt, and seen, in relation to the environment, giving rise to patterns of neural activity. The goggles alter the way neural patterns are coordinated by action on and in the world. An intramodal external rerouting in vision generates an intermodal conflict between vision and proprioception, which in turn induces an intramodal change in proprioceptive qualitative expression that resolves the conflict, relative to the power of vision.²³ The illusion is thus compounded, as proprioception inherits the illusion the goggles have perpetrated on vision. If Harris is right, deference here cannot be deference to the true position of distal objects. But nevertheless, what would drive the change in qualitative expression of LP-cortex, and could eventually give rise to the secondary illusion that visual experience is re-inverted, would still be a change in the dynamic patterns of sensorimotor contingency among intentional movement, proprioception and vision.

We are skeptical about Harris' view. We will consider his view as a rival to Taylor's in order to spell out our misgivings about it. Consider why proprioception might bend to vision, as Harris argues his experiments show it to do. Why might the new co-stimulation relation induce illusory proprioceptive deference rather than veridical visual deference? Suppose, for the sake of argument, that there are two distinct qualitative possibilities, veridical Taylor-type visual adaptation and illusory Harris-type proprioceptive adaptation with the secondary adaptation such that vision again seems normal. Then it is an empirical question which form of adaptation in fact occurs. Perhaps it is even possible that adaptation occurs one way under certain conditions and the other way under other conditions. Either way, cortical activity somewhere changes its qualitative expression intramodally, driven by the changes in dynamic patterns of sensorimotor contingency. So either way, adaptation to reversing goggles illustrates intramodal deference.

It would be nice to stop there. But we have yet to pin down, in sensorimotor terms, *which* intramodal change, in vision or in proprioception, has occurred in response to changes in dynamic patterns of sensorimotor contingency. What does a dynamic sensorimotor approach predict about which

modality will, as it were, host the adaptive change? Does it predict the kind of deference illustrated by Taylor's view or that illustrated by Harris' view?

As we saw in the previous section, the dynamic sensorimotor approach gives a plausible characterization of intermodal differences. Vision and hearing and touch each have distinctive sensorimotor relations to bodily movement and to each other. But for one animal, the dynamic patterns of sensorimotor contingency governing different simultaneously operative modalities are interpenetrating, superposed on the same organism and its neural system operating in a given environment. We propose that qualities of experience reflect practical knowledge of higher-order sensorimotor patterns. But when a change occurs in the higher-order sensorimotor patterns in which several modalities participate and the animal comes to acquire know-how in relation to the new pattern, how does a dynamic sensorimotor approach attribute the corresponding qualitative adaptation to one modality or another? For example, when a change occurs in the normal relations between visual and proprioceptive inputs and motor outputs, on what basis does a dynamic sensorimotor approach predict that visual or proprioceptive experience will adapt with renormalization?

To answer this question, we need to consider what is needed for sensorimotor know-how to be regained by someone wearing left-right reversing goggles. First, you need, presumably, to be in intermodal sensorimotor harmony, so that vision, hearing, touch, proprioception, and motor intentions are integrated and not in conflict. Intermodal harmony may demand intramodal adaptation, but it may be possible to satisfy this demand in more than one way, as the difference between Taylor's and Harris's views suggests. It can be satisfied so that intermodal harmony is veridical, as in Taylor's view – your right hand looks and feels rightward – or is illusory, as in Harris' view – you right hand both looks and feels leftward.

Second, however, sensorimotor know-how also requires you to be able to negotiate your public environment successfully. Now recall the two parts of Harris' view. The primary aspect of adaptation of proprioception to vision means that your right hand comes at the same time to feel as well as look leftward. But there is also a secondary aspect, in that universal mirror reversal comes at the same time to look normal and familiar. If adaptation only had the primary aspect, your environmental know-how would be compromised by the primary illusion. Here, for example, is how things might go wrong with only the primary adaptation: When your wedding ring is on your left hand, it doesn't look right. So you put your wedding ring on your right hand, which looks and feels leftward. When someone you don't like at all indicates his romantic interest in you, you hold your right hand, stylishly, near to his face, hoping to stop him in his tracks. His response is not at all what you intended.

Can the secondary aspect of adaptation save you from such blunders? While it is not completely clear how Harris understands this secondary aspect, it could be interpreted as a kind of higher-order illusion: because *everything* rightward seems leftward, you come to interpret seeming leftward, wrongly, as seeming rightward. Whether or not this is the correct interpretation of Harris, it is worth spelling out why we are skeptical of this dual-illusion hypothesis.

The dual illusion hypothesis under scrutiny is that while rightward things *really* look and feel leftward to you, they come to *seem* to look and feel rightward. So the true qualities of your experience are no longer self-evident to you. Both the primary and secondary adaptations generate illusions. By canceling out the primary illusion, such a secondary illusion would at least restore your know-how in relation to your environment. It would, for example, save you from the kind of wedding ring blunder just described, as your right hand would come with ‘normalization’ to seem to you to look and feel rightward, even though it still really looks and feels leftward. So your wedding ring would stay where it belongs (see Figure 6).

The primary aspect of adaptation postulated by Harris without such a secondary illusion would not restore dynamic sensorimotor know-how; it would leave you open to blunders. So a dynamic sensorimotor view will not predict the primary illusion without the secondary illusion to cancel it out for practical purposes. But both the doubly illusory adaptation and Taylor’s veridical adaptation will equally restore intermodal and environmental know-how in the context of the new pattern of contingencies imposed by the goggles. Does our account favor one or the other?

