THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Perception, action, and experience: Unraveling the golden braid

Citation for published version:

Clark, A 2009, 'Perception, action, and experience: Unraveling the golden braid', *Neuropsychologia*, vol. 47, no. 6, pp. 1460-1468. <https://doi.org/10.1016/j.neuropsychologia.2008.10.020>

Digital Object Identifier (DOI):

[10.1016/j.neuropsychologia.2008.10.020](https://doi.org/10.1016/j.neuropsychologia.2008.10.020)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Neuropsychologia

Publisher Rights Statement:

This is an author's Accepted Manuscript of the following article: Clark, A. (2009), "Perception, action, and experience: Unraveling the golden braid", *Neuropsychologia*. 47 (6) p. 1460-1468. The final publication is available at: <http://www.sciencedirect.com/science/article/pii/S0028393208004259>

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Perception, Action, and Experience: Unraveling the Golden Braid*?

Andy Clark

Abstract

Much of our human mental life looks to involve a seamless unfolding of perception, action and experience: a golden braid in which each element twines intimately with the rest. We see the very world we act in and we act in the world we see. But more than this, visual experience presents us with the world in a way apt for the control and fine guidance of action. Or so it seems. Milner and Goodale's ((1995) (2006)) influential work on the dual visual systems hypothesis casts doubt on certain versions of this intuitive vision. It does so by prising apart the twining strands of conscious visual perception and the fine control of visuomotor action. Such a bold proposal is of major interest both to cognitive science and philosophy. In what follows I first clarify the major claims that the bold proposal involves, then examine three sets of worries and objections. The first set concern some important matters of detail. The second set concern a certain kind of conceptual or philosophical worry to the effect that the perception/action model unfairly equates visual experience itself with what are in fact certain *elements* within visual experience. The third set concern

the very idea of conscious experience as a well-defined conceptual or experimental target.

Keywords: Vision, Dual Visual Systems, Consciousness, Action, Visual Experience

1. A Bold Proposal: Vision for Perception, Vision for Action.

Milner and Goodale's bold proposal offers a certain functional gloss on the anatomical complexity of the visual system. This gloss takes as its starting point the existence of two major processing streams (the ventral and the dorsal) projecting from early visual areas to the rest of the human brain. Ungerleider and Mishkin (1982) famously depicted these as the 'what' and 'where' streams, each specialized to perceive different aspects of the visual world. On this view, both streams contributed contents (though different ones) to human visual experience. Milner and Goodale (1995) described and defended an alternative functional gloss, according to which one stream supports 'vision-for-action' and the other 'vision-for-perception'. Thus the dorsal stream, projecting to the posterior parietal cortex, is said to support the kinds of visuomotor transformation in which visual input leads to fluent actions such as reaching and grasping, while the ventral stream, projecting to the temporal lobe, seems

to be especially implicated in the recognition and identification of objects and events. Milner and Goodale (1995) and Goodale and Milner (2004) further suggest that the dorsal stream computes 'vision-for-action' in a way that is fast, transient, and unconscious, and that the ventral stream computes 'vision-for-perception' in a way that is slower, more enduring, and at least sometimes sometimes conscious. The contents of conscious visual experience, for Milner and Goodale, are thus strongly associated with the coding and processing operations carried out by the ventral stream.

To illustrate the way these two streams are then meant to interact, Goodale and Milner (2004) develop an analogy with tele-assistance approaches to the control of distant robots in distant or hostile environments. Here, a conscious human operator and a non-conscious semi-intelligent distal robot combine forces so as to perform actions in some environment (for example, in the control of a Mars rover, where the human operator reviews images on a screen in Texas, flagging items of interest that the robot can locate and retrieve using its own on-board sensory systems and sensorimotor routines. Such approaches are contrasted with tele-operation solutions, in which the conscious human operator controls all the spatial and temporal aspects of the robots movements (perhaps via a joystick or a set of sensors that allow the operators own arm and hand movements to be relayed to the robot). The tele-assistance analogy thus

identifies the role of the conscious human operator with the role of the ventral stream (working in concert with stored memory and various 'executive control' systems). The task of this coalition, the analogy suggests, is to identify objects and to select types of action that are appropriate given the agent's current goals, background knowledge, and currently attended perceptual input. The task of the dorsal stream (and associated structures) is then to turn these high-level specifications into metrically accurate, egocentrically specified forms of world-engaging action. The dorsal stream (plus associated structures) thus plays the non-conscious robotic Mars Rover to the ventral coalition's conscious human operator, such that:

"Both systems have to work together in the production of purposive behavior- one system to select the goal object from the visual array, the other to carry out the required metrical computations for the goal-directed action." (Goodale and Milner (2004) p.100)

This picture of the distinct but interlocking contributions of the two visual streams explains some of the coarse dissociations and psychophysical effects described in the literature. In particular, this account of the division of labour offers a neat explanation of the differing pattern of deficits seen in patients

with lesions affecting one or other of the two streams. Thus ventrally compromised patients (such as DF) display severe impairments in recognizing shapes, orientations, and objects. Not only can DF not recognize most everyday objects, or faces, she cannot distinguish between squares, rectangles, triangles and circles. In the famous 'mailbox' task DF was pretty well at chance for turning a handheld card to match the perceived orientation of a posting slot. Yet asked to actually post the card through the slot she was almost indistinguishable from normal controls (see Milner and Goodale (1995) pp 128-133). On an intuitive model, DF's apparent lack of conscious visual awareness of features such as shape and orientation might suggest that she has (for whatever reason) simply failed to compute the very information needed to guide the relevant visuomotor actions. Yet her fluent performance belies this. The dual streams/tele-assistance model accounts for this, since the visual information required to support the world-engaging action is computed independently (though on the basis of the same retinal input) of the information required to support conscious identification¹.

Moving to the case of normal, unimpaired subjects, the same story neatly accounts for work by Aglioti et al (1995). In this work, the experimenters set up a graspable version of the famous Ebbinghaus or 'Titchener Circles' visual illusion in which two central circles are presented, surrounded by a ring of

other circles. In one case, the surrounding circles are larger than the central one. In the other, they are smaller. This leads to the well-known illusions in which subjects misjudge the relative size of the two central circles. Such mistaken estimates of relative size do not, however, affect subjects' abilities (in the physical, poker chip version) to form precision grips that perfectly anticipate the true size of the centre discs. The explanation, according to Goodale and Milner (2004, pp 88-89) is that the conscious scene is computed by the ventral stream in ways that are at liberty to make a variety of assumptions on the basis of visual cues (e.g. attempting to preserve size constancy by treating the smaller circles as probably further away than the larger ones). The dorsal stream, by contrast, uses only the kinds of information that are metrically reliable and exploit specific opportunities for elegant, fast, metrically accurate diagnosis. For example, the dorsal stream may make great use of binocular depth information. These differences in processing, combined with the quasi-independent modes of operation of the two streams, account for the illusion's ability to impact conscious visual experience while leaving our visuomotor engagements intact.

