Brit. J. Phil. Sci. **58** (2007), 563–594

What Reaching Teaches: Consciousness, Control, and the Inner Zombie

Andy Clark

ABSTRACT

What is the role of conscious visual experience in the control and guidance of human behaviour? According to some recent treatments, the role is surprisingly indirect. Conscious visual experience, on these accounts, serves the formation of plans and the selection of action types and targets, while the control of 'online' visually guided action proceeds via a quasi-independent non-conscious route. In response to such claims, critics such as (Wallhagen [2007], pp. 539–61) have suggested that the notions of control and guidance invoked are unacceptably vague, and that that the image of 'zombie systems' guiding action fails to take account of the possibility that there is genuine but unconceptualized, unnoticed, and/or unreportable experience taking place and guiding or controlling the actions. I address both sets of concerns. I try to show that refining and clarifying the key notions of control and guidance leaves the original argument intact, as does the appeal to unconceptualized, unnoticed, or unreportable experiences. The exercise serves, however, to highlight an important complex of considerations concerning the relations between control, agency, and experience. Better understanding these relations is, I suggest, an important source of insights concerning the nature of phenomenal experience.

1 Introduction: The Dual Visual Systems Challenge

In recent years, there has emerged something of a consensus (at least among those working in the neuroscience of vision) concerning the existence of an important division of labor between the two main processing streams projecting from early visual areas to the rest of the human brain. One of these streams, the dorsal stream, projects to the posterior parietal cortex and seems to be especially implicated in the fast online control of motor actions. The other, the ventral stream, projects to the temporal lobe and seems to

doi:10.1093/bjps/axm030

be especially implicated in the recognition and identification of objects and events.

That distinct aspects of visual processing are handled by each of the streams seems beyond reasonable doubt. But just how best to *further* characterize this division of labour remains a matter of some dispute. In particular, the question arises whether the representations processed by each stream are conscious (but of differing types), or line up with the conscious/non-conscious distinction itself. The putative link with conscious visual experience is suggested by (Milner and Goodale [1995]; Goodale and Milner [2004]). The dorsal and ventral streams, on this model, deal respectively with 'vision-for-action' (fast, automatic, unconscious) and 'vision-for-perception' (slow, conscious). The contents of conscious visual experience, Goodale and Milner (see e.g., their [2004]) continue to believe, are thus associated pretty well exclusively with the coding and processing operations carried out by the ventral stream. Others, such as (Jacob and Jeannerod [2003]; Jeannerod and Jacob [2005]), lay greater stress on the complex interactions that bind the activity of the streams into successful practical action, while still endorsing the general 'dual visual systems' picture and (hence) the 'basic bifurcation between vision-for-perception and vision-for-action' (Jeannerod and Jacob [2005], p. x). Regarding the nature and degree of dissociation and segregation of labour between the two pathways, it has also been suggested (Decety and Grèzes [1999]), that this may be a context-modulated property that depends crucially on the nature of the activity involved. Kanwisher ([2001]), challenges the idea that conscious visual perception corresponds preferentially to ventral stream activity, offering instead a picture where both streams play a major role; while fully opposing the dual visual systems view are those (such as Shull and Bingham [2000]), who argue that the degree, density, and complexity of the interactions required render the whole idea of substantial division of labor, at least along anything like the lines suggested by Milner and Goodale, misleading and unhelpful.

The dual visual systems picture itself is supported by a large and impressively varied body of evidence. Such evidence includes the original neuropsychological studies of patient DF (Milner *et al.* [1991]; Milner and Goodale [1995]; Goodale and Milner [2004]), 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]; Aglioti *et al.* [1995]; Fecteau *et al.* [2001]). Of these, the Aglioti *et al.* results are both the best known and the most controversial. They involved normal subjects engaging in reaching and grasping movements whose target was subject to a perceptual size illusion, and seemed to show that grip aperture was sensitive to the real size of the target, whereas judgments based on conscious visual experience were influenced by the illusion. The explanation, according to (Milner and Goodale [1995]), was that grip aperture was determined by the dorsal 'vision-for-action' system operating independently of the ventral 'vision-for-perception' stream whose codings, on this interpretation, informed both conscious experience and verbal response.[1]

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, because of 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 in which 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, some of which we shall touch upon later in this treatment (for some balanced reviews of the topic, see Carey [2001]; Jacob and Jeannerod [2003], Chapter 4).

There are, finally, compelling computational and information-processing considerations that speak in favour of a dual visual system architecture (though just how this in turn lines up with the issues concerning conscious and non-conscious vision is open to question, as we shall see later). 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, and exquisitely distance and orientation-sensitive encoding of the visual array. The latter requires the computation of object-constancy (objects do not 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 looks to preclude the use of the very same

[1] It is worth noting that the data here show only [at best] a divergence between verbal report and visuomotor action. The claim that the representations that inform verbal report exhaust what is given to perceptual consciousness is an additional claim whose support and plausibility will be at issue throughout the present treatment. Thanks to an anonymous referee for suggesting the need to flag this issue at the very outset.

encoding for the other (Milner and Goodale [1995], pp. 25–66). Different uses of visual information thus impose quite different computational demands on the brain and the common sense idea that a *single* visual encoding might both inform recognition, verbal judgment, etc. *and* guide fluent world-engaging action turns out to be unworkable in the natural order.

It remains to be seen, however, whether we have good reason to identify conscious visual experience preferentially with processing in one of these streams. In (Clark [2001]) I argued that the dual visual systems hypothesis, even when modified in the light of various results suggesting greater interaction between the streams, presented a computationally compelling, empirically well-supported, and philosophically important challenge to the view that conscious visual experience really does control daily, world-engaging action. The present paper aims to defend that view against some recent critical challenges (Sections 3 and 4 following). It then uses those challenges to introduce a complex of questions concerning the nature of perceptually grounded intentional action (Section 5) and the relations between planning, reasoning, agency, and visual experience (Section 7). I begin, though, by briefly rehearsing the main arguments and considerations offered in (Clark [2001]) as a reason for rejecting the view that conscious visual experience controls fluent world-engaging action.

2 Experience-based Selection

The dual visual system findings cast doubt, or so I argued in (Clark [2001]), on the following rather natural assumption concerning the nature and functional role of conscious visual experience in the management of human action:

> **Assumption of experience-based control (EBC)**
>
> Conscious visual experience presents the world to the subject in a richly textured way; a way that presents fine detail (detail that may, perhaps, exceed our conceptual or propositional grasp) and that is, in virtue of this richness, especially apt for, and typically utilized in, the control and guidance of fine-tuned, real-world activity. (Clark [2001], p. 496)

The point of this formulation was not to suggest that, in fact, we do *not* experience the world in a richly textured way (though how best to account for such experience is another, and itself contested, matter: see (Dennett [1991]; Noë and O'Regan [2000]; Noë [2001], [2002]; Clark [2002]). Rather, it was to suggest that our visual experience, although it may indeed present a world of fine and rich detail, is not in the business of controlling our subtle physical engagements with that very world. We shall return, in Section 3, to the question of just what that word 'control' needs to mean if this is to express a truth about the nature of human visuomotor activity.

The ([2001]) treatment then contrasted experience-based control (EBC) with a formulation meant better to capture the actual role of visual experience in relation to human action. This was:

> **Hypothesis of Experience-based Selection (EBS)**
>
> Conscious visual experience presents the world to a subject in a form appropriate for the reason-and-memory based selection of actions. (Clark [2001], p. 512)

The contrast was thus between the claim that visual experience was apt for and utilized in the fine control of human action and the claim that visual experience was (rather) apt for and utilized in the gross selection of actions in ways that respect stored knowledge and may invoke chains of conscious reasoning.

A good (though potentially somewhat blunt and over-dramatic, as we shall see) way to focus the intended difference is by means of an analogy (Goodale [1998]; Goodale and Milner [2004]), with tele-assistance approaches to the control of distant robots in distant or hostile environments. In a typical tele-assistance set-up, a human operator and a semi-intelligent distal robot combine forces so as to perform actions in some environment. A familiar example might be a Mars rover, where the human operator reviews images on a screen in Texas, flagging items of interest (such as a strangely shaped rock in the top left of the screen). The operator commands the robot to retrieve the flagged item, perhaps adding commands that specify the use of one of several retrieval modes (according to estimated weight, fragility, etc.). The robot rover then does the rest, locomoting to the spot and calculating the local commands needed to deploy the robot body and gripper so as to achieve the goal.