Note that the dual-illusion hypothesis conflicts with the claim that qualities of experience are self-evident. In effect, the subject loses knowledge of the qualities of his own experience, in regaining sensorimotor know-how. So any independent arguments for self-evidence would favor Taylor’s hypothesis over the dual-illusion hypothesis.²⁴

Our view leads to the same conclusion by a different route. It predicts that the two postulated illusions would cancel one another out qualitatively as well as practically. The dual-illusion hypothesis relies on the bare idea that *this* is what it is really like qualitatively for the subject, even though it does not seem to him that *this* is what is really like. But our view predicts that if there were no difference in sensorimotor know-how between the dual-illusion view and Taylor’s view, then there would be no qualitative difference either.

It may be natural to presuppose that experience is either visual or not visual intrinsically, or at least in virtue of something other than dynamic sensorimotor know-how. But to do so in effect begs what is an empirical question against our account of intermodal differences in qualities of experience.

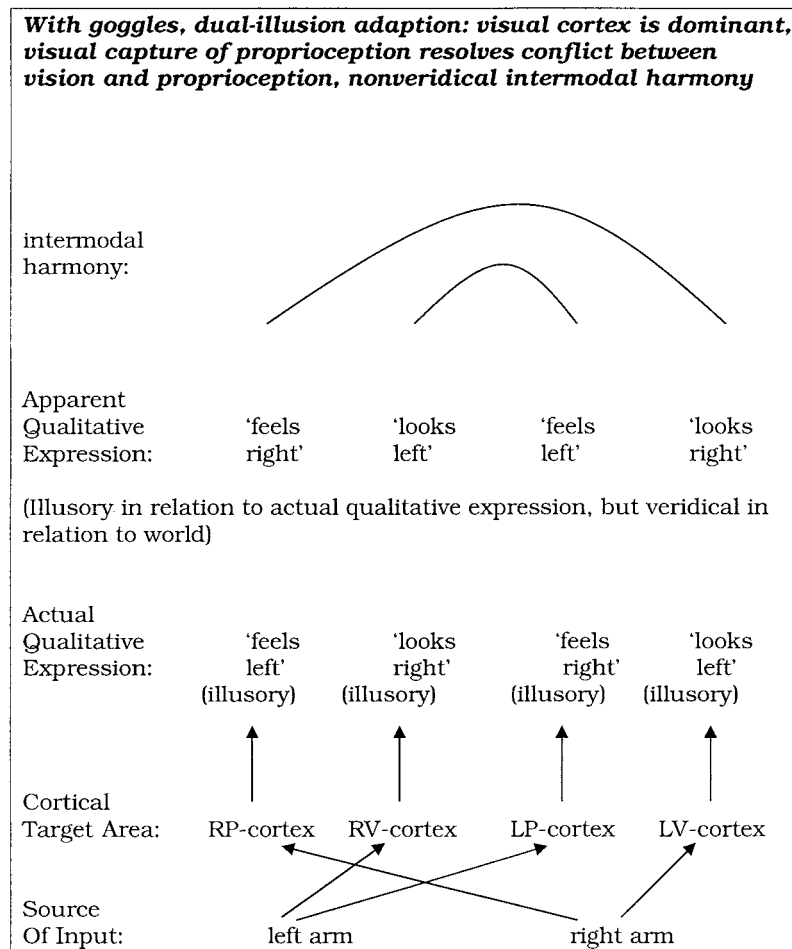


Figure 6.

Why is this presupposition natural? Perhaps because it is natural to assume intramodal cortical dominance: if RV-cortex is still active, then visual experience must still be of something rightward. But certainly Harris would not be entitled to this assumption, by his own lights, since if his view is correct then proprioceptive cortex defers intramodally, even if visual cortex dominates, as explained above. Moreover, we have argued that cortical dominance cannot in general be assumed. It is an empirical question whether cortex dominates or defers, and the evidence from cases of neural plasticity suggests it can do either.

To sum up our response to Harris' position: If adaptation involves illusory proprioception, as Harris suggests, then a canceling secondary adaptation

or illusion is required to restore full know-how. If there is a qualitative difference between adaptation on the dual-illusion interpretation and adaptation on Taylor's veridical interpretation, then the former implies intramodal cortical deference just as much as the latter does. Nevertheless, we suggest that adaptation on the dual-illusion view would not be qualitatively different from adaptation on Taylor's view, since there is no difference in sensorimotor know-how. The two illusions would cancel out qualitatively as well as practically.

Note that this response is not put forward on a priori verificationist grounds, along the lines of denying a priori the qualitative possibility of an inverted spectrum despite complete practical adaptation. Rather, it is an empirical prediction based on a theory supported by evidence in other cases (see also Cole 1990). Indeed, it could prove difficult to verify this particular prediction after adaptation, although the dynamic sensorimotor theory is in general open to empirical assessment, as we have explained.²⁵

The dominance/deference distinction explained

We will now review in general terms how the cases we have considered constrain explanation of the dominance/deference distinction and move us toward a dynamic sensorimotor account.

Deference resulting from neural rerouting, in our original cases, shows that qualitative expression cannot be explained just in terms of the area of cortex activated. Deference resulting from external rather than neural rerouting, as with adaptation to TVSS and reversing goggles, shows that it is not enough to appeal in addition to the character of peripheral sources of input, since there is no rerouting from peripheral inputs to cortex in these cases. Rather, rerouting, whether internal or external, changes the pattern of dynamic sensorimotor contingencies in which given areas of cortex participate. Recall the way changes in qualitative expression in TVSS and goggle adaptation depend on the agents' active control and exploratory movement in their environment. As a result of *rerouting plus a subject's activity*, a global dynamic pattern characteristic of a specific modality, or a more local pattern characteristic of a specific quality within a modality, may be newly established or relocated to new neural pathways. And a given area of cortex may find itself newly participating in and integrated into such dynamic patterns of sensorimotor contingency. Such characteristic patterns govern agents' skillful perceptual activities in their environments, their perceptual know-how. Changes in the neural paths of such characteristic sensorimotor patterns after rerouting can disrupt agents' perceptual know-how and with it the qualitative character of experience, but with practice such know-

how can be reacquired. Deference reflects agents' know-how in relation to dynamic patterns of sensorimotor contingencies that are characteristic of specific modalities or qualities, but which use nonstandard neural paths that include areas of cortex that would normally participate in different sensorimotor patterns. We suggest that the dominance/deference distinction can be explained in terms of such skill-governing sensorimotor patterns in both intermodal and intramodal cases.