More recently, a similar effect has been shown using the so-called 'hollow face illusion'. In this illusion a concave model of a human face appears convex, due to the influence of top-down knowledge concerning normal human faces. This

suggests it is a purely ventral stream based illusion. Kroliczak et al (2006) showed that in a task where subjects were asked to flick small targets off the actually hollow (though visually convex) face, the flicking movements found the real (non-illusory) locations of the targets. According to Milner and Goodale:

"This demonstrates that the visuomotor system can use bottom-up sensory inputs...to guide movements to the veridical locations of targets in the real world, even when the perceived positions of the targets are influenced, or even reversed, by top-down processing"

Milner and Goodale (2006) p.245

Such demonstrations, and the more general issue of perception-action dissociations and visual illusions, have spawned a large and complex literature devoted to the search for counter-examples, alternative explanations, exceptions, refinements, and additional support (for some useful reviews, see Carey (2001), Clark (2001), Goodale and Westwood (2004))ⁱⁱ.

Aspects of the dual visual systems picture are supported by a large and impressively varied body of evidence including a swathe of neuroimaging experiments (e.g. Le et al (2002), James et al (2003)), a wide variety of single cell

recordings and other experimental interventions in monkeys (Taira et al (1990) Sakata (2003)), and psychophysical experiments involving normal human subjects (Bridgeman et al (1979), Goodale et al (1986), Fecteau et al (2001), and (as we just saw) Aglioti et al (1995)). More generally, it may be observed that the online control of motor action requires the extraction and use of radically different kinds of information (from the incoming visual signal) than do the tasks of recognition, recall and reasoning. The former requires a constantly updated, (multiply) egocentrically-specified, exquisitely distance and orientation sensitive encoding of the visual array. The latter requires the computation of object-constancy (objects don't change their identity every time they move in space) and the recognition of items by category and significance irrespective of the fine detail of location, viewpoint and retinal image size. A computationally efficient coding for either task thus looks to preclude the use of the very same encoding for the other (Milner and Goodale (1995) 25-66). Different uses of visual information impose quite different computational demands on the brain, so there are compelling computational and information-processing considerations that speak in favor of a dual (or at any rate, multiple) visual systems architecture. But just how all this in turn lines up with issues concerning conscious and non-conscious vision remains open to question, as we shall see.

2. Complications

The very strongest claim in this area, that the contents of conscious visual experience are *exclusively* determined by the coding and processing operations carried out by the ventral stream, now looks to be empirically suspectⁱⁱⁱ. Schenk and Milner (2006) show that form discrimination, in the visual form agnostic DF, can be improved if she engages in a concurrent visuomotor task. Specifically, when asked to name the shape of a visually presented object while reaching for the object, performance on the naming task was significantly improved. Further experiments ruled out the suggestion that this effect is due to DF using proprioceptive or efferent cues, suggesting instead that she is able directly to access shape-relevant information (in fact, it turns out to be information about width) being processed by the dorsal stream. This means that DF's shape judgements can be influenced by ongoing dorsal stream activity. Does this mean that processing in the dorsal stream can contribute to the contents of DF's visual experience? It is not clear. For the possibility remains that the dorsal influence provides only some kind of elusive, non-visual cue. As the authors put it:

"better 'discrimination' does not necessarily imply better 'perception', and instead Df could have been employing some indefinable 'implicit'

cues of the kind that enable blindsight patients to perform above chance in discrimination tasks." Schenk and Milner (2006) p. 1502

It has also been suggested (Matthen (2005)) that dorsal processing may contribute a kind of non-specific 'feeling of presence' to our contact with the visual scene. To bring this idea into focus, it helps first to distinguish, following Matthen, 'descriptive' from 'motion-guiding' vision. Descriptive vision, as Matthen uses the term, corresponds rather closely to the kind of vision supported by ventral stream coding. It is the kind of vision that supports sensory classification, allowing us to experience the scene as composed of objects and elements that might be classified as similar or different in such-and-such sensory (visual) respects. Descriptive vision is thus, as we might say, 'epistemically pregnant': it presents the visual world in ways that are apt for reasoning about that world (see Clark (2001)). Matthen contrasts descriptive vision and 'motion-guiding vision', identifying the latter as supported by both non-cortical routes and the dorsal (cortical) stream. According to Matthen (op cit p 296) "the link between motion-guiding vision and bodily motion is direct; it is not routed through consciousness". Conscious seeing may, of course, provide the information that leads us to choose a certain target and a certain type of action. But it does not provide moment-by-moment guidance of fine visuomotor action. This picture thus comports nicely with Milner and

Goodale's view concerning the non-conscious nature of the processing and control operations carried out by the dorsal stream. Nonetheless, Matthen suggests, there is at least one way in which that picture may be incomplete. For it fails to recognize the dorsal processing origins of the 'feeling of presence' that accompanies many of our visual encounters with the world. The idea here is that the kinds of metrically accurate depth and location information computed by the dorsal stream allow us to feel that we are really in the presence of, say, a cup and not merely seeing a 2D drawing or photograph of a cup. This feeling of presence (which can be at least partially duped by some kinds of two dimensional depiction) is said to form a genuine part of our visual experience. Even though it is not part of 'descriptive vision', it "makes a difference to the quality of one's visual awareness of an object" (op cit p.301).

Rizzolatti and Matelli (2003), and Gallese (2007), suggest that the dorsal/ventral distinction itself is too coarse-grained, and should be replaced by a tripartite distinction between dorso-dorsal, ventro-dorsal, and ventral streams (for the anatomical details, see Gallese (op cit) sections 2-4). The dorso-dorsal stream is said to support fast, non-conscious visuomotor transformations, just as envisaged by Milner and Goodale. While the ventral stream, again as Milner and Goodale suggest, supports object identification and classification. But the

ventro-dorsal stream, though itself anatomically part of the dorsal stream, nonetheless contributes, these authors suggest, to the determination of conscious experience. One major way it does so, it is argued, is by supporting the conscious visual experience of others' intentional actions. Such a suggestion gains plausibility if one way in which we come to grasp the nature and goals of others' observed actions is by activating our own motor systems, perhaps under the influence of mirror system (Rizzolatti et al (1996)) processing. The ventro-dorsal stream is also said to be essential for the conscious visual awareness of space as an arena for our own motor activity, as when we see objects as reachable, or as manipulable in such and such ways. Holding both these strands together is the guiding idea that parieto pre-motor circuits support various forms of 'embodied simulation' (Gallese (2005)) such that "side-by-side with the sensory description of the observed phenomena, the motor schemata associated with these actions, or objects, are evoked in the observer" (Gallese (2007) p.14). These active motor schemas are said to be the mechanism by which "conscious awareness of actions and spatial locations are generated" (op cit).

In this section we have scouted what are perhaps best seen as a series of complications: refinements and additions to the strong dual visual systems view. Though outstanding questions remain^{iv}, accommodating these

refinements and additions leaves intact many of the central tenets of the account. Other challenges, however, go deeper than this, some of them threatening even the claims concerning the non-conscious status of the basic (dorso-dorsal) visuomotor transformation itself. It is to these challenges that we now turn.

3. The 'Narrow Vision of Conscious Vision' Worry

Milner and Goodale's bold proposal is, as we have already begun to see, subject to an interesting and important (though initially somewhat elusive) kind of worry. We can dub it the 'Narrow Vision of Conscious Vision' worry. The worry takes many shapes and forms, some of them more plausible, and some of them less elusive, (and some both), than others. What they all have in common is the thought that the bold proposal illicitly identifies conscious visual experience with one or more of the components, forms, or styles, of conscious visual experience. That is to say, conscious visual experience may involve many different 'things', supported by many different processing streams and neural coalitions, and it would be a mistake to identify conscious vision with just one of them.

