Cognitive science and epistemic openness

MICHAEL L. ANDERSON
Institute for Advanced Computer Studies, University of Maryland, College Park, USA (E-mail: anderson@cs.umd.edu)

Abstract. Recent findings in cognitive science suggest that the epistemic subject is more complex and epistemically porous than is generally pictured. Human knowers are open to the world via multiple channels, each operating for particular purposes and according to its own logic. These findings need to be understood and addressed by the philosophical community. The current essay argues that one consequence of the new findings is to invalidate certain arguments for epistemic anti-realism.

Keywords. knowledge, vision, action, proprioception, realism

Introduction

Here is a common view about how, in general, our epistemic cognition is structured: our senses deliver impulses of some kind that are arranged, or synthesized, or are the basis for inference, or are otherwise interpreted by large-scale cognitive structures to produce conceptually ordered beliefs. Or, as Richard Rorty has put the matter: "Since Kant, we find it almost impossible not to think of the mind as divided into active and passive faculties, the former using concepts to 'interpret' what 'the world' imposes on the latter." (Rorty 1979 p.3) Whatever the merits of this view as history, it certainly seems an accurate characterization of the contemporary orthodoxy. Thus, Steven Pinker writes,

> When [organisms] apprehend the world by sight, they have to use the splash of light reflected off its objects, projected as a two dimensional kaleidoscope of throbbing, heaving streaks on each retina. The brain somehow analyzes the moving collages and arrives at an impressively accurate sense of the objects out there that give rise to them. . . . The selective advantage is obvious: animals that know where the food, the predators, and the cliffs are can put food in their stomachs, keep themselves out of the stomachs of others, and stay on the right side of the clifftop. (Pinker 1997 p.212-3)

Likewise, David Milner and Melvyn Goodale suggest: "It is commonly assumed that vision in humans has a single function: to provide a unified internal representation of the external world which can then serve as the perceptual foundation for visually based thought and action" (Milner and Goodale 1995 p.5), and Alva Nöe and Evan Thompson introduce a recent collection of essays on vision by summarizing: "*The* central puzzle for traditional visual science has been to explain how the brain bridges the gap between what is given to the visual system and what is actually experienced by the perceiver." (Nöe and Thompson 2002b p. 4-5)

More generally, we have Quine's insistence that "whatever evidence there is for science is sensory evidence . . . the stimulation of sensory receptors is all the evidence anyone has to go on, ultimately, in arriving at his picture of the world" (Quine 1969 p. 75), a sentiment clearly echoed in Richard Gregory's more recent formulation:

> The key notion of cognitive psychology since the collapse of behaviourism is that we build brain descriptions of the world of objects, which give perception and intelligent behaviour. Perceptions are not regarded as internal pictures or sounds, but rather as language-like descriptions coded, we suppose, by brain structures of what may be out there. We carry in our heads predictive hypotheses of the external world of objects and of ourselves. . . . From patterns of stimulation at the eyes and ears and other organs of senses, including touch, we project sensations of consciousness into the external world. (Gregory 1998 p. 1693)

These are but a few such statements among the numerous available, but the picture they paint of our epistemic predicament could hardly be clearer, or more in tune with the neo-Kantian orthodoxy: a passively received sensory stimulation (apparently a single sort of thing, a "given"), insufficient by itself to qualify as experience (and misbehaving besides, "heaving" and "throbbing" about), is "somehow" analyzed by the brain, which overcomes the insufficiency of the inputs, builds a predictive hypothesis, or otherwise

"bridges the gap", and thereby transforms sensory stimulation into an experienced description of the world by which we hope to fruitfully act and survive.[1]

Thus stated, this orthodox view of the nature of perception—which I shall call the received view—clearly involves a number of distinct, implicit assumptions. It is worth trying to explicate these:

1. The brain or mind provides the active element of perception, while the senses are the passive recipients of "stimulations".

2. All perception works by analogy with vision. From the epistemic standpoint, "sensory stimulation" describes a single class of event, regardless of the sensory modality involved. The raw characteristics of the signal and low-level details of the processing may be different for different sense modalities, but the overall logic of the processing, and therefore the epistemic import of the sensory stimulation, is identical in each case.

3. In order to become a perceptual experience, or to sustain a cognitive attitude like belief, "sensory stimulation" must be synthesized by inferential, conceptual, or other cognitive structures.

4. Conceptually synthesized sensory stimulation is our only epistemically relevant mode of access to the world.

5. Perception has a single, unified function: to build a (conceptually informed) representation or description of the world.

6. This single, perceptually generated description of the world is the basis for all perceptually guided thought and action.

Relying heavily on some recent work in cognitive science, as well as on phenomenological considerations, I shall argue: (a) that although the received view may accurately characterize one epistemically relevant mode of access to the world, human beings in fact have many; (b) that each operates according to its own logic and for its particular purposes; and (c) that each thereby produces or contributes to some element of our overall (but not necessarily centrally represented) set of beliefs about the world. By a

mode of access, I shall mean a certain, discriminable pathway between the world and an agent's beliefs, representations, and other intentional states and information-carrying inner structures, which opens these to the world's influence. A "logic of operation", in this context, means a high-level functional description of the processing of the world's causal impact, within or by such a pathway, such as would be involved in an explanation of the effect of a given event on the content of an information-carrying structure. The term "representation" should be read in light of the action-grounded theory of representation given in (Anderson 2005; Rosenberg and Anderson 2004; forthcoming), and "beliefs, representations, and other intentional states and information-carrying inner structures" should be broadly construed to include symbolically rendered and conceptually informed language-like mental entities, non-conceptual cognitive contents and situation-relative bodily dispositions, explicit "know-that" and implicit "know-how". I cannot claim to identify here *all* our epistemically relevant modes of access to the world, nor even to fully characterize those I will identify; I can hope only to identify enough, and characterize these sufficiently, to motivate the criticisms of (and alternative to) the received view that I offer.

As this essay relies largely on the same work by Milner and Goodale cited by (Clark 2001), my arguments bear some similarity to his. We are both attracted by the idea that a single perceptual system (in this case vision) can have multiple functions, and thus we both insist on the inadequacy on thesis (5), above. However, whereas Clark's interest is to question (5) in the service of undermining (6), I intend instead to use Milner and Goodale's critique of (5) to focus attention directly on thesis (4), and indirectly on (2) and

(3). For, as I have already indicated, I believe that human beings are more epistemically porous than the received view allows.

Of course, in gathering these various elements—gleaned from different theories and theorists—and making of them one view, one runs the risk of assembling a theory that no one in fact holds. Although it is not clear to me that it is useless to show something false that no one will admit to believing, there is in this case little reason for such concerns. For, to put it bluntly, *every* element of the received view is questionable. By way of illustration, consider just three examples: first of all, *contra* thesis (1), it has long been a staple of phenomenology (Merleau-Ponty 1962), and has more recently been the focus of the enactive and embodied schools of cognitive science (Anderson 2003a; Clark 1997; Varela, Thompson and Rosch 1991) that the sense organs can be active tools of exploration: the hand reaches out to grasp objects and brush surfaces, and likewise the eye does not wait for stimulation, but moves selectively and proactively about the scene. Central to the process of sensing is the participation of the sense organs in an internally guided, need-driven search for further stimulation.[2] This perspective further suggests, *contra* theses (2) and (3), that part of the meaning of this actively sought sensory stimulation can be found, not in the conceptual structures by which these stimulations might be interpreted or synthesized, but in the nature and purpose of the activities the sense organs participate in to sense the world. Sensing generally occurs in a context of activity that can itself confer meaning on what might, abstracting away from this context, seem to be mere "stimulation".

Second, against the notion that perceptual experience requires conceptual synthesis (thesis 3, above), proponents of nonconceptual content suggest that there may be an element in the experienced deliverances of perception that is not constrained, guided or informed by conceptual structures. This view appears to get some support from naïve introspection, as indeed it seems that the experienced content of perception is richer than any description one could make of it. Theorists of nonconceptual content typically analyze this extra richness in terms of the skills, dispositions or nondiscursive knowledge an agent has in virtue of the perception in question. (Bermúdez 1995a; 1995b; 1998; Chrisley 1995; Evans 1985; Gunther 2003; Peacocke 1998)

Finally, Andy Clark questions what he calls the Assumption of Experience-Based Control (thesis 6): "that conscious visual *experience* provides the very information continuously used for visually based motor control." In contrast, he argues for the "deep and abiding dissociation between the contents of conscious seeing, on the one hand, and the resources used for the on-line guidance of action, on the other." (Clark 2001)

These are but three examples, for the list of critiques of the received view is long and growing.[3] The current essay adds one more that growing list.

In addition to being interesting and worth consideration in its own right, the multiple modes theory of our epistemic openness to the world has an immediate philosophical benefit: it undercuts an important premise of a popular, relatively simple and apparently convincing argument for epistemic anti-realism. The work of this essay, then, is three-

fold: first, to criticize the received view along the lines I have indicated; second, to begin to sketch (and in the current essay I can do no more than this) a substantial alternative picture; and third, to show how this alternative picture of our epistemic openness undermines at least one version of epistemic anti-realism. For the most part, these tasks will be pursued in parallel; but before the real work can begin, we must first present the argument for anti-realism it is my hope to undermine. Thus it is to this task that we now turn.