Such approaches should be contrasted with *tele-operation* solutions, in which the human operator controls all the spatial and temporal aspects of the robot's movements (perhaps via a joystick or a set of sensors that allow the operator's own arm and hand movements to be relayed to the robot). In a tele-assistance solution:

> the human operator doesn't have to worry about the real metrics of the workspace or the timing of the movements made by the robot; instead, the human operator has the job of identifying a goal and specifying an action towards that goal in general terms. [The robot then uses] its on-board range-finders and other sensing devices to work out the required movements for achieving the specific goal (Goodale and Milner [2004], p. 99).

The tele-assistance analogy identifies the human operator with the ventral stream (working 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 (+) thus plays the robotic Mars Rover to the ventral stream (+)'s human operator. In this way:

> Both systems have to work together in the production of purposive behaviour—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 contributions of the two visual streams fits quite neatly with some of the coarse dissociations described in the literature. In particular, as touched upon in Section 1 above, 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. 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 she was pretty well at chance for turning a hand-held 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. 128–33). Optic ataxics, with dorsal stream damage, display a very different profile, succeeding at the recognition task while proving unable to actually post the card through the slot.

One intuitive (but apparently incorrect) way to account for optic ataxia is to embrace the Hypothesis of EBC, but suggest that the perceptual information, that is clearly consciously available to the subject, is somehow blocked from getting through to the motor systems. On that view 'there is only one kind of visual processing, which not only provides our conscious perception, but also the visual guidance for all our actions' (Goodale and Milner [2004], p. 38). But this intuitive model, Goodale and Milner argue, cannot easily accommodate the converse impairment seen in DF. On the intuitive model, DF's lack of conscious perceptual awareness should suggest that she has (for whatever reason) failed to compute the information needed to guide visuomotor action. Yet her fluent performance belies this. The dual streams/tele-assistance model accounts for both cases, 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 identification.

Moving to the case of normal, unimpaired subjects, the same story neatly accounts for the (Aglioti *et al.* [1995]) results mentioned in Section 1. 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. 88–9) 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 (information that, they claim, 'makes only a small contribution to our [conscious] *perception* of depth' (*op. cit.*, p. 91). 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 splendidly intact.

The claims concerning the immunity of visuomotor action to visual illusions have spawned a large industry 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]). 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 impacting[2] visumotor acts of grasp scaling and

[2] Some evidence suggests that, in the Ponzo and Müller-Leyer cases, conscious visual illusions *do* influence grasp, although only to a small degree. Thus according to (Ellis *et al.* [1999]) the visual illusion, in these cases, *does* influence the action systems, but the action [grasp] system

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. 170–3).

The Assumption of EBC cannot, I conclude, be tolerated by an empirically informed account of the role of experience in human action. Conscious visual experience is not, despite appearances, in the business of providing the information used to control most daily visuomotor action. Instead, it is in the business of delivering representations optimized for quite different purposes. These are the purposes of recognizing objects, selecting targets for action, and determining what kinds of action, broadly speaking, to perform (should the Mars Rover fetch the rock, or push it away?). Visual experience is systemically tied to capacities of recall, reason, and planning, not to the support of finely tuned world-engaging motor activity.[3]

Matthen ([2005]), captures the situation well. Asking us to imagine a case where we are told, by our host, to 'take the blue coffee cup', the systems underlying visual experience enable us to identify the cup, to draw on stored knowledge concerning proper etiquette and the dangers of hot liquids so as to select a particular way to grip the cup (carefully, by the handle), and to launch the whole reach-and-grasp routine at roughly the right time. But:

> ... once the object has been identified and the type of grip chosen—once the action has been, as it were, launched—the guidance of the hand is taken over by a system that is independent of the one that provides experiential information about shape, size, and colour. (Matthen [2005], p. 295)

Such a view, as Matthen (*op. cit.*, p. 294) notes, seems to go against the intuitive view that I dubbed EBC—here, the view that 'throughout this process

also has access to more veridical information. The results obtained then reflect the interaction between the two. In a similar fashion (Jeannerod [1997]; Jacob and Jeannerod [2003]; and Jeannerod and Jacob [2005]) offer a variety of evidence favouring a greater degree of interaction between vision-for-perception and vision-for-action, but without casting doubt on the general correctness the dual systems perspective.

3 This general picture also looks to be a good fit with the so-called 'Theory of Event Coding' (Hommel *et al.* [2001]), according to which conscious perception and action-planning share resources, working together as a kind of 'distal-event system' that 'cares about' the overall effects of action rather than the specifics of the action itself; see (Jordan [2003]) for further discussion.

you were guided by the blue shape in your visual field'. Conscious processes of seeing, willing, and intending, if the alternative (Experience-Based Selection, EBS) view I am defending is correct, are at work only at one remove from most cases (pantomime acts aside) of world-engaging action. This fact (if it is a fact) should have important consequences for how best to think about the relations between conscious experience and intentional states, and perhaps even for our notions of agency and responsibility. But is the contrast (between EBC and EBS) really as clear as I have made it sound? And (assuming there is indeed a contrast) do the facts really support the EBS alternative?

3 Selection, Programming, and Control

One natural way to probe the story on offer is to focus on the meaning and scope of the key concepts of 'control' versus 'selection'. This is exactly the strategy employed by (Wallhagen [2007], pp. 539–61), who argues that the evidence I adduced 'does not in fact threaten the view that conscious states play an important and intimate role in the control of much everyday behavior' (*op. cit.*, pp. 539–61).

It is important, before proceeding, to remove one misunderstanding that in various ways seems to colour Wallhagen's argument. It is to suppose that my account was meant to lend support to any general epiphenomenalism about the role of conscious states in the control of everyday behaviour. Wallhagen, apparently taking this to be the endpoint of my argument, sets out to show that 'conscious states play an important and intimate role in the control of much everyday behavior' (*op. cit.*, pp. 539–61). The trouble here lies, indeed, with this slippery notion of control. For the point of my argument was to show that certain conscious states (the ones that seem to present a rich and detailed visual scene laid out in space) do not, in fact, play the *kind of role* in the control of behaviour that one might pre-theoretically expect. Instead of the conscious visual states providing the information that the motor systems use to engage the world, they rather (or so I claimed, in line with the dual systems view) support various acts of reasoning, selection, and choice. These acts of reasoning, selection, and choice identify worldly targets and broad types of intentional action, whose motoric unfolding is then controlled by scene layout information computed by a quite different suite of circuits. Notice that as far as this picture is concerned, there is nothing whatsoever 'epiphenomenal' about the role of the conscious states that present a visual scene in a form apt for reasoning and action selection. For the acts of selection that ensue form a vital part of the causal chain that leads to world-engaging behaviour.

To see this, consider once again the Aglioti *et al.* experiments using a poker chip version of the Titchener Circles illusion. Here, according to the dual systems view, the information made available in conscious awareness (which is subject to the size illusion) causes the agent to initiate a motor routine that picks up one of the two centre chips. The surprise is that the agent's finger grip is precisely matched to the real size of the chip. Thus, if asked to pick the other chip, she will form exactly the same precision grip, despite her illusory experience of a difference in size between the two. The conclusion to be drawn is *not* that the information presented in conscious awareness has no impact on action, for the whole choice of which chip to pick up and the selection of (let us say) a right-chip-retrieving motor routine depend on it. Rather, it is that the impact is less direct that we might pre-theoretically expect, with the consciously presented information serving only as a basis for the selection of targets and action types, rather than as the information used to actually compute the details of the reach and grasp. Hence my suggestion that we speak of 'experience-based selection' rather than 'experience-based control of fine-tuned real-world activity'. It should be clear that the appeal to 'selection rather than control', in the technical sense thus intended, does not render the consciously presented information impotent, or even 'mildly epiphenomenal' (to use Wallhagen's phrase), in the chain leading to world-engaging action.[4]

Wallhagen's subsequent critique begins by attempting to pin down the notions of 'control and guidance' and 'fine-tuned activity'. Concerning control and guidance, Wallhagen notes, correctly, that one could not cast doubt on EBC simply by showing that non-conscious mental states play some role in the generation of behaviour. For this is something that all sensible theorists should allow. So the question is not whether non-conscious stuff gets into the act at all. Rather, it concerns the exact role that conscious visual experience gets to play in the complex web of causation that leads to some real-world behaviour. After some discussion, Wallhagen offers the following positive suggestion:

> ... we can say that a mental state—particularly a sensory/perceptual state—plays a role in the control/guidance of some behavior when the state provides the information used by the motor system in developing motor instructions (e.g., information about joint movements) that bring about the behavior in question. (Wallhagen [2007], pp. 539–61)

I think this is an acceptable formulation, though it is not itself without ambiguity, since that notion of 'the information used [. . .] in developing motor

[4] In this respect, the view I defend is far less radical than, say, the views of Libet ([1985]) and others who thought to show that even the acts of (what I am calling) selection occur too late to count as the instigators of the 'selected' motions. I believe the arguments meant to support this more radical view are flawed, but to consider them is far beyond the scope of the present treatment.

instructions' might, or might not, be thought to include the 'types and targets' style information mentioned above.[5] We shall return to this later.