An analogy^v may be helpful here. Visual and auditory experience are different phenomena, and depend on different mechanisms and forms of internal processing. In the light of this mundane fact, we would be wrong to identify what we perceptually experience (in general) with what we *visually* experience. Correlatively, we would be wrong to simply identify the mechanisms of perceptual experience with the mechanisms of visual experience. Now suppose that within the realm of visual experience itself there exist multiple quite different varieties of experiential phenomena, each supported by different types or forms of mechanism and processing. Suppose too that one or more such elements turns out to depend preferentially, just as Milner and Goodale suggest, upon processing in the ventral visual stream. Still it would not follow that conscious visual experience itself (in all its varieties) depends preferentially upon processing in that stream. To do so would be to fall prey to what I am calling a "Narrow Vision of Conscious Vision" error.

Here are two (distinct but related) versions of the 'narrow vision' worry drawn from the recent philosophical and cognitive neuroscientific literature:

3.1 Comparing Like with Like

Jeannerod and Jacob (2005) make a strong case that simple comparisons between the contributions of the dorsal and ventral streams are misleading insofar as they fail to control for three factors. The factors concern direction of fit, direction of causation, and level of processing. Let's take the first two first. The underlying model here is the philosophical distinction (Anscombe (1957), Searle (1983)) between states that are belief-like and states that are desire-like. Beliefs have a mind-to-world direction of fit, and a world-to-mind direction of causation. That is, a belief is a mental state that aims to fit the way the world is (thus, it exhibits a mind-to-world direction of fit). One good way to ensure that kind of fit is for it to be *caused* by the way the world is (thus also exhibiting world-to-mind direction of causation). Desires, by contrast describe ways the world should be (they exhibit world-to-mind direction of fit), and they may function so as to help bring it about that the world be that way (thus displaying a mind-to-world direction of causation). Notice, then, that different attitudes (belief versus desire) are here distinguished in part in terms of their direction of fit. Thus an agent may be said to be in states that share a content (eg that the shops be open) with differences in attitude (one may believe it, another hope it, another fear it).

Apart from this question about direction of it, there is a question concerning nature and level of processing. Thus consider next the various ways an active agent may need to use visually presented information. The information may be used to ensure that they know how things are out in the world. Jeannerod and Jacob (op cit p.3) call this the 'semantic processing of visual information' Or (or in addition) it may be used to help act upon, and alter, that world. They call this the 'pragmatic processing of visual information'. Semantic processing has a belief-like profile, exhibiting a mind-to-world direction of fit and a world-to-mind direction of causation. Pragmatic processing has a more desire-like profile^{vi}, exhibiting a world-to-mind direction of fit and a mind-to-world direction of causation. Moreover, each kind of processing involves many levels of abstraction. Early stages of semantic processing yield representations with 'pictorial non-conceptual contents' (for present purposes, we can think of these as essentially 'iconic', non-propositionally structured contents, in the sense of Coltheart (1983), but see Clark (2001) for further discussion). Later stages yield representations whose contents are more structured and highly processed, presenting a world of objects, relations, and attributes, in ways apt to inform processes of explicit reasoning and planning. Similarly, early stages of pragmatic processing are said to represent basic geometrical properties (width, as it figures in the Schenk and Milner results discussed in section 2, would be one such property) of objects in ways apt for the guidance of actions such as fluent

reaching and grasping, while later stages yield more abstract, or 'conceptual' representations. Crucially, according to Jeannerod and Jacob:

"The scope of pragmatic processing... is not limited to the visuomotor transformation, since pragmatic processing is involved in conceptually more complex operations like evaluating the feasibility of an action, anticipating its consequences, planning further steps, and learning the skilled use of tools by observation"

Jeannerod and Jacob (2005) p 5

The upshot is that:

"The visuomotor transformation is but a first, lower level component, of the human 'pragmatic processing' of objects. [We should] contrast this lower level pragmatic processing with a higher level of pragmatic processing of objects involved in the skilled use and manipulation of complex cultural tools and artifacts"

Jacob and Jeannerod (2003) xviii

These higher levels of pragmatic processing reach, the authors argue, all the way up to representations with consciously accessible contents, as in the case

where we experience motor imagery involving the manipulation of these complex tools and artifacts. Such uses go well beyond simple grasping and involve the retrieval of what Jacob and Jeannerod ((2003) p. 216) call 'stored scripts' for the manipulation and use of cultural objects. At this point, the authors argue "the distinction between action and perception loses much of its significance" (Jacob and Jeannerod (2003) p 253).

Let's return now to Milner and Goodale's bold proposal. The deepest worry raised by Jeannerod and Jacob may be put like this: we should be wary of conclusions (concerning the functional roles of the dorsal and ventral streams, and the alignments of those roles with conscious visual experience) that do not compare like with like. In particular, we should be wary of conclusions reached by comparing *early levels of pragmatic processing* with *later levels of semantic processing*. Early stages of pragmatic processing (those devoted to the basic visuomotor transform) involve representations with little or no conceptual content, and these are indeed not the kind of content that normally figure in our own experience. But early levels of semantic processing, which fall short of identifying or classifying objects, are similarly silent, phenomenally speaking. While later stages of both semantic and pragmatic processing involve contents that are not silent: that are present to phenomenal awareness. If we keep the nature and levels of processing matched, what we find are thus not differences

in phenomenal status but rather differences in direction of causation and of fit. That is to say, we find that there are indeed (at least^{vii}) two distinct roles that visual information may be serving in the agent economy. But neither role lines up exclusively with conscious visual experience.

3.2 Directive Content?

But what about the basic visuomotor transformation itself? Here, Jeannerod and Jacob seem to accept that the processing (that involved in e.g. transforming visual input into fluent acts of basic reaching and grasping) fails to reach or inform conscious visual experience. But even this may be called into question.

Thus consider once again the case of DF, the ventrally compromised carbon monoxide poisoned subject studied by Milner and Goodale. In an interesting paper, Wallhagen (2007) suggests that DF may really experience visually presented shape, but be unable to report that experience due to some problem with bringing her experience 'under concepts'. The idea is that the intact dorsal stream processing associated with the basic visuomotor transformation has its own attendant phenomenology, and that it is this attendant phenomenology

that explains DF's remarkably fluent visuomotor behaviour. Thus DF sees (visually experiences) shapes and orientations, but when asked to *report* on what she sees, she is unable (due to her compromised ventral processing) to do so. Not only can she not report what she sees, in a real sense she does not *know* what she sees. Just as a non-human animal, lacking the concept of 'chair', might well visually experience a chair yet not in any sense know it is experiencing a chair, so too DF might experience a world of oriented lines, shapes and forms but not know (neither recognize nor be able to report) that she is doing so. According to Wallhagen:

"Aspects of form may well be phenomenally present to D.F., she may well consciously sense, and hence represent, the shapes, sizes, and orientations of things, even if she cannot properly conceptualize these aspects of form" Wallhagen (2007)

O'Regan and Noë (2001) seem tempted by a similar thought. They describe DF as a case of 'partial awareness' in which 'she is unable to describe what she sees but is otherwise able to use it for the purpose of guiding action' (op cit p.969).

