Cell PRESS

How rich is consciousness? The partial awareness hypothesis

Sid Kouider, Vincent de Gardelle, Jérôme Sackur and Emmanuel Dupoux

Laboratoire de Sciences Cognitives & Psycholinguistique, Ecole Normale Supérieure - CNRS, 29 rue d'Ulm, 75005, Paris, France

Current theories of consciousness posit a dissociation between 'phenomenal' consciousness (rich) and 'access' consciousness (limited). Here, we argue that the empirical evidence for phenomenal consciousness without access is equivocal, resulting either from a confusion between phenomenal and unconscious contents, or from an impression of phenomenally rich experiences arising from illusory contents. We propose a refined account of access that relies on a hierarchy of representational levels and on the notion of partial awareness, whereby lower and higher levels are accessed independently. Reframing of the issue of dissociable forms of consciousness into dissociable levels of access provides a more parsimonious account of the existing evidence. In addition, the rich phenomenology illusion can be studied and described in terms of testable cognitive mechanisms.

One or two types of consciousness?

Understanding the psychological and neurobiological determinants of consciousness constitutes a major challenge in contemporary science [1,2]. Following a long period of neglect, the study of consciousness has gained respect by acknowledging the importance of data collection, empirical support, and all of the core principles of scientific investigation. This research field now offers functional descriptions and testable predictions regarding conscious processing [3–7].

However, critics of this approach to consciousness argue that functional explanations come at the price of sacrificing the phenomenal aspects of consciousness: functional explanations are restricted to the cognitive mechanisms (i.e. attention, working memory, etc.) underlying access to conscious contents, ignoring the problem of how these contents arise in the first place [8,9]. From this perspective, consciousness should be dissociated into two components, namely access and phenomenal consciousness, following a popular dichotomy introduced by Block [9]. Importantly, the contents of phenomenal experiences are assumed to be much richer than the limited representations we can access at a given time: in Block's terms, 'phenomenology overflows access' (here phenomenology is synonymous with phenomenal consciousness). This hypothesis has become so popular that most neurobiological theories of consciousness now endorse it explicitly (Box 1).

In this paper, we claim that despite its appeal, this dissociation raises more problems than it solves and leads, in fact, to an epistemological impasse. We show that previous efforts to demonstrate phenomenal consciousness without access not only are at odds with the subjects' own reports, but also that they have been misled by a restrictive account of access and by a confusion between phenomenal contents and unconscious processing (i.e. information is processed but remains below the threshold of consciousness; see [10] for a review). We argue that functional descriptions that rely on the notion of partial awareness along a hierarchy of levels of representation, including, crucially, lower levels of representation, provide a more parsimonious explanation of previous evidence. Finally, we propose that a key aim for future research on consciousness is to focus specifically on the cognitive mechanisms (e.g. top-down expectation, confidence monitoring, probabilistic inferences, etc.) that might underlie the illusion of a rich phenomenology. We hope to show that by developing a refined description of the mechanisms underlying conscious access, one can not only obtain phenomenal consciousness for free, but also explain its illusory nature.

The arguments: phenomenal overflow and neural purity

Two main empirical arguments, the overflow argument and the purity argument, have been offered by proponents of the access–phenomenal consciousness dissociation. The overflow argument is rooted in the intuition that we are conscious of much more than we can describe and manipulate. For instance, when observing a complex visual scene, we feel that we have a rich visual experience even if we can report only a few elements. This phenomenon was operationalized in Sperling's now famous study over half a century ago [11], which used brief presentations to quantify the number of letters available regardless of report (Figure 1). However, it remains controversial whether the large amount of information that is available without being reportable reflects phenomenal or unconscious processing [12–14].

The neural purity argument follows from the assumption that specific neural mechanisms exist for phenomenal experience (e.g. local neural recurrence; Box 1). Such mechanisms supposedly constitute a pure index of consciousness that is more reliable than subjective reports, which are limited by verbal abilities [14,15]. For instance, Block and Lamme argue that in paradigms in which subjects cannot report the presence of a stimulus because of inattention (e.g. change blindness, inattentional blindness, attentional blink), subjects might still be phenomenally conscious of the stimulus because it induces local recurrence in perceptual brain regions [14,15]. Notably, contrary to the overflow

Co o d : Kouider, S. (sid.kouider@ens.fr).