Next in line is the notion of 'fine-tuned activity'. Here, Wallhagen offers two interpretations, only one of which (he suggests) can serve the purposes of my argument. The readings are:

> EBC (gen): Conscious visual experiences are typically utilized in the control and guidance of voluntary/intentional behaviors.
>
> EBC (res): Conscious visual experiences are typically utilized in the control and guidance of the fine/minor adjustments made to behaviors during their execution (i.e., 'on the fly'). (Wallhagen [2007], pp. 539–61)

Rejecting EBC (res), Wallhagen argues, would be relatively uninteresting, since it is unclear whether either common sense or philosophy has much to say about the most minute adjustments involved in world-engaging action. What matters are, rather, the larger scale aspects of apparently consciously controlled behaviour. For example:

> In bringing one's finger down upon the center of a printed cross, one not only needs to make minor adjustments to the position of one's finger, one needs to move one's finger in the general direction of the cross, and the appropriate distance, so as to get one's finger in the general vicinity of the cross. These latter, grosser, aspects of behavior need to be controlled every bit as much as the minor adjustments one makes to bring one's finger down on the very center of the cross. (Wallhagen [2007], pp. 539–61)

I agree. These grosser aspects need to be (in my terms) both selected *and* controlled. We need to first choose (in the poker chip case) which chip to reach for, and what kind of reach to utilize, and we need to generate motor commands apt for the purpose. In the normal case of fluent, online reaching, my claim was that even the gross heading and kinematics of the reach were programmed (though not selected) by the distinct representational resources proper to the dorsal stream. Thus my target was (and remains) a version of EBC (gen).

In this regard, my frequent stress on 'fine-tuned motor activity' may have been misleading. That stress was rooted in two concerns. The first was as part of a critical response to (Cussins [1998]), claim that the conscious (but non-conceptual—see next section) contents of a motorcyclist's visual experience are what accounts for her 'knowingly making micro-adjustments of [. . .] speed

[5] An anonymous referee usefully points out that the more general distinction between 'what we try to do' and 'how we try to do it' is itself rather elusive. In trying to pick up the cup, I do not try to flex my muscles thus and so. But I do perhaps mean to pick the cup up by its handle. How much of this detail is properly speaking part of the conscious content of my trying is unclear.

[...] in response to changing road conditions' (Cussins [1998], p. 10). Cussins' claim was thus about the micro-details of the visuomotor response, and I was concerned to stress that these were pretty much solely the work of the (putatively) non-conscious dorsal stream. The second was to recognize the wide range of experiments (surveyed and discussed in the original paper) showing our surprising behavioural sensitivities to minor alterations in the actual local visual layout, even when those alterations go unmarked in conscious experience.[6] Wallhagen is right to suggest, however, that there is nothing here to really challenge common sense, for all that common sense requires is some grosser role for conscious visual perception in guiding motor action, that is, what he dubs EBC (gen).

Nonetheless, the thrust of my argument (as Wallhagen correctly infers) was indeed to cast doubt on EBC (gen). In the case of the finger and cross, for example, my claim was that *if* the Milner and Goodale story were correct, then even (most of—see below) the grosser aspects of the motor engagement turn out to be controlled by scene layout information that has been computed quite independently of the information, computed by the ventral stream, that (arguably) supports recognition, reasoning, and conscious visual experience. One way to see this is to return to the analogy with tele-assistance (Goodale [1998]; Goodale and Milner [2004], Chapter 7; Milner and Goodale [2006], pp. 231–4) introduced in Section 2 above. Goodale and Milner ([2004], p. 98) describe this as a 'close analogy' devoting many pages to its pursuit and fine-tuning. We can use their discussion to further clarify both the scope (EBC (gen)) and the complexities (inter-stream interactions complicate, but do not undermine, EBC (gen)) of a nuanced dual stream story.

Recall that in the tele-assistance model, the ventral stream (working in concert with memory and other key cognitive systems) is said to 'perform the role of the human operator in recognizing a potentially important target' (Milner and Goodale [2006], p. 232). These systems work, according to the dual visual systems account, with a representational format that is 'rich and detailed but not metrically precise' (*op. cit.*). Once an action target and action type have (on this basis) been selected, the actual motor programming is left to the dorsal stream (again, working in concert with other key systems such as the basal ganglia). One account of how the ventral stream might 'tell' the dorsal stream what the target should be is via their joint access (by recurrent projections) to the low-level retinatopic maps in primary visual cortex. Thus Goodale and Milner ([2004], p. 102) speculate that (perhaps working also with the lateral intraparietal sulcus area, LIP) 'it would be possible for the ventral

6 We may count the Aglioti *et al.* demonstrations as belonging to this broad category, even though in that case it is not unseen alterations that matter, but the illusion-induced discrepancies between what we see and fine motor response.

stream to flag a target for the dorsal stream using this common frame of reference', adding that 'once a target has been highlighted on the retinal map, it can be converted into any other coordinate system that the dorsal stream might need to use' (*op. cit.*).

Suppose, if only for the sake of argument, that some such account proves correct. In that case, I suggest, the ventral stream has indeed merely selected a target (and, of course, an action type: a reach, a punch, fast, slow, etc.). For the full job of actually programming the reach, *including* motor programming the initial heading and gross kinematics, is here left to the dorsally dominated, metrically precise resource. To clarify the suggestion, it helps to now distinguish three elements in the genesis of an intentional reaching action, in place of my original choice of two ('selection' and 'control'). The three elements are selection (of targets and actions types), initial motor programming of the heading and gross kinematics, and online subtle control and error correction. Milner and Goodale's position, and the one I invoked in the ([2001]) treatment, was that of these three elements, only the first (selection) was carried out by the awareness-supporting ventral stream.

Now, it is true that in recent years this claim has been challenged. Thus (Rossetti *et al.* [2003]; Glover [2004]) suggest that both selection *and* the initial motor programming of heading may be carried out by the ventral stream. The place where dorsal stream activity seems most securely implicated is, they suggest, in the small online corrections, often carried out in the complete absence of conscious awareness that the target has moved. These corrections are what seem most impaired in the cases of optic ataxia, often presented as the 'complementary' deficit to that of visual form agnosia. Milner and Goodale continue to reject this suggestion, arguing that optic ataxics also show deficits in the initial heading of the reach, and in their initial gross directional movements in obstacle avoidance tasks etc (see Milner *et al.* [2003]; Schindler *et al.* [2004]; and discussion in Milner and Goodale [2006], p. 238). Milner and Goodale (*op. cit.*) conclude that 'the weight of current evidence therefore confirms that the dorsal stream is implicated in both the programming and the online control of reaching'.[7]

The image is thus one of a kind of parallel unfolding of visual awareness and visumotor control, where 'control' really does include not just fine-grained online error correction, but also the programming of the gross heading and kinematics of the intentional act itself. As they more recently put it:

[7] Milner and Goodale allow, however, that the initial programming of the heading and the online correction and control may depend on different subsystems within the dorsal stream, and that this may explain why one patient showed a deficit only in online control (see Milner and Goodale [2006], p. 238).

> It is our claim [. . .] that all of this *visual* awareness will be mediated by processing in the ventral stream and that this processing will typically unfold independently of, and in parallel with, the visual control of the action itself [. . .] The compelling nature of visual consciousness makes it hard to resist the intuition that it is one's [conscious] perception of the goal object that is guiding the action. But sometimes the truth can be counterintuitive. (Milner and Goodale [2006], p. 222, their emphasis).