The lure of anti-realism (and how to resist it)

A common and apparently convincing argument for epistemic anti-realism starts precisely from one of the elements of the received view: our epistemic contact with the world is mediated by concepts or conceptual structures (cognitive, social, ethical, linguistic), which interpret or structure the deliverances of our senses so as to provide our picture of the world. It is but a short step from here to the notion that it is these conceptual structures—and not the world itself—which determine the shape of experienced reality.[4] After all, isn't it *obvious* that people approaching the world with significantly different sets of conceptual schema will thereby understand the world differently? And does it not follow that the shape of this understanding is determined by the conceptual schema here employed, and not—most importantly—by the world (even assuming that "the world" has an unambiguous referent)? For according to this view, all of our epistemic contact with the world is accounted for by sensation, and yet epistemic *content* is determined by concepts.[5]

More formally, the anti-realist claim under consideration is that for no belief that *P* can it be known whether *P* accurately reflects a fact or state-of-affairs in an independently existing world. For the only available method of verification (or falsification) of *P* is to derive from it (and/or from –*P*) some set of observable facts about the world $\{F_1, F_2, \ldots F_n\}$ such that $P \rightarrow F_i$ & $-F_i \rightarrow -P$ (or, alternately $-P \rightarrow F_i$ & $-F_i \rightarrow P$).[6] The trouble with this method, according to the claim, lies in the notion of an *observable fact*. In so far as perceptual experience operates by the logic of the received view, no experience (observation) *E* describes or establishes a fact about the world *F*, but only a "fact" about how the world appears to a being with conceptual scheme *C*. Every *E* is a belief-like mental state made contentful only in virtue of interpretation by *C*—the same scheme of concepts by which all sensory stimulation is interpreted (and, not incidentally, in terms of which *P* is expressed). What follows from this is either that no experience *E* can establish any fact *F* about the world independent of its interpretation, or, more typically, that insofar as *E*'s do pertain to or establish *F*'s, the *F*'s in question are likewise "facts" concerning what a being with scheme *C* can expect by way of experience. One is never in the position of verifying the belief *P* or its implications directly against the world, or against some pure experience of that world, but only, at best, of judging the internal consistency of some set of beliefs. Considerations like this have led to the truism that one cannot "step outside" one's concepts to see the world as it is; yet without the possibility of verification this would provide, there seems to be no basis for the judgment that our concepts accurately reflect—or can be made through experience to conform more closely to—the actual structure of the world. Thus, in the absence of some other account

of how the world can limit, direct, and change these concepts, it must follow that the phenomenological order—the "world" of our experience—is epistemically closed.

Certainly, this conclusion has seemed obvious, or at least inevitable, to many thinkers.[7] And I find myself agreeing: if *we* were the way the received view supposes, this would be the way things were for us. But we are not (I will argue), and, therefore . . . here I am tempted to say, "things are not", but of course I am—or will be—entitled only to the conclusion that we are not forced, by our best judgment of how we are, to adopt this particular anti-realist picture of how things are.[8]

In my view, the way to defuse this kind of argument for epistemic anti-realism is to question an apparently innocuous element of the premise: that we have but one epistemically relevant mode of openness to the world (conceptually mediated sensory input), or, to put it differently, that our epistemic access to the world operates by this single logic.[9] For although an important aspect of our epistemic contact with the world is indeed conceptually mediated sensory input, there are additional modes of epistemic openness to the world, which, in not operating by the same logic, may offer avenues for epistemically meaningful contact between mind and world precluded by the received view.[10] In particular I will argue that there are other modes of experience or kinds of knowing that indeed allow the possibility of establishing certain facts about the physical world, and not just about the appearance of that world relative to a given conceptual scheme. In so far as this is right, these epistemically relevant modes of openness would allow the opportunity to verify (at least some) beliefs against, and conform them more

closely to, a conceptually uninterpreted world, potentially justifying the claim that conceptual contents are ultimately grounded in (and guided by) the structure of (or facts about) the physical world.

In (O'Donovan-Anderson 1996; 1997; 2002) I argued that proprioception—the bodily awareness of the position and motion of our own limbs—was a candidate for one such mode of conceptually unmediated epistemic openness to the world, and therefore presents, along the lines sketched above, a direct challenge to the epistemic anti-realist.[11] One reason for focusing on the epistemic importance of the active body was my own (unwarranted and unconsidered) assumption that vision, at least, worked entirely by the logic of the received view.[12] However, recent work in cognitive science suggests, in contrast, that vision in fact has multiple epistemic modes—that it consists of several distinct systems each serving a different function—and that there is (at least) one mode of vision that meets the criteria for being conceptually unmediated in its epistemic contact with the world. Let us begin, then, with this.

Two kinds of knowing in vision: the Milner-Goodale hypothesis

Although (Milner and Goodale 1995) point out that the retina in fact projects to *many* different regions of the brain—they identify ten distinct areas, including the dorsal part of the lateral geniculate nucleus of the thalamus (LGNd), the superior colliculus (SC) and the superchiasmatic nucleus (SCN), indicating that the visual system has a multitude of functions (e.g. the SCN is thought to govern the synchronization of one's circadian rhythm)—they are mostly interested in establishing the existence and function of two

major visual pathways. The first consists of a dorsal projection from primary visual cortex (V1) to the posterior parietal cortex (towards the region of the brain implicated in motor control), and the second of a ventral projection from V1 to the infero-temporal cortex (towards the regions of the brain involved with conceptual processing and, in humans, with language).[13] The ventral stream (VS) is a specialized visual processing system optimized for representing information to allow the recognition, identification and classification of objects in terms of those categories most useful to such cognitive tasks as modeling the stable features of the environment and planning. It is this pathway that is most closely identified with the experience of seeing (and the experience of thinking about what is seen) and it has therefore been, illegitimately albeit understandably, identified with the entirety of vision.

In contrast, the dorsal stream (DS) is a specialized perceptual processing system that represents information in a form optimized for calculating and directing motor responses. Information from this pathway is used to guide such things as the orientation not just of the sense organs for optimal perception or perceptual tracking (e.g. the visual grasp reflex), but also of the whole body to facilitate actions like reaching, grasping, and snapping at prey, which actions it also governs. The natural way to characterize what the agent discerns in virtue of DS representations[14] is the location, size, shape and orientation of an object, and how to get the agent-object relation into a preferred state. One might say that the DS places the object in an egocentric (and task-specific, see below) visuomotor space, or an egocentric action field, and the object is (re-)presented to the agent in these terms.

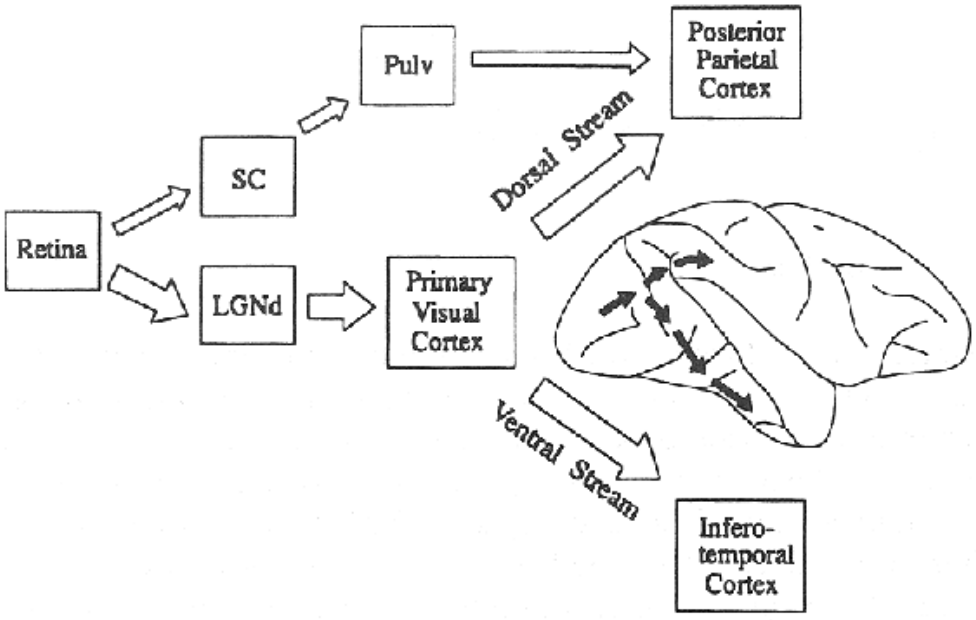

Fig. 1 Schematic diagram illustrating the two major visual processing streams. From (Milner and Goodale 1995).

Milner and Goodale are at pains to distinguish this view from similar claims made for the existence of "what" and "where" pathways (Schneider 1969; Ungerleider and Mishkin 1982). According to this latter view, the two visual pathways are specialized for identifying objects by category or features, and for locating objects in space, respectively. While not denying that the DS plays a role in fixing the location of objects, Milner and Goodale emphasize that this localization of objects is only one of the many sub-functions which the DS serves in virtue of its role in visuomotor control. This clarification is important for two reasons: first, because it underlines the greater scope of the Milner-Goodale position, and second because the what/where pathway hypothesis leaves the impression that there is a single visual system for representing space, and another single system dedicated to categorization. In contrast, Milner and Goodale suggest that

> While all visually guided actions take place in space, the spatial coding required will vary according to the action performed. In other words, there is no single representation of space in the brain, but instead multiple effector-specific coordinate systems. (p. 45)

Furthermore, in a system specialized for visuomotor control of action, one would expect encoding of *both* "what" and "where" information about the target object, albeit specified in a way appropriate to action rather than high-level cognition. And evidence suggests that this is indeed the case:

> Many of the cells studied by (Taira et al. 1990) were sensitive not only for the appearance of an object, but also for the appropriate motor act performed in grasping it. This motor information may reasonably be assumed to come from the rostral sector of inferior area 6 (area F5), a part of the frontal lobe known to be intimately connected with this part of the parietal lobe (Goldschalk et al. 1984; Matelli et al. 1986) and whence neurons project directly to the motor cortex (area 4). Cells in this premotor area include ones with motor and visuomotor properties associated with grasping behaviour (Rizzolatti et al. 1988) and (Sakata et al. 1992) proposed that they provide a feedback message to enable a 'matching' process between sensory input and motor output to take place in the parietal cortex. (p. 53)

The immediate point is that it is a mistake to imagine that one processing system (VS) is accounting for all the discrimination and classification ("what"-related) tasks, while the other (DS) in engaged only in spatial tracking ("where"-related) tasks. Rather each stream is identifying and classifying the objects in the world according to a set of categories most appropriate to its function, and in each case, therefore, it is legitimate to say that the objects are understood in these terms.[15] This further underlines the main message of the current essay, that there is more than one epistemically relevant mode of openness to the world, each operating according to its own logic and utilizing its preferred epistemic (sensory) materials. It is worth emphasizing, however, that although these two processing streams are functionally distinct, there is a great deal of anatomic interconnection and informational feedback between them. I will close this essay with

some suggestions as to the epistemic import of this cooperation between the two processing streams.