Box 1. Dissociating consciousness: from philosophy to neurobiology

Easy vs. hard problem. The easy problem of consciousness consists of a set of issues about the informational properties of conscious states that are tractable with the standard tools of cognitive science: we can use objective measures of consciousness to explore its relationship with the integration of sensory information, attention, working memory, etc. The hard problem of consciousness consists of explaining the experiential dimension of consciousness: the first-order subjective nature of qualia and phenomenal states, the 'what is it like to be conscious', and how and why we experience consciousness at all [8,34].

Access vs. phenomenal consciousness. The epistemic distinction between easy and hard problems maps, according to Block [9], on two forms of consciousness. Phenomenal consciousness is related to the private first-person experience (i.e. qualia). Understanding this constitutes the hard problem. Access consciousness corresponds to the fact that some representations are 'poised for direct control of thought and action' [9]. Block further links conscious-accessed contents with global broadcasting [35], similar to workspace theories of consciousness [3,6]; that is 'contents information about which is made available to the brain's ''consumer'' systems: systems of memory, perceptual categorization, reasoning, planning, evaluation of alternatives, decision-making, voluntary direction of attention, and more generally, rational control of action.' Arguably, this property can be explained in terms of computational mechanisms that, through attention, amplify transitory information, maintain it in short-term memory and exploit it in controlled cognitive operations, eventually leading to long-term memory storage and report. Phenomenal consciousness presumably occurs without attention (Box 3) and reflects rich-capacity contents in sensory memory (e.g. iconic buffer), whereas access consciousness necessitates attention and reflects a limited set of elements in working memory.

Dissociative theories in neuroscience. Several theories have adopted Block's dissociation and explicitly distinguish between two neural correlates of consciousness. For instance, the duplex vision theory of Milner and Goodale [36] has recently been updated to associate sustained ventral stream activity with phenomenal consciousness, whereas only the involvement of more anterior (e.g. prefrontal) regions supports conscious access [37]. Similarly, Zeki [38] has recently linked micro- and macro-consciousness in his original theory [39] with the phenomenal consciousness of specific attributes (colors, contrasts, etc.) and bound objects, respectively, whereas unified consciousness is somewhat analogous to access consciousness. In the local recurrence theory of Lamme [15,40], phenomenal experience is explicitly associated with any recurrent neuronal activity (i.e. local or global loops), whereas conscious access occurs only for global recurrence. Although all these theories diverge in many respects, they all link phenomenal consciousness with posterior (i.e. occipitotemporal) regions, whereas anterior (i.e. prefrontal, workspace) areas are linked to conscious access (see [2] for a review). Notably, the original motivation underlying Block's distinction has been somewhat lost. Although it was primarily intended to stress the non-functional, non-mechanistic nature of phenomenal contents, neurobiological accounts actually treat phenomenological aspects in terms of functional (i.e. neurocomputational) mechanisms and are rather driven by the motivation to probe consciousness in the absence of subjective reports. In any case, both neurobiological and philosophical dissociative accounts share the assumption that a fundamentally inaccessible form of consciousness exists.

argument for which subjects' reports matter, according to the neural purity argument, reports are not to be trusted: subjects could be conscious of stimuli, even when they themselves deny it.

The limits: paradoxes and circularities

Here we assert that arguments favoring dissociative approaches to consciousness suffer from serious flaws, as previously argued in the context of higher-order theories

