



Topics in Cognitive Science 2 (2010) 643–657

Copyright © 2010 Cognitive Science Society, Inc. All rights reserved.

ISSN: 1756-8757 print / 1756-8765 online

DOI: 10.1111/j.1756-8765.2010.01084.x

Control Consciousness

Pete Mandik

Department of Philosophy, William Paterson University

Received 24 April 2009; received in revised form 5 January 2010; accepted 12 January 2010

Abstract

Control consciousness is the awareness or experience of seeming to be in control of one's actions. One view, which I will be arguing *against* in the present paper, is that control consciousness is a form of sensory consciousness. In such a view, control consciousness is exhausted by sensory elements such as tactile and proprioceptive information. An opposing view, which I will be arguing *for*, is that sensory elements cannot be the whole story and must be supplemented by direct contributions of non-sensory, motor elements. More specifically, I will be arguing for the view that the neural basis of control consciousness is constituted by states of recurrent activation in relatively intermediate levels of the motor hierarchy.

Keywords: Consciousness; Motor control; Neurophilosophy; Imagery; Will; Agency; Phenomenology

1. Introduction

Control consciousness is the awareness or experience of seeming to be in control of one's actions. In the sense that I use the term “control consciousness,” states of control consciousness can occur both when the appearance is accurate (one really is in control) and when it is inaccurate (one really is not in control).¹ Key questions in the present investigation are the following: Are these appearances exclusively sensory? Are states of control consciousness a kind of exclusively sensory state explicable either as a kind of sensory perception or as a kind of sensory imagery? Or does control consciousness instead involve, at least partially, distinctively nonsensory elements such as motor commands or states of the will?

My main interest in the current paper is to address these questions by developing an extension of my Allocentric-Egocentric Interface theory of consciousness (AEI) (Mandik,

Correspondence should be sent to Pete Mandik, Department of Philosophy, William Paterson University, 300 Pompton Road, Wayne, NJ 07470. E-mail: mandikp@wpunj.edu

2005, 2008, 2009). My previous expositions of AEI have focused on sensory consciousness, in particular, visual consciousness. However, my proposal for extending AEI to control consciousness will *not* proceed by assimilating control consciousness to a kind of sensory consciousness. I will instead be developing the application of AEI to control consciousness by arguing for the direct involvement of distinctively nonsensory control signals.

Now, I do not deny that sensory consciousness is often part of the story. When I consciously experience flipping a fried egg without breaking the yolk or consciously experience attaching a delicate component to a scale model, much of my complex conscious experience integrates what I *see* arrayed before me as well as what I *feel* in my skin and muscles. Nonetheless, despite acknowledging the role that perceptual input plays in contributing to control consciousness, I will argue that *some* aspect of the consciousness involves (in a *direct* way) nonsensory signals.

Before continuing, it will be useful to make further clarifying remarks concerning the definition of “control consciousness.” The following pair of distinctions concerning mental states will be useful in clarifying what I intend by “control consciousness”: The first is a distinction between states that are sensory states and states that are control states. The second is a distinction between states that are conscious and states that are unconscious.

Many mental states and processes may be put into one of two categories: Those that are more closely involved with sensory or input systems and those that are more closely involved with control or output systems. (There may additionally be mental states and processes that belong in neither group, but they are of minor concern to the present project.) It is uncontroversial that sensations and sensory perceptions belong in the first category. It is slightly more controversial to hold that sensory images do too. According to some accounts (such as Kosslyn, 1994), sensory images are endogenously triggered sensory representations whose exogenous triggering are crucially involved in sensory perception. (The controversy concerning the degree to which imagery is a sensory state will be discussed in section 3.) Mental states more closely involved with output systems include, for example, an intention to raise one’s hand and the hypothesized motor plans and motor commands that are the causal antecedents of one’s hand being raised. Into this latter group belong various states of will or “willings”—states via which our movements are controlled. For ease of exposition, I shall refer to the distinction between more input-related states and more output-related states as a distinction between sensory states and control states.

I turn now to the distinction between conscious and unconscious mental states. We are perhaps most familiar with states of sensory consciousness. Many readers of the current sentence are undergoing a conscious visual state of words on a page or screen. In addition to conscious sensory states, there is evidence of unconscious sensory states, as in states of subliminal perception. Just as various sensory states come in conscious and unconscious varieties, so too do various output-related states. One may consciously decide to go back and reread the previous sentence or to close one’s eyes and count to three out loud. In addition to such conscious control states there are control states that are unconscious. Examples of unconscious control states include the control states that are the causal antecedents of actions performed when one is “on autopilot” or doing something “absent-mindedly.” It should be noted here that the claim that there are conscious control states that are not a kind

of conscious sensory state is a matter of some controversy. It is one of the main aims of the present paper to address this controversy. One side of the controversy—a side that I will be opposing—holds that so-called conscious control states are really just a variety of conscious sensory state. Much more on this will be discussed in sections 2 and 3.