Wallhagen's own suggestion concerning what it would take for a conscious visual experience to play a role in the 'control and guidance' of behaviour is, as mentioned earlier, somewhat hard to interpret. He suggests that:

> your visual experience of [a] cup will have been utilized in the control and guidance of [cup-reaching] behavior if it provides the information your motor system uses in causing your hand to move in the relevant direction and distance, to take on the appropriate orientation, and in configuring your fingers appropriately (Wallhagen [2007], pp. 539–61)

But what does it mean for the conscious percept to 'provide the information used by the motor system in developing motor instructions' (*op. cit.*, pp. 539–61)? If this simply means flagging the target and selecting an action type, leaving the dorsal servants to do the rest, the claim would, I suggest, be consistent with the rejection of EBC (gen) and hence with my own preferred EBS. It seems, then, that Wallhagen must have something stronger in mind. The idea, I think, is that the conscious experience provides the information that the motor system actually uses to compute the initial heading and gross aspects of the action, even though small online corrections rely on information from other sources. The required distinction is that between providing the information used to compute the gross aspects of the reach and merely providing a specification of the target and action type required. It is the difference between telling the waiter to get the apples from the pantry and providing at least a rough sketch of the motion that would get the waiter to the pantry. If this is the right way to read Wallhagen's claim, then we have now isolated the key point of disagreement. For according to EBS, and to Milner and Goodale, the consciously available information is used only for the specification of action types ('don't bruise the apples, they are delicate') and targets, and is not used even to compute a rough sketch of the trajectory itself.

We can now turn to the promised complications. They emerge once we admit that the tele-assistance model can, as Goodale and Milner recently note, make it seem as if the ventral stream (and associated resources) 'plays only a very distant role in the implementation of action, rather like a chief executive officer in a corporation, setting goals and writing mission statements, and then delegating the real work to others' (Goodale and Milner [2004], p. 103). For *some* aspects of the motor programming, however, the ventral system does rather more than this. As we already noted, it provides information about the action type (if the object is a screwdriver, it provides information about which end to grip, and in what way, so as to actually use the tool—see e.g. Goodale and Milner [2004], pp. 105–7). The patient DF, with ventral damage, though able to scale her grip to the shape of the (to her, unidentifiable) tool, will often grip the tool from the wrong end. Moreover, the programming of grip force required to take account of an object's likely weight turns out to require ventral involvement too, and not just in respect of familiar objects whose weight is known and stored in memory. In effect, the dorsal resource has nothing beyond visual attributes, processed 'bottom up' on the spot, to go on. So even to use knowledge of material (lead versus plastic) to determine grip force requires information beyond its means. Grip force turns out to be left to the ventral stream, so much so that illusions of size (in this case using a special version of the Ponzo or railway lines illusion) have been shown (Jackson and Shaw [2000]) to affect the scaling of grip force even in cases where grip size is computed correctly, presumably by the dorsal stream. This suggests that the computation of size that determines the applied grip force is the one carried out by the ventral (illusion-prone) resource, even though the computation of size that determines the precision grip itself is not.

The ventral stream, it seems, is therefore involved in the actual programming of some aspects of intentional action, though (arguably at least) not those that determine even the grosser aspects of the motor trajectory itself. With regard to that trajectory, even the initial heading and gross structure are programmed by (though not selected by: recall our earlier discussion) the 'zombie' dorsal resource whose task is to 'convert visual information directly into action'.

4 Noticing What You See

Wallhagen's next move is to press a series of tricky questions concerning the relation between seeing (in the sense of being consciously aware) something and noticing that you see that very thing. Wallhagen suggests (correctly in my view) that:

> it is implausible to suppose that being consciously aware of something is no different from noticing that something is the case. Instances of awareness without noticing abound. For example, suppose you see a friend along the street, and your friend has had her hair cut. You might not *notice* this for some time. Nevertheless, you were *aware* of her hair, and the properties in virtue of which it looks different to you (which explains the thought you're likely to have when you *do* notice her haircut: 'I had a sense there was something different about her today')... (Wallhagen [2007], pp. 539–61).

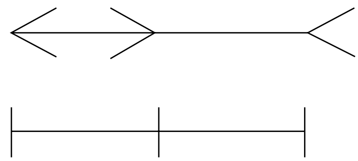

Figure 1. Augmented Müller–Lyer Illusion. Figure courtesy of Jeroen Smeets.

This possibility matters when considering various empirical results (e.g., Bridgeman *et al*. [1979]), in which our actual eye movements accurately track the location of stimuli that, owing to various illusion-inducing effects,[8] we (falsely) judge to be differently located. In (Clark [2001]) I take this as evidence that the accurate eye movements were controlled by non-conscious visual states. But perhaps, Wallhagen suggests, the actual location of the target *is* accurately represented in conscious awareness, but the subject simply fails to notice this due to the presence of other (illusory, conflicting) information also present in the conscious visual array? Such a possibility requires, of course, that illusory and veridical contents can both be simultaneously present to conscious awareness. But this, Wallhagen argues (and see also the similar comments by Davies and Peacocke mentioned in Clark [2001]) is both possible and actual. Thus consider Wallhagen's account of the familiar Müller–Lyer illusion. What is given in conscious visual experience, Wallhagen suggests, is *both* a veridical and an illusory presentation of length. To show this, he draws the reader's attention to the cleverly revised version (for useful discussion see Smeets and Brenner [2001]) shown in Figure 1.

Here it seems that we can see (literally see) both that the lines look different in length and that they look the same. According to Wallhagen, this clever manipulation merely reveals what was already present to visual experience even in the standard unaugmented version. It is just that in the standard case we do not *notice* that our own conscious perception has the veridical content too. We notice this only when the new cues draw it to our attention. According to Wallhagen then:

> ... something similar happens in the eye movement experiments. In particular, one's perceptual constancy mechanisms give rise to an illusion of movement, and cause subjects to *think* that the stimulus has moved. As such, when asked to look to the target's original position, they do so in accordance with what they believe. Nevertheless, the target's true movement is still represented in conscious vision. The subjects simply

[8] Typically, frame-ground effects, in which a target point and the frame that contains it both disappear, then reappear slightly over to one side, but with the frame displaced more greatly than the target. This causes subjects to judge (falsely) that the target has moved more than it actually has.

> do not notice that their visual states have this content—just as someone subject to the Müller-Lyer illusion does not notice that her visual states accurately present the two lines as being of equal length. But since their states do contain the veridical content, there is no reason to suppose that conscious visual states do not provide the information about the target's location that eye-control centers use to generate a saccade to the target. (Wallhagen [2007], pp. 539–61).

This is an interesting challenge, but one that needs to be handled with care if it is not to prove too much. For once we so firmly distinguish what the subject currently thinks and judges from what they are aware of, it becomes hard to see what kind of evidence could ever show that some information, assuming it is demonstrably picked up by some part of the organism's response-determining apparatus, failed to be in some manner present to conscious awareness.[9] For it might always be claimed that the information made it into conscious awareness but was simply not noticed by the agent. Any unrestricted version of the argument therefore requires a kind of division of the agent into what might be termed the *bare experiencing* agent and the *noticing* agent. On this model the bare experiencing agent can be fully consciously aware of much that the noticing agent misses. The augmented Müller–Lyer illusion does not, it seems to me, lend support to any unrestricted version of this claim. For here, the augmentations create a state of conflict in the noticing agent, and this is our best evidence that under those special conditions multiple inconsistent contents are indeed present. It is unclear what, if anything, we should then infer about the unaugmented case.

There is an interesting comparison to be made with Johnson and Haggard's ([2005]) suggestion that subjects can exhibit a kind of 'motor awareness' in which they are both (a) unaware that a perceptually presented target has been displaced during some visuomotor routine (the subjects fluently accommodate the displacement without reporting seeing or noticing it) yet (b) aware, at least as evidenced by their subsequent on-demand reproduction of the spatial path, of a displacement (a 'visuomotor adjustment') in the motor routine itself. This is somewhat surprising and is said to demonstrate:

> . . . the novel finding that motor awareness can dissociate from perceptual awareness. That is, despite reporting that the target did *not* shift, subjects were nevertheless aware that their movement contained a lateral deviation in the appropriate direction. They are able to report their motor awareness of their movement despite having no conscious knowledge of the target event that drove it. (Johnson and Haggard [2005], p. 232).