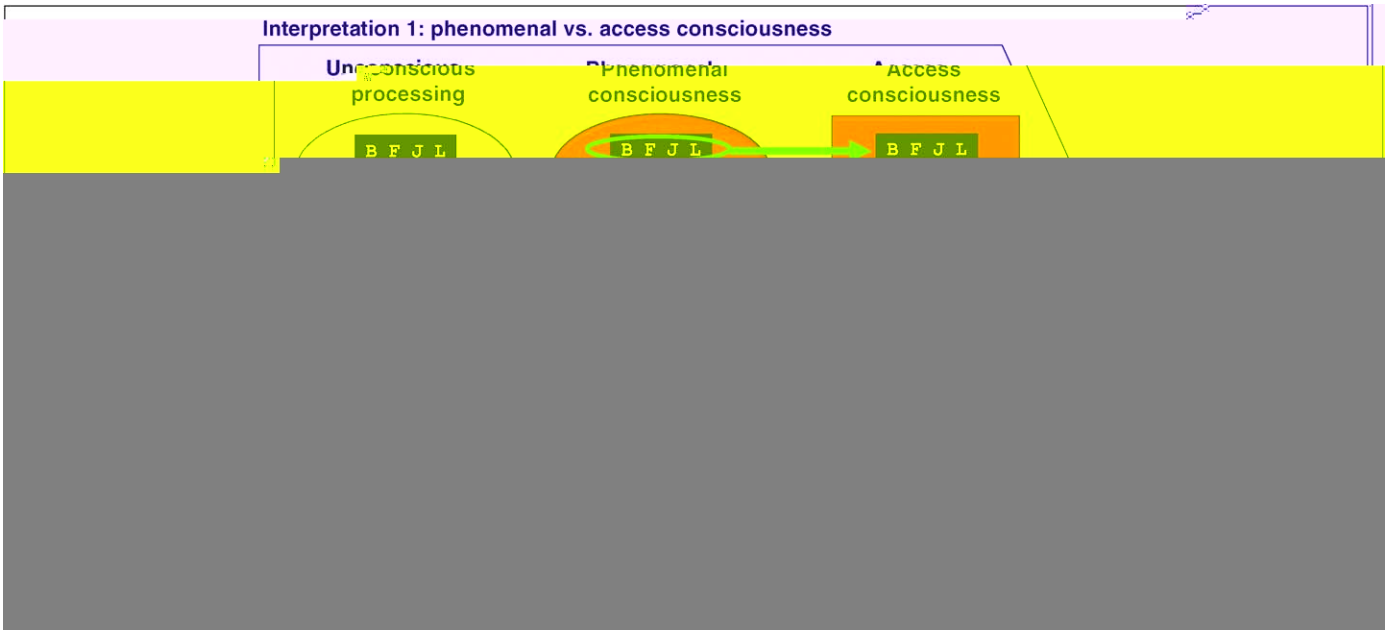

Figure 1. Two interpretations of Sperling's classical study on the availability of information in brief presentations [11]. In the original study, observers were presented with flashed arrays of 12 letters. Although they could report only three or four items, they had the impression of 'seeing all the letters'. When Sperling presented an auditory cue following the array, instructing subjects to report only one of the three rows, they reported nearly all items, indicating that there were indeed more contents available than the three or four items they could initially report. The green ellipse represents attentional capture and amplification. According to dissociative approaches (Interpretation 1), subjects are phenomenally conscious of all the items in the array, but then, because of time pressure, can only access a few elements before they fade away from the iconic buffer. According to the partial awareness hypothesis (Interpretation 2), observers consciously access a large quantity of low-level information (i.e. fragments) that is transiently activated and are also heavily biased towards the presence of letters in the entire array. Consequently, when the information is only fragmentary, subjects automatically fill in the array either with reconstructed letters (perceptual illusions) or with letter tags (cognitive illusions), leading to the impression of a rich and complete visual experience (Box 2). According to this proposal, the impression of richness is not basic and primary, but is actually a late construct [12].

Box 2. Partial awareness and perceptual illusions

Cognitive vs. perceptual illusions. When a complex stimulus is degraded (e.g. brief, masked, peripheral), observers cannot report its total content but nonetheless experience a rich visual experience. This has been described as a form of cognitive illusion, whereby subjects have a feeling of seeing that does not correspond to what they can access [14,27]. This cognitive illusion occurs because a scene can normally be inspected at will, and hence temporarily missing information is not a cause of alarm. In other words, the perceptual system uses the external world as a memory buffer. However, recent research has shown that perceptual contents at a given level of representation reflect the merging of bottom-up stimulus-related information with information already present at that level. For instance, observers can rapidly extract the gist from a brief visual presentation and use this information as prior information for identification of the details [41]. Furthermore, perceptual interpretation of a visual stimulus can be biased by means of expectations [42] or prior exposure [43]. Although cognitive illusions are probably real, it seems more appropriate to account for overflow phenomena in terms of perceptual illusions whereby partially represented sensory signals contribute, along with top-down expectations, to the reconstruction of perceptual contents.

Perceptual illusions under partial awareness. Partial awareness situations are those in which the subject accesses the stimulus information at some but not all representational levels. Information at other levels can remain inaccessible or, in some situations, can be accessed by filling in plausible content. This occurs when the signal is weak or degraded and reliance on prior information is high. We studied such illusory contents in two previous studies (Figure I). In a modified Stroop priming paradigm with visible but degraded stimuli, subjects treated nonwords as if they were real color words (e.g., GEREN perceived as GREEN) only under conditions that combined 1) strong expectations that there were real color words and 2) the possibility of detecting letters without identifying words accurately [44]. In a modified Sperling paradigm, we intermixed classical trials (Figure 1) with trials containing nonletters (e.g. rotated letters) in the uncued part of the array that shared the same features as letters (Figure I below). Furthermore, at the end of some trials, subjects were asked to decide which of several alternatives were actually present in the uncued parts of the array. We found that not only did subjects fail to detect nonletters, but also actually tended to perceive them as real letters [12].