Simple illustrations of the Milner-Goodale hypothesis

There is a large and growing body of evidence for the two visual systems hypothesis, which in some form or other has become a standard account of visual processing. However, as the purpose of the current essay is neither to establish nor to argue for the existence of separate visual processing streams, but rather to explore one of its implications, evidence for the hypothesis will be recounted only partially and briefly, and primarily with an eye to illustrating its significance. As is often the case in cognitive science, some of the most striking illustrations come from experiments with brain-damaged patients.

Cortical blindness

Cortical blindness occurs when, although the eye and optic nerve are intact and functional, damage to V1 prevents its processing of visual information. Cortically blind patients report no visual experience, and are therefore unable to do such things as describe visual scenes. Given that V1 is the sole input source for VS processing, this deficit is to be expected. However, since the posterior parietal cortex has an alternate input pathway, through the superior colliculus and pulvinar (see fig. 1), one might expect some DS functionality to remain intact despite V1 lesions. And indeed, there is evidence for such 'blindsight' (Sanders et al. 1974). Patient DB, for instance, can accurately point to a visual target, despite his insistence that he is unable to see it. (Weiskrantz et al. 1974)

Optic ataxia

In contrast, patients with lesions in the posterior parietal cortex, while able to describe the orientation and location of objects in space (Perenin and Vighetto 1988), nevertheless show a number of impairments in reaching and grasping. For instance (Perenin and Vighetto 1983; 1988) recorded inaccuracies in the direction of reaching movements, and also a failure to correctly orient the hand in order to pass it through a slot. Similarly, the patient RV (Goodale et al. 1993) shows impairments in 'grip scaling'—the ability to pre-form the hand to the proper size and shape in order to grasp an object—as well as in judgments regarding where to position one's fingers so as to avoid having the object slip from one's grip (Goodale et al. 1994). This despite the fact that the relevant objects in each case were in plain view and could be accurately described.

Visual form agnosia

Visual form agnosia is a deficit in the ability to recognize and interpret shapes, resulting in difficulty in object and face recognition. Lower level visual abilities, such as color and texture recognition, are often preserved, allowing for the identification of some objects based on these characteristics (Humphrey et al. 1994). The condition is linked to occipital lobe damage, often following an anoxic episode, caused for instance by carbon monoxide poisoning. Patient DF, the most extensively studied visual form agnosic, shows bilateral necrosis in her lateral occipital cortex, the result of asphyxiation by a faulty water heater.

DF is not able to recognize line drawings of common objects, nor is she able to copy them. She *is* able to draw simplified versions of common objects from memory, although when later shown her own drawings, she cannot name the objects (Humphrey et al. 1994).[16] Despite these perceptual deficits, DF shows no impairment of typical visuomotor skills. She can accurately reach for and grasp objects, even those with complex shapes, and can catch a ball or a short wooden stick. In a particularly striking experiment, it was shown that DF is easily able to post a letter through differently oriented mail slots, despite being unable to report on the orientation either verbally or by matching the orientation—for instance by manually setting a comparison slot or by orienting the letter without moving to post it (Milner and Goodale 1995; Milner et al. 1991; Goodale et al. 1991).[17] More recently, DF's cognitive impairments, and their correlation with specific functional deficits in her brain, were confirmed by high-resolution fMRI (James et al. 2003). These results further support the Milner-Goodale hypothesis regarding the isolation and functional specialization of the two visual streams.

Epistemic openness and the received view

It does appear that the received view may accurately describe the operation of ventral stream visual processing, for in this case it appears that sensory impulses are interpreted in terms of a set of explicitly conceptual structures which together determine the meaning and significance of the visual input, and allow us to perform such critical cognitive tasks as modeling the stable features of our environment.[18] Were this the whole story, it might be that the structure of our epistemic cognition would land us in epistemological quandaries of the sort earlier described. But, as we have seen, this is not the whole

epistemic story, not even for vision. Instead, there is an additional visual system, which interprets and places objects not in an allocentric, conceptually structured cognitive space, but rather in an egocentric action field, in which objects are understood and presented in terms allowing the agent to appropriately select and guide actions taken with, or with respect to, the objects in question. As the examples of catching a ball or posting a letter indicate, this information involves not just static representations of the current scene, but also predictions of future states and expectations for the outcome of directed interventions.

This fact alone is enough to present a direct challenge to the anti-realist; at the very least he (and indeed, any philosopher interested in human knowledge) is called upon to assimilate this new information and alter his arguments accordingly. It may well be that he will be able to once again make a case for anti-realist conclusions.[19] But there are reasons to suspect this task will be difficult. For unlike the conceptual structures targeted by the ventral stream, the action-guiding structures of the dorsal stream, in terms of which objects are also understood, can be directly compared with the actions it is indeed possible to take in the world, and the success or failure of these actions can be used to change the relevant representations.[20] This ability to directly measure the appropriateness of interpretive structures to the world they represent is precisely what is missing in (or denied by) the received view, and it is exactly (and only) this that the realist needs to make the case that the world can indeed provide epistemic friction sufficient to limit and guide our conceptions of it.

In the sections that remain, I would like to *briefly* elaborate this hypothesis regarding the structure of our epistemic cognition, focusing in particular on how the cooperation of different epistemically relevant modes of openness allows for the possibility that the concepts by and through which we understand the world are themselves open to, and able to be influenced by, the world.[21] Since this possibility is intimately tied up with the epistemic import of acting in the world, we need first to have in place at least a provisional account of the nature of our awareness of, and the means by which we monitor, our own actions.

Bodily awareness: two more kinds of knowing

Somatoception, our awareness of the state of our own bodies, breaks into three categories of two different epistemic types. First there is exteroceptive information about the state of the world and its objects (exemplified by feelings of texture). Next we have interoceptive information about the state of one's body or its parts (exemplified by hunger or pain). And finally we have relational information about the relative state of the body and its environment (exemplified by the vestibular system).[22] The difference between the three *categories* of somatoception is in their object: awareness of the body's interior in interoception, perception of the external world in exteroception, and awareness of the relation between the two. These kinds of information can be further classified, following (O'Shaughnessy 1980; 1995), by their sensory or epistemic mode. The first kind consists of perceptions that are mediated by explicit sensory feelings. In touch, for instance, the information in question is cast in terms of, or perhaps it is better to say *carried by*, such sensory feelings as pain, heat, cold and roughness. Perceptual seemings—perceived

matters of fact—in this case centrally depend upon the quality or characteristics of the feelings in question. In contrast, there are some aspects of somatoception that provide information unmediated by sensory feelings, and without the need for conceptual interpretation. Proprioceptive information about the position and motion of one's limbs appears to be of this sort. Here one is presented simply with the fact of the arrangement; in the proprioceptive seeming there is nothing else to notice but the fact that one's limbs are arranged thus. Given that the primary function of proprioceptive feedback is to guide and monitor physical actions, the simplicity and immediacy of the processing is sensible. If one is suddenly thrown a ball, one doesn't need to attend to one's limb position, or be in a position evaluate its characteristics, before moving to catch the ball, rather, one's bodily position is constantly and unqualitatively given, and this informational state, whether attended to or not, is used in the calculation of the motion. Of course, one *can* become aware of position, and express it in conceptual terms (bent at 90°), but even so there is an important distinction between this, and feeling heat or seeing an apple. In the latter cases, it seems we should say that the information in question is *carried by*, or *relies upon*, the quality of the feeling, or the application of the concept; without these there *is* no information, no belief. In proprioception, however, while the information it provides is of course used to guide its conceptual expression, it does not consist in this expression, nor does this information rely on the conscious identification of the quality of any sensuous feeling.[23] Still, as (O'Shaughnessy 1995) is at pains to establish, proprioception can indeed be a kind of perceiving:[24]

> [I]t is an experience, of the type attending, whose content is caused nondeviantly by its object, and it can form the basis of an inference to the existence of its object. Above all, it is no kind of cognitive attitude; even though, as befits perception, it causally sustains such, since one could in principle have this

> experience when one knows irrefutably that the bodily facts are other than they seem in the experience. In short, we have here an attentive experience in which a small sector of physical reality appears one way, which is to be sharply distinguished from cognitive attitudes of all kinds, even though it naturally sustains such. In a word, a perceiving. (p.176)

Even given these distinctions, it is clear that the types of somatoception must cooperate in various ways, and also with other categories of perception. Thus, the position of one's limbs can be given in non-qualitative awareness, but also by touch (the feeling of the desk pressing against one's knees) and by vision (Ghez, Gordon and Ghilardi 1995). Indeed, vision can sufficiently confound one's sense of limb position that it is apparently possible to locate—to feel—the touch of a feather in a clearly visible and strategically placed rubber arm, instead of in one's actual arm, which is being simultaneously touched but is hidden from view (Botvinick and Cohen 1998). Likewise, a single touch can simultaneously give interoceptive information—a heat in one's finger—and exteroceptive information—the heat of the stove one is touching. And finally, it seems that certain kinds of tactile perception, e.g. feelings of texture, insofar as they involve not just contact between the sensing organ and the object, but also the *motion* of that organ, require both proprioceptive and tactile awareness. The fact of this cooperation is important to the current essay, but as the details are not, and are quite complex besides, further inquiry will be left for a future work.