9 There is an interesting discussion of these issues in (Dretske [2006]). Dretske's conclusions, though arrived at by an alternative route, seem generally consistent with the line taken in the present treatment.

Strikingly, subjects did not find anything odd in their combination of a report of no target displacement with their awareness of a deviated trajectory. This leads the authors to speculate that:

> . . . we think the mind does not even attempt to resolve lack of perceptual awareness with positive motor awareness. Rather, motor awareness and perceptual awareness might be quite different and independent processes in the brain [. . .] In some respects, this position recalls the concept of multiple 'micro-consciousnesses' within the visual system. (Zeki and Bartels [1999]). (Johnson and Haggard [2005], p. 236)

The point I want to stress, though, is that to support this counter-intuitive view, Johnson and Haggard had to create a situation in which there was at least some kind of report (and hence some evidence of 'noticing') available, distinct from the actual performance itself.[10] Is this a reasonable demand, or does it too force us to conflate conscious experience with something like self-conscious (because *in some way* noticed or reportable) experience? My own view, that I shall try to defend as the discussion progresses, is that there is a perfectly proper underlying demand here, but one has more to do with deep ties between experience and deliberate response and agency than with any shallow conflation of experience and reportability. It is only when in some sense the agent has access to perceptually transduced information that that information can be counted as given to the agent in her experience. This demand is compatible with all manner of multiplicities and inconsistencies[11] in experience: for experience, on this view, can be multiple, fractionated, and inconsistent to whatever degree this access may be multiple, fractionated and inconsistent. But subtract that access entirely, leaving nothing but a fully isolated islet, and the notion of experience fails (I want to say) to get a grip.

Certainly, though, Wallhagen is right to draw our attention to this issue. For the literature is replete with the kinds of inference (from failures of noticing or report to the absence of conscious awareness) he seeks to question. In a lovely case study of obstacle avoidance during reaching in a patient (VE)

10 More generally, it seems to me that the putative experiences of any 'bare experiencing agent' must always be parasitic on at least some *possible* performances of the noticing agent. For example, I did not notice my friend's haircut but I certainly could have, had my attention been so drawn.

11 The psychologist of vision Jeroen Smeets has written extensively on such cases (see e.g., Smeets and Brenner [2001]) On the basis of a wide range of compelling empirical demonstrations, Smeets now suggests that perceptual experience does not even attempt to construct any kind of 'visual or perceptual space'. Instead, the brain merely detects information about various attributes, and uses the most effective attributes (for some task) to guide behaviours. As I understand it, Smeets' view is thus that the question of what is or is not given in conscious awareness is neither a useful nor a tractable one. That might be true. But if so, things are much worse than Wallhagen expects. For far from establishing EBC against EBS, such a view amounts to a kind of eliminativism concerning the appeal (whether by EBC or EBS) to the contents of the consciously presented scene in explanations of human behaviour.

with left visual extinction,[12] McIntosh *et al.* ([2004]), found that the patient avoided obstacles during a reaching movement 'to exactly the same degree whether or not he was able to report their presence' (*op. cit.*, p. 15). That is to say, obstacle avoidance for reaches conducted in the extinguished field was indistinguishable from obstacle avoidance in the intact (reportable) field. According to McIntosh *et al.*, this shows (*op. cit.*) that 'conscious awareness is not necessary for normal obstacle avoidance'. In their subsequent discussion, the authors accept that their method equates conscious awareness with verbal report, writing that:

> It could conceivably be argued that V.E. was aware of extinguished stimuli but at a level that was insufficient for verbalization. Ultimately, this objection may be unanswerable.

They add, however, that:

> If awareness were necessary for obstacle avoidance and V.E. had been only dimly aware of the [obstacles] we might have expected their influence to have been less than that of non-extinguished [obstacles], and this was not observed. (McIntosh *et al.* [2004], p. 19).

Perhaps Wallhagen's suggestion, at least if it is to cope with this kind of response, needs to be not that the agent may be very dimly conscious of that which they fail to notice or report, but that they may be fully conscious of it, whilst nonetheless failing to be able to report on it: much as, one might imagine, the cat is not dimly but vibrantly aware of the mouse, even while failing to be able to issue a verbal report. (The cat, though, has certainly *noticed* the mouse).

We here enter a new phase of the discussion, and one that invites us to tiptoe into some murky waters. Troubling those waters are two large and intertwined issues. The first concerns the possible impact, on the kinds of evidence and arguments I have presented, of making a distinction between what we *fully consciously experience* and what we currently, as advanced verbal beings, *conceptualize* ourselves as experiencing. The second concerns the relations, if any, between the very notion of conscious perceptual experience and the availability of transduced information to inform agency and choice.

[12] This is an attentional disorder, caused by unilateral brain damage, in which a patient can detect single stimuli on either side of their visual space, but fail to do so on the side opposite the lesion when a stimulus there is paired with one on the other side. The condition is akin to visual neglect but (unlike neglect) is elicited only under conditions of bilateral stimulation.

5 Conceptualization versus Experience

We should admit at the outset that the empirical evidence thought to support the fully non-conscious ('zombie') status of *all* dorsal stream contributions is at least questionable. For example, it seems quite possible that more neglected elements of experience, ones other than those concerning the typical suspects such as visually perceived shape, may depend more directly on dorsal stream activity. Thus Matthen ([2005], p. 301) argues that the 'feeling of presence' may depend on dorsal stream activity even if the other more descriptive elements do not. The dorsal stream contribution, thus understood, would be what determines the difference *in experience* between seeing a painting of a cup and a cup. In the latter case, but not the former, the dorsal stream is able to compute the details of a cup-engaging reach, and this arguably lends a feeling of real presence to the cup. The painted scene, by contrast, fails to provide cues as to the precise metric distance of the cup from the perceiver, and hence fails to engage the dorsal machinery in the same way (it can compute the metric distance of the painting, but that is, of course, another matter).

That such extra elements in experience may be dependent on the dorsal computations does not, however, threaten the claim that we should (where conscious shape-presenting 'descriptive vision' is concerned) prefer EBS over EBC. Indeed, it is not clear to me to what extent this 'feeling of presence' should even be said to be properly *visual* in nature.[13] What *would* threaten the argument for EBS is evidence that there really is conscious experience of visually presented shape and form[14] in the range of cases offered as demonstrations of successful, but dorsally controlled, behaviour. The flagship case here is, of course, the carbon monoxide poisoned subject DF. DF, it was claimed, is sometimes able to accurately engage the very shapes that she says she cannot see. This was taken as strong evidence for a dissociation between vision-for-perception (i.e., for conscious experience) and vision-for-action (i.e., for the programming and control of visuomotor action), and was certainly

[13] It does not seem to me sufficient, as an argument for the visual nature of the feeling of presence, to point out (as does an anonymous referee) that the cues used to compute presence involve visual features such as precise binocular depth. The question is whether there is here anything that we might properly regard as visual experience of the dimension of presence. I am not convinced there is, though I concede that an alternative explanation of my lack of conviction might be the dominance of descriptive visual experience (reportable experience of visual features) over other forms of visual experience. The relative ineffability of such putative non-descriptive-yet-genuinely-visual-experiential elements should, however, give us pause for thought.

[14] There is important emerging evidence that the dorsal stream can make width information available to control conscious response. The information thus provided falls radically short, however, of constituting awareness of visual shape and form (thus, objects of the same width but different forms fail to be distinguished). See (Schenk and Milner [2006]). In addition, the dorsal stream, and parietal lobe functions more generally, seem to play an important role in making information about higher level actions, such as other's tool-using behaviour, available to conscious awareness (see Jeannerod [2006]).

a major part of the machinery meant to favour EBS over EBC. Indeed, Wallhagen ([2007], pp. 539–61) concedes that if such an interpretation of DF is correct, this would provide 'a powerful case against [. . .] EBC (gen), since the movements that non-conscious states supposedly control are not mere minor adjustments to behavior'. According to Wallhagen, however, this description of DF as lacking conscious experience of shape is too hasty. For it is possible that DF experiences visually presented shape, but is unable to *report* that experience because of some problem with conceptualizing aspects of the forms of the objects experienced. If this were so, DF might indeed be making use of conscious shape information to control her successful actions, even though she reports no visual awareness of shape as such.

