

phenomenal experience but these representations might be too aberrant to be integrated into the wider cognitive workspace.

In the target article Block aims to identify circumstances in which stimuli might elicit phenomenal experience but not elicit a representation which is cognitively accessible. In the examples Block considers, cognitive accessibility is impeded through deficits in attention (the extinction example) or through limitations in processing time (the partial report examples). Both mechanisms depend upon some constraint in processing the stimulus. There must be a concern that this might affect sensory processing, and hence that phenomenal experience of the stimulus is affected along with cognitive access. We know, for example, that attention affects visual sensitivity (Solomon 2004). Changes in the response gain of neurons in sensory areas of cortex are as likely to be affecting phenomenal experience as cognitive access (Carrasco et al. 2004; Treue & Martínez Trujillo 1999). Are there circumstances in which the object of sensory processing can be examined at leisure and can be fully attended, and yet cognitive access is lost when phenomenal experience survives?

What does it mean for a representation to be cognitively accessible? As Block notes, “mechanisms of reporting, reasoning, evaluating, deciding, and remembering” (sect. 11, para. 2) should be able to make use of such a representation – a key aspect of cognitive access in models such as Baars’ Global Workspace Theory (Baars 1997) is that information about a stimulus becomes widely available. The representation must therefore take a form that permits interaction with memories or plans about other stimuli (in Piagetian terms, for a cognitive system to assimilate a representation, the former must be accommodated to the latter). If stimuli no longer engage representations that can be integrated in any sense with the rest of a person’s representational framework, then surely they cannot be said to be cognitively accessible.

How can representations become isolated from cognition? One might erroneously assume that sensations are “raw” – that they do not need to be prepared in a fashion to make them accessible to cognition. Sensations do not, however, correspond simply to the activity of sensory receptors. In color vision, for example, color appearance is far more closely related to a distal property of surfaces (things in the world about which we have cognitions), their spectral reflectance, than to the activations of cone photoreceptors in the retina (Hofer et al. 2005). Color sensation is the endpoint of a complex process. If this process is prevented from running to its conclusion we may be left with an incomplete signal that gives rise to sensation but cannot be integrated with cognition. Stimuli may therefore potentially elicit phenomenal experiences and these experiences may be discriminable yet they remain cognitively isolated. There is no framework within which to reason about or evaluate their differences, nor can they be remembered (I am not sure I see mere discrimination or decision as an act of cognition, but neither does Block when he discusses “guessing” in blindsight).

Do such circumstances exist? I will argue there are neurological patients who have all the time in the world to process stimuli, who can attend to them, but who have sensations divorced from cognition.

Cerebral achromatopsia is a neurological condition in which color vision is lost as a consequence of damage to ventromedial occipital cortex, usually in the vicinity of the fusiform and lingual gyri (Meadows 1974). Unlike more usual forms of color blindness there are no deficits or abnormalities in the retinal cone photoreceptors which form the starting point of normal color perception. Patients with cerebral achromatopsia do not usually make spontaneous comments about color. Questions about color sensation seem quite alien to them. They are unable to name the colors of stimuli presented to them or to perform nonverbal tests of color perception such as sorting or odd-one-out tasks. Although they may remember some semantic associations of color words (e.g., that bananas are yellow), they appear to have no sensation of color or means of remembering

any aspects of the nature of color sensation (see, e.g., Heywood & Kenridge [2003] for a recent review). It is, however, well known, that cerebral achromatopsics do see (quite consciously) the border formed between regions of equiluminant color (Heywood et al. 1998). If a red and a green that a patient cannot tell apart are used to construct a red square against a green background, then the patient will effortlessly see the square although they will be unable to explain how the square and its background differ. My colleagues and I (Kenridge et al. 2004) have shown that not only can these patients see such color borders but they can discriminate between borders formed from different colors (and that this discrimination cannot be accounted for on the basis of chromatic contrast salience) even though they cannot see the colors of the adjoining surfaces that form the borders. Again, the discrimination is conscious. The borders somehow look different from one another. The patients cannot, however, explain in any sense how or why the borders differ (they have no accessible representations of color) – they just know that the borders look different. The patient is surely having differing phenomenal experience of these borders, yet these experiences are not cognitively accessible. It is true that they know of the existence of these borders, but subjects in a partial report experiment know of the existence of items they cannot describe. I suggest that this situation, in which sensory representations simply cannot be integrated into the global workspace, provides a better example of phenomenology in the absence of cognitive access than cases in which the integration of representations into the workspace is possible but temporarily unachieved.