What I do wish to draw attention to is the apparent existence of two more epistemically relevant modes of access to the world, which do not seem to operate by the logic of the received view.[25] The first mode—and the one most important to the current essay—is of course non-qualitative proprioception. Again quoting O'Shaughnessy:

> We stand to our limbs in a relation of *awareness*. A concrete or *intuitional* awareness. And since that awareness ceases when feeling ceases, and in despite of the fact that feeling is not its evidential ground, it must be a *sensuous* intuition. Then what information do we glean in being thus aware? We learn of *the presence of* the limb, and of such properties as that it is straight or moving away from one's body, i.e. of certain *spatial properties* of the limb. (O'Shaughnessy 1980, p. 167)

For O'Shaughnessy, our awareness of our bodies is *immediate* and *sensuous*; he supposes the latter because bodily awareness evidently requires the cooperation of our nervous system (severing the appropriate nerves will disrupt proprioceptive feedback), but it should be clear that, given the immediacy of the awareness, the necessary sensory impulses play only a causal, and not an epistemic role in bodily awareness. It is because of, but not *through* bodily sensations that one is aware of the limbs and their posture; there is, in such awareness, *nothing else to notice* but the limb and its posture. No *qualia* intervene. "This immediate awareness, which is possible only because of feeling in the limb, we *non-sense-perceptually* characterize as 'being able to feel that it is flexed'." (O'Shaughnessy 1980 p.217)[26] Likewise, no conceptual mastery is required to represent and make use of the information in question. Although one may require the mastery of certain concepts—bent, flexed, straight, degree, etc.—to *express* the state of one's limbs, it cannot be the case that conceptual mastery of this sort is required for an agent to know how his limbs are arranged and make use of this information in action. Walking, reaching, grasping, eating all require proprioceptive awareness, and none of these require conceptual mastery. What follows from this is that one important way of experiencing and monitoring our own actions is through this immediate non-qualitative awareness; we have immediate, non-qualitiative, and unsynthesized access to the current positional state of our bodies, and hence to the progress of ongoing activity. As will be explained in

more detail below, this gives us the ability to compare expectations for action—generated, for instance, as a result of one's identification of an object as being in a particular place and of a particular type, from which it follows that a given intervention will produce a certain result—with the actual outcome of action. Success or failure of action, then, can be used to question or confirm the validity of one's perception. Indeed—and this speculation is at the center of the next section—such feedback can perhaps even, by allowing one to make adjustments to the concepts involved in perceptual organization, actually change the character of one's perception (Gauthier et al. 2003).

The second mode of access is that identified as the qualitative element of somatoception. This kind of perception at first blush seems to operate identically to conceptually mediated perception: in touch, for instance, elements of the world can be understood and categorized in terms of publicly available concepts. Indeed, one can certainly perceive there to be an apple or a sphere in one's hand, and there is not likely to be a *quale* for an apple or a sphere; rather the experience is organized under, and experienced in terms of, these conceptual structures. At the same time, however, heat, cold, texture, and pain seem to be distinguishable from each other, and divisible into internal categories, not in terms of conceptual structures like these, but rather in terms of *qualitative* features of the experience.

If I see a dagger before me I do indeed perceive *the dagger*. I do not *feel* that it is a dagger, nor do I judge that it is so based on *qualitiative* evidence.[27] I perform no post-perceptual manipulation or arrangement of *qualia* into a dagger-like mental object; this is

not how perception works. Rather, perception makes the dagger present to mind, and that it is a dagger is a given feature of my experience. On the other side of the coin, I *do* feel its sharpness, and the coldness of its steel, and these (along with, perhaps, its silvery sheen[28]) are identifiable qualities that are likewise part of my experience of the dagger. What this suggests is that experience has not just *conceptual* or *categorical*, but *qualitative* features; and these features represent different aspects of our consciousness of objects. Further, it may be that each such feature of experience is a sign of a different facet of our epistemic sensitivity to the world.

Note that what has been said above does *not* amount to a restatement of classical empiricism. That doctrine envisioned the existence of a set of basic sensual *qualia*, from which a picture of the world would be constructed. I make no claim that there is any level of pure experience consisting of raw feels, pure *qualia*, or the like, that is epistemically basic or the primal, pre-conceptual product of perception. Nor should what I have said here be construed to imply even that there are some sensory modalities that deliver, in isolation and without cognitive processing, a stream of pure *qualia*. First, it may well be that it is only in cooperation with other elements of perception, and the context they provide, that qualities can be experienced and understood; thus, for instance, it appears that the experience of a pain or a tickle requires the proprioceptively rooted sense of the spatial characteristics of one's body. It is evidently not possible to feel an itch in a place that does not at least seem to be part of one's body.[29] Likewise, we may not, in general, be able to experience a quality—red—without it seeming to be the quality *of* something—the fire engine. In this case the experience of red would require the

cooperation of qualitative *and* conceptual processing.[30] Further, qualitative elements of perception may be the result of high-level processing mechanisms every bit as complex as those involved in conceptual or categorical perception. The suggestion that the qualitative aspect of experience is different from the conceptual, and exploits different features of sensory input, does not imply that it is simpler, purer, or more accurately reflects the basic nature of that sensory input.[31] And finally, the notion of pure *qualia* generally implies a nonintentional component of experience, a pure "feel" that has no intrinsic connection to any property in or element of the world. What is being posited here, in contrast, is that what is given in perception is, in some cases *conceptually* mediated information about the state world, and in other cases *qualitatively* mediated information. Likewise, perception can provide *action-oriented* information, and in at least one case, *unmediated* awareness of the state of one very special part of the world: one's own body. There is simply no suggestion in any of these cases that perceptual experience is essentially (or even largely) nonintentional.

Because it is a central purpose of this essay to suggest that human knowers possess more than one, and perhaps many, epistemically relevant modes of openness, I have tried to suggest a number of possibilities for what these modes might be. I have identified four candidates so far: cognitive-conceptual perception, exemplified by ventral stream visual processing; qualitative perception, exemplified by touch; action-guiding perception, exemplified by dorsal stream visual processing; and a special kind of unmediated awareness of bodily position and motion, provided by proprioception.[32] I am aware that some of these suggestions may be controversial, and that establishing all of them would

require far more time and careful argument than I have provided here. Still, I think the general *idea* of multiple modes of epistemic openness to the world is sound, even if one or another of my particular suggestions must be abandoned. And in any case, it is only to the existence of the latter two epistemic modes that the central thesis of this essay is committed. Thus it is to these modes of epistemic access to the world, and their postulated significance, that I turn in the next, concluding, section.

Conclusion: the multiple modes theory of epistemic openness

It is of course obvious that most of our information about the state of, and goings-on in the world are the result of the complex cooperation of our various sensory modalities. What I have suggested, in addition, is that we are possessed of multiple epistemic modes—I have proposed four candidates—and that cooperation between the various epistemic modes can account for our openness to the world, and the ability of the world to limit and guide our conceptions of it. As I mentioned already above, there mere existence of non-conceptually mediated epistemic modes is enough to block, at least temporarily, the anti-realist arguments with which we began this essay. Insofar as this is correct, the main argumentative responsibility of the essay has already been discharged. Nevertheless, the case would undoubtedly be stronger were I able to provide a concrete illustration of how such cooperation might work, and thereby begin to construct a substantial alternative to the received view, and its implication of our cognitive confinement. Thus I will end with one particular, highly speculative, suggestion as to how cooperation between VS visual processing, DS visual processing, and proprioception

might allow the possibility of determining the appropriateness of a concept or set of concepts to the world.

Let it first be stipulated that it is a central role of (at least some) concepts to provide guidance for behavior. I take this to be relatively uncontroversial, so far as it goes, and it is not meant to contradict any claim that their role is to structure perceptual experience, participate in inference, or model the environment. Indeed, on my view, it is largely *in virtue* of their role in guiding behavior that concepts structure perceptual experience, participate in inference, and model the environment (Anderson 2003a; 2003b; 2005; O'Donovan-Anderson 1996; 1997; Rosenberg and Anderson 2004; forthcoming). In any case, the concepts that *do* provide guidance for interacting with the physical objects[33] that fall under them, do so in virtue of such connections with or information about the repertoire of actions available to the agent, and the relevant action-related properties of the object, as will allow the agent to make appropriate decisions about which actions to take under the circumstances to produce the desired outcomes. Again, I take the sentiment (if not the details of its expression) to be obvious and uncontroversial. Something like this must be the case, insofar as actions are guided by concepts. But we have yet to put this observation in the context of the current essay; doing so marks the beginning of admittedly controversial speculation.

To this point, we have been focusing primarily on the evidence for the differences in function between DS-related and VS-related representations. However, there is evidence to the effect that certain actions—for instance posting a complex shape through a

matching slot—require that *both* dorsal and ventral stream representations of the relevant objects cooperate (Goodale et al. 1994). Indeed, it seems that conscious, intentional action would in general require cooperation between the two streams, precisely as a result of their different functions. Thus, for instance, we know the cognitive/conceptual system is involved in decisions regarding what actions to take with respect to which object (kick the ball), and even helps determine some elements of the character of the action—for instance, how much force is expected to be required—in light of the goals of the action, and conceptually-encoded knowledge about the nature of the object (Ellis and Lederman 1998). But in order to effectively guide action, this high-level intention must be passed to on-line behavioral control systems, capable of translating or interpreting the abstract and high-level intention into the specific mechanical requirements of the action. Given the different roles served by the two visual streams, it is hypothesized that the action-guiding role of concepts—that is, the fact that conceptually rendered high-level intentions can exert a specific influence on on-line action-guidance—is facilitated by a set of specific, and presumably flexible, anatomical associations and correspondences between these concepts and the dorsal-stream-targeted (posterior parietal cortex) structures implicated in on-line, real-time action guidance. It might be reasonable to hypothesize further that these associations would be sensitive to—and, indeed, track or co-vary with—learning and concept change, and moreover that feedback gathered in the course of action could influence such learning.[34] After all, knowledge influences interaction; but interaction also influences knowledge.

It is important to keep in mind that how concepts are represented in the brain is not well understood, and I am not committing myself to any particular model here. The basic thought is that the high-level cognitive-conceptual structures of the occipital cortex have use or task-related information associated with them—perhaps as a result of their own construction from more basic conceptual elements such as mental models, feature sets, semantic relations, and the like—and that this information could be reflected in, copied by, or even *stored as* a set of correspondent (and task relative) associations among the basic motor schemas stored in the parietal cortex and motor areas (Jeannerod 1997). In this way, inferencing with the concept "kick" could also prime or otherwise trigger the relevant associated motor schemas (and, perhaps, kicking might prime or otherwise trigger the relevant concept "kick"). In point of fact, I know of no explicit evidence either for or against this hypothesis,[35] but it seems that, *if* concepts do provide guidance for action, and *if* acquiring and adjusting our concepts involves changes (say) both to a concept's inferential role or cognitive contents and to its action-guiding role, and *if*, further, changes to its inferential role or cognitive contents can have implications for its action guidance (and vice versa), and *if*, finally, on-line action guidance is handled by (and in terms of) dorsal stream processes and representations, and inference by (and in terms of) ventral stream processes and representations, *then* one reasonable mechanism to account for the coordination of all these elements, is that concepts are stored as distributed representations, with dorsal stream (action-related) and ventral stream (cognition-related) elements, which elements are closely connected and co-varying.[36] In any case, I offer this as a speculative but testable hypothesis as to one kind of cooperation between dorsal and ventral stream processing—the maintenance of distributed conceptual

structures with both dorsal and ventral elements, supposed to be active in each case in the structuring of perceptual information as processed by these two perceptual streams.