The organization of the remainder of the paper is as follows: In the next section, I unpack AEI and its application to control consciousness. The remaining sections concern objections to my account.

2. The Allocentric-Egocentric Interface theory of consciousness

The general aims and methods of AEI are *neurophilosophical*. Neurophilosophy is an interdisciplinary enterprise that brings neuroscientific results and methodologies to bear on philosophical concerns, especially as they arise in the philosophy of mind (Bickle, Mandik, & Landreth, 2006; Brook & Mandik, 2007). The main aim of neurophilosophical theories of consciousness is to bring neuroscience to bear on the following central questions of consciousness (Mandik, 2007, p. 420):

The Question of State Consciousness:

In what consists the difference between mental states that are conscious and mental states that are unconscious?

The Question of Transitive Consciousness:

When one has a conscious mental state, what is one thereby conscious of?

The Question of Phenomenal Character:

When one has a conscious state, in what consists the properties in virtue of which there is something it's like for one to be in that state?

For a review of neurophilosophical theories of consciousness, see Mandik (2007).

2.1. AEI and sensory consciousness

I begin my exposition of the application of AEI to sensory consciousness by focusing on the answer to the question of state consciousness. Not just any input to sensory systems gives rise to a conscious percept. Instances of subliminal perception and blindsight are two kinds of example. The solution I advocate for distinguishing conscious from unconscious perception is twofold. I shall label the two parts of the solution “intermediacy” and “recurrence.” The first part, intermediacy, involves identifying conscious perceptual states with states at intermediate levels of sensory-processing hierarchies. The second part, recurrence, restricts consciousness to intermediate-level states involved in recurrent interaction between representations at high and at low levels of sensory-processing hierarchies.

The “what” and “why” of intermediacy. Sensory processing, as in vision, for example, is hierarchical. In this processing hierarchy the lowest levels are neural activations close to the sensory periphery that represent local and egocentric visible features and the highest levels are abstract, invariant, and allocentric representations employed in categorization and recognition.

It is natural to ask *where* in a sensory-processing hierarchy conscious states reside. It is crucial to any account of consciousness that it connect the *reality* accessible from the third-person point of view (e.g., states of activation in neural circuits) with the *appearance* of what it’s like from the first-person point of view. Further, both introspective methods (first-person methods) and observational methods (third-person methods) converge to indicate that conscious states are *intermediate* between the highest and lowest levels of the hierarchy. My visual perception of a coffee cup represents the cup as having a specific orientation relative to my point of view and a specific location in my visual field. However, the percept is not so high level as to merely indicate the presence of a cup in a way abstracting from all observer-relative information. Nor is it so low level as to register every change in irradiation of various regions of my two retinas (the lowest levels are prior to even the integration of information from the disparate retinas). The intermediacy criterion on sensory consciousness means that not just any neural response to a sensory input will count as a conscious percept. This “Goldilocks criterion” will exclude from consciousness those neural activations that are too high or too low.²

The “what” and “why” of recurrence. While intermediacy is necessary, it does not seem to alone suffice for consciousness. Strictly feed-forward activation of representations in sensory-processing hierarchies can occur without consciousness. Pascual-Leone and Walsh (Pascual-Leone & Walsh, 2001) showed, with precisely timed pulses of transcranial magnetic stimulation, that visual consciousness was suppressed if recurrent activation was suppressed and only feed-forward activation was allowed. Additionally, Lamme et al. (Lamme, Supèr, & Spekreijse, 1998) suggest that responses to stimuli in animals under general anesthetic are feed-forward activations without accompanying recurrence.

Allocentric-Egocentric Interface theory incorporates both intermediacy and recurrence in the following manner: Conscious states are intermediate-level states in processing hierarchies that are constituted by pairs of recurrently interacting allocentric and egocentric representations. Thus, do we have an answer to the question of state consciousness. As for the question of transitive consciousness (when one has a conscious state, what is one conscious of?) the answer delivered by AEI is that what one is conscious of is one and the same as what the reciprocally interacting egocentric and allocentric representations are representations of. As for the question of phenomenal character, the answer AEI delivers is that phenomenal character is to be identified with the representational contents of the reciprocally interacting egocentric and allocentric representations.

In capsule form, the AEI account of sensory consciousness identifies states of sensory consciousness with states of recurrent activation in intermediate levels of sensory-processing hierarchies. In section 2.3, I will spell out the case for identifying states of control consciousness with states of recurrent activation in intermediate levels of motor-processing hierarchies. But before spelling that out, I briefly address, in section 2.2, the question

of whether control consciousness should simply be regarded as a species of sensory consciousness.

2.2. Why not a sensory theory of control consciousness?

There are various reasons why it might be appealing to assimilate control consciousness to a form of sensory consciousness. First, of our various mental states, our states of sensory consciousness are the ones most vividly present to us. Second, sensory consciousness, especially visual consciousness, is, from the point of view of science, perhaps our best-understood form of consciousness. Third, the institution of science itself is influenced by a long tradition of empiricism, an early motto of which is that there is nothing in the mind that is not first in the senses.