These views are potentially more radical than that defended by Jeannerod and Jacob, since they call into question the identification (accepted by Jeannerod

and Jacob) of conscious perceptual contents with the more highly conceptualized products of later stages of (both semantic and pragmatic) processing. One immediate problem is that these less highly processed (more 'unconceptualized') contents, if they are indeed consciously experienced, will nonetheless be hard if not impossible to bring into focus for report and description. For whenever we do so, we in effect move up the processing ladder, calling on our grasp of the scene as a structured, attended entity populated by nameable objects, shapes, and relations. But for the moment, let's bracket that worry and try to put a little more flesh on the suggestion itself.

One way to think about this kind of proposal is further developed in Nudds (2007). Nudds, like Jeannerod and Jacob, urges us to distinguish between two kinds of content that visual perceptual experiences may possess. He dubs these the 'presentational' and the 'directive' contents of visual experience. Presentational contents correspond to what Matthen (2005) termed 'descriptive vision' and what Jeannerod and Jacob (2005) termed 'semantic processing': they are contents that depict how things are in the world. Directive contents correspond to what Matthen termed 'motion-guiding vision' and what Jeannerod and Jacob called 'pragmatic processing': they are concerned with how to guide actions so as to bring about desired results. Nudds suggestion, rather like that of Wallhagen, is that these contents figure in visual experience

and help to determine the shape and nature of our world-engaging activity. It is this latter claim that marks the point of departure from both Milner and Goodale's model and (at least for the basic visuomotor transformation) that of Jeannerod and Jacob)^{viii}.

Why should we believe that directive contents, associated with basic visuomotor transformations, form part of our conscious visual experience? Nudds reasoning is best appreciated in contrast to a more standard model that depicts consciously visually guided action as guided by conscious intentions (eg Peacocke (1992)). On this model visual experience presents a world relative to which I may form an intention (eg to pick up the coffee cup I see over there) that then determines the action. Such a model is consistent with (though it does not commit you to) the tele-assistance image favoured by Milner and Goodale (see section 1 above). For it leaves room for the detailed kinematics of the action to be determined by something other than the content of the visual experience (which need only allow us to form apt intentions). But such a model, Nudds argues, is inadequate to explain visually-guided action. Instead, Nudds argues that there is a distinctive kinematics (as Milner and Goodale showed) to reaching and grasping performed while in visual contact with the scene. Remove the contact and the kinematics alter. This alteration -and this is crucial for his argument- is not exhaustively accounted for by an agent's

intentions (eg to pick up the cup, or even, moving to the higher levels of pragmatic processing, to use the screwdriver so as to screw in the screw), which may often be the same in both cases. But whereas Milner and Goodale depict these further precise kinematic details as determined by a non-conscious stream of low-level pragmatic processing, Nudds thinks this fails fully to capture the phenomenon. The reason he gives is that the precise way we perform the action seems to be something for which *we* (the agent, rather than just some subsystem of the agent's brain) are responsible. The precise way the action unfolds is, he suggests, something *we* do. That I tie my shoelaces *like this* is not something that just happens to me, or that I just find myself doing. The explanation, Nudds claims, is that (e.g.) the detailed lace-tying kinematics are driven by conscious but directive contents actually given in visual experience. Nonetheless we will not necessarily be well-placed to *report* on those contents, and it will not 'seem like' anything very specific to be guided by them. This is because both reporting and (more generally) knowing what we are experiencing depend on content-monitoring capacities informed by the other ('presentational') dimensions of conscious visual content.

This is perhaps the most difficult suggestion we will examine in the present treatment. On this model, DF (to take one striking example) may have visual experiences with rich conscious directive contents that help explain her

successful and often self-initiated behaviors^{ix} even though [she] “will not be under the impression that anything is any way, nor have any basis for judging that anything is any way” (Nudds (2007)). This also opens up an alternative interpretation of the Aglioti et al experiments rehearsed in section 1. For on Nudds' account visual experience guides *both* the action (the well-calibrated grasp) and the illusion-prone verbal response, but it is the directive content of the experience that guides the visuomotor action and the presentational or descriptive content that guides the verbal response. Both the judgment and the action are thus guided by (different forms of) visual experience^x. Moreover there is no inconsistency in the content of visual experience here, since these different kinds of content share no 'common currency' in which to frame a disagreement.

4. The Argument from Agency

In this section I aim to challenge the claim that these directive contents (both in DF and in neurally uncompromised subjects) are properly depicted as forming part of our conscious visual experience. As already noted, however, we cannot challenge such a claim simply by pointing to the agent's honest reports, since these will only reveal what is present or absent to descriptive vision. What

we can do, however, is attempt to gain some leverage from a simple, but I think plausible, observation. The observation (which may of course be doubted, as we'll later see) is that conscious experience must always be the experience *of* some agent. The putative directive contents of visual experience, being available only for one kind of purpose, and then only in highly transient form (that is, only while in actual contact with the visual scene), seem to fail to meet this requirement. They are more like encapsulated pockets of processing than genuine contents of visual (or any) experience. In very much this vein Gareth Evans once argued that an informational state may underpin a conscious experience only if it (the informational state) is in some sense input to a reasoning subject. According to Evans, to count as a conscious experience an informational state must:

"[serve] as the input to a thinking, concept-applying and reasoning system: so that the subject's thoughts, plans, and deliberations are also systematically dependent on the informational properties of the input. When there is such a link we can say that the person, rather than some part of his or her brain, receives and processes the information"

Evans (1982) p.158

I think this is almost right. But the real point here is (or should be) quite independent of Evan's appeal to the subject as concept-using. What matters, rather, is that the information must be available to the agent qua 'reasoning subject', where this may be unpacked in many different ways, not all of them requiring full-blown concept-use on the part of the agent (see e.g. Hurley (1997), Bermudez and McPherson (1998)). As long as an animal can form (nonconceptualized) goals, and can become directly and non-inferentially aware of specific environmental opportunities that allow, or that block, the fulfillment of those (limited) goals and projects, then transduced information can be, or fail to be, input to this kind of minimal reasoning subject (for more on this slightly less demanding conception, see Dretske (2006), Hurley (1998), Clark (2007)). A positive suggestion^{xi} thus emerges according to which conscious perceptual experience occurs when, and only when, information is poised, however briefly, for direct and non-inferential use in the guidance of (at least minimal) rational action. By contrast, the sensory transduction of visual information can sometimes (as in the case of blindsight) simply channel information so as to guide response, without providing the agent herself with any reasons, justifications, or rationales, for her action (see Campbell (2002), Dretske (2006)). In such cases, behaviour whose success depends on that very information will (all other things being equal) surprise the agent herself. In such cases, information is transduced, impacts behaviour, but is never poised so as

to provide me with a reason (visible to myself as an agent) for my actions or choices. At other times, however, something about the form or nature of the processing poises specific items of transduced information in a way apt (if my attention is so directed) to provide me with reasons or motivations for my own actions and choices (to provide what Dretske (2006, p.168) calls 'justifying reasons' for my actions. Importantly, it may be (see discussion below) that even elements that we don't yet attend to or notice in our own visual experience can be thus poised, as might elements that are not yet fully conceptualized (in this way, as Fodor (2007) notes, even the contents of 'iconic' encodings may provide justifying reasons for a subsequent act or judgement).