It is at this point that we are in a position to see how different epistemic modes might cooperate to transmit epistemic friction from the world to the relevant conceptual structures. For let us suppose that a given agent has, in his acquisition of the concept "lead", somehow not gathered that it is very dense. This agent, seeing a lead ball (say 28 inches in diameter) and, being an avid soccer player, knowing that balls are good for kicking, may well decide to kick this ball into the nearby (conveniently placed) net. At the same time, the semantic information associated with "lead" gives a misleading estimate of the feasibility of the action, and the approximate amount of force to apply. Thus having decided on an action, and a target, the continuing real-time guidance of the action will be handled by DS processes. The relevant motor representations for "ball" direct his kick in such a way that a certain spot of the top of his foot will contact a particular spot between the ball's equator and the ground. At the same time, the force expectations passed on in virtue of the conceptual associations of "lead" are translated into the muscle force expected for the ball to reach the goal. Thus the agent kicks the ball, hard.[37] Let us assume, out of compassion for our agent, that he does not break his foot. Nevertheless, the action will of course fail; his foot will stop before he expects, and the ball will not move. The agent will be immediately, non-qualitatively aware of the failure of his bodily motion; the actual trajectory and ending position of his kick can be directly compared with that expected in virtue of the chosen motor schema.[38] This failure will suggest the inappropriateness of the chosen motor schema,[39] which in turn will change

the action-guiding content of the relevant concepts, which (by earlier hypothesis) will alter its inferential content. "Lead", that is, will come to imply dense, or heavy. Now, of course, this process, described here so briefly, is likely to be quite complex, and may well involve further explorations or the testing of provisional hypotheses by our surprised agent (he may try to pick up or roll the ball, or kick a different ball, or some such). But the point here is not to suggest a general theory of learning, nor to provide any details about how, exactly, learning in this case would take place—although it is certainly worth mentioning the fact that quite complex learning can take place guided *just* by positive and negative feedback, in the form, for instance, of task success and failure (Sutton and Barto 1998).

Rather, the point is to establish the possibility, against many years and pages of theorizing to the contrary, that such concept change can indeed not only take place (the various coherentist and internalist theories of course allow for that possibility, and the story told above is not incompatible with those) but that this change can be directly attributed to the (not conceptually mediated) influence of the (not conceptually interpreted) world. Returning to the somewhat more formal terms offered early in this essay, theory (or concept, or belief) *change*, and thereby, indirectly, theory (or concept or belief) *content*, can be attributed to the world itself just in case it is possible to test and establish whether some set of observable facts $\{F_1, F_2, \ldots F_n\}$, which follow from the theory (or concept, or belief) in question *P*, do indeed obtain. The claim I am making is that some members of this set of observable facts specify possible actions, and the expected, immediate, bodily outcomes of those actions. Let us represent this set: $\{F_{1a},$

$F_{2a}, \ldots F_{na}\}$, such that that $P \rightarrow F_{ia}$ & $-F_{ia} \rightarrow -P$ (or, alternately $-P \rightarrow F_{ia}$ & $-F_{ia} \rightarrow P$). Proprioceptive feedback offers the possibility of verifying (or falsifying) *these* facts—it can establish whether F_{ia} or $-F_{ia}$—without any need for conceptually interpreted, mediated, or structured experience; *proprioceptive* experience, then, can indeed establish facts about the world (regarding the possibility, and immediate bodily outcomes of action). Thus, insofar as the claim that mental content cannot ultimately be attributed to the world is based on a theory of the content of experience that has been shown to be incomplete and inadequate to the complexity and diversity of human epistemic ability, this generic anti-realist argument fails.

It may indeed not be possible to "step outside" our conceptual structures to see the world as it "really is"; but neither does this imagined direct confrontation between concepts and the bare structure of the world appear to be necessary to make the claim that we know the world, and that the concepts in terms of which we express this knowledge are grounded in our contact with that world. Naturally, acting in the world is not seeing it, and the proprioceptive experience arising from our physical encounter with that world is not conceptually or cognitively rich. But insofar as it provides even the minimal information required to determine the success or failure of an action (and that it does *at least* this is not speculation; the role of proprioceptive monitoring in governing action adjustment and repetition in light of failure is well documented) it is nevertheless sufficient to drive concept correction and change and, in doing this, to establish our epistemic openness to the physical structure of the world.

Acknowledgments

Thanks are due to Peter Carruthers, Andy Clark, David Milner and Gregg Rosenberg for helpful comments on earlier drafts of this essay.

Notes

[1] Nöe and Thompson comment: "What is given to us, one might suppose, is not the world itself, but the pattern of light on the retina, and *that* pattern does not supply enough information to determine how things are in the environment. For example, from the retinal image of a table alone, it may not be possible to tell whether it is large and far away, or small and nearby. . . .In the face of this puzzle an orthodox or 'Establishment View' of perception (Fodor and Pylyshyn 2002) has taken shape over the last fifty years. According to this orthodoxy, perception is a process whereby the brain, or a functionally dedicated subsystem of the brain, builds up representations of relevant features of the environment on the basis of information encoded by sensory receptors. As David Marr (2002) puts it: 'Vision is a process of discovering from images what is present in the world, and where it is.' Because the patterns on the retina are not sufficient by themselves to determine the layout of the surrounding environment, perception must be thought of as a process of inductive inference. Perceptions are, as Richard Gregory (2002) suggests, hypotheses concerning the distal causes of proximal stimulation. In the famous phrase of Helmholtz, perception is unconscious inference." (Nöe and Thompson 2002b p.4-5)

[2] Following out this general line of reasoning, Kevin O'Regan and Alva Nöe make an interesting case for the role (and necessity) of the eye's activity in the processing of vision. According to their view, the output of the visual system is not a static picture or scene description, but rather a set of 'sensory-motor contingents' describing the relation between changes in sensory stimulation and the movements of the eye, head, and perceived object. "Under the present view of what seeing is, the visual experience of a red color patch depends on the *structure of the changes* in sensory input that occur when you move your eyes around relative to the patch, or when you move the patch around relative to yourself." (O'Regan and Nöe 2001 sec. 5.7)

[3] See (Nöe and Thompson 2002a) for a nice collection of essays on the subject

[4] Rorty again: "Since Kant, we find it almost impossible not to think of the mind as divided into active and passive faculties, the former using concepts to "interpret" what "the world" imposes on the latter. . . . But as soon as we have this picture in mind it occurs to us, as it did to Hegel, that those all important *a priori* concepts, those which determine what our experience or our morals will be, might have been different. . . . The possibility of different conceptual schemes highlights the fact that a Kantian unsynthesized intuition can exert no influence on how it is to be synthesized—or, at best, can exert an influence we shall have to describe in a way . . . relative to a chosen conceptual scheme. . . . Insofar as a Kantian intuition is effable it is just a perceptual judgment, and thus not merely "intuitive." Insofar as it is ineffable, it is incapable of having an explanatory function. This dilemma . . . casts doubt on the notion of a faculty of "receptivity." There seems no need to postulate an intermediary between the physical thrust of the stimulus upon the organ and the full-fledged conscious judgment that the properly programmed organism forms in consequence." (Rorty 1979 p.3-4)

[5] This is the upshot of (Quine 1985) and (Davidson 1985). What Quine and Davidson deny is that, prior to conceptual synthesis, sensation has qualities which can be experienced non-cognitively as such; it is to deny that (again, prior to cognitive interpretation) experience has what Rorty calls "raw feels" in terms of which we can access, or get some phenomenological handle on, our sensory experience. Quine and Davidson do not claim that our experience has *no* phenomenological content (that there is nothing it "feels like" to touch wet grass); instead they argue that whatever wet grass "feels like"—whatever the content of this experience—no stimulation of our sensory receptors warrants the name "experience" except that which is conceptually structured.

[6] Note that putting the problem of verification this way points in the direction of some familiar problems of induction and scientific confirmation. These are not of concern in the present essay, the aim of which is to

question the anti-realist claim that the verification process—whatever its other difficulties—cannot even get off the ground, for facts themselves cannot be verified.

[7] Naturally, there have been numerous attempts to avoid this conclusion. Among the most influential have been (Davidson 1985) and (McDowell 1994), who decry the scheme/content division this picture still seems to imply, and insist that, contrary to the received view, there are *no* epistemic mediators active in our contact with the world; our knowledge of the world, while conceptually structured, is not thereby conceptually *mediated*. Davidson argues that the very structure of interpretation—its required maximization of local and global coherence—guarantees that most of our beliefs are true. This, Davidson suggests, when combined with the fact that world of our experience is always already conceptually structured, shows that there is no room for the notion that we could be incorrectly interpreting some independently defined entity. Once we give up on this idea that we can identify a set of concepts on the one hand, and an uninterpreted world on the other, it simply follows that skepticism (and anti-realism) is false. But few have found this convincing, and Davidson is generally understood to support some kind of internalism or coherentism, neither of which is a strong ally to realism. I myself find it hard to make sense of this distinction between conceptual *structuring* and conceptual *mediation*, and argued extensively in (O'Donovan-Anderson 1997) that, while it may avoid one particular version of skeptical relativism, in so far as it fails to provide an account of our epistemic openness, it must also fail to provide grounds for epistemic realism. Even after Davidson's refutation of the third dogma of empiricism, the attraction to anti-realism remains.