What, then, would it mean to assimilate control consciousness to sensory consciousness? It is natural among researchers to take deliberate bodily motion as a basic and paradigmatic case of an action. And the most natural way to assimilate the consciousness of deliberately moving parts of one's own body to a form of sensory consciousness is to do so in terms of sensory feedback from the muscles, tendons, and skin as the body parts in question are moved. Part of the view that control consciousness is sensory is the denial of any *direct* contribution of a motor-command signal to the subjective aspect of control consciousness. It allows *indirect* contributions as when, for example, motor commands trigger musculo-skeletal activity, which in turn triggers sensory feedback. However, the denial of direct contributions leads to the key weakness of the sensory theory of control consciousness.

The sensory theory has the implausible consequence that there could be an arm movement that results from a subject's issuing a motor command but, due to effects of anesthesia, is unaccompanied by sensory feedback and, lacking sensory feedback, the subject would be completely unaware of having either moved or even having tried to move his or her own arm. However, as pointed out by Prinz (2007a, p. 344) and Peacocke (2007, p. 359), a subject may be quite aware of moving a body part even while not perceiving the part due to either local anesthetic or severing of afferent nerves.

Sensory input alone seems insufficient for control consciousness. Something more than sensory input is needed to account for such cases. I turn now to examine how AEI can be extended to show what, in addition to sensory input, is needed for control consciousness.

2.3. AEI and control consciousness

I turn now to sketch a nonsensory theory of control consciousness. In its most basic form, the theory is comprised of a pair of theses, one negative and one positive. The negative thesis is that sensory inputs are alone insufficient for distinguishing between the conscious states associated with controlled movements and the conscious states associated with *mere* movements (movements that are not controlled, or are not the result of control signals). The positive thesis is that control signals such as motor commands make a direct contribution to states of control consciousness.

The notion of *directness* employed in the positive thesis requires further clarification. The main idea of “directness” here means something like “not mediated by sensory inputs or imagery thereof.” To illustrate, even on the sensory theory, motor commands have an *indirect* influence: I turn my head and thus see something other than what I was looking at. But here the changes exerted on the conscious state by the motor command are mediated by changes in sensory input. If motor commands themselves or copies of motor commands—so-called efference copies—can make a difference on conscious states without the difference being mediated by changes caused to sensory inputs, then this would be an instance of the influence of motor commands being *direct*.

As stated, the most basic form of the theory is inadequate for distinguishing conscious from unconscious control states. Not just any contribution from motor commands will make a contribution to consciousness. This is especially evident in the case of unconscious control states. What is needed, then, is a means for distinguishing conscious control states from unconscious control states. And here an extension of AEI can do the trick.

The applicability of AEI to motor systems looks to be a straightforward affair. First, motor systems are arranged hierarchically in much the same manner as sensory systems, with the highest levels being most remote (in terms of intervening connections) from the periphery of the nervous system. Focusing here just on cortex, the highest level is in the prefrontal cortex, the lowest level is in the primary motor cortex, and intermediate is in the premotor cortex. Further, there exist both forward projections and back projections between successive levels of the motor hierarchy (Churchland, 2002, p. 72). We may further characterize levels in the motor hierarchy as differing along an allocentric-egocentric dimension. To illustrate, a high-level motor representation may have an abstract content such as *Grab a coffee mug*—a content that abstracts away from low-level egocentric details about the precise sets of muscular contractions required to get the job done. The most specific details will be the task of motor representations at the lowest levels of the motor hierarchy.

The neuroanatomical features of the motor system make it quite natural to suppose that both intermediacy and recurrence can apply to motor processing. The basic suggestion here is twofold.

First, *conscious* control states are states consisting in reciprocally interacting pairs of intermediate-level motor representations where one member of the pair is more allocentric than the other. Second, *unconscious* control states are control states that are either not intermediate (they are too high or too low) or are intermediate but lack the requisite recurrence.

One question worth considering about the current proposal is the question of what function recurrence is playing in control consciousness. Here I think insight can be gained from a certain interpretation of some ideas from control theory.