For present purposes, what matters is that even these kinds of (weak but important) link between experience and agency do not seem to be present in the case of (merely) directive contents. Nudds himself comments, as we saw, that "since directive content doesn't present an object as being some way, in having a visual experience with directive content the subject will not be under the impression that anything is any way, nor have any basis for judging that anything is any way" (Nudds (2007)). As a result, this information will not be available for use in any form of practical reasoning, and will not provide the agent herself with any reasons (visible to herself) for her own actions. But such total fractionation, of the putative conscious content from what the agent

knows or (more generally) has reason to do, seems in conflict (I suggest) with the image of genuine experience being experience *of* the agent.

5. The Most Radical Challenges

Block (2007) asks a question that reaches to the very roots of current scientific attempts to study conscious experience. He asks how we can know whether some content actually forms part of our conscious experience. The answer, he notes, is only straightforward as long as we take reportability of content (for example, the broadcasting of content to a 'global workspace'^{xii}) as a requirement on conscious presence.

But suppose we don't. Suppose we ask, instead, whether conscious experience might go beyond that which we can access or report? A natural worry is that such a question is simply unanswerable. Block's response to this (in part) is to appeal to work by Sperling (1960) and Landman et al (2003), arguing that these studies suggest that "in a certain sense phenomenal consciousness overflows cognitive accessibility" (Block (2007) p.481). In the Sperling experiments^{xiii} subjects are briefly (50ms) shown a 3x3 grid of letters, such as:

T D A

S R N

F Z B

The stimulus grid is removed after 50 ms. Many subjects then claim to have had and to have briefly retained a conscious experience of all the letters, even though they could reliably report (in the so-called ‘full report’ condition) only about 4 of them. It may seem that there is no way to empirically test whether the subjects actually *saw* (properly saw, in distinct individual detail) more letters than they can *report*. But Sperling then tested subjects in a further ‘partial report’ condition. This showed that if rapidly asked instead for the letters in any given row (the top, middle, or bottom) subjects could quite accurately respond, regardless of which row was chosen. What this (at least taken in the context of the subjects’ own experiential reports) suggests to Block (and see also Dretske (2006), Fodor (2007)) is that detailed and consciously encoded information about each and every letter was temporarily available to drive report and noticing (if attention were rapidly to be so directed) even though the subsequent selection of some letters to thus report renders the rest unavailable. That is to say, the initial experience contained more phenomenal information than (perhaps for reasons having to do with limitations on working memory) any full report can subsequently display.

In more recent studies (Landman et al (2003)) subjects were shown 8 oriented rectangles for half a second, then a gray screen, then the array of 8 but were informed that one rectangle may have changed orientation. Subjects were able to keep track of the orientation of about four rectangles from the group of eight (so their 'capacity measure' was 4) . Yet, much like the Sperling subjects, many of them reported seeing the specific orientation of all eight rectangles. Once again, a partial report condition seems to bear out the subject's claim. If the experimenter adds a pointer on the gray screen to ask the orientation question of any given rectangle, subjects can track almost all the rectangles (they display a capacity measure up to 6 or 7). The explanation once again, according to Block, is that the initial phenomenal experience contains more information than the full report condition can display^{xiv}.

With these results in mind, Block's (2007) strategy is then to display a neuroscientific story concerning strong back-of-the-head neural coalitions (involving pockets of recurrent processing - see Lamme (2006)) that nonetheless just fail to win a winner-take-all competition for broadcast to the 'global workspace' (recall note 11) and hence for reportability. Such a story (which I shall not attempt properly to rehearse here) is meant to make sense of the claim that in these (and in many other) cases phenomenal consciousness

'overflows' cognitive accessibility and thus that we can (and do) have experiences even in cases where we lack the kind of access that would yield some form of report that such and such an experience had occurred.

The point I want to notice is that this argument (which is actually a form of inference to the best explanation) takes as its starting point the assertion that the *only grounds* we have for treating the just-losing coalitions as non-conscious is the unreportability of the putative perceptual experiences. But perhaps this is premature. For underlying the appeal to reportability is, I suggest, a deeper and perhaps more compelling access-oriented concern. It is the concern, raised in the previous section, that any putative conscious experience should be the experience *of an agent*.

Can we really make sense of the image of free-floating experiences, of little isolated islets of experience that are not even potentially more widely available to act as fodder for a creature's rational choices and considered actions? Evans' insight was that the notions of conscious experience and reasoned agency (here very broadly construed) are deeply intertwined: that there are non-negotiable links between what is given in conscious awareness and the enabled sweep of deliberate actions and choices available to a reasoning subject. One way to begin to flesh this out (see e.g. Dretske (2006)) is, as we saw, to depict

conscious perceptual experience as providing an agent with self-transparent reasons for her own actions. Such a story opens up a different way of interpreting the Sperling and the Landman et al results. In these cases subjects report phenomenally registering all the items because information concerning each item was, at that moment, available to be deployed in the service of deliberate, reasoned, goal-directed action. Responses selected on the basis of this information would meet the key condition of being self-transparently grounded in the agent's perceptual connection to the world (for more on this idea, see Clark (2001)). Such momentary potentiality is not undermined by the (interesting and important) fact that the selection of a few items to actually play that role then precludes the selection of the rest. Contrariwise, Block argues that a subject such as GK (Rees et al (2000) (2002)) suffering from visuo-spatial extinction may be having an experience of a face and yet it be impossible for him, qua agent, to know anything of this experience. This is because Block takes GK's phenomenal experience to consist in recurrent processing in the fusiform face area. My suggestion, following Evans' would be that GK can be consciously experiencing a face only if the information given in the putative experience is at least momentarily poised in a way that makes it apt for use (though it need not actually be used) in the agent's personal level reasoning, planning and for the deliberate and goal-driven selection action. In that way the link to agency is maintained. Recurrent processing in the fusiform

area will no doubt prove to be among the many conditions necessary (but not sufficient) for realizing a state that plays this distinctive causal role. Block's just-losing coalitions fail to trigger winning frontal coalitions and hence fail to be in a position to contribute their contents in this manner to the full sweep of the agent's deliberate acts and choices. It is this fact (rather than the related but admittedly more superficial and unreliable indicator of mere non-reportability) that should motivate our treating the contents of the just-losing coalitions as non-conscious.

Perhaps, though, the very idea of clean facts hereabouts (facts concerning what does and does not form part of our current perceptual experience) is itself a deeply mistaken one. For it increasingly seems (and see Schwitzgebel (2008) for some compelling arguments to this effect) our grip on what it *means* for something to form part of our current conscious experience is tenuous at best. It is a grip compromised by our own congenital inability to know what we are experiencing without turning attention to it, or attempting to recollect it at a later moment, or introspecting upon it right now. Each such act alters the set of cognitive mechanisms in play, yet to eschew reliance on such methods *tout court* is to leave us with no anchor points at all. It leaves us with no means by which to decide, for example, what subsets of behaviour and response to look for in other (e.g. non-verbal, non-human, or impaired) cases. In other words, it

seems we must make some antecedent decisions to get the experimental balls rolling, but that these decisions themselves cannot be checked because the experiments that do so must be interpreted according to some closely related, equally unverified, set of assumptions. From this one might conclude that there are facts here that are terminally resistant to scientific resolution. But a better conclusion, it seems to me, is that there are no such finer-grained facts here at all^{xv}.