This is not an easy suggestion to bring into clear focus, but a concrete example may help. DF is adept at avoiding furniture and other obstacles as she crosses a room, even one of which she has no previous experience and hence no stored knowledge of layout, of the heights of tables and chairs, etc. It was my contention, following Milner and Goodale, that in so doing she was not consciously aware of the shapes and forms of the very objects she was successfully avoiding. More precisely, she was not (it was claimed) consciously aware of the objects' categories (tables versus chairs versus rocking horses), shapes or orientations.[15] That DF cannot, on the basis of her conscious experience, bring the perceived objects under categories ('table' 'chair', etc.) is, Wallhagen seems to allow, well supported by the behavioural evidence. DF is unable to name most of the objects she perceptually encounters, despite having knowledge of the names of objects and fully intact linguistic skills. But on what basis do we in addition deny that DF experiences the shapes and orientations of those very objects (that she fails to name)? Might DF in fact *experience* shapes, sizes, and orientations, but simply fail at the task of forming judgments (concerning shape, size, and orientation) on that basis? In short, might DF's problem be 'one of conceptualization, not one of consciousness per se' (Wallhagen [2007], pp. 539–61). The evidence for her supposed lack of experience of shape is that (1) DF denies experiencing shape information, and (2) she fails at tasks requiring her to identify objects by shape (she cannot position her hand so as to indicate, rather than grasp, the shape, or to match the orientation). But such evidence is, Wallhagen suggests, equally compatible with DF's experiencing shapes and orientations in some non-conceptual manner. Wallhagen thus suggests that:

> Aspects of form may well be phenomenally present to D.F., she may well consciously sense, and hence represent, the shapes, sizes, and orientations

[15] DF is, however, known to be aware of some colour and texture information, and where further information (e.g., concerning category or orientation) can be inferred from that, she is usually able to do so. See, for example, (Goodale and Milner [2006])

> of things, even if she cannot properly conceptualize these aspects of form. Similarly, if one is distracted, one may be aware of a sound outside without recognizing it, without conceptualizing it, as a particular type of sound (e.g., a car idling). Therefore, the data derived from the study of D.F., while very interesting, do not show that she lacks conscious awareness *of* form. At best, it shows that she lacks awareness *that* things have particular forms. (Wallhagen [2007], pp. 539–61)

This distinction, between bare awareness of form and the conceptualization of form is not, it seems to me, as clear as one might wish. But the general idea seems to be that DF might have some kind of non-conceptual but fully conscious experience of form, perhaps in the way we imagine a non-linguistic animal (a dog or a cat, for example) might experience the shapes and forms of tables and chairs without bringing the relevant form concepts (such as square, tall, narrow, slanted, etc.) to bear.

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). O'Regan and Noë here opt for 'partial awareness' since visual information is still playing an action guiding role in, for example, DF's successful posting (on demand) of a card through a visually presented slot. On the face of it this simply conflates visual awareness with the use of visual information, precisely the knot that Milner and Goodale were trying to untie. For this reason, Goodale ([2001]) rejects O'Regan and Noë's account of DF, claiming that DF 'shows almost perfect visuomotor control in the absence of any evidence that she actually "sees" the form of the object she is grasping'. But this now seems too hasty. For were it not for DF's verbal reports of lack of shape experience, surely we would treat her behavioural success as itself evidence for the ascription of experience of shape. The possibility that Wallhagen is raising is that these verbal reports really reflect only her failure to conceptualize shapes that she does, in some way, experience. To request a report is, after all, simply to demand that the subject bring her experiences under some concepts or other!

Wallhagen goes on to argue that not only is this 'deflationary' interpretation of DF possible, it is actually to be preferred. Two reasons are adduced. The first is that DF has excellent abilities to deal with objects according to their shapes. He contrasts this with the case of the blindsighter whose more limited abilities lead Wallhagen ([2007], pp. 539–61) to conclude they are genuinely unaware of form. The second concerns the difficulty in imagining DF's actual phenomenology. The worry here is that since DF can consciously see colours and textures, we must try to imagine such seeings combining with the lack of awareness of shaped and oriented regions bounding those colours and textures (*op. cit.*, pp. 539–61). Given these facts, and given the availability of the alternative deflationary explanation, the proper conclusion (Wallhagen argues) that what we find in DF is 'a failure of knowledge about shape, not a failure to [enjoy] conscious sensation of shape' (*op. cit.*, pp. 539–61). The upshot of all this is meant to be that:

> Since D.F. (probably) does consciously represent such aspects of form as size, shape, and orientation, we cannot rule out the possibility that these representations enter into the control and guidance of her motor behaviors. As such, Clark's appeal to D.F. does not provide compelling evidence against EBC—restricted or general (Wallhagen [2007], pp. 539–61).

6 Responding to the Non-conceptualist Challenge

How might a proponent of the 'zombic' interpretation of DF's shape-related performances respond to this non-conceptualist challenge? A simple response would be to deny outright the intelligibility of the notion of non-conceptualized visual experience. I shall not pursue such a response. Given that I am willing to accept the existence of genuine but non-conceptualized forms of experience, why not depict DF as enjoying genuine but non-conceptualized visual experience of shape and form?

First off, the positive reasons adduced by Wallhagen in favour of such a view seem unimpressive. It is not clear what follows from the difficulties in imaginability that Wallhagen describes. Given that our experience is not structured in the same way as DF's, such difficulties are surely to be expected. Nor should we simply assume that fluent abilities to deal with shaped objects imply conscious experience of shape. If this were so, then all manner of assembly line robots and other apparently non-conscious devices, not to mention a wide variety of insects and simple organisms, would need to be admitted, post-haste and without further argument, to the ranks of the conscious perceivers. It seems there must be more, even to non-conceptualized visual experience, than effective visuomotor response. Wallhagen's complaint thus raises a very serious and important challenge. What exactly is it about DF that should make even a non-conceptualist (but genuinely shape-experiential) reading unappealing?

One place to look for clues is by asking what else (if anything) DF cannot do. That is, suppose we try to put aside everything to do with report and with conceptualization and ask in what other ways (if any) DF's actual performance is compromised? In this regard it seems striking that DF, despite her putative non-conceptual experience of shape and form, is unable to put that experience into certain kinds of contact with practical action. Here are some examples. I will try to bring some unity to them after a brief review:

1. DF is quite adept at drawing objects from memory, but she is totally unable to draw or copy objects visually present in the here and now. Nor is she able to recognize the objects she has herself drawn as displaying the visual forms they do (see Servos and Goodale [1995]). It is not obvious (to me at least) that copying a visually presented form requires conceptualizing the form as such. Rather (assuming one understands the instruction and has the required skills) it seems like just the kind of thing even a non-conceptualized visual experience should enable one (assuming one can draw at all) to take a shot at.
2. DF is famously unable to self-prompt the appropriate use of her own form-accommodating visuomotor skills. For example, she would not spontaneously attempt to post the letter through the slot, even though she had previously been told that a large prize would be given to the first successful posting. Again, it is not clear to me that this is a deficit only in conceptualization. Rather, her experience does not (I want to say) present her with a world (not even a non-conceptualized world) in which visual form information can provide the basis for deliberate, reasoned, goal-and-project responsive action.
3. DF, though perfectly able (if asked) to grasp a screwdriver with a form-appropriate grip, will not typically choose the grip that is most apt given her needs and purposes (e.g., to use the screwdriver to tighten a screw). Thus it is not just that the experience fails to be brought under the linguistic concept 'screwdriver'. Rather, it is not an experience (not even a non-conceptualized experience) as of an object whose visually presented shape affords a certain kind of activity.
4. (Perhaps merely a variant on 2 and 3?). Suppose DF were to visually inspect a room. She would not (I claim) be able to see that, for example, the non-conceptualized object *over there* would fit into the non-conceptualized space *over here*, yet such judgments are surely paradigms of the kinds of judgment meant to be available to an agent in virtue of the non-conceptual contents of her experience. Of course, we cannot ask DF a question about what might fit where, on pain of requiring her to bring her experience under concepts. But if we could somehow set things up so that she had a goal of, say, blocking up the space that is (as it happens) between the sofa and the table, my strong bet is that she would nonetheless fail to experience the visual scene in such a way as to make the potential fit visible to herself. She would then not, for example, proceed to move the furniture even on a 'hunch' that the goal would be achieved.
5. Finally, even where DF manages to produce a behaviour that is indeed apt for the fulfillment of some goal, plan or project, she is not

able to assess her own unfolding activity as well suited to success. Thus she cannot see that her own unfolding posting behaviour is apt to the target. For if she could, why would she be surprised that it succeeds? It is not obvious (again, to me at least) that anything here turns on bringing the experience under concepts. Rather, what is missing is the visual awareness itself. (For more on this condition, see Campbell [2002]).