Many philosophers are aware of control theory via the work of Rick Grush (e.g., Grush, 2001) and I here rely on his exposition of its basic ideas. In the simplest kind of control system, *open-loop control*, a desired goal signal is fed into a controller, which sends control signals to a target system or “plant.” Applying these concepts to motor control involves viewing parts of the musculoskeletal system as plants and the neural systems generating motor commands as controllers. The controller implements a mapping, the *inverse mapping*, of goal states onto command sequences. The plant implements a mapping, the *forward*

mapping, of command sequences onto goal states (Grush, 2001, pp. 352–353). A slightly more complex control system, *closed-loop control*, has all of the same components as in open-loop control, along with the addition of feedback signals from the plant to the controller. While for many control purposes, closed-loop control is superior to open-loop control, closed-loop-control is not without certain problems. If, for example, there are significant delays in the receipt of the feedback signal due to slow signal speeds and/or a distant plant, then the system can oscillate wildly through potentially destructive cycles of overshooting and overcompensation. A slightly more complex control system that potentially overcomes such problems is *pseudo-closed-loop control*. One way of conceiving of pseudo-closed loop control is by thinking of it as built by adding features to open-loop control. The first addition involves a second signal being sent by the controller: an *efferent copy*, which is a duplicate of the signal sent to the plant. This duplicate signal, however, is not sent to the plant, but instead to an *emulator* or *forward model*, which in turn sends signals back to the controller.

With these key concepts of control theory in hand, we can appreciate the following proposal for the function that recurrence is playing in the AEI account of control consciousness: Recurrent signaling implements a pseudo-closed-loop control architecture. Outgoing signals from the highest levels of the hierarchy may be identified with the specification of a goal state. The next lowest level receives the goal state and sends on the inverse mapping. This inverse mapping may be sent to the lowest levels, eventuating in command signals. But it or, more precisely, a copy of it, may be sent down to intermediate areas wherein activation is utilized as a forward model, with results that may be propagated back up to higher levels.

We may view the relation of pseudo-closed loop control to control consciousness in the following manner: Open-loop control would implement a kind of “shooting in the dark” where one is acting but has no awareness that one is acting. Both closed-loop control and pseudo-closed-loop control introduce elements that circumvent this problem, and pseudo-closed-loop control does so in a manner that has certain advantages over closed-loop control (such as avoiding harmful feed-back-induced oscillations).

3. Why not a sensory imagery theory of control consciousness?

In the present section, I want to examine and argue against an alleged alternative to the nonsensory account of control consciousness. Such an alternative tries to account for control consciousness in control-theoretic terms while at the same time regarding the hypothesized pseudo-closed-loop control architecture as involving a kind of sensory imagery. On such a view, the forward model is interpreted as in some way being sensory.

Now, it is tempting to follow Grush in calling the signal from the forward model “mock *sensory* information about what the real target system would do under various conditions” (p. 356, emphasis added), but I will want to resist such temptation. Yielding to such a temptation would result in an account of control consciousness wherein control consciousness turns out to be a form of sensory consciousness after all. Prinz (2007a) has recently defended a sensory account of control consciousness that has as its key thesis the view that forward models are a kind of sensory image.

Prinz (2007a) supplies a concise statement of his view, captured here in the following quotation:

The feeling of agency could be explained by a kind of prediction that the brain makes when we are about to act. If you elect to move your arm, you will be able to anticipate its movement. According to some leading neurobiological theories, when a plan is generated in the premotor cortex, a representation is sent to the somatosensory cortex corresponding to what the bodily senses should perceive when that action is executed. That representation is called a “forward model.” A forward model is an anticipatory somatosensory image. When our bodies carry out motor plans, the forward model is compared with the actual changes that take place in our body as we move. The feeling of agency may arise from this matching process. If a match occurs, we feel we are in control. If a match doesn’t occur, it’s because our bodies didn’t move as we predicted they would, and that results in an experience of being passively moved by an external source. (p. 342)

One way of appreciating a problem with Prinz’s view involves the way it combines a concept from control theory—that of a forward model—with the concept of a sensory image. That forward models are involved in the control of bodily movement is a highly plausible suggestion. That they be regarded as sensory images is somewhat less plausible.

It is useful here to consider the following two questions: First, what is involved in something’s being *sensory* in the sense of the term relevant to the current discussion? Second, do we have adequate reason for thinking that a forward model is *sensory*?

Starting with the first question, it is useful to look at Prinz’s own account of what makes something sensory. Prinz writes:

I will define a perceptually conscious mental state as a mental state that is couched in a perceptual format. A perceptual format is a representational system that is proprietary to a sense modality. To say that phenomenal states are perceptual is to say that their representational vehicles always belong to one of the senses: touch, vision, audition, olfaction, and so on. (Prinz, 2007a, p. 336)

Further elaboration comes from what Prinz takes the negative aspects of his key thesis to be: “We do not have conscious states couched in non-perceptual formats. If I am right, we never have conscious states in our motor systems, and no conscious experiences are constituted by amodal representations...” (Prinz, 2007a, p. 337).

In an earlier work dedicated to elaborating Prinz’s brand of empiricism, he spells out his view that “the senses are *dedicated input systems*” (Prinz, 2002, p. 115). Crucial to Prinz’s characterization is that each sense has both a proprietary class of inputs (physical magnitudes) and a proprietary representational format (thus denying that separate senses share a “common code” [Prinz, 2002, p. 117]).

It is worth noting that in this earlier work Prinz endorses a view of imagery whereby “we can form mental images by willfully reactivating our input systems” (Prinz, 2002, p. 115).