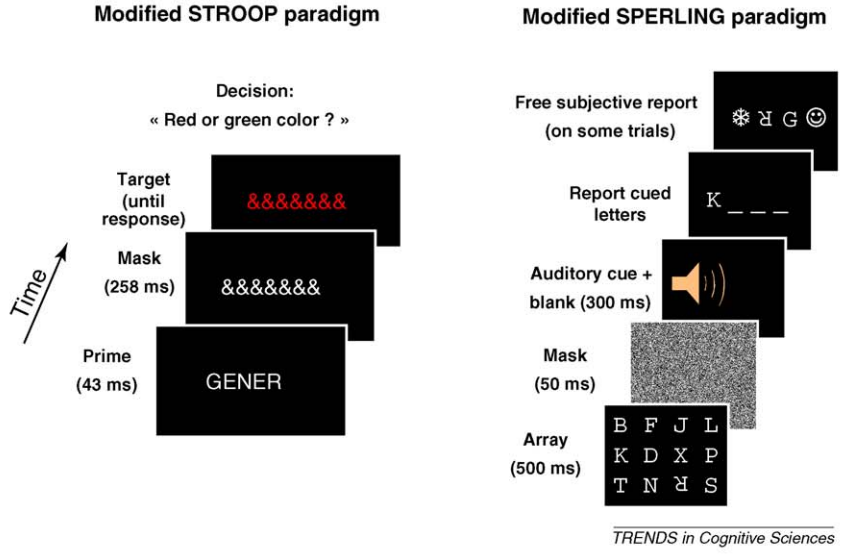

Figure I. Modified versions of the Stroop (left) and Sperling (right) paradigms (based on Refs [12,44]).

[16] and workspace theories [17] of consciousness. Our counter-arguments center on the fact that the phenomenal overflow argument is confounded with partial awareness situations, whereas the neural purity argument reflects the confusion between phenomenal consciousness and unconscious perception. We start with the overflow argument.

First, it is important to stress that limits in (verbal) reportability should not be equated with limits in access. Perhaps our visual experience seems rich because we lack the conceptual representations and the words to describe it. Perception involves non-conceptual contents that are difficult to verbalize, such as shades of colors, smells, etc. However, the relative poverty of verbal reports in these domains should not be equated with poverty in access. Indeed, the hallmark of psychophysics is precisely to uncover the rich, graded and multidimensional aspects of domains such as color and smell perception using indirect measures such as similarity judgments [18]. Furthermore, because verbal reports take time and are performed in a sequential manner, accessible information might have disappeared prior to verbalization. Nonetheless, subjects' performance in non-verbal tasks such as detection (e.g. 'Is there something or nothing on the screen'?) or discrimination (e.g. 'Is it X or Y?') shows that information can be accessed before it fades away. In other words, the overflow argument might only show that access overflows verbal report.

Second, the intuition of a rich phenomenal experience on which the overflow argument is built might be overstated. Indeed, observers might overestimate both the quantity and accuracy of the information they experience at a given moment, lured either by a 'cognitive illusion of seeing' [19] or by perceptual illusions (Box 2).

Third, the possibility of inaccessible consciousness is methodologically dubious: if subjects do not have access to their experience, how can we determine that they are conscious of it? What difference would it make to them if

Box 3. Consciousness without attention?

Following from the assumption that phenomenology overflows access [14], it has recently been proposed that consciousness overflows attention [40,45]. In this perspective, consciousness without access is roughly equivalent to consciousness without attention [46]. Whether consciousness is independent of attention remains highly controversial, notably given the difficulty of demonstrating consciousness without attention [13,45,47]. Koch and colleagues base their claims for consciousness without attention on two types of evidence [45,48,49]. First, they showed that under dual-task conditions, which they term the near absence of attention, although subjects are performing a main task on a target at a central location, they can still perceive, at least indistinctly, a stimulus in the periphery. Second, they rely on situations in which attention and consciousness, presumably, have opposite effects (see [45] for a review).