[8] Note in particular that the current essay is not intended to make serious headway against brain-in-a-vat or Matrix-style arguments for radical skepticism. It is only meant to establish that, according to our best understanding of the nature of our epistemic access to the world, we can justify the claim that we are cognitively "in touch" with the world we inhabit. That is not by itself enough to establish the truth of any arbitrarily chosen belief about that world, whether physical, metaphysical, ontological or micro-structural. It is nevertheless a significant claim, and one that might be used to leverage further realist arguments. For one useful approach to Matrix-style skeptical arguments, see (Chalmers 2003).

[9] For reasons obvious to those who have followed this debate, I call this the *fourth* dogma of empiricism, and in what follows I will argue that it, too, has to go.

[10] At this point in the argument, Peter Carruthers has raised the following questions: Is it really necessary to postulate another mode of epistemic access to the world to avoid anti-realism? Why can't one point, for example, to the evidence of innate conceptual structure, and to the argument from evolutionary psychology that such structures will have been selected for precisely because they at least approximate to the causal structure of the world in which the organism operates? Well, perhaps one can. What I am claiming is first, that without an account of how the world can influence and direct *whatever* structures are implicated in the production of perceptual content, the anti-realist argument has all it needs to move forward, and second, that no such account is available on the received view. In this light, it is quite true that the various naturalistic/evolutionary accounts of epistemology can be construed as offering the missing account. Note, however, that it is not clear that in doing so they preserve the uni-modal epistemic assumption of the received view. For selection pressure can perhaps be construed as another kind of epistemic friction—another mode of epistemic openness to the world—the mechanisms of which work directly and slowly on the (for instance) innate conceptual structures implicated in generating perceptual content from sensory stimulations. Of course, this interpretation assumes that the mechanism of natural selection can be characterized as truth-seeking (tending to result in more accurate cognitive systems), and there are serious questions about this assumption (on this point see (Akins 1996)). But whatever the case, the matter does not, so far as I can tell, significantly affect the arguments of the current essay. If selection pressure emerges as another mode of epistemic openness to the world, so much the better for my overall project, to argue for our possession of multiple—and perhaps even very many—epistemic modes. On the other hand, if selection pressure fails on some grounds to qualify as a mode of epistemic openness to the world, this failure does not necessarily undermine the status of any other proposed mode of epistemic openness to the world.

[11] I argued, in essence, that bodily activity, tracked and known through proprioception, allowed the possibility that concepts, in so far as they standardly provide guidance for acting in the world, could be refined and modified in virtue of the success or failure of the actions in question. See also (Anderson 2003b; 2005; Rosenberg and Anderson 2004; forthcoming).

[12] Closer attention to (Gibson 1966; 1977) might have helped avoid this error, as he suggests that vision, as with the other senses, has both exteroceptive and proprioceptive elements.
[13] The anatomical differentiation is of course much more complex and interesting than this characterization suggests. For the details the reader is encouraged to consult (Milner and Goodale 1995), chapter 2.
[14] Here, and throughout this paper, the term "representation" and its cognates should be understood in accordance with the theory of representation presented in (Anderson 2005; Rosenberg and Anderson 2004; forthcoming). Roughly speaking, a state **R** in agent **A** represents entity **E** for **A** in circumstances **C** just in case **A** has an enduring conscious preference or conditioned reflex to use **R** to guide its behavior with respect to **E** in **C**. This definition abstracts away from the possible instantiations of **R**, and does not require that **R** can be cashed out semantically or conceptually.
[15] Likewise, each stream is involved in *spatial* processing. Thus, the specialization of the two systems must be understood in terms of their functional roles, and not in terms of the "sort" of processing they do, abstractly defined. For a recent example of the "what" and "where" systems hypothesis, which sits uneasily with the Milner and Goodale position, see (Hurford 2003). Hurford argues that the "where" pathway might function to fix a spatially-coded reference (or deictic pointer) to an object, which could serve as the foundation of the variable *(x)* in *predicate(x)* structure. One question worth asking Hurford in light of the Milner-Goodale position is: "Which 'where' pathway"? For a different criticism of Hurford, but driven by similar considerations, see (Anderson and Oates 2003).
[16] Presumably, creating more complex renditions of these objects—for instance, including a leaf on the stem of an apple—would involve form recognition. In order to know where to put the leaf, one must know what one is looking at and how it is oriented, which DF does not.
[17] Interestingly, although she performed well on the mail slot task, when posting a 'T'-shaped object into a matching slot, her performance deteriorated to 50%, with the failed attempts almost always off by 90 degrees from the correct orientation (Goodale et al. 1994). This appears to indicate that she is using only one element of the 'T' to guide her posting behavior, and is unable to combine the two elements into an entire oriented shape. One reason this is interesting is it suggests the importance of cooperation between dorsal and ventral stream processing in some visuomotor tasks. In contrast, DF's ability to accurately *grasp* complex shapes is unimpaired, indicating that grasping ability is isolated to the dorsal stream.
[18] Note, once again, that this modeling includes spatial as well as conceptual elements. It is not just that objects are recognized, but they are seen as oriented and in relation to other objects and to the self.
[19] There is, of course, a movement called 'the new scepticism' (Nöe 2002b), which, on the basis of some interesting findings in the psychology of vision (e.g. change blindness) and our susceptibility to visual illusions, argues for a strong constructivist (and therefore anti-realist) account of perception. I believe that the considerations I advance here go some distance in answering the challenge of the new skepticism, but I will not give that argument in the current essay. For a review of the subject see (Nöe 2002a), and for some arguments against the new scepticism compatible (so far as I can see) with my own, see (Nöe 2002b).
[20] This general thesis that acting in the world can guide conceptual change is hardly new. See, for instance, (James 1912; 1981; Peirce 1955; 1958; Hacking 1983). However, the pragmatist notion was that action increased experience, and their concept of experience was still too closely tied to the classical empiricist notion of sense data to allow them to side-step anti-realist arguments. For the details of this account of Pragmatism see (O'Donovan-Anderson 1997, ch. 3). What I am proposing, in contrast, is that there is more than one kind of experience, and that acting in the world exploits these other kinds.
[21] Some elements of this hypothesis were introduced in (O'Donovan-Anderson 1997) and (Anderson 2003b).
[22] Interestingly enough, while the thermoceptive and the vestibular systems both provide relational information *in fact* (we sense thermal differences between our body and the environment, and not absolute temperature), this information is not always experienced *as* relational. Apparent motion and apparent temperature, for instance, are generally imputed either to the self or to the experienced object.
[23] Likewise with bodily extent and shape—these are given as facts, not feelings, unmediated by quality, and are the basis for such things as the calculation of motion. A good account of the nature and origin of the body schema is given in (O'Shaughnessy 1980; 1995) Note that the body schema is not the same thing as the body image, which is cognitive in nature, with this kind of self-perception mediated by concepts and qualities. The body image can play a role in conscious choices (eating behaviors, for instance), and is subject to a unique set of pathologies. Phantom limb phenomena are pathologies of the body schema, anorexia of the body image (Gallagher 1986).

[24] It is not clear that, in what might be called its most typical, attentively recessive form, proprioception *is* a kind of perception, strictly speaking. See (Gallagher 2003) for a good account of the issue, and an argument that proprioception is typically a form of non-perceptual awareness. None of the arguments in the current essay depend on the resolution of this debate. All that is required for the argument of the current essay is that the awareness of the state of the body it provides is not a *conceptually or qualitatively mediated* form of awareness.

[25] For those who find anatomical evidence convincing, it turns out that not only are there different sets of receptors for heat/cold, texture, and limb position and motion, but what I am calling qualitiative somatoception (heat/cold, texture) is processed by a different pathway from non-qualitiative somatoception (proprioception). Among other differences, touch is importantly cortical, while proprioception is importantly spinocerebellar, although it integrates with other sensory modalities at a cortical level (Sommer and Wurtz 2002; Nicolelis et al. 1998; Craig and Rolman 1999; Bosco and Poppele 2001).

[26] In case the claim is not apparent from experience, consider the following: It is characteristic for sensory feelings to differ both in type and intensity. Pain of the same qualitative type (ache) can be mild or intense and everything in between. There seems to be no axis of intensity for proprioceptive seemings. Without an axis of intensity, it would seem that each sensed position would require a difference in experienced quality. But there are innumerable bodily positions which differ from each other only in very subtle ways, whereas it is characteristic of different sensory *qualities* to be radically different, if not incommensurable (is this more rough than that is red?). Indeed, the subtle differentiations of bodily position are characteristic precisely of the subtle differences in intensity allowed by most sensory modalities. This is not surprising, as the perception of subtle difference here depends on a shared basic aspect of the sensors involved—their ability to fire more or less rapidly, that is, their capacity for gradations in stimulation. In proprioception, this gradation of stimulation is used to signify gradations of position, whereas in pain, for instance, it is translated directly into intensity of experience. This is perfectly reasonable for a sense like pain or thermoception, where the intensity of the stimulation might well reflect the strength of its cause, and one would want to react accordingly (pulling quickly away from the hot stove). In contrast, given the function for which this characteristic of the sensors is used in proprioception, there is no need for the output to be translated into intensity of experience. As with the qualitative difference in other sense modalities, what it reflects instead is a different matter of fact.

[27] A word about vocabulary: when I talk about feelings, I mean those features of our perceptual experience that are possessed of quality. I make no claim for the autonomy of *qualia*; they are not the pre-existing building blocks of experience, nor are they experience *per se*, nor the pure "pre-conceptual" products of our sense organs. I am claiming only that they are one identifiable feature of—one class of object to be found in—our experience. I use "perception" to refer to the entire range of our epistemic and information-gathering interaction with the world. Or, to put it differently, perception includes all sensory experience/input that generates or influences representations, where "representation" is defined in terms of the theory of representation described in (Anderson 2005; Rosenberg and Anderson 2004; forthcoming). (See also fn. 14.) Thus understood, the product of perception is not *qualia*, pure or otherwise; it is our experience—our *awareness*—of the world. However, I believe that each such feature of experience is a sign of a different facet of our epistemic sensitivity to the world.