Our account predicts deference where two general conditions are met: first, where perceptual experience of the kinds in question normally arises out of distinct patterns of sensorimotor contingency, which are systematically transformed by rerouting, and second, where the agent is able to explore and learn the new operative contingencies and their relations to the old ones. Thus, dominance should result when the second condition is not met because the agent is relatively passive, or where the first condition is not met because particular kinds of rerouted input give rise to 'dangling' cortical activity, not substantially tied into a dynamic pattern of sensorimotor contingencies, of cross-modal and feedback relationships. Such dangling activity latter could nevertheless generate a limiting case of perceptual experience. There may be a spectrum of degrees of richness and complexity in patterns of sensorimotor contingency, with a kind of nearly-null case at one extreme whereby a source of input stimulates only one cortical area and is unaffected by motor activity.²⁶ These two conditions for dominance are connected: inactivity by the subject may leave a new input dangling, until activity ties it in to the network of sensorimotor contingencies through co-stimulation and feedback.

Mappings from different input sources to cortex are affected in different ways and to different degrees by motor activity, and the way and degree to which they are affected can be altered by rerouting. Rerouting of inputs that are affected by motor activity should generally produce changes in dynamic patterns of sensorimotor interdependence. Thus, on this approach, such rerouting should generally produce deference. Deference is the norm, and dominance is the exception that needs to be explained, as a kind of limiting case.

How do these predictions play back onto our original cases? Intermodal deference in the ferret and Braille cases are straightforwardly predicted, since neither condition predicting dominance is met: the agents are active, and the rerouted input does not dangle but is tied into a network of sensorimotor contingencies. It would be interesting to know what would happen if TMS were applied to visual cortex when no Braille reading activity by the subject is taking place: would it still generate a tactile sensation?

By contrast, dominance could be predicted in the phantom referral case on the basis that the rerouted input from face-stroking to cortex that used

to signal touch to arm is dangling as a result of inactivity in relation to this specific input. Why? Because the experimenter, not the subject himself, does the face stroking while the subject is blindfolded, so no feedback is set up. If the subject were to stroke his or her own face, while also watching in a mirror, rather than having the experimenter stroke it, the dynamic sensorimotor hypothesis would predict that qualitative expression would defer. Such self-stroking would make available a new set of relevant contingencies.

This prediction is in line with Ramachandran's well-known mirror-box results (Ramachandran and Blakeslee 1998: 47ff; Ramachandran and Rogers-Ramachandran 1996; Ramachandran et al. 1995). Ramachandran's patient had an immobilized phantom hand, paralyzed in a painful position for ten years since he had lost his limb. Ramachandran used a box in which mirrors had been positioned to create an illusion of the patient's intact hand in the felt position of his phantom hand. The patient was asked to try to move both his hands simultaneously. When he moved his intact hand and saw it move in the mirrors, in the felt position of his phantom hand, he felt his phantom hand move as well. Moreover, the movement in his phantom relieved the pain in his phantom. The mirrors created illusory visual feedback of phantom movement, harking back to the sensorimotor contingencies familiar to the subject from before the loss of the limb. Experience changed accordingly. Ramachandran suggests that when the brain sends out motor commands for movement, and copies of these commands, but gets no corresponding feedback of actual arm movement because the arm is missing, it learns that the arm does not move, that it is paralyzed. The illusory feedback created by the mirror box allows it temporarily to unlearn paralysis.

How then, Ramachandran asks, can we understand the persisting experience of phantom limb movements in congenital phantoms? He suggests that a normal adult has a lifetime of practical familiarity with what in our terminology are dynamic patterns of sensorimotor contingency. These are missing after amputation; the brain's normal 'expectations' of sensorimotor feedback are 'disappointed', so adaptation is needed. As a result the phantom may freeze or even disappear over time. But movement in a congenital phantom may persist indefinitely because the congenital absence of a limb to provide co-stimulation and feedback relationships between various modalities and motor activity means that there are no normal expectations of sensorimotor feedback from such a limb to be disappointed. So no adaptation is called for. In effect, the part of the innate body image corresponding to the phantom is never overwritten but dangles, disconnected from the network of sensorimotor contingencies into which it would normally be integrated (see Ramachandran and Blakeslee 1998: 57).

A challenge for the dynamic sensorimotor view is to explain apparent intermodal dominance in synaesthesia. We do not have an account of this worked out. But here it is interesting to note that while synaesthetes are like normal subjects in displaying cross-modal priming effects for consciously perceived colors, they do not display the *covert* cross-modal priming effects shown by normal subjects. This may provide a clue of use to the dynamic sensorimotor hypothesis, suggesting a degree of ‘dangle’, or disconnection of synaesthetic color perceptions from the usual network of cross modal contingencies.²⁷ To the extent color perceptions resulting from synaesthetic rerouting do dangle, the dynamic sensorimotor approach would predict dominance.

Conclusion

Our main aims in this article have been to draw the dominance/deference distinction, to indicate its relationship to the comparative explanatory gaps, and to raise the question of how the distinction should be explained. In our view, the NCC perspective does not *by itself* have the right kind of resources to explain the dominance/deference distinction. We have proposed a dynamic sensorimotor hypothesis as a way of explaining the dominance/deference distinction and have suggested further experiments.²⁸ If our dynamic sensorimotor hypothesis is successful – and we emphasize that its success or otherwise turns among other things on empirical issues – the approach should by the same token go some way to bridge the comparative explanatory gaps.