It seems natural to suppose that what is responsible for these reactivations counting as *sensory* imagery is that they are *input* systems that are reactivated.

With these remarks about what the “sensory” in “sensory imagery” consists in, let us return to the question of whether forward models need be conceived of as sensory imagery. In the basic outlines of pseudo-closed-loop control, there is nothing that makes compulsory a sensory-imagery interpretation of the forward model. The forward model is not receiving sensory inputs and thus cannot count as a *sensory* system as characterized by Prinz. A fortiori, it cannot count as sensory *imagery*, as it does not count as the *reactivation* of an input system.

Of course, it should be noted that there may be alternate architectures that incorporate forward models satisfying criteria for being sensory.³ However, the core idea of a forward model does not alone satisfy such criteria.

The criticism of Prinz’s view that I have offered so far is that while it may be plausible to try to account for control consciousness in terms of pseudo-closed-loop control, this does not by itself suffice to make the resultant theory a sensory-imagery theory of control consciousness. I turn now to present a different criticism of Prinz’s view—namely, that even sensory imagery involves a direct contribution of control signals to control consciousness. The key idea here is that given a characterization of imagery as the willful reactivation of input systems, what distinguishes sensory imagery from sensory perception is the presence, in imagery, of the control signals that reactivate input systems.

Even though there are various commonalities between sensory perception and sensory imagery, the main way in which we are able to distinguish an image from a percept with similar content is by the differential degrees of direct control that we have over the image (Kosslyn, 1994, pp. 102–104). For example, in imagining an apple, I can easily rotate, enlarge, or distort the shape of the *imagined* apple. But I cannot enact such transformations on a *perceived* apple unless I can move my body or parts thereof.

Because of similarities between percepts and images, subjects do sometimes confuse the two (Perky, 1910). However, the degree to which subjects confuse a percept and an image can be manipulated experimentally by introducing factors that either vary how difficult the imagery task is (Finke, Johnson, & Shyi, 1988) or whether the images are created intentionally rather than incidentally (Durso & Johnson, 1980; Intraub & Hoffman, 1992). An intentionally formed and difficult to manipulate image (say, an image of a rotating, complex three-dimensional figure) is less likely to be mistaken in memory for a percept than a comparatively less difficult image.

It is worth here spelling out the superiority of the AEI account of control consciousness over the sensory-imagery theory of control consciousness with regards to the above points concerning the distinguishability (and occasional lack thereof) of imagery and perception. The key distinction between the imagery theory and the motor theory is that the imagery theory does not allow for any direct awareness of the contributions of control signals. According to the imagery theory, any influence on consciousness that control signals have will always be indirect—that is, mediated by the reactivation of sensory representations. But problems arise when we realize what the difference between imagery and perception amounts to. If (a) the only difference between conscious perception and conscious imagery

is that the former has exogenous causes and the latter has endogenous causes and (b) the subject has no direct consciousness of the causes but only the sensory effects (the activations of sensory representations), then there should be no phenomenological difference between perception and imagery. If an adherent of the imagery theory endorses a view, as Prinz does, that imagery is the willful reactivation of sensory elements, and also endorses the empiricist view that the actual presence of willfulness does not show up in consciousness, then such an adherent seems thereby committed to there being no phenomenological difference between imagery and perception. That there is in fact such a difference thus seems to favor a theory such as AEI, which allows for a direct phenomenological contribution of nonsensory control signals.

Before proceeding, I pause here to summarize the remarks so far. Because one may undergo states of control consciousness with respect to motions of body parts that are anesthetized, a pure sensory theory of control consciousness seems a nonstarter. This motivates an account that relies, at least in part, on nonsensory control signals that make a direct contribution to control consciousness. One such theory is an extension of AEI that identifies states of control consciousness with states of intermediate-level recurrent activation in motor-processing hierarchies. I then addressed whether a different alternative to a pure sensory theory is a sensory-imagery theory such as Prinz's. I argued that Prinz's sensory-imagery theory seems not to constitute a viable alternative to a nonsensory theory for two reasons. The first is that Prinz's classification of the control-theoretic forward model as being *sensory* is unjustified. The second is that insofar as sensory imagery is first-person distinguishable from sensory perception due to imagery's involving the *endogenous* activation of sensory states, we have grounds for holding that there is a direct involvement of control signals in control consciousness.

4. Epiphenomenal conscious will?

I turn now to address what might seem to be a potentially troubling consequence of the AEI account of control consciousness (AEICC), namely that it offers an account of control consciousness as an inefficacious state with regards to the production of action. Arguably it is a part of our common-sense view of ourselves that we have states of conscious will that are efficacious with respect to the production of our bodily movements. And if AEICC is correct, then the conscious state that precedes action (an intermediate-level state in a motor hierarchy) is causally preceded by an unconscious state (a higher-level state that is the genuine initial cause of an instance of bodily movement). This view that an unconscious state of will precedes the conscious appearance of being in control of one's action is a similar result to the much discussed work of Benjamin Libet (1999).