[28] Actually, I have no considered judgment regarding whether vision has, in addition to the modes already discussed, yet another, qualitative mode—exemplified, for instance, by color vision. For a discussion see (Thompson 1995). The main thesis of the current essay does not require the hypothesis that vision has a qualitative mode—nor, for that matter, does it require that *any* sensory modality possesses a qualitative mode. What *does* matter is that there is at least one epistemic mode other than that described by the received view, and that it be conceptually unmediated in its contact with the world. It is further necessary to my particular speculations that this mode be operative in bodily activity.

[29] "Bodily sensations cause an awareness of themselves *as* set in a specific position in a determinately postured limb, and simultaneously those same sensations cause awareness of the very limb, and *as* determinately postured, in which they themselves come as seemingly set." (O'Shaughnessy 1980 p.204) Given that we can experience such sensations at any, but only at some, place on a seeming body (and recalling that it is not by and through having such sensations at given points on our bodies that we are aware of the body) we can postulate the existence of an always already present seeming body, which provides the framework or substrate for bodily perception, and is the object of bodily awareness. This is the *body schema*. In its short-term manifestation, it consists of an awareness of one's current posture; in its

long-term manifestation it consists of a sense of the persisting spatio-structural features of one's body, thus not current postures, but possible ones.

[30] In general, as Merleau-Ponty puts it: sensation is always "a formation already bound up with a larger whole, already endowed with a meaning . . ." (Merleau-Ponty 1962 p.9; see also Thompson 1995 ch. 5)

[31] Indeed, in my view, qualitative perception is just an older, simpler form of categorical/conceptual perception, developed for the same purpose: to allow for differential action in light of different sensed features of the world.

[32] These different epistemic modes cut across different sensory modalities; in addition to possible cases of one-to-one correspondence between epistemic mode and sensory modality, a given sensory modality may employ more than one epistemic mode, and a given epistemic mode may require the cooperation of more than one sensory modality.

[33] Concepts can also provide guidance for dealing with abstract entities, e.g. numbers. This possibility is left aside here, but it is not thereby ruled out.

[34] Note that (Goodale et al. 1994) does not establish this connection nor any degree of covariance.

[35] An account of the form concepts must take to direct both action and inference must wait for some future paper.

[36] I hinted earlier that such cooperation might be effected via the anatomic connections that exist between dorsal and ventral streams. That suggestion cannot be ruled out based on current evidence, but David Milner is doubtful, and suggests instead that the coordination between the two streams is more likely to be a function of the frontal lobe (personal communication). I am willing to place an admittedly risky bet on the *existence* of cooperation, with the function I outlined, between the two streams, but I'll make no bets as to how this cooperation is anatomically instantiated.

[37] The similarity of this example to Dr. Johnson's refutation of Berkeley is intentional. I've often thought that this refutation was unjustly maligned—for the notion that Johnson's performative argument is a non-starter depends on the assumption that the epistemic import or result of the kick is the production of more sense data. But this needn't be the case. Kicking is indeed a kind of touching, which is indeed a kind of perception; but not all perception operates in the manner supposed by Berkeley.

[38] In addition, of course, there will be unexpected qualitative experience (pain), and visual experience (the stationary ball), which each might contribute in their own way to a reconsideration of the relevant concepts.

[39] In this instance in a rather radical way, but we have all experienced cases where our expectation for the weight of something (expressed in terms of the initial force applied to lift it) had to be quickly adjusted during the act of lifting, lest we throw it through the ceiling. Such weight illusions can apparently be caused by different factors: estimates of weight based on visual perception of size (Flanagan and Beltzner 2000); estimates of weight based on haptic perception of size (Kawai 2000); and estimates of weight based on knowledge of the objects or materials to be lifted (Ellis and Lederman 1998). Interestingly, although subjects quickly adjust the forces required for lifting objects to their actual weight, when asked to verbally judge or estimate weight, the original expectations can continue to influence judgments, and the verbal estimates can prove, in the short term at least, recalcitrant to kinesthetic experience. Thus, for instance, in the classic case of the size-weight illusion, a small object is judged to feel heavier than a large object of the same weight, even when, after repeated lifting of both, identical force is in fact applied to lift each (Flanagan and Beltzner 2000). Likewise, in the "golf-ball" illusion, experienced golfers who *expect* there to be a weight difference between practice balls and regular golf-balls in fact judge practice balls to be heavier after lifting them, even though in the experiment the two kinds of balls weigh the same. Non-golfers, who have no such expectations, accurately judge the two kinds of golf-balls to weigh the same amount (Ellis and Lederman 1998). In part, these experiments further underline the distinctions between systems concerned with conceptual or semantic knowledge and expectations, and those involved in on-line behavioral control. Indeed, in a particularly striking experiment with a patient with ventral stream lesions, it was shown that the visual estimates of size processed by the dorsal stream cannot be used to produce a size-weight illusion at all (Dijkerman, et al. 2004). The experiments also show that the effect of kinesthetic feedback on semantic knowledge and conceptual content is, at least in the case of weight, not always immediate (since the illusion of a weight difference persists even when the same force is applied to lift both objects). On the other side of the coin, the golf-ball experiment *does* suggest the existence of a feedback mechanism (whatever the time scale of its operation) whereby kinesthetic experience with real and practice golf-balls leads to the semantic or conceptual knowledge of their relative weight, which in turn influences the expectations brought to bear when interacting with them, and the patients in (Dijkerman, et al. 2004)

did experience the size-weight illusion when allowed to gather haptic and kinesthetic information about the size of the object.

References

Akins, K. 1996. Of sensory systems and the "aboutness" of mental states. *Journal of Philosophy*, 93: 337-72.

Alston, W. 1997. *A Realist Conception of Truth*. Ithaca, NY: Cornell University Press.

Alston, W. 2001. *A Sensible Metaphysical Realism*. Milwaukee: Marquette University Press.

Anderson, M. L. 2003a. Embodied cognition: a field guide. *Artificial Intelligence*, 149(1): 91-130.

Anderson, M. L. 2003b. Representations, symbols, and embodiment. *Artificial Intelligence*, 149(1): 151-6.

Anderson, M. L. 2005. Representation, evolution and embodiment. In: D. Smith (ed.), *Evolutionary Biology and the Central Problems of Cognitive Science*, special issue of *Theoria et Historia Scientarum*, 9 (1).

Anderson, M. L. and Oates, T. 2003. Prelinguistic agents will form only egocentric predicates. *Behavioral and Brain Sciences* 26(3): 284-5.

Bermúdez, J. L. 1995a. Ecological perception and the notion of a non-conceptual point of view. In: J. L. Bermúdez,, A. Marcel and N. Eilan (eds) *The Body and the Self*. MIT Press. pp 153-174.

Bermúdez, J. L. 1995b. Non-conceptual content: From perceptual experience to subpersonal computational states. *Mind and Language*, 10(4): 333-69.

Bermúdez, J. L. 1998. *The Paradox of Self-Consciousness*. Cambridge, MA: MIT Press.

Bosco, G. and Poppele, R. 2001. Proprioception from a spinocerebellar perspective. *Physiological Reviews*, 81: 539-68.

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

Chalmers, D. 2003. The Matrix as metaphysics. http://whatisthematrix.warnerbros.com/rl_cmp/new_phil_fr_chalmers.html

Chrisley, R. 1995. Non-conceptual content and robotics: Taking embodiment seriously. In: K. Ford, C. Glymour, and P. Hayes (eds) *Android Epistemology*. Cambridge, MA: AAAI/MIT Press, pp 141-66.

Clark, A. 1997. *Being There: Putting Brain, Brain and World Together Again*. Cambridge, MA: MIT Press.

Clark, A. 2001. Visual experience and motor action: are the bonds too tight? *Philosophical Review*, 110.

Craig, J. and Rollman, G. 1999. Somethesis. *Annual Review of Psychology*, 50: 305-31.

Davidson, D. 1985. On the very idea of a conceptual scheme. In: *Inquiries into Truth and Interpretation*. Oxford: Clarendon Press.

Devitt, M. 1991. *Realism and Truth*, 2d ed. Oxford: Basil Blackwell.

Dijkerman, H.C., Lê, S., Démonet, J.F., and Milner, A.D. 2004. Visuomotor performance in a patient with visual agnosia due to an early lesion. *Cognitive Brain Research*, in press.

Ellis, R.R. and Lederman, S.J. 1998. The golf-ball illusion: evidence for top-down processing in weight perception. *Perception*, 27(2): 193-201.

Evans, G. 1982. *The Varieties of Reference.* Oxford: Oxford University Press.

Evans, G. 1985. Molyneux's question. In: G. Evans (ed) *The Collected Papers of Gareth Evans*. London: Oxford University Press.

Flanagan, J.R. and Beltzner, M.A. 2000. Independence of perceptual and sensorimotor predictions in the size-weight illusion. *Nature Neuroscience*, 3(7): 737-41.

Fodor, J. and Pylyshyn, Z. 2002. How direct is visual perception?: Some reflections on Gibson's "Ecological Approach". In: Nöe, A. and Thompson, E. (eds) *Vision and Mind: Selected Writings in the Philosophy of Perception*. Cambridge, MA: MIT Press. pp. 167-228.

Gallagher, S. 2003. Bodily self-awareness and object perception. *Theoria et Historia Scientarum: International Journal for Interdisciplinary Studies*, 7(1).

Gallagher, S. 1986. Body image and body schema: a conceptual clarification. *Journal of Mind and Behavior*, 7: 541-54.

Gauthier, I., James, T.W., Curby, K., and Tarr, M.J. 2003. The influence of conceptual knowledge on visual discrimination. *Cognitive Neuropsychology*, 20(3-6): 507-23.

Ghez, C., Gordon, J., and Ghilardi, M.F. 1995. Impairments of reaching movements in patients without proprioception. II. Effects of visual information on accuracy. *Journal of Neurophysiology*, 73: 361-72.

Gibson, J.J. 1966. *The Senses Considered as Perceptual Systems*. Boston: Houghton Mifflin.

Gibson, J.J. 1977. On the analysis of change in the optic array in cotemporary research in visual space and motion perception. *Scandinavian Journal of Psychology*, 18: 161-3.