We are thus suggesting that an empirical account can in principle scratch an explanatory gap itch – in particular, the comparative gap itches. Understanding the way certain dynamic patterns of sensorimotor contingency are characteristic of particular modalities and qualities provides a kind of insight and intuitive illumination into what they are like, which does not leave us asking at once, “OK, but *why* does that characteristic sensorimotor pattern go with what it is like to *see*?” When the sensorimotor pattern characteristic of vision is explained, we have an “aha!” reaction; we see *through* the dynamic pattern of sensorimotor contingency to what vision in particular is like.²⁹ Again, keep in mind that these claims relate to the comparative explanatory gaps, and not the absolute gap.

By contrast, understanding that activity in certain brain areas is characteristic of particular modalities and qualities leaves us wanting to rephrase the question immediately. Finding neural correlates of consciousness, or NCCs, is a splendid thing to do, but it does not by itself scratch explanatory gap itches. We still want to know, for example, *why* the qualitative expression of activity in a given brain area is like seeing instead of hearing, or like one

quality rather than another. Viewed out of the contexts of the dynamic patterns of sensorimotor contingency in which they function, NCCs are qualitatively inscrutable; we do not see through them to what their qualitative expressions are like.

This contrast is a matter of degree. Characteristic patterns of sensorimotor contingency may be qualitatively translucent rather than transparent. This may stem in part from differences in the degree to which patterns of sensorimotor contingency are accessible at the personal level, a thought we hope to pursue in further work. But even if sensorimotor patterns are not always completely qualitatively transparent, they are a lot more qualitatively scrutable than NCCs.

The itch-scratching/scrutability contrast just drawn between dynamic sensorimotor and NCC approaches is at a psychological level. But meta-questions arise about this contrast itself. If the expanded gaze strategy has the potential to provide more satisfying answers than the inward neural scrutiny strategy, why is the latter so prevalent? What assumptions orient us inwardly this way? This may be related to the question of the logical relationship of our account to the claim that qualitative character supervenes on neural properties. It may seem that our account is incompatible with a claim of neural supervenience, but we deny this. Both are empirical claims, and they are logically compatible. As an empirical matter, both claims may be true. Qualitative character may supervene on neural properties even if the qualitative expression of neural activity is determined, as we have argued, by dynamic sensorimotor considerations. This is so because the rerouting (whether neural or extraneural) that changes patterns of sensorimotor contingency may well also induce changes in neural properties at a given locus. In work in progress we explain the compatibility claim further. But if both claims are true, we hold that our account is explanatory in a way that the neural supervenience claim is not. The neural supervenience claim may be true, but may nevertheless encourage us to look in the wrong place for an explanation of qualitative character. The explanatory promise of our approach, we suggest, is closely connected to the fact that it expands our scrutiny both spatially and temporally, to the dynamic relations between brain, body, and world.

Notes

¹ For helpful discussions and comments, we are grateful to David Chalmers, Patricia Churchland, Paul Churchland, Jonathan Cohen, Alan Cowey, Jeffrey Gray, Mark Greenberg, Rick Grush, Robert Hanna, Colin McGinn, Erik Myin, Philip Pettit, Kim Plunkett, Nicholas Rawlins, Eric Schwitzgebel, Evan Thompson, Michael Tooley, and Larry Weiskrantz. We are also grateful to participants in the 2002 NEH Summer Institute on Consciousness and Intentionality in Santa Cruz, California. 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 Fellowship in the Humanities.

² Compare Chalmers 1996, p. 5, who in effect distinguishes the absolute and comparative gaps, but not the intermodal and intramodal comparative gaps. The latter distinction is useful in setting out our argument. But our argument does not require there to be a sharp distinction between intramedal and intermedal comparisons.

³ "How pulses of water in pipes might give rise to toothaches is indeed entirely incomprehensible, but no less so than how electro-chemical impulses along neurons can." Maudlin 1989: 413.

⁴ Compare related discussions in von Melcher et al. (2000), Merzenich (2000), Pallas (2000). Note that the dominance/deference distinction is not necessarily tied to cortical activity, but can also arise for activity in other parts of the brain; the more general distinction could be labelled 'neural dominance vs. neural deference'. Thanks here to Pat Churchland.

⁵ The distinction raises further philosophical issues as well, about how supervenience claims relate to explanatory gaps. We address these in another article, in progress. In particular, we claim that our dynamic sensorimotor account is compatible with neural supervenience claims, but does more to address explanatory issues.

⁶ Teller and Pugh (1983) introduced the term "bridge locus." They write: "Most visual scientists probably believe that there exists a set of neurons with visual system input, whose activities form the immediate substrate of visual perception. We single out this one particular neural stage, with a name: the bridge locus. The occurrence of a particular activity pattern in these bridge locus neurons is necessary for the occurrence of a particular perceptual state; neural activity elsewhere in the visual system is not necessary. The physical location of these neurons in the brain is of course unknown. However, we feel that most visual scientists would agree that they are certainly not in the retina. For if one could set up conditions for properly stimulating them in the absence of the retina, the correlated perceptual state would presumably occur" (Teller and Pugh 1983: 581). For discussion of this conception, see Noë and Thompson, forthcoming.

⁷ It will sometimes be convenient to refer to the A-feeling as 'A', as in the looks-yellow-feeling and 'looks yellow', in the discussion of synaesthesia in section 3 below.

⁸ Such referral of touches to the face to the phantom arm can occur less than one day after amputation, suggesting that the referral may be due to the unmasking of ordinarily silent inputs rather than the sprouting of new axon terminals (Borsook et al. 1998; Ramachandran and Rogers-Ramachandran 1996: 385). However, the precise topographic mapping from facial stimulation to phantom arm can become extremely disorganized and may be highly unstable over time, suggesting that the relevant alterations in sensory processing may not be hardwired but rather be mediated by an extensive and interconnected neural network with fluctuating synaptic strengths (Knecht et al. 1998; see also Halligan et al. 1994).