Thus take, for example, the currently extremely 'live' question of the phenomenal status of currently unattended visual stuff. Is such unattended stuff phenomenally experienced at the time, or is it at best poised to feed phenomenal experience (or perhaps merely to inform memory) at some later moment? Wright (2006) argues, convincingly it seems to me, that this question is unanswerable in the present state of the science. But more importantly, he also suggests that the question itself is relatively scientifically uninteresting. What matters, Wright argues, is getting a firm grip on what contribution is made by various systems and sub-systems, and how those contributions enable us to maintain fluent contact with, and interact successfully with, the world. In getting such a grip, we are not forced to resolve questions concerning the phenomenal status of unattended visual stuff.

The final view that I want to display, then, is one that we may call the Mere Motley model of conscious perceptual experience. According to this model the phrase 'conscious visual experience' is just a rough and ready label for a typically integrated, but potentially highly dissociable, complex of capacities. Some of these involve recall and report, some involve attention and noticing, others (if Block and Lamme are right) involve only various forms of recurrent-processing amplified neural activity. Such a model would be an instance of what Sloman (2007) calls a "labyrinthine" theory according to which visual experience is itself highly structured and multiply layered, such that different combinations of the many bits of the labyrinth determine different (often dissociable) aspects and nuances of what we have come to think of as 'our visual experience'. Much the same picture, again based upon theoretical apparatus and insights from work in artificial intelligence and robotics, is endorsed in Ballard (2007). Conscious visual experience, if such views are correct, is not usefully understood via the metaphor of a single inner light that is either on or off (compare: one leading voice) but consists instead in a motley swathe of surprisingly dissociable elements and effects, relative to which pressing the simple binary question ("is conscious visual experience occurring or not?") is just a recipe for trouble and confusion. The most famous defence of such a view is probably Dennett's (1991) 'multiple drafts' model of conscious experience, according to which the only real facts hereabouts concern the ways

the system would respond to various kinds of probe made at various points in the ongoing cycle of processing. But the essential core of the view (which I take to be the assertion of motley processing with no simple facts of the matter concerning the presence or absence of conscious experience in many cases) may be developed in many different ways.

7. Conclusions: Still Revealing After All These Years?

In the light of all this, what is most clearly right and important in the strong dual visual systems model is the claim that ventral stream processing (along with some of the highest levels of what Jeannerod and Jacob call 'pragmatic processing) preferentially determine what might be called our 'reflective take on our own visual experience'. That is to say, such processing is preferentially involved in the way our own visual experience presents itself *to us*: to rational, reflective agents motivated by a variety of plans, goals, and projects. But step outside that self-reflective arena, however, and the landscape changes dramatically. If experience, or some varieties of experience, outrun report, and dissociate from processes of top-down attention and consolidation into agent-memory (perhaps occurring simply courtesy of recurrent processing in encapsulated pockets of the cognitive economy) then all bets are off. Worse still, we may be forced to embrace what I have called a Mere Motley model of

conscious experience, according to which there is simply no answer to questions concerning (for example) the phenomenal status ('seen or unseen') of what Wright (2006) calls 'unattended visual stuff'. Instead, 'visual experience' would depend on a messy, multi-faceted web of processing that links us variously to the world: a web of processing that we probe in various ways and on various time-scales, some of which inevitably recruit processing in ways tied up with report, memory, attention, and noticing, and others of which do not.

Milner and Goodale have done more than just about anyone else to bring these foundational issues into focus, involving neuroscientists, cognitive psychologists, AI researchers, philosophers, and many others in what has become one of the most exciting, important, and productive debates in recent decades. On those increasingly elusive questions concerning the nature and neural underpinnings of conscious visual experience itself, the jury (it seems to me) remains out. But whatever the outcome, there is no doubting the value and impact of the dual visual systems model itself: still revealing after all these years.

* Thanks to Robert McIntosh, Tim Bayne, Thomas Schenk, Michael Tye, Matthew Nudds, Mog Stapleton, Julian Kiverstein, Simon Garrod, Steve

Draper, Liz Irvine, John Henderson, Ned Block, and Tillmann Vierkant for stimulating exchanges on these topics. Thanks too to the two referees, to all the participants in the *Perception, Action and Consciousness: Sensorimotor Dynamics and Dual Vision* conference held at Bristol in July 2007, and to all the members of the Edinburgh Philosophy, Psychology and Informatics Reading Group for stimulating discussions. This paper was prepared thanks to support from the AHRC, under the ESF Eurocores CNCC scheme, for the CONTACT (Consciousness in Interaction) project, AH/E511139/1.

References

Aglioti, S., Goodale, M. and DeSouza, J.F.X. (1995) 'Size contrast illusions deceive the eye but not the hand'. *Curr. Biol.*, 5, 679-85.

Anscombe, G. E. (1957). *Intention*. Oxford: Blackwell.

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

Baars, B.J. (1997). *In the Theater of Consciousness: The Workspace of*

the Mind. Oxford University Press.

Bach, Kent (1978) "A Representational Theory of Action," *Philosophical Studies* 34:361-79.

Ballard, D (2007) 'Putting Consciousness in its Place' presented at the 'Perception, Action and Consciousness: Sensorimotor Dynamics and Dual Vision' Conference, Bristol July 2007

Bayne, T (ms) "Deconstructing the Global Workspace"

Bermúdez, J and Macpherson, F (1998) Nonconceptual content and the nature of perceptual experience. *Electronic Journal of Analytical Philosophy* 6
Archived at: <http://ejap.louisiana.edu/archives.html>

Block, N (2007) Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences* 30(5-6): 481-499.

Bridgeman, B., Lewis, S., Heit, G., and Nagle, M. (1979) 'Relation between cognitive and motor-oriented systems of visual position perception' *Journal of Experimental Psychology (Human Perception)* 5:692-700

Byrne, A., Hilbert, D., and Siegel, S. (2008) Do we see more than we can access? Commentary on Block, "Consciousness, accessibility, and the mesh between psychology and neuroscience", *Behavioral and Brain Sciences* (2008)

Campbell, J. [2002]: *Reference and Consciousness*, Oxford: Oxford University Press.

Carey, D. (2001) "Do Action Systems Resist Visual Illusions?" *Trends in Cognitive Sciences* 5:3: 109-113

Clark, A (2001) 'Visual Experience and Motor Action: Are The Bonds Too Tight?' *Philosophical Review* 110:4: 495-519

Clark, A (2007) "What Reaching Teaches: Consciousness, Control, and the Inner Zombie" *British Journal for the Philosophy of Science* 58: 563-594

Coltheart, M. (1983). Iconic memory. *Philosophical Transactions of the Royal Society, B302*, 283-294.

Dehaene, S. & Naccache, L. (2001). Towards a Cognitive Neuroscience of Consciousness: Basic Evidence and a Workspace Framework. *Cognition* 79, 1–37.