### Phenomenology without conscious access is a form of consciousness without top-down attention

doi: 10.1017/S0140525X07002907

Christof Koch<sup>a</sup> and Naotsugu Tsuchiya<sup>b</sup>

<sup>a</sup>Division of Biology, California Institute of Technology, Pasadena, CA 91125;

<sup>b</sup>Division of Humanities and Social Sciences, Psychology and Neuroscience, California Institute of Technology, Pasadena, CA, 91125.

koch.christof@gmail.com

<http://klab.caltech.edu/~koch/naotsu@gmail.com>

<http://www.emotion.caltech.edu/~naotsu/Site/index.html>

**Abstract:** We agree with Block’s basic hypothesis postulating the existence of phenomenal consciousness without cognitive access. We explain such states in terms of consciousness without top-down, endogenous attention and speculate that their correlates may be a coalition of neurons that are consigned to the back of cortex, without access to working memory and planning in frontal cortex.

We agree with Block’s hypothesis that phenomenally conscious states may sometimes not be cognitively accessible. Partial report and dual-tasks paradigms show that we have only limited access to some aspects of phenomenal experience. Or, to adopt Block’s earlier language, phenomenal consciousness can occur without access consciousness, a revision of our earlier position (Crick & Koch 1998a; Koch 2004). We argue here that sensory psychology has a long-standing framework to consider such cases, involving attentional selection processes.

We recently (Koch & Tsuchiya 2007; Tsuchiya & Koch 2008) summarized the empirical evidence that consciousness and top-down, volitionally controlled endogenous attention are distinct neurobiological processes with distinct functions (see also, Iwasaki 1993; Lamme 2003). In particular, sensitive psychophysical techniques can dissociate these two. Indeed, a range of phenomena exists in which subjects are conscious of certain stimuli attributes without top-down attention. This list includes

the pop-out target in a visual search task, gist perception of a scene, categorization of peripheral targets when a very intensive cognitive-resource task has to be performed at fixation, and iconic memory.

Consider Sperling's original iconic memory experiment (Sperling 1960) or Landman et al.'s (2003) variant. Subjects report that they clearly, vividly, and consciously see a field of letters or a bunch of bars arranged on a circle. This is also what we experience when we look at such displays. However, it is well known that subjects have only very limited access to the detailed properties of the individual elements, unless top-down attention is directed to a subset of stimuli using appropriately timed cues. Our basic point is that phenomenology without conscious access is an example of consciousness without top-down attention processing, though the converse is not true; that is, not every example of conscious perception in the absence of top-down attention is cognitively non-accessible. For example, the gender of a briefly presented face can be accurately reported even if subjects are engaged in a highly demanding task at the fixation (Reddy et al. 2004).

So what is the story at the level of the brain? Decades of electrophysiological recordings in monkeys have proven that the spiking response of neurons in the ventral visual stream (e.g., in areas V4 and IT) representing attended stimuli is boosted at the expense of the response to non-attended items. According to Crick and Koch (1995), this enables these neurons to establish a reciprocal relationship with neurons in the dorsolateral prefrontal cortex and related regions that are involved in working memory and planning (and language in humans), leading to reverberatory neuronal activity that outlasts the initial stimulus duration. Critical to the formation of such a single and integrated coalition of neurons are the long-range axons of pyramidal neurons that project from the back to the front of cortex and their targets in the front that project back to the upper stages of the ventral pathway (possibly involving stages of the thalamus, such as the pulvinar [Crick & Koch 1998b], and the claustrum [Crick & Koch 2005]). The subject now consciously sees these stimuli and can report on their character (e.g., identify the letter [Sperling 1960] or the orientation of the square [Landman et al. 2003]). Furthermore, the subject also has a strong conscious sense of the entire scene ("I see an array of letters") that is likewise mediated by a loop that involves the inferior temporal cortex and the frontal lobes half-way across the brain.