⁹ Other imaging work has shown that visual cortex of blind subjects is activated by sound changes, when the task is to detect these changes (Kujala et al. 2000).

¹⁰ Speech was unaffected by TMS, and subjects given a chance to correct their reports after TMS had ended did not do so, suggesting that errors were not due to interference with speech output.

¹¹ The perception of signing by the congenitally deaf provides another example suggestive of intermodal deference. Auditory cortex lights up when some congenitally deaf persons receive inputs from the visual periphery (where manual signing is perceived by those fluent at sign language, who focus on faces; Elman et al. 1996: 299). It would be interesting to discover whether in these cases TMS to auditory cortex produces visual distortions.

¹² Thanks here to Jeffrey Gray.

¹³ See also Kujala et al. (2000) for further evidence on cross-modal cortical reorganization in the mature brain.

¹⁴ Typical patterns of afference in development may be critical for the emergence of typical patterns of cortical organization; thus neural activity at a given locus can have different neural properties, as we have emphasized in explaining why neural supervenience is compatible with deference. The question then becomes whether dominance vs. deference is best explained in terms of differences in local neural properties induced by different patterns of afference in development.

Neural supervenience may well be true, and neural correlates research is important for understanding the mechanisms by which conscious states are achieved. Nevertheless, we claim in work in progress that our dynamic sensorimotor approach provides a deeper explanation, in the sense that it does more to bridge the comparative explanatory gaps. There is more than one explanatory project, of course, and our account does not challenge the value of identifying neural process correlated with consciousness.

¹⁵ See also Arho et al. (1999) on the possibility of an auditory-visual substitution system, as well as Meijer (1992), and also http://ourworld.compuserve.com/homepages/Peter_Meijer/.

¹⁶ This is also argued in Hurley 1998: ch. 9, and O'Regan and Noë (2001, a, b, c; Noë 2002).

¹⁷ Here we draw on Noë and O'Regan (2002), O'Regan and Noë (2001a, b, c) and on Hurley 1998a, especially chapter 9.

¹⁸ A more satisfyingly detailed characterization of the distinctive sensorimotor patterns characteristic of different modalities is given in O'Regan and Noë (2001a, c).

¹⁹ We do not accept the constraint that explanatory gaps concerning consciousness can only be bridged by conceptual truths. See and cf. Levine 1993; van Gulick 1993: 146; Chalmers 1996. For a different kind of attempt to bridge the explanatory gaps in sensorimotor terms, see O'Regan, Myin and Noë, submitted and O'Regan and Noë (2001c).

²⁰ This may hold even for differences in color, which we do not address here. But see O'Regan and Noë 2001a; O'Regan, Myin & Noë submitted; Noë, forthcoming; and see also Pettit in press for an account of specific color qualities very much in harmony with our approach here. The fact that there are asymmetries in color relationships that prevent inversions from working smoothly is relevant here. See Thompson 1995; Myin 2001; van Gulick 1993, 144–145; Hardin 1997.

²¹ Note that this terminology can mislead. RV-cortex is the part of visual cortex that normally subserves the experience as some something being visually on the right. It is not the anatomically right part of the visual cortex.

²² Taylor reports that one of his long-term subjects experienced no aftereffect when removing or reinstating the goggles (while riding a bicycle!). This suggests a very striking variability in qualitative expression. With goggles on, left arm stimulates RV-cortex, and looks leftward. With goggles off, right arm stimulates RV-cortex, and looks rightward. Here, the qualitative expression of activity in RV-cortex would vary between 'looks rightward' and 'looks leftward'. The subject has acquired know-how in relation to both sets of patterns of sensorimotor contingencies, with and without goggles, and switches between them seamlessly. See Hurley 1998a: ch. 9, for further discussion.

²³ Visual capture is a widespread phenomenon. It is exhibited in such phenomena as the ventriloquist effect, in which we hear the voice coming from the mouth of the dummy because we see things that way, and in the McGurk effect, in which visual information about lip-movements trumps acoustic information in determining heard phonemes. An interesting case has recently been reported (Botvinik and Cohen 1998). Subjects were "seated with their left arm resting upon a small table. A standing screen was positioned beside the arm to hide it from the subject's view and a life-sized rubber model of a left hand and arm was placed on

the table directly in front of the subject. The subject sat with eyes fixed on the artificial hand while we used two small paintbrushes to stroke the rubber hand and the subject's hidden hand, synchronising the timing of the brushing as closely as possible" (Botvinik and Cohen 1998: 756). After a short interval subjects have the distinct and unmistakable feeling that they sense the stroking and tapping in the visible rubber hand and not in the hand which is in fact being touched. Further tests show that if you ask subjects with eyes closed now to point to the left hand with the hidden hand, their pointings, after experience of the illusion, are displaced toward the rubber hand.

²⁴ See Hurley 1998: ch. 4; cf. Chalmers 1996: ch. 5, on the paradox of phenomenal judgment. While we cannot pursue the details here, we note that two distinctions are critical in assessing the relevance of Chalmers' arguments to present concerns.

First, self-evidence should be explicitly distinguished from incorrigibility: if experience must be self-evident, there is an entailment from the character of experience to judgment. But the reverse entailment does not follow: judgment about the character need not be incorrigible, even if his experience is self-evident. Hurley 1998a denies incorrigibility, but defends a version of self-evidence. Self-evidence is all we need to rule out Harris' dual-illusion adaptation, not incorrigibility. Chalmers also denies incorrigibility, insisting that our judgements about our own consciousness do not entail the truths of such judgments; a zombie could judge himself to be conscious. His position on self-evidence is more complex; see pp. 97, 205ff.