Dehaene, S., Changeux, J.P., Naccache, L., Sackur, J. and Sergent, C (2006) Conscious, preconscious, and subliminal processing: a testable taxonomy *Trends in Cognitive Sciences* 10:5: 204-210

Dehaene, S., Kerszberg, M. & Changeux, J.-P. (1998). A Neuronal Model of a Global Workspace in Effortful Cognitive Tasks. Proceedings of the National Academy of Science 95, 14529–14534.

Dennett, D 1991 *Consciousness Explained* New York, Little Brown

Searle, J. R. (1983). *Intentionality, an essay in the philosophy of mind*. Cambridge: Cambridge University Press.

Dennett, D. (2001). Are We Explaining Consciousness Yet? *Cognition* 79, 221–

Dretske, F. [2006]: ‘Perception without Awareness’, in T. Gendler and J. Hawthorne (eds), *Perceptual Experience*, New York: Oxford University Press, pp. 147 – 80.

Dyde, R. T. and Milner, A. D. (2002): ‘Two Illusions of Perceived Orientation: One Fools All of the People Some of the Time; the other Fools All of the People All of the Time’, *Experimental Brain Research*, 144, pp. 518 – 27.

Evans, G. (1982) *The Varieties of Reference*. J. McDowell (ed.), Oxford: Oxford University Press.

Fecteau JH, Chua R, Franks I, Enns JT (2001) Visual awareness and the on-line modification of action. *Canadian Journal of Experimental Psychology* 55:104-110.

Fodor, J (1983) *The Modularity of Mind* (MIT Press, Camb. MA)

Fodor, Jerry A. (2007). The revenge of the given. In Brian P. McLaughlin & Jonathan D. Cohen (eds.), *Contemporary Debates in Philosophy of Mind*. Blackwell, Oxford

Gallese, V (2007) The “Conscious” Dorsal Stream: Embodied Simulation and its Role in Space and Action Conscious Awareness *Psyche* 13/1 (archived electronic journal: <http://psyche.cs.monash.edu.au/>)

Gallese, V. (2005). Embodied simulation: from neurons to phenomenal experience. *Phenomenology and the Cognitive Sciences*, 4: 23-48.

Gazzaniga, M. 1998: *The Mind's Past*. Berkeley: University of California Press.

Sperling, G. (1960) The information available in brief visual presentations. *Psychological Monographs* 74(498): whole issue.

Goodale MA, Pelisson D, Prablanc C (1986) "Large adjustments in visually guided reaching do not depend on vision of the hand or perception of target displacement". *Nature* 320:748-750.

Goodale, M and Milner, D (2004) *Sight Unseen: An Exploration of Conscious and Unconscious Vision* (Oxford University Press, Oxford)

Goodale, M and Westwood, D (2004) An Evolving View of Duplex Vision: Separate but Interacting Cortical Pathways for Perception and Action *Current Opinion in Neurobiology* 14: 203-221

He, S. et al. (1996) Attentional resolution and the locus of visual awareness. *Nature* 383, 334–337

Hurley, S. 1998: *Consciousness in Action*. Cambridge, Mass.: Harvard University Press.

Hurley, S.L (1997) Non-Conceptual Self-Consciousness and Agency: Perspective and Access. *Communication and Cognition* Vol. 30, (Part 1 of Special Issue: Approaching Consciousness) Nr. 3/4 207-248

Hurley, S.L (1997) Non-Conceptual Self-Consciousness and Agency: Perspective and Access. *Communication and Cognition* Vol. 30, (Part 1 of Special Issue: *Approaching Consciousness*) Nr. 3/4 207-248

James TW, Culham J, Humphrey GK, Milner AD, and Goodale MA (2003)
"Ventral occipital lesions impair object recognition but not object-directed grasping: an fMRI study" *Brain* 126: 2463-2475.

Jeannerod, M. and Jacob, P. [2005]: 'Visual Cognition. A New Look at the Two Visual Systems Model', *Neuropsychologia*, 43, pp. 301 – 12.

Jiang, Y. et al. (2006) A gender- and sexual orientation-dependent spatial attentional effect of invisible images. *Proc. Natl. Acad. Sci. U.S.A.*, DOI: 10.1073/pnas.0605678103 (<http://www.pnas.org>)

Kelly, S (ms) "Merleau-Ponty on the Body: The Logic of Motor Intentionality"

Kent, B (1978), "A Representational Theory of Action," *Philosophical Studies* 34:361-79.

Koch, C and Tsuchiya, N (2006) Attention and consciousness: two distinct brain processes *Trends in Cognitive Sciences* 11: 1: 16-21

Kouider, S., Dehaene, S., Jobert, A. & Le Bihan, D. (2006) Cerebral bases of subliminal and supraliminal priming during reading. *Cerebral Cortex*.

Kroliczak, G, Heard, P.F., Goodale, M.A. and Gregory R.L. (2006) Dissociation of perception and action unmasked by the hollow-face illusion
Brain Research 1080, 1, 1-16

Lamme, V (2006) Towards a true neural stance on consciousness *Trends in Cognitive Sciences* 10: 11 : 494-500

Landman, R., Spekreijse, H. & Lamme, V. A. F. (2003) Large capacity storage of integrated objects before change blindness. *Vision Research* 43:149–64.

Le S, Cardebat D, Boulanouar K, Henaff MA, Michel F, Milner D, Dijkerman C, Puel M, Demonet JF (2002) Seeing, since childhood, without ventral stream: a behavioural study. *Brain* 125:58-74.

Legrand, D. (2006). The Bodily Self: The Sensori-motor Roots of Pre- reflexive Self-Consciousness. *Phenomenology and the Cognitive Sciences* 5, 89–118.

Li, F.F. et al. (2002) Rapid natural scene categorization in the near absence of attention. *Proc. Natl. Acad. Sci. U. S. A.* 99, 9596–9601

Marcel, A.J. (1983) Conscious and unconscious perception: An approach to the relations between phenomenal experience and perceptual processes. *Cognitive Psychology*, **15**, 238-300.

Matthen, M. (2005). *Seeing, Doing and Knowing*, Oxford: Oxford University Press

McIntosh RD, Dijkerman HC, Mon-Williams M, and Milner AD (2004) Grasping what is graspable: evidence from visual form agnosia. *Cortex* 40 (4-5):695-702

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

Milner, A and Goodale, M (2006) "Epilogue: twelve years on" in Milner, A and Goodale, M *The Visual Brain in Action (2nd Edition)* Oxford University Press, Oxford p. 207-252

Milner, A. D. and Dyde, R. [2003]: ‘Why do Some Perceptual Illusions Affect Visually Guided Action, When Others Don’t?’, *Trends in Cognitive Sciences*, 7, pp. 10 – 1.

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

Montaser-Kouhsari, L. and Rajimehr, R. (2004) Attentional modulation of adaptation to illusory lines. *J. Vis.* 4, 434–444

Nudds, M (2007) "Seeing How to Move: Visually Guided Action and the ‘Directive’ Content of Visual Experience” draft manuscript available at: <http://homepages.ed.ac.uk/mnudds/papers/shtm.pdf>

O’Regan, J. K. and Noë, A. [2001]: ‘A Sensorimotor Approach to Vision and Visual Consciousness’, *Behavioral and Brain Sciences*, 24(5), pp. 939 – 73.

Peacocke, C (1992) *A Study of Concepts* (MIT Press: Camb. Ma.)