Libet's experiment involved having experimental subjects note, while looking at a clock, at what time they made the conscious decision to flick their wrist. Libet found, by noting EEG recordings of a *readiness potential* (a marked increase of neural-electrical activity preceding the wrist-flick), that there was a delay of 300–500 ms between the readiness potential and the reported time of the conscious decision (the subjective time or time in

which the decision seemed to the subject to be made). One implication, perhaps troubling, of Libet's result is that control consciousness is an illusion. We do not consciously will anything. Willing occurs prior to a conscious state that itself is a by-product of the act of willing, not the willing itself.

This general sort of result is to be expected according to AEICC. The highest levels of activation in a motor-processing hierarchy occur unconsciously and prior to the recurrent signaling in intermediate levels that constitute the conscious state. However, I will argue that we should not regard such a result as puzzling or troubling.

Central to the line I want to develop here is an analogy between conscious states of perception and conscious states of will. Conscious states of perception have causal antecedents that, being causal antecedents, occur prior to the state of causal perception, and unless we apply overly stringent criteria on what counts as consciously perceiving something, there is no illusion of conscious perception thereby generated. Similarly, conscious states of will have causal antecedents that temporally precede the occurrence of the conscious states, and by analogy we should not apply overly stringent criteria on what counts as a conscious willing.

It is no more an illusion that we will consciously than an illusion that we perceive consciously. The external-world events that we perceive are perceived consciously even though *they*, the external events, are causal antecedents of our states of perceptual consciousness. If we find such a view nonparadoxical and nonpuzzling, then we should be able to come to a similarly nontroubling view of the implications of Libet's results for control consciousness. Just as external events are consciously perceived even though they are causal antecedents of states of consciousness, certain inner events are conscious willings or consciously willed even though they are causal antecedents of states of consciousness.

We could, if we wanted, apply overly stringent criteria to perception to generate a "puzzle of conscious perception" that parallels the puzzle of conscious will that many see raised by Libet's results. One overly stringent criterion is a *time-of-occurrence-criterion*, whereby in order to be distinct from a *memory*, a perception of an event has to occur at the same time as the event perceived. Another overly stringent criterion is a *factivity criterion*, whereby in order to be distinct from an *illusion*, a perception of the time of occurrence of an event as "now" has to be accurate (the perception that now is noon cannot occur a little after noon without counting as an illusion). If we applied these stringent criteria, then we could derive that we never have accurate perceptions and, instead, we either have accurate memories of what was happening or inaccurate illusions of what is currently happening. More natural, however, is to avoid such overly stringent criteria and thus go on, just as common sense does, saying that we frequently perceive events at their actual time of occurrence.

5. Explaining the apparent lack of control phenomenology

Some have claimed that introspection of control consciousness reveals no distinctively nonsensory component (Prinz, 2007a). Others make a contrary claim about the relevant phenomenology and hold it to be obvious that there is a distinctive phenomenology associated with control (Mandik, 1999). Phenomenological disputes are notoriously difficult to

adjudicate, thus leading some researchers to be quite skeptical of both the phenomenological enterprise and the reliability of introspection. Nonetheless, claims about consciousness should be made to square with the self-reports of subjects, if not to *explain*, then to *explain away*. If there's controversy regarding some point of phenomenology, it can be quite satisfying to discover an explanation of *why* such a controversy arises.

While I think careful reflection reveals a distinctively nonsensory component to control consciousness, I do think there is something worth taking seriously in various claims against nonsensory control phenomenology. In the remainder of this section, I offer two possible explanations why it may have seemed obvious to some that there would be no such thing. The first concerns the differential bandwidth between prototypical instances of sensory inputs and motor outputs. The second concerns the degree to which introspection is itself an act.

Sensory inputs may be compared with one another and with motor outputs in terms of bandwidth. Estimates of the bandwidth of the human eye for color vision range from 4.32×10^7 bits/s (Jacobson, 1950, 1951) to a more recent estimate of 10^6 bits/s (Koch et al., 2006). It is perhaps not surprising that hearing has a significantly lower bandwidth than vision (this may be what underlies the common assertion that a picture is worth a thousand words). Jacobson (1950, 1951) gives a bandwidth estimate of 4.32×10^6 bits/s for the eye for black-and-white vision and an estimate of 9,900 bits/s for the bandwidth of the human ear. These differences in bandwidth perhaps account for widespread intuitions such as the intuition that visual "qualia" are ineffable, the intuition that a person blind from birth can never be told what it's like to see (Hume, Locke), and the intuition that a person reared in a black-and-white environment would not know what it's like to see red (Jackson, 1982). The auditory channel is impoverished compared to the visual channel, and the black-and-white visual channel is impoverished compared to the color channel.

So what happens when we turn our attention to motor systems? Bandwidth estimates for motor output systems are far lower than either vision or hearing. Fitts (1992) estimates motor output bandwidth at 10–12 bits/s.