Second, we should distinguish absolute from comparative issues: self-evidence and incorrigibility may be more or less plausible for absolute or comparative issues, and our concern here is with comparative issues, not zombies.

²⁵ Moreover, it is difficult to see how Harris's view could explain for the visual doubling that Taylor's subject reported before adaptation was complete: "the simultaneous perception of an object and its mirror image, although . . . the chair on the right [its true position] was rather ghost-like" (1962: 202, 206).

²⁶ For example, qualities of taste experiences may seem to be close to this extreme of the spectrum, since taste may appear to participate in few sensorimotor or co-stimulation regularities. But this case may be harder to make out than first appears. Olfactory qualities are well tied in to the network of action and co-stimulation (from hound dogs to Proust), and strongly influence taste experience. And consider the way coffee tastes when you expect it to be tea, or vice versa. We hope to consider how our account would handle taste and smell elsewhere.

²⁷ In particular, colored graphemic synaesthetic perception does not have all the properties of normal color perception. Synaesthetically induced colors give rise to cross-modal priming effects, as do colors perceived by normal subjects. However, synaesthetically induced colors do not give rise to covert cross-modal priming effects, while colors perceived by normal subjects do (Mattingly et al. 2001).

Here is what this means, operationally. Normal subjects asked to name the color of ink in which a word is written show longer reaction times when the word spells the name of an incongruent color: for example, if the word 'red' is printed in blue ink, they will take longer to say 'blue' than they will if the word 'blue' is printed in blue ink (the Stroop effect). Synaesthetes display a synaesthetic version of a Stroop effect: if they are asked to judge the physical color of a letter that induces a different synaesthetic color, their reaction times are slowed (whereas that of normal subjects are not). If they are shown a letter prime that induces experience of a certain synaesthetic color, and then shown a color patch and asked to name its color, they are slower when the induced synaesthetic color of the letter prime is different from the color of the color patch. Synaesthetic and normal colors are thus similar in respect of Stroop priming effects, where subjects are conscious of the colors in question.

However, normal subjects also display covert priming effects: a letter is briefly presented and masked so that subjects are not consciously aware of having seen it, and then asked to name a subsequently presented letter. Normal subjects are slower to name the subsequent letter when the masked prime was a different letter, even though they are not consciously aware of the masked prime. And synaesthetes show this same covert priming effect on letter recognition, displaying unconscious processing of the letter prime. Note that this task does not involve synaesthetic perceptions.

However, if the synaesthetes are asked to name not a letter but the color of a color-patch, preceded by a masked letter prime which they are not aware of having seen, there is no covert priming effect: reaction times are no longer when the masked letter prime would, if unmasked, induce a different synaesthetic color from the color of the presented patch. Thus, synaesthetically induced colors in particular appear to generate distinctively synaesthetic intermodal priming effects only when they are consciously perceived, even though synaesthetes show normal covert intramodal priming effects for letter recognition.

This suggests that synaesthetic color perception lacks some of the links that normal color perception has. Mattingley et al. (2001) suggests that synaesthetic interactions occur after initial processes of recognition in the inducing modality are complete. See also Rich and Mattingley (2002).

²⁸ Namely, we have suggested investigating whether applying TMS to auditory cortex of the congenitally deaf produces visual (or other sensory) distortions; whether Sadato's later-blind subject would experience tactile distortions when TMS is applied to visual cortex; whether TMS produces tactile illusions when no Braille reading is in progress when applied to visual cortex in the congenitally and early blind; whether referred sensations are found in congenital phantoms; and whether referred sensations are found in late-acquired phantoms who stroke themselves while watching in a mirror.

²⁹ Again, more detailed sensorimotor characterizations of the modalities are given in O'Regan and Noë (2001a, c). We recommend attention to these for the most robust 'aha!' reaction.

References

- Arho, P., Capelle, C., Wanet-Defalque, M.C., Catalan-Ahumada, M. and Veraart, C.: 1999, 'Auditory Coding of Visual Patterns for the Blind', *Perception* **28**(8), 1013–1029.
- Bach-y-Rita, P.: 1972, *Brain Mechanisms in Sensory Substitution*, Academic Press, New York.
- Bach-y-Rita, P.: 1995, *Nonsynaptic Diffusion of Neurotransmission and Late Brain Reorganization*, Demos-Vermande, New York.
- Bach-y-Rita, P.: 1984, 'The Relationship between Motor Processes and Cognition in Tactile Visual Substitution', in W. Prinz and A.F. Sanders (eds), *Cognition and Motor Processes*, Springer-Verlag, Berlin, pp. 149–160.
- Bach-y-Rita, P.: 1996, 'Sustitucion sensorielle et qualia', in J. Proust (ed.), *Perception et Intermodalité*, Presses Universitaires de France, Paris). Reprinted in English translation in Noë, A. and Thompson, E.: 2002, *Vision and Mind: Selected Readings in the Philosophy of Perception*, MIT Press, Cambridge, MA.
- Bach-y-Rita, P., Collins, C.C., Saunders, F., White, B., and Scadden, L. 1969, *Nature* **221**, 963–964.
- Borsook, D., Becerra, L., Fishman, S., Edwards, A., Jennings, C.L., Stonjanovic, M., Papinicholas, L., Ramachandran, V.S., Gonzalez, R.G. and Breiter, H.: 1998, 'Acute Plasticity in the Human Somatosensory Cortex Following Amputation', *Neuroreport* **9**(6), 1013–1017.