Uniting these various shortfalls, I want to suggest, is a common feature that might be described as the lack of a certain kind of link to personal agency. The suggestion that DF enjoys non-conceptual experience of visual form is undermined, it seems to me, by the lack of integration between this (putative) experience and DF's goal- and project reflecting choices and attempted actions. (Optic ataxics, by contrast, exhibit full integration—they are aware when visually presented information is apt for their goals and purposes—and are compromised instead in putting this information to effective use).

The resolute skeptic, however, may still be unconvinced. For at best I may have shown that:

1. There is a complex of skills missing in DF, roughly groupable as the 'link to personal agency'

   and

2. That the absence of these skills is not simply due to her failure to conceptualize her own experiences of shape and form.

Even granting this much, what remains missing, of course, is some independent reason to believe that the absence of these very skills (and more generally of the link to personal agency) undermines the claim that DF nonetheless enjoys rich but non-conceptualized experiences of visually presented shape and form. Why link phenomenal experience to integration with (even non-conceptual versions of) agent-level goals, plans and projects? With this question we enter. . .

7 The Murk

Why suppose there is some deep connection between visually presented information being linked to personal agency (i.e., to an agent's choice and understanding of her own actions, in the rather specific, sub-linguistic, ways just described) and that information's being present in first person phenomenal experience? One reason is simply the empirical association between verbal reports of (lack of) visual experience of form and the unavailability of this coherent functional complex. The skeptic, however, may well concede this link while insisting that it in no way bears on the questions about visual experience. The cost of this move, it seems to me, is to drive an unwelcome wedge between the agent and her own (putative) experience. It is to allow experience to occur in something like informationally isolated islands within the overall agent-economy. We should ask ourselves whether this idea really makes sense.

Gareth Evans famously rejects the very idea of such informationally isolated islands of experience. According to Evans, an informational state may underpin a conscious experience only if it (the informational state) is in some sense input to a reasoning subject. 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 the real point here is (or should be) independent of Evans' appeal to conceptualization. Thus consider an Advanced Non-conceptual Animal (Ana for short). Ana, by hypothesis, lacks all forms of conceptualized thought. But Ana can form non-conceptualized goals, and can spot environmental opportunities that allow the fulfillment of those goals and projects. Ana can grasp visually presented objects in ways dictated by her intended future use. She can choose between courses of action on the basis of her conscious experience of what the world, as currently visually presented, offers. She can execute her motor routines with extra, and appropriate, caution when she recalls hidden dangers. She can also perceptually judge when her ongoing action is falling short of its target and take remedial action. If she could talk (which she cannot) she would not in general be surprised by the success of her own actions in the world. She can, in short, put visual information to use in the service of her personal (albeit, limited and non-conceptualized) goals and projects, and she can see and monitor the aptness of her own actions to their goals. Ana, I want to say, meets the plausible requirement at the heart of Evans' account, viz, to put it brutally, that *visual experience should be the experience of an agent*, where that is cashed in terms of the direct availability of that information for use in the planning and selection of deliberate, stored knowledge-exploiting, goal-reflecting, and goal-responsive actions. This is what is lacking (I claim) in DF. It may even be (I am strongly tempted to argue) that this lack is ultimately *constitutive* of DF's lack of conscious visual experience of shape and form. Were shape and form information directly available to DF for use in these kinds of ways, she simply *would* consciously experience shape and form.

This further reductive step is by no means compulsory, but I think it is worth considering (for some early exploration, see Roberts, Ward and Clark [unpublished]).

I do not suppose that any of this will be enough to satisfy the critic who doubts the existence of any link between experience and those features of personal agency I have been describing. Such a critic, by endorsing the idea of experience as that which is somehow phenomenally present to what I dubbed the 'bare agent', places the discursive goalposts in a most inaccessible position. But I hope to have at least gestured at a possible rationale for refusing such a vision of bare experience. For the bare-experience-laden agent is not really an agent at all, but rather an unanchored islet of experience isolated within a larger economy.[16] Surely experience should, as Evans recognized, have more to do with agency than this. By tying conscious visual experience to the complex concerning deliberate action and personal agency, we begin to reveal a functional role distinctive of the information-processing episodes underlying conscious seeing. Conscious seeing, if this is correct, is about putting currently transduced information into contact with stored knowledge and expectations, and doing so in ways that support just the kinds of reasoned, self-intelligible, and flexibly goal-responsive behaviours we have come to expect of those who genuinely experience, and hence can knowingly engage, their worlds.

8 Conclusions

Wallhagen's arguments, I have tried to show, fail to undermine the reasons for rejecting EBC (gen) in favour of EBS. It is not simply minor adjustments and fine motor corrections that are handled by the visual zombie systems. Nonetheless, we should also be aware of the potential role of the dorsal stream in at least some forms of conscious experience (e.g., the experience of presence) and of the many complex interactions between zombie systems and the rest. We should also be aware of the possibility that conscious experience itself presents multiple inconsistent contents. Such complexities, I have tried to show, merely enrich our understanding of what still seems to be a real and important division of labour within the neural economy.

Second, and more importantly, the claim that the evidence and arguments for such a view reveal only failures of noticing and conceptualization, rather than true failures of conscious awareness, is not convincing. It is not convincing

[16] An anonymous referee suggests that information about visual form, in DF, might be counted (contra the account on offer) as poised for agent-level use by DF but simply blocked, by the neural damage, from performing that role in regard to certain tasks. But the considerations in the closing sections aim to motivate the thought that actual availability for use in planning, reasoning, and stored-knowledge exploiting response is required for information to thus count as 'given to the agent'. So it is not clear to me in what way we might then depict the form information as (in the sense at issue) nonetheless 'there for DF'.

if we are skeptical of the truly radical (and perhaps empirically intractable) vision of bare islets of awareness, and instead insist on preserving what are arguably non-negotiable links between what is given in conscious awareness and the enabled sweep of deliberate actions and choices by a thinking subject. Such links can be made intelligible, I have tried to show, even where the notion of the thinking subject is stripped of its conceptualist baggage. The guiding idea here is that there exist essential, though still not fully understood, links between phenomenal experience and the capacity to factor information into (potentially non-linguistic, limited, but still perfectly real) forms of reasoning about practical action. Making good on such a link or links would go a long way towards revealing just why conscious awareness is (as it seems to be) so closely associated with neural subsystems that specialize in the categorization and classification of environmental stimuli, and in putting incoming information into contact with long-term stored knowledge.

Notice finally[17] that, even though the story on offer is one of quite significant divisions of labour and fragmentation among the neural origins of our own visuomotor actions, there is at least one place (though not in the brain) where it does all finally 'come together'. That place is, of course, the world itself, whose constant presence binds together the unfolding activity of the two streams. Courtesy of the stable presence of the world, conscious seeing and acting are tightly bound and synchronized, linked by multiple re-entrant loops. That, it seems reasonable to assert, is why we find ourselves acting in the very world we experience, and why experience has the world-presenting character it does. We are wrong if we think that conscious visual experience is itself providing the moment-by-moment information used to guide many key aspects of our fast, fluent, world-engaging action. But we are right to think that perceptual experience informs us about the world we act in and that we act in the very world we perceive.

Acknowledgements

Thanks are due to Rob McIntosh, Julian Kiverstein, Matthew Nudds, Susan Hurley, Tom Roberts, Dave Ward, Tim Hodgson, and the members of the Edinburgh CONTACT team. Thanks are also due to the anonymous referees whose helpful and constructive comments did much to improve the final version of this paper. This work was funded by the Arts and Humanities Research Council, under the European Science Foundation's Eurocores: Consciousness

[17] This point is due to Tim Hodgson, whose comments on an earlier version were invaluable in preparing the present paper.

in a Natural and Cultural Context (CNCC) scheme, to A. Clark as part of the CONTACT (Consciousness in Interaction) project, AH/E511139/1.