I offer that bandwidth differences between various sensory systems and output systems can serve as a basis for explaining why many may have the intuition that there is no distinctive phenomenology for control consciousness. Such an explanation proceeds as follows: It is obvious that vision is phenomenologically richer than hearing, and color vision is richer than black-and-white vision. It is a very natural explanation of this differential richness to cite the bandwidth differences of the relevant input systems. Further, the very low bandwidth of motor output predicts that control phenomenology is comparatively impoverished. Indeed, it is so impoverished that it is no surprise that some people come to the mistaken conclusion that it is totally impoverished.

Another explanation of why some may have supposed that there is no control phenomenology—an explanation that may work together with the bandwidth-based explanation—hinges on the fact that introspecting is itself an act. As such, it is reasonable to suppose that a greater load on cognitive resources is presented in introspecting control consciousness than in introspecting sensory consciousness. To spell this out further, let us

assume, for purposes of illustration, that motor systems and sensory systems have the same bandwidth. If so, bandwidth alone would not serve to account for an apparent difference in phenomenological richness. If, however, there were some additional factor present that inhibited the ability to introspectively attend to motor systems but not sensory systems, then that factor would serve to explain a difference in apparent richness.

What could such a factor be? It is well known that attempting simultaneous multiple control tasks diminishes the capacity one would otherwise have to do them singly. If introspection is itself an act, then introspecting motor control is a doubling of tasks in a way that introspecting otherwise passive sensory input is not. The doubling introduced in introspecting control consciousness thus serves as the sought-after factor that can explain a comparative lack of richness between control and sensory systems.

It is worth stressing that the point of these remarks about bandwidth and control is *not* that it is impossible to notice any control phenomenology, only that it is predictably difficult. Some people have not noticed that there is control phenomenology because (a) informationally speaking, there is very little to it in the first place, making it difficult (though not impossible) to notice and (b) given that introspection itself puts a load on control processes, control phenomenology is even more difficult to notice (although not impossible).

Notes

1. I thus interpret “control consciousness” as being what philosophers call *intensional* or *nonfactive*—applying alike to cases wherein one *is* in control and when one *merely seems* to be in control. Synonymous and near-synonymous expressions for control consciousness include “the sense of agency,” “consciousness of action,” “the phenomenology of agency,” “agentive experiences,” and “the phenomenology distinctive of first-person agency.” For excellent recent reviews of current work on control consciousness, see Bayne (2008), Gallagher (2007), and Jeannerod (2007).
2. Thus, the current account is an example of an intermediate-level theory of consciousness. There are several besides mine (Jackendoff, 1987; Koch, 2004; Prinz, 2000, 2001, 2005, 2007b). As Prinz’s work is of central interest to the current project, I here briefly mention the main similarities and differences between his account of consciousness and mine. We both agree that conscious states are intermediate-level representations. We disagree over the roles of recurrence and attention. Prinz requires and I do not require that intermediate-level representations be modulated by attention to count as conscious. I require and Prinz does not that intermediate-level representation be involved in recurrent processing to count as conscious.
3. Grush (2007) discusses what he calls “modal emulators” and “amodal emulators,” though it is not clear that Grush’s “modal emulators” satisfy the criteria for sensory systems that Prinz lays out.

Acknowledgments

I am incredibly grateful for detailed comments on this paper from Rick Cooper, William Bechtel, and two anonymous referees for this journal. For helpful discussion of an earlier version of the paper, I am grateful to Fred Adams, Myrto Mylopoulos, Eric Thomson, Arnold Trehub, and Josh Weisberg. For invaluable copyediting assistance, I thank Rachelle Mandik.