- Botvinik, M., and Cohen, J.: 1998, 'Rubber Hands 'Feel' Touch that Eyes See', *Nature* **391**(6669), 756.
- Buchel, C.: 1998, 'Functional Neuroimaging Studies of Braille Reading: Cross-modal Reorganization and Its Implications', *Brain* **121**, 1193–1194.
- Buchel, C., Price, C., Frackowiak, R.S.J. and Friston, K.: 1998, 'Different Activation Patterns in the Visual Cortex of Late and Congenitally Blind Subjects', *Brain* **121**, 409–419.
- Carman, L.S., Pallas, S.L. and Sur, M.: 1992, 'Visual Inputs Routed to Auditory Pathway in Ferrets', *Society for Neuroscience Abstracts* **18**, 593.
- Chalmers, D.: 1996, *The Conscious Mind*, Oxford University Press, New York.
- Cohen, L.G., Celnik, P., Pascual-Leone, A., Corwell, B., Faiz, L., Dambrosia, J., Honda, M., Sadato, N., Gerloff, C., Catala, M.D. and Hallett, M.: 1997a, 'Functional Relevance of Cross-modal Plasticity in Blind Humans', *Nature* **389**, 180–183.
- Cohen L.G., Weeks, R., Celnik, P. and Hallett, M.: 1997b, 'Role of the Occipital Cortex During Braille Reading (Cross-Modal Plasticity) in Subjects with Blindness Acquired Late in Life', *Society for Neuroscience Abstracts* (92.1).
- Cohen, L.G., Weeks, R.A., Sadato, N., Celnik, P., Ishii, K. and Hallett, M.: 1999, 'Period of Susceptibility of Cross-Modal Plasticity in the Blind', *Annals of Neurology* **45**(4), 451–460.
- Cole, D.: 1990, 'Functionalism and Inverted Spectra', *Synthese* **82**, 207–222.
- Elman, J., Bates, E., Johnson, M., Karmiloff-Smith, A., Parisi, D. and Plunkett, K.: 1996, *Rethinking Innateness*, MIT Press, Cambridge.
- Hadjikhani, M., Liu, A.K., Dale, A., Cavanagh, P. and Tootell, R.B.H.: 1998, 'Retinotopy and Color Sensitivity in Human Visual Cortical Area V8', *Nature Neuroscience*, 235–241.
- Halligan, P.W., Marshall, J.C. and Wade, D.T.: 1994, 'Sensory Disorganization and Perceptual Plasticity after Limb Amputation: A Follow-up Study', *Neuroreport* **5**(11), 1341–1345.
- Hardin, C.L.: 1997, 'Reinverting the Spectrum', in A. Byrne and D.R. Hilbert (eds), *Readings on Color: Volume 1, The Philosophy of Color*, MIT Press, Cambridge, pp. 289–301.
- Harris, C.: 1965, 'Perceptual Adaptation to Inverted, Reversed, and Displaced Vision', *Psychological Review* **72**(6), 419–444.
- Harris, C.: 1980, 'Insight or Out of Sight?: Two Examples of Perceptual Plasticity in the Human Adult', in Charles S. Harris (ed.), *Visual Coding and Adaptability*, Erlbaum, Hillsdale, NJ, pp. 95–149.
- Hurley, S.L.: 1998a, *Consciousness in Action*, Harvard University Press, Cambridge.
- Hurley, S.L.: 1998b, 'Vehicles, Contents, Conceptual Structure, and Externalism', *Analysis* **58**(1), 1–6.
- Knecht, S., Henningsen, H., Hohling, C., Elbert, T., Flor, H., Pantev, C. and Taub, E.: 1998, 'Plasticity of Plasticity? Changes in the Pattern of Perceptual Correlates of Reorganization after Amputation', *Brain* **121**(4), 717–724.
- Kohler, I.: 1951, 'Über Aufbau und Wandlungen der Wahrnehmungswelt', *Österreichische Akademie der Wissenschaften. Sitzungsberichte, philosophisch-historische Klasse* **227**, 1–118.
- Kohler, I.: 1964, *The Formation and Transformation of the Perceptual World*. Published as a monograph in *Psychological Issues* vol. 3 (monograph 12), New York International University Press. [This is a translation of Kohler 1951.]
- Kujala, T., Alho, K. and Naatenen, R.: 2000, 'Cross-modal Reorganization of Human Cortical Functions', *Trends in Neuroscience* **23**(3), 115–120.
- Levine, J.: 1993, 'On Leaving out What It's Like', in Martin Davies and Glyn Humphreys (eds), *Consciousness*, Blackwell, Cambridge, pp. 121–136.