Goldschalk, M., Lemon, R.N., Kuypers, H.G.J.M., and Ronday, H.K. 1984. Cortical afferents and efferents of monkey postarcuate area: an anatomical and electrophysiological study. *Experimental Brain Research*, 56: 410-24.

Goodale, M.A., Milner, A.D., Jakobson, L.S., and Carey, D.P. 1991. A neurological dissociation between perceiving objects and grasping them. *Nature*, 349: 154-6.

Goodale, M.A., Murphy, K.J., Meenan, J.-P., Racicot, C.I. and Nicolle, D.A. 1993. Spared object perception but poor object-calibrated grasping in a patient with optic ataxia. *Society for Neuroscience Abstracts*, 19: 775.

Goodale, M.A., Meenan, J.-P., Bülthoff, H.H., Nicolle, D.A., Murphy, K.J. and Racicot, C.I. 1994. Separate neural pathways for the visual analysis of object shape in perception and prehension. *Current Biology*, 4: 604-10.

Goodman, N. 1978. *Ways of Worldmaking.* Indianapolis: Hackett Publishers.

Gregory, R. 1998. Brainy mind. *British Medical Journal*, 317: 1693-5.

Gregory, R. 2002. Perceptions as hypotheses. In: A. Nöe, A. and E. Thompson, E. (eds) *Vision and Mind: Selected Writings in the Philosophy of Perception*. Cambridge, MA: MIT Press. pp. 111-34.

Gunther, Y. (ed) 2003. *Essays on Nonconceptual Content*. Cambridge MA: Bradford Books.

Hacking, I. 1983. *Representing and Intervening*. Cambridge: Cambridge University Press.

Humphrey, G.K., Goodale, M.A., Jakobson, L.S., and Servos, P. 1994. The role of surface information in object recognition: studies of a visual form agnosic and normal subjects. *Perception*, 23: 1457-81.

Hurford, J. 2003. The neural basis of predicate-argument structure. *Behavioral and Brain Sciences*.

James, T.W., Culham, J., Humphrey, G.K., Milner, A.D., and Goodale, M.A. 2003. Ventral occipital lesions impair object recognition but not object-directed grasping: an fMRI study. *Brain*, 126: 2463-75.

James, W. 1912. *Essays in Radical Empiricism*. New York: Longmans, Green.

James, W. 1981. *Pragmatism*. Indianapolis: Hackett.

Jeannerod, M. 1997. *The Cognitive Neuroscience of Action*. Oxford: Blackwell Publishers.

Kawai, S. 2002. Constant involvement of haptically perceived size in weight discrimination. *Experimental Brain Research*, 147(1): 16-22.

McDowell, J. 1994. *Mind and World*. Cambridge, MA: Harvard University Press.

Marr, D. 2002. Selections from *Vision*. In: A. Nöe, A. and E. Thompson, E. (eds) *Vision and Mind: Selected Writings in the Philosophy of Perception*. Cambridge, MA: MIT Press. pp. 229-66.

Matelli, M., Camarda, R., Glickstein M., and Rizzolatti, G. 1986. Afferent and efferent projections of the inferior area 6 in the macaque monkey. *Journal of Comparative Neurology*, 251: 281-98.

Merleau-Ponty, M. 1962. *The Phenomenology of Perception*. C. Smith (tr) London: Routledge.

Milner, A.D., Perrett, D.I., Johnston, R.S., Benson, P.J., Jordan, T.R., Heeley, D.W. et al. 1991. Perception and action in "visual form agnosia". *Brain*, 114: 405-28.

Milner, A. D. and Goodale, M. 1995. *The Visual Brain in Action*. Oxford: Oxford University Press.

Nicolelis, M., Ghazanfar, A., Stambaugh, C., Olivera, L., Laubach, M., Chapin, J., Nelson, R., and Kass, J. 1998. Simultaneous encoding of tactile information by three primate cortical areas. *Nature Neuroscience*, 1: 621-631.

Nöe, A. (ed) 2002a. *Is the Visual World a Grand Illusion?* Special issue of the *Journal of Consciousness Studies*, 9.

Nöe, A. 2002b. Is the visual world a grand illusion? In: A. Nöe (ed) *Is the Visual World a Grand Illusion?* Special issue of the *Journal of Consciousness Studies*, 9.

Nöe, A. and Thompson, E. (eds) 2002a. *Vision and Mind: Selected Writings in the Philosophy of Perception*. Cambridge, MA: MIT Press.

Nöe, A. and Thompson, E. 2002b. Introduction. In: A. Noe and E. Thompson (eds) *Vision and Mind: Selected Writings in the Philosophy of Perception*. Cambridge, MA: MIT Press.

O'Donovan-Anderson, M. 1996. Science & things: Scientific method as embodied access to the world. In: M. O'Donovan-Anderson (ed) *The Incorporates Self: Interdisciplinary Perspectives on Embodiment*. Lanham, MD: Rowman and Littlefield Publishers.

O'Donovan-Anderson, M. 1997. *Content and Comportment: On Embodiment and the Epistemic Availability of the World*. Lanham, MD: Rowman and Littlefield Publishers.

O'Donovan-Anderson, M. 2002. Understanding bodily awareness. *Proceedings of Tucson 2002:Towards a Science of Consciousness*. Tucson, AZ.

O'Regan, J.K. and Nöe, A. 2001. A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24.

O'Shaughnessy, B. 1995. Proprioception and the body image. In: J.L. Bermúdez, A. Marcel and N. Elian (eds) *The Body and the Self*. Cambridge, MA: MIT Press. pp.175-203.

O'Shaughnessy, B. 1980. *The Will: A Dual Aspect Theory*. Cambridge: Cambridge University Press.

Peacocke, C. 1998. Nonconceptual content defined. *Philosophy and Phenomenological Research*, 58:2: 381-8.

Peirce, C.S. 1955. *The Philosophical Writings of Peirce*. J. Buchler (ed) New York: Dover.

Peirce, C.S. 1958. *The Collected Papers of Charles Sanders Peirce*. Cambridge, MA: Harvard University Press.

Perenin, M.T. and Vighetto, A. 1983. Optic ataxia: a specific disorder in visuomotor coordination. In: A. Hein and M. Jeannerod (eds) *Spatially oriented behavior*. New York: Springer-Verlag.

Perenin, M.T. and Vighetto, A. 1988. Optic ataxia: a specific disruption in visuomotor mechanisms. I. Different aspects of the deficit in reaching for objects. *Brain*, 111: 643-74.

Pinker, S. 1997. *How the Mind Works*. New York: Norton.

Power, S. 2003. The AIDS rebel. *The New Yorker*, May 19: 54-67.

Putnam, H. 1981. *Reason, Truth and History*. Cambridge: Cambridge University Press.

Quine, W. V. O. 1969. *Ontological Relativity* New York: Columbia University Press.

Quine, W. V. O. 1985. Two dogmas of empiricism. In: A. Martinich (ed), *The Philosophy of Language*, pp. 26-39. Oxford: Oxford University Press.

Rizzolatti, G., Camarda, R., Fogassi, L., Gentilucci, M., Luppino, G. and Matelli, M. 1988. Functional organization of inferior area 6 in the macaque monkey. II. Area F5 and the control of distal movements. *Experimental Brain Research*, 71: 491-507.

Rorty, R. 1979. The world well lost. In: *The Consequences of Pragmatism*. Minneapolis: University of Minnesota Press.

Rorty, R. 1991. *Objectivity, Relativism and Truth*. Cambridge: Cambridge University Press.

Rosenberg, G. and Anderson, M. L. 2004. A brief introduction to the guidance theory of representation. *Proceedings of the 26th Annual Conference of the Cognitive Science Society*.

Rosenberg, G. and Anderson, M. L. *forthcoming*. Content and action: the guidance theory of representation.

Sakata, H., Taira, M., Mine, S., and Murata A. 1992. Hand-movement related neurons of the posterior parietal cortex of the monkey: their role in visual guidance of hand movements. In: R. Caminiti, P.B. Johnson and Y. Burnod (eds) *Control of Arm Movement in Space*. Berlin: Springer-Verlag.

Sanders, M.D., Warrington, E.K., Marshall, J., and Weiskrantz, L. 1974. "Blindsight": vision in a field defect. *Lancet*, 20:707-8.

Schneider, G.E. 1969. Two visual systems: brain mechanisms for localization and discrimination are dissociated by tectal and cortical lesions. *Science*, 163: 895-902.

Sommer, M. and Wurtz, R. 2002. A pathway in primate brain for internal monitoring of movements. *Science*, 296: 1480-1482.

Stroud, B. 1984. *The Significance of Philosophical Skepticism*. Oxford: Clarendon Press.

Stroud, B. 2000. *Understanding Human Knowledge: Philosophical Essays*. Oxford: Oxford University Press.

Sutton, R. and Barto, A. 1998. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press.

Taira, M., Mine, S., Georgopoulos, A.P., Mutara, A., and Sakata, H. 1990. Parietal cortex neurons of the monkey related to the visual guidance of hand movements. *Experimental Brain Research*, 83: 29-36.

Thompson, E. 1995. *Colour Vision*. New York: Routledge.

Ungerleider, L.G. and Mishkin, M. 1982. Two cortical visual systems. In: D.J. Ingle, M.A. Goodale and R.J.W. Mansfield (eds) *Analysis of Visual Behavior*. Cambridge, MA: MIT Press.

Varela, F. J., Thompson, E. and Rosch E. 1991. *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, MA: MIT Press.

Weiskrantz, L., Warrington, E.K., Sanders, M.D., and Marshall, J. 1974. Visual capacity in the hemianopic field following a restricted occipital ablation. *Brain*, 97: 709-28.

Williams, M. 1996. *Unnatural Doubts*. Princeton: Princeton University Press.

Wright, C. 1988. Realism, antirealism, irrealism, quasirealism. In: P. French, T. Uehling, and H. Wettstein (eds), *Midwest Studies in Philosophy XII: Realism and Antirealism*. Minneapolis: University of Minnesota Press.

Wright, C. 1994. *Truth and Objectivity*. Cambridge, MA: Harvard University Press.