References

- Bayne, T. (2008). The phenomenology of agency. *Philosophy Compass*, 3(1), 182–202.
- Bickle, J., Mandik, P., & Landreth, A. (2006). The philosophy of neuroscience [Electronic Version]. *The Stanford Encyclopedia of Philosophy (Spring 2006 Edition)*, available at: <http://plato.stanford.edu/archives/spr2006/entries/neuroscience/>.
- Brook, A., & Mandik, P. (2007). The philosophy and neuroscience movement. *Analyse & Kritik*, 29(1), 3–23.
- Churchland, P. S. (2002). *Brain-wise: Studies in neurophilosophy*. Cambridge, MA: MIT Press.
- Durso, F. T., & Johnson, M. K. (1980). The effects of orienting tasks on recognition, recall, and modality confusion of pictures and words. *Journal of Verbal Learning and Verbal Behavior*, 19(4), 416–429.
- Finke, R. A., Johnson, M. K., & Shyi, G. C. W. (1988). Memory confusions for real and imagined completions of symmetrical visual patterns. *Memory & Cognition*, 16(2), 133–137.
- Fitts, P. M. (1992). The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology: General*, 121(3), 262–269.
- Gallagher, S. (2007). The natural philosophy of agency. *Philosophy Compass*, 2(2), 347–357.
- Grush, R. (2001). The architecture of representation. In W. Bechtel, P. Mandik, J. Mundale, & R. Stufflebeam (Eds.), *Philosophy and the neurosciences: A reader* (pp. 349–368). Oxford, England: Blackwell.
- Grush, R. (2007). Skill Theory v2.0: dispositions, emulation, and spatial perception. *Synthese*, 159(3), 389–416.
- Intraub, H., & Hoffman, J. E. (1992). Reading and visual memory: Remembering scenes that were never seen. *American Journal of Psychology*, 105(1), 101–114.
- Jackendoff, R. (1987). *Consciousness and the computational mind*. Cambridge, MA: MIT Press.
- Jackson, F. (1982). Epiphenomenal qualia. *Philosophical Quarterly*, 32, 127–136.
- Jacobson, H. (1950). The informational capacity of the human ear. *Science*, 112(2901), 143–144.
- Jacobson, H. (1951). The informational capacity of the human eye. *Science*, 113(2933), 292–293.
- Jeanerod, M. (2007). Consciousness of action. In M. Velmans & S. Schneider (Eds.), *The Blackwell companion to consciousness* (pp. 540–550). Oxford, England: Blackwell.
- Koch, C. (2004). *The quest for consciousness*. Englewood, CO: Roberts & Company.
- Koch, K., McLean, J., Segev, R., Freed, M. A., Berry, M. J., Balasubramanian, V. & Sterling, P. (2006). How much the eye tells the brain. *Current Biology*, 16(14), 1428–1434.
- Kosslyn, S. M. (1994). *Image and brain: The resolution of the imagery debate*. Cambridge, MA: MIT Press.
- Lamme, V. A. F., Supér, H., & Spekreijse, H. (1998). Feedforward, horizontal, and feedback processing in the visual cortex. *Current Opinion in Neurobiology*, 8(4), 529–535.
- Libet, B. (1999). Do we have free will? *Journal of Consciousness Studies*, 6, 47–57.
- Mandik, P. (1999). Qualia, space, and control. *Philosophical Psychology*, 12(1), 47–60.
- Mandik, P. (2005). Phenomenal consciousness and the allocentric-egocentric interface. In R. Buccheri, A. C. Elitzur, & M. Saniga (Eds.), *Endophysics, time, quantum and the subjective* (pp. 463–485). Singapore: World Scientific.
- Mandik, P. (2007). The neurophilosophy of consciousness. In M. Velmans & S. Schneider (Eds.), *The Blackwell companion to consciousness* (pp. 418–430). Oxford, England: Basil Blackwell.

- Mandik, P. (2008). An epistemological theory of consciousness? In A. Plebe & V. M. D. L. Cruz (Eds.), *Philosophy in the neuroscience era* (pp. 136–158). Rome: Squilibri.
- Mandik, P. (2009). The neurophilosophy of subjectivity. In J. Bickle (Ed.), *The Oxford handbook of philosophy and neuroscience* (pp. 601–618). Oxford, England: Oxford University Press.
- Pascual-Leone, A., & Walsh, V. (2001). Fast back projections from the motion to the primary visual area necessary for visual awareness. *Science*, 292, 510–512.
- Peacocke, C. (2007). Mental action and self-awareness (1). In B. McLaughlin & J. Cohen (Eds.), *Contemporary debates in philosophy of mind* (pp. 358–376). Oxford, England: Blackwell.
- Perky, C. W. (1910). An experimental study of imagination. *American Journal of Psychology*, 21, 422–452.
- Prinz, J. (2000). A neurofunctional theory of visual consciousness. *Consciousness and Cognition*, 9(2), 243–259.
- Prinz, J. (2001). Functionalism, dualism, and the neural correlates of consciousness. In W. Bechtel, P. Mandik, J. Mundale, & R. Stufflebeam (Eds.), *Philosophy and the neurosciences: A reader* (pp. 278–294). Oxford, England: Blackwell.
- Prinz, J. (2002). *Furnishing the mind: Concepts and their perceptual basis*. Cambridge, MA: MIT Press.
- Prinz, J. (2005). A neurofunctional theory of consciousness. In A. Brook & K. Akins (Eds.), *Cognition and the brain: The philosophy and neuroscience movement* (pp. 381–396). Cambridge, England: Cambridge University Press.
- Prinz, J. (2007a). All consciousness is perceptual. In B. McLaughlin & J. Cohen (Eds.), *Contemporary debates in philosophy of mind* (pp. 335–357). Oxford, England: Blackwell.
- Prinz, J. (2007b). The intermediate level theory of consciousness. In M. Velmans & S. Schneider (Eds.), *The Blackwell companion to consciousness* (pp. 247–260). Oxford, England: Basil Blackwell.

Copyright of Topics in Cognitive Science is the property of Wiley-Blackwell and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.