- Marks, L.E.: 1987, 'On Cross-modal Similarity: Auditory-visual Interactions in Speeded Discriminations', *Journal of Experimental Psychology: Human Perception and Performance* **13**(3), 384–394.
- Mattingley, J.B., Rich, A.N., Yelland, G. and Bradshaw, J.L.: 2001, 'Unconscious Priming Eliminates Automatic Binding of Colour and Alphanumeric form in Synaesthesia', *Nature* **410**, 580–582.
- Maudlin, T.: 1989, 'Computation and Consciousness', *Journal of Philosophy*, 407–432.
- Meijer, P.B.L.: 1992, 'An Experimental System for Auditory Image Representations', *IEEE Transactions on Biomedical Engineering* **39**(2), 112–121. Reprinted in the 1993 *IMIA Yearbook of Medical Informatics*, pp. 291–300.
- Merzenich, M.: 2000, 'Seeing in the Sound Zone', *Nature* **404**, 820–821.
- Mueller, J.: 1838, *Handbuch der Physiologie des Menschen* (Vol. V), Hölscher, Coblenz.
- Myin, E.: 2001, 'Color and the Duplication Assumption', *Synthese* **129**, 61–77.
- Noë, A.: 2002, 'On What We See', *Pacific Philosophical Quarterly* **83**(1), 57–80.
- Noë, A.: forthcoming, *Action in Perception*, The MIT Press, Cambridge, MA.
- Noë, A. and O'Regan, 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. and O'Regan, K.: 2002, 'On the Brain-basis of Visual Consciousness: A Sensorimotor Account', in A. Noë and E. Thompson (eds), *Vision and Mind: Selected Readings in the Philosophy of Perception*, MIT Press, Cambridge, MA.
- Noë, A. and Thompson, E.: forthcoming, 'Are there Neural Correlates of Consciousness?' *Journal of Consciousness Studies*.
- Nunn, J.A., Gregory, L.J., Brammer, M., Williams, S.C.R., Parslow, D.M., Morgan, M.J., Morris, R., Bullmore, E., Baron-Cohen, S. and Gray, J.A.: 2002, 'Functional Magnetic Resonance Imaging of Synesthesia: Activation of Color Vision Area V4/V8 by Spoken Words', *Nature Neuroscience* **5**(4), 371–374.
- O'Regan, K. and Noë, A.: 2001a, 'A Sensorimotor Account of Vision and Visual Consciousness', *Behavioral and Brain Sciences* **24**(5), 883–917.
- O'Regan, K. and Noë, A.: 2001b, 'Acting out Our Sensory Experience', *Behavioral and Brain Sciences* **24**(5), 955–975.
- O'Regan, K. and Noë, A.: 2001c, 'What It Is Like to See: A Sensorimotor Theory of Perceptual Experience', *Synthese* **129**(1): 79–103.
- O'Regan, J.K., Myin, E. and Noë, A.: submitted, 'What It's Like to Feel: An Explanation in Terms of Bodiliness and Grabbiness', *Trends in Cognitive Sciences*.
- Pallas, S.L.: 2001, 'Intrinsic and Extrinsic Factors that Shape Neocortical Specification', *Trends in Neuroscience* **24**(7), 411–423.
- Pallas, S.L. and Sur, M.: 1993, 'Visual Projections Induced into the Auditory Pathway of Ferrets: II. Corticocortical Connections of Primary Auditory Cortex', *Journal of Comparative Neurology* **337**(2), 317–333.
- Pettit, P.: in press, 'Looks Red', *Philosophical Issues* **13**.
- Ramachandran, V.S., Rogers-Ramachandran, D. and Cobb, S.: 1995, 'Touching the Phantom Limb', *Nature* **377**, 490–491.
- Ramachandran, V.S. and Rogers-Ramachandran, D.: 1996, 'Synaesthesia in Phantom Limbs Induced with Mirrors', *Proceedings of the Royal Society* **263**, 377–386.
- Ramachandran, V.S. and Blakeslee, S.: 1998, *Phantoms in the Brain*, Fourth Estate, London.
- Ramachandran, V.S. and Hirstein, W.: 1998, 'The Perception of Phantom Limbs', *Brain* **121**(9), 1603–1630.
- Ramachandran, V.S. and Hubbard, E.M.: 2001, 'Synaesthesia – A Window into Perception, Thought and Language', *Journal of Consciousness Studies* **8**(12), 3–34.

- Rauschecker, J.P.: 1996, 'Substitution of Visual by Auditory Inputs in the Cat's Anterior Ectosylvian Cortex', *Prog. Brain Research* **112**, 313–323.
- Rich, A. N. and Mattingley, J.: 2002, 'Anomalous Perception in Synaesthesia', *Nature Reviews: Neuroscience* **3**, 43–52.
- Roe, A.W., Pallas, S.L., Hahm, J. and Sur, M.: 1990, 'A Map of Visual Space Induced in Primary Auditory Cortex', *Science* **250**(4982), 818–820.
- Roe, A.W., Pallas, S.L., Kwon, Y.H. and Sur, M.: 1992, 'Visual Projections Routed to the Auditory Pathway in Ferrets', *Journal of Neuroscience* **12**(9), 3651–3664.
- Sadato, N., Pascual-Leone, A., Grafman, J., Ibanez, V., Deiber, M.P., Dold, G. and Hallett, M.: 1996, 'Activation of the Primary Visual Cortex by Braille Reading in Blind Subjects', *Nature* **380**(6574), 526–528.
- Sadato, N., Pascual-Leone, A., Grafman, J., Deiber, M.P., Ibanez, V. and Hallett, M.: 1998, 'Neural Networks for Braille Reading by the Blind', *Brain* **121**(7), 1213–1229.
- 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.
- Sharma, J., Angelucci, A. and Sur, M.: 2000, 'Induction of Visual Orientation Modules in Auditory Cortex', *Nature* **404**(6780), 841–847.
- Stratton, G.M.: 1897, 'Vision Without Inversion of the Retinal Image', *Psychological Review* **4**, 341–360, 463–481.
- Sur, M., Angelucci, A. and Sharma, J.: 1999, 'Rewiring Cortex: The Role of Patterned Activity in Development and Plasticity of Neocortical Circuits', *Journal of Neurobiology* **41**(1), 33–43.
- Taylor, J.: 1962, *The Behavioral Basis of Perception*, Yale University Press, New Haven and London.
- Teller, D.Y. and Pugh, E.N. Jr.: 1983, 'Linking Propositions in Color Vision', in J.D. Mollon and L.T. Sharpe (eds), *Colour Vision*, Academic Press, London, pp. 577–589.
- Tootell, R.B.H. and Hadjikhani, N.: 1998, 'Has a New Color Area Been Discovered? Reply to Zeki', *Nature Neuroscience* **1**(5), 335–336.
- Thompson, E.: 1995, *Color Vision*, Routledge, London.
- van Gulick, R.: 1993, 'Understanding the Phenomenal Mind', in M. Davies and G. Humphreys (eds), *Consciousness*, Blackwell, Oxford, pp. 137–154.
- von Melcher, L., Pallas, S.L. and Sur, M.: 2000, 'Visual Behaviour Mediated by Retinal Projections Directed to the Auditory Pathway', *Nature* **404**(6780), 871–876.
- Zeki, S., McKeefry, D.J., Bartels, A. and Frackowiak, R.S.J.: 1998, 'Has a New Color Area Been Discovered?' *Nature Neuroscience* **1**(5), 335–336.