Reddy, L. et al. (2006) Face identification in the near-absence of focal attention.
Vision Research 46, 2336–2343

Rees, G., Wojciulik, E., Clarke, K., Husain, M., Frith, C. & Driver, J. (2002)
Neural correlates of conscious and unconscious vision in parietal extinction.
Neurocase 8:387–93.

Rizzolatti G, Fadiga L, Gallese V, Fogassi L. (1996). Premotor cortex and the
recognition of motor actions. *Cognitive Brain Research*, 3: 131-141.

Rizzolatti, G and Matelli, M (2003) "Two different streams form the dorsal
visual system: anatomy and functions" *Experimental Brain Research* 153:2:146-157

Sakata H (2003) The role of the parietal cortex in grasping. *Adv Neurol* 93:121-
139.

Schenk, T and Milner, A.D. (2006) Concurrent visuomotor behaviour improves form discrimination in a patient with visual form agnosia *European Journal of Neuroscience* 24: 1495–1503

Schwitzgebel, E (2008) The Unreliability of Naïve Introspection *Philosophical Review* 117: 245-273

Schwitzgebel, Eric and Michael S. Gordon. 2000. How Well Do We Know Our Own Conscious Experience? The Case of Human Echolocation. *Philosophical Topics* 28:235-46.

Sloman A (2007) "Consciousness in a Multi-Layered Multi-functional Mind" poster presented at the 'Perception, Action and Consciousness: Sensorimotor Dynamics and Dual Vision' Conference, Bristol July 2007 and available online at <http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#pac07>

Taira M, Mine S, Georgopoulos AP, Murata A, Sakata H (1990) Parietal cortex neurons of the monkey related to the visual guidance of hand movement. *Exp Brain Res* 83:29-36.

Tse, P. U., Martinez-Conde, S., Schlegel, A. A. & Macknik, S. L. (2005) Visibility, visual awareness and visual masking of simple unattended targets are confined to areas in the occipital cortex beyond human V1/V2. *Proceedings of the National Academy of Sciences* 102:17,178–83.

Ungerleider, L and Mishkin, M (1982) Two Cortical Visual Streams in D. Ingle., M. Goodale., and R. Mansfield (eds) *Analysis of Visual Behavior* MIT Press: Camb. Ma p 549-586

Wallhagen, M. (2007) ‘Consciousness and Action: Does Cognitive Science Support (Mild) Epiphenomenalism?’, *British Journal for the Philosophy of Science* 58:3:539 – 61.

Wright, W (2006) Visual Stuff and Active Vision *Philosophical Psychology* 19: 129-149

ⁱ Conversely, optic ataxics, with dorsal stream lesions, are adept at visually identifying objects that they cannot fluidly reach and grasp. Optic ataxics: 'have little trouble seeing [i.e. identifying] objects in a visual scene, but a lot of trouble reaching for objects they can see. It is as though they cannot use the spatial information inherent in any visual scene' Gazzaniga (1998) p. 109.

ⁱⁱ For example, it has been shown that some visual illusions do affect visuomotor engagement. Importantly, however, this seems to be the case only when the illusion is rooted in very early stages of visual processing (in primary visual cortex) and is thus 'passed on' to both streams when they subsequently diverge (Dyde and Milner (2002) Milner and Dyde (2003)). This is, of course, fully compatible with the strong dual systems view. Moreover, several other perceptual illusions have subsequently been shown to affect conscious experience without impactingⁱⁱ visumotor acts of grasp scaling and reaching including the Ponzo ('railway lines') and Müller-Lyer illusions (see Goodale and Milner (2004) p. 89). In such cases, motor effects *are* observed when delays are introduced between viewing the illusion and producing the motor response. But this is as predicted by the model, which treats time-delayed actions as 'pantomimed' in that they cannot rely on the here-and-now computations of the dorsal stream and are instead driven by the illusion-prone deliverances of the ventral stream (see Milner and Goodale (1995) pp. 170-173).

ⁱⁱⁱ In addition, some aspects of *online object engagement* may require ongoing ventral stream effort (grip force (Jackson and Shaw), functionally informed grips (Goodale and Milner (2004)), complex object engagements (McIntosh et al (2004))

^{iv} One residual issue hereabouts, which I won't attempt to resolve in this treatment, concerns the nature of these projected additional elements of conscious experience. Is the 'feeling of presence' truly part of my visual experience? Is the width information available to DF for form discrimination really given to her as some (perhaps weak or indeterminate?) form of *visual* experience? Does the meaningfulness of a visually presented action really belong to the visual experience itself?

^v This analogy was originally suggested to me by an anonymous referee for Clark (2007).

^{vi} The actual story is a little more complex, since 'pragmatic' visuomotor representations are said (op cit p.3) to have a hybrid direction of fit, insofar as they also provide motor intentions with information about what actions the world affords.

^{vii} Jeannerod and Jacob go on to suggest a further decomposition of function within the parietal lobe/dorsal pathway, with varying admixtures of directions of fit and causation associated with the various functions. The superior parietal lobule, they suggest, is concerned with visuomotor processing, while the right inferior parietal lobule deals with the perception of spatial relationships, and the left inferior lobule with representations of goal-directed actions. Both these latter roles, they argue, are plausibly associated with certain contents of conscious visual experience.

^{viii} Similarly Mathen's comments on the 'feeling of presence' are probably best seen as a kind of restricted version of Nudds' claim (restricted insofar as the only contribution the directive contents make to experience is, on Matthen's account, the addition of that sense of presence).

^{ix} It is worth noting here that DF retains descriptive visual experience of colour and texture, and thus knows when there is an object out there, and what kind of surface (shiny, dull etc) it has. This is what is usually taken to explain her ability to self-initiate actions, and thus already distinguishes her quite sharply (in terms of practical action) from blindsight patients.

^x Tim Bayne (personal communication) suggests that the directive content described by Nudds, insofar as it is indeed consciously experienced, may be best thought of as 'motor intentional' rather than genuinely visual. This is an interesting suggestion, but one that I shall not pursue in the present context.

^{xi} Versions of this suggestion are found in Evans (1982), Marcel (1983), Milner and Goodale (1995), Goodale and Milner (2004), Hurley (1998), Clark (2001), Jacob and Jeannerod (2003), Dretske (2006), Clark (2007)

^{xii} Thus according to 'global workspace' theory (Barrs (1988) (1997), Dehaene et al (1998), Dennett (2001), Metzinger (2003) etc) information becomes conscious when it is poised for dissemination to many cortical areas (perhaps via long-range white matter pathways linking cortical areas- see Dehaene and Naccache (2001)). Information poised for such widespread dissemination (information 'in the global workspace') will ipso facto be poised for the control of an open-ended variety of rational responses, including report where available, so the global workspace model can be seen as providing one mechanism by means of which a strong link to personal-level agency may be implemented in the brain.

^{xiii} These experiments are also discussed in Dretske (2006) and Fodor (2007).

^{xiv} For some interesting worries, see Byrne, Hilbert and Siegel (2008).

^{xv} Schwitgebel's position (personal communication) is that although he can see how the arguments might be taken this way, he himself finds it compelling that there must be a clean fact of the matter concerning e.g. the richness, or lack of it, of ongoing visual experience, and that this will be so even if that fact turns out to be permanently resistant to scientific resolution.