School of Philosophy, Psychology, and Language Sciences
University of Edinburgh
George Square
Edinburgh EH8 9JX
UK
andy.clark@ed.ac.uk

References

Aglioti, S., Goodale, M. and DeSouza, J. F. X. [1995]: 'Size Contrast Illusions Deceive the Eye But Not the Hand', *Current Biology*, **5**, pp. 679–85.

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**, pp. 692–700.

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), pp. 109–13.

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

Clark, A. [2002]: 'Is Seeing All It Seems? Action, Reason and the Grand Illusion', *Journal of Consciousness Studies*, **9**(5/6), pp. 181–202.

Cussins, A. [1998]: 'Nonconceptual Content, Frames of Reference and Trails of Information', Manuscript of talk presented to the Symposium on Nonconceptual Content, American Philosophical Association, Eastern Division Meeting, Washington DC, January 1998.

Decety, J. and Grezes, J. [1999]: 'Neural Mechanisms Subserving the Perception of Human Actions', *Trends in Cognitive Sciences*, **3**(5), pp. 172–8.

Dennett, D. [1991]: *Consciousness Explained*, New York: Little Brown.

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.

Ellis, R., Flanagan, J. and Lederman, S. [1999]: 'The Influence of Visual Illusions on Grasp Position', *Experimental Brain Research*, **125**, pp. 109–14.

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

Fecteau, J. H., Chua, R., Franks, I. and Enns, J. T. [2001]: 'Visual Awareness and the On-Line Modification of Action', *Canadian Journal of Experimental Psychology*, **55**, pp. 104–10.

Glover, S. [2004]: 'Separate Visual Representations in the Planning and Control of Action', *Behavioral and Brain Sciences*, **27**, pp. 3–78.

Goodale, M. [1998]: 'Where Does Vision End and Action Begin?', *Current Biology*, **8**, pp. R489–91.

Goodale, M. [2001]: 'Real Action in a Virtual World: Commentary on O'Regan and Noë', *Behavioral and Brain Sciences*, **24**(5), pp. 984–5.

Goodale, M. and Milner, A. D. [2004]: *Sight Unseen: An Exploration of Conscious and Unconscious Vision*, Oxford: Oxford University Press.

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**, pp. 203–21.

Goodale, M. A., Pelisson, D. and Prablanc, C. [1986]: 'Large Adjustments in Visually Guided Reaching do not Depend On Vision of the Hand or Perception of Target Displacement', *Nature*, **320**, pp. 748–50.

Hommel, B., Müsseler, J., Aschersleben, G. and Prinz, W. [2001]: 'The Theory of Event Coding (TEC): A Framework for Perception and Action Planning', *Behavioral and Brain Sciences*, **24**(5), pp. 849–78.

Jackson, S. and Shaw, A. [2000]: 'The Ponzo Illusion Affects Grip Force but not Grip-Aperture Scaling During Prehension Movements', *Journal of Experimental Psychology: Human Perception & Performance*, **26**, pp. 418–23.

Jacob, P. and Jeannerod, M. [2003]: *Ways of Seeing: The Scope and Limits of Visual Cognition*, Oxford: Oxford University Press.

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**, pp. 2463–75.

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

Jeannerod, M. [2006]: *Motor Cognition: What Actions Tell the Self*, Oxford: Oxford University Press.

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

Johnson, H. and Haggard, P. [2005]: 'Motor Awareness without Perceptual Awareness', *Neuropsychologia*, **43**, pp. 227–37.

Jordan, S. [2003]: 'Emergence of Self and Other in Perception and Action: An Event-Control Approach', *Consciousness and Cognition*, **12**, pp. 633–46.

Kanwisher, N. [2001]: 'Neural Events and Perceptual Awareness', *Cognition*, **79**, pp. 89–113.

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), pp. 1–16.

Le, S., Cardebat, D., Boulanouar, K., Henaff, M. A., Michel, F., Milner, A. D., Dijkerman, C., Puel, M. and Demonet, J. F. [2002]: 'Seeing, Since Childhood, Without Ventral Stream: A Behavioural Study', *Brain*, **125**, pp. 58–74.

Libet, B. [1985]: 'Unconscious Cerebral Initiative and the Role of Conscious Will in Voluntary Action', *Behavioral and Brain Sciences*, **8**, pp. 529–66.

Matthen, M. [2005]: *Seeing, Doing and Knowing*, Oxford: Oxford University Press.

McIntosh, R. D., McClements, K. I., Schindler, I., Cassidy, T. P., Birchall, D. and Milner, A. D. [2004]: 'Avoidance of Obstacles in the Absence of Visual Awareness', *Proceedings of the Royal Society Series B-Biological Sciences*, **271**, pp. 15–20.

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

Milner, A. D. and Goodale, M. [2006]: 'Epilogue: Twelve Years On', in A. D. Milner and M. Goodale (*eds*), *The Visual Brain in Action*, (2nd edn), Oxford: Oxford University Press, pp. 207–52.

Milner, A. D., Dijkerman, H. C., McIntosh, R. D., Pisella, L. and Rossetti, Y. [2003]: 'Delayed Reaching and Grasping in Patients with Optic Ataxia', *Progress in Brain Research*, **142**, pp. 225–42.

Milner, A. D., Perrett, D. I., Johnston, R. S., Benson, P. J., Jordan, T. R., Heeley, D. W., Bettucci, D., Mortara, F., Mutani, R., Terazzi, E. and Davidson, D. L. W. [1991]: 'Perception and Action in "Visual form Agnosia" ', *Brain*, **114**, pp. 405–28.

Noë, A. [2001]: 'Experience and the Active Mind', *Synthese*, **129**, pp. 41–60.

Noë, A. [2002]: 'Is the Visual World a Grand Illusion?', *Journal of Consciousness Studies*, **9**(5–6), pp. 1–12. Also available in A. Noë (*ed.*), *Is the Visual World a Grand Illusion?*, Thorverton, Imprint Academic, pp. 1–12.

Noë, A. and O'Regan, J. K. [2000]: 'Perception, Attention and the Grand Illusion', *Psyche*, **6**(15), <psyche.cs.monash.edu.au/v6/psyche-6-15-noe.html>.

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.

Roberts, T., Ward, D. and Clark, A. [unpublished]: 'Knowing What We Can Do'.

Rossetti, Y., Pisella, L. and Vighetto, A. [2003]: 'Optic Ataxia Revisited: Visually Guided Action Versus Immediate Visuomotor Control', *Experimental Brain Research*, **153**, pp. 171–9.

Sakata, H. [2003]: 'The Role of the Parietal Cortex in Grasping', *Advances in Neurology*, **93**, pp. 121–39.

Schindler, I., Rice, N. J., McIntosh, R. D., Rossetti, Y., Vighetto, A. and Milner, A. D. [2004]: 'Automatic Avoidance of Obstacles is a Dorsal Stream Function: Evidence from Optic Ataxia', *Nature Neuroscience*, **7**, pp. 779–84.

Servos, P. and Goodale, M. [1995]: 'Preserved Visual Imagery in Visual form Agnosia', *Neuropsychologia*, **33**, pp. 1383–94.

Schenk, T. and Milner, A. D. [2006]: 'Concurrent Visuomotor Behavior Improves form Discrimination in a Visual form Agnosic', *European Journal of Neuroscience*, **24**, pp. 1495–503.

Shull, J. A. and Bingham, G. P. [2000]: 'Two Visual Systems Must Still Perceive Events', *Behavioral and Brain Sciences*, **25**, pp. 118–9.

Smeets, J. and Brenner, E. [2001]: 'Action Beyond Our Grasp', *Trends in Cognitive Sciences*, **5**(7), p. 287.

Taira, M., Mine, S., Georgopoulos, A. P., Murata, A. and Sakata, H. [1990]: 'Parietal Cortex Neurons of the Monkey Related to the Visual Guidance of Hand Movement', *Experimental Brain Research*, **83**, pp. 29–36.

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

Zeki, S. and Bartels, A. [1999]: 'Towards a Theory of Visual Consciousness', *Consciousness and Cognition*, **8**, pp. 225–59.