But what happens to those stimuli that do not benefit from attentional boosting? Depending on the exact circumstances (visual clutter in the scene, contrast, stimulus duration) these stimuli may likewise establish coalitions of neurons, aided by local (i.e., within the cortical area) and semi-local feedback (i.e., feedback projections that remain consigned to visual cortex) loops. However, as these coalitions of neurons lack coordinated support from feedback axons from neurons situated in the prefrontal cortex, thalamus, and claustrum, their firing activity is less vigorous and may decay much more quickly. Yet, aided by the neuronal representation of the entire scene, these weaker and more local coalitions may still be sufficient for some phenomenal percepts.

Block cites functional magnetic resonance imaging (fMRI) studies of patients with visuospatial hemi-neglect (Driver & Vuilleumier 2001; Rees et al. 2000; 2002b) that offer evidence of enhanced hemodynamic activity in the fusiform face area contralateral to a face that the patient is not aware of. For Block, this raises the question of whether this is likewise an example of phenomenal consciousness without cognitive access. We answer this question clearly in the negative. First, one should trust the first-person perspective: That is, in the absence of compelling, empirical evidence to the contrary (such as Anton's blindness, also known as hysterical blindness; Sackeim et al. 1979), if the subject denies any phenomenal experience, this should be accepted as a brute fact. If we take the existence of mere recurrent, strong neuronal activation as evidence for consciousness, why not argue

that the spinal cord or the enteric nervous system is conscious but is not telling me (Fearing 1970)? Second, the relationship between neuronal firing activity and the associated hemodynamic BOLD response is a very complex one. In particular, there are well-documented cases where a vigorous fMRI signal is observed in the absence of any spiking from the principal neurons in that area (Harrison et al. 2002; Logothetis 2003; Logothetis & Wandell 2004; Mathiesen et al. 1998). Synaptic activity is a much larger driver of hemodynamic activity than are action potentials. Therefore, a much more cautious reading of these studies is that they demonstrate synaptic input into the fusiform face area in these patients; however, whether or not this input is vigorous enough to establish a sustained coalition of neurons is totally up in the air and requires further investigations.

In conclusion, the quiddity of the neuronal correlates of conscious access are long-range loops between the back and the front of cortex and its associated satellites (thalamus, basal ganglia and claustrum), enabled by top-down attention. Without this amplification step, most coalitions in the back are fated to die; however, given the right conditions, a few may survive and may be consciously experienced by the subject. Yet, as the informational content of these coalitions are not accessible to working memory and planning circuits in the front, the subject cannot consciously access the detailed stimulus attributes. Our explanation provides a plausible account of how phenomenal consciousness can occur without cognitive access.

## Partial awareness and the illusion of phenomenal consciousness

doi: 10.1017/S0140525X07002919

Sid Kouider, Vincent de Gardelle, and Emmanuel Dupoux

*Laboratoire de Sciences Cognitives et Psycholinguistique, Ecole Normale Supérieure, 75005 Paris, France, and CNRS/EHESS/DEC-ENS, 75005 Paris, France.*

sid.kouider@ens.fr    gardelle@ens.fr  
dupoux@lscp.echess.fr    www.lscp.net

**Abstract:** The dissociation Block provides between phenomenal and access consciousness (P-consciousness and A-consciousness) captures much of our intuition about conscious experience. However, it raises a major methodological puzzle, and is not uniquely supported by the empirical evidence. We provide an alternative interpretation based on the notion of levels of representation and partial awareness.

In his target article, Ned Block is dealing with a difficult problem: how to empirically demonstrate that phenomenal consciousness (hereafter P-consciousness) is dissociable from access consciousness (hereafter A-consciousness). An a priori argument in favor of this dissociation is the common intuition that the representational content of phenomenal experience is much richer than the limited content we can access at a given time. In Block's words, "phenomenology overflows cognitive accessibility" (sect. 8, para. 6). This intuition is so strong that it appears very easy, at first glance, to show how much richer P-consciousness is, compared with A-consciousness.

However, providing an empirical demonstration of this dissociation leads to a major methodological difficulty: any measure of consciousness seems inevitably to require the involvement of A-consciousness. From there on, it seems impossible to show evidence for P-consciousness without A-consciousness. This methodological puzzle arises also in dissociating consciousness and top-down attention. Demonstrating consciousness without attention seems impossible for similar reasons: To assess consciousness of the stimulus, one needs to direct the subject's attention on the stimulus! Although there is converging evidence that attention can affect both conscious and unconscious