However, as for inaccessible consciousness, demonstration of a psychological state of consciousness without attention is plagued by the observer effect described above. Indeed, probing an individual's consciousness of a stimulus necessitates directing the observer's attention to the stimulus. Thus, it seems to be extremely difficult, if not impossible, to assess whether subjects consciously perceive objects in the periphery without relying on some form of access, as doing so inevitably requires observers to engage their attention on the stimulus. We contend that, similar to consciousness without access, the possibility of consciousness without attention is usually based on a restrictive definition that does not take into account the possibility of residual attention at lower (i.e. sensory, non-conceptual) levels of processing. For instance, a peripheral stimulus in Koch's paradigm can be considered as both conscious and unattended when spatial attention is defined as a focal, all-or-none component of the cognitive system. However, the same data can be interpreted quite differently if we consider that there are residual or non-focal components resulting from the division of attention [50,51]. In this view, although a large part of attentional resources is indeed engaged on the central stimulus, it is arguable that subjects can also attend to lower levels of information in the periphery.

they did not have this conscious experience at all? By this hypothesis, subjects should not be able to distinguish between situations with and without this conscious experience. In fact, someone experiencing phenomenology without access should not be able to talk about it. If subjects can tell 'something' about their experience (even summarily), then this can no longer be considered as phenomenal experience without access. In other words, reporting a rich visual experience demonstrates that we have access to some kind of information.

Fourth, if, nevertheless, one maintains that phenomenal experiences can arise in the absence of access, one reach an epistemological impasse. Indeed, to prove that particular content is phenomenal, one has to ask the subject about it. But if the subject is attempting to report on her experience, it also means that she is attempting to access it. In other words, we face an observer effect, according to which any observation of the internal states of a system changes the state of the system [2]. Thus, any attempts to observe internal states prior to access will necessarily be contaminated by access mechanisms themselves.

A potential solution to the problems outlined above might be to follow the neural purity argument, according to which phenomenological consciousness can be probed regardless of reportability. However, this strategy leads to circularity because validation of the neural index in the first place necessarily requires reliance on access mechanisms. Indeed, demonstration that a specific neural mechanism (e.g. local recurrence) is sufficient for consciousness initially requires assessment of neural events while probing whether the subject is conscious. As the best way to probe consciousness in the first place is reliance on access mechanisms, it seems impossible to map neural and phenomenal states without relying on access. Thus, although neural indexes offer interesting possibilities when reports are impossible, for instance for patients with locked-in syndrome or for prelinguistic babies, they cannot be assumed to reflect more than conscious access.

Finally, the neural purity argument often only reflects a theoretical confusion: it merely shows that the brain processes information without consciousness, but not that there is phenomenal experience associated with these processes. A supposed neural index of inaccessible consciousness might thus simply reflect an unconscious form of processing [13,20]. However, because we cannot demonstrate whether phenomenal experience is involved or not, the neural purity argument becomes unfalsifiable: if, say, local recurrence is observed in the absence of conscious access, stipulation of alternative forms of consciousness, instead of unconscious processing, cannot be verified and becomes a matter of faith.

Basic assumptions: a hierarchical view of conscious access

Much discussion has centered on the definition of phenomenal consciousness while assuming that access consciousness poses no difficulty. We believe, however, that some of the problems raised above actually stem from an ill-defined view of access mechanisms and a lack of consideration for the possibility of partial awareness. Here, we adopt the standard definition of access consciousness, according to which a mental content is conscious if it is broadcast to cognitive subsystems, notably working memory and control mechanisms (Box 1). However, in addition to this standard notion of access, we make specific claims about the nature of what can be accessed. In particular, we propose a few simple but fundamental assumptions to reframe the issue of dissociable forms of consciousness into dissociable levels of conscious access.

The first assumption is that most stimuli are processed through a hierarchy of representations, ranging from lower levels (e.g. visual energy, simple geometric elements, etc.) to higher levels (e.g. letters, word forms and meaning), as assumed in classical approaches to human cognitive architecture [21]. Our second and crucial assumption is that representations at each level can be accessed independently from each other (Figure 2). When observing a complex visual scene, a subject can access low-level representational contents (e.g. colors, textures, forms) over the whole visual field and thus have an impression of a rich, yet impossible to verbalize, visual experience, while being unable to access representations at higher levels (e.g. recognition of objects) over the entire visual field. Conversely, the subject could access the global meaning of the scene without accessing specific details [22,23].

Stimulus perception can thus fall into three possible categories. (i) Complete awareness in which all the levels of processing relevant to the task at hand are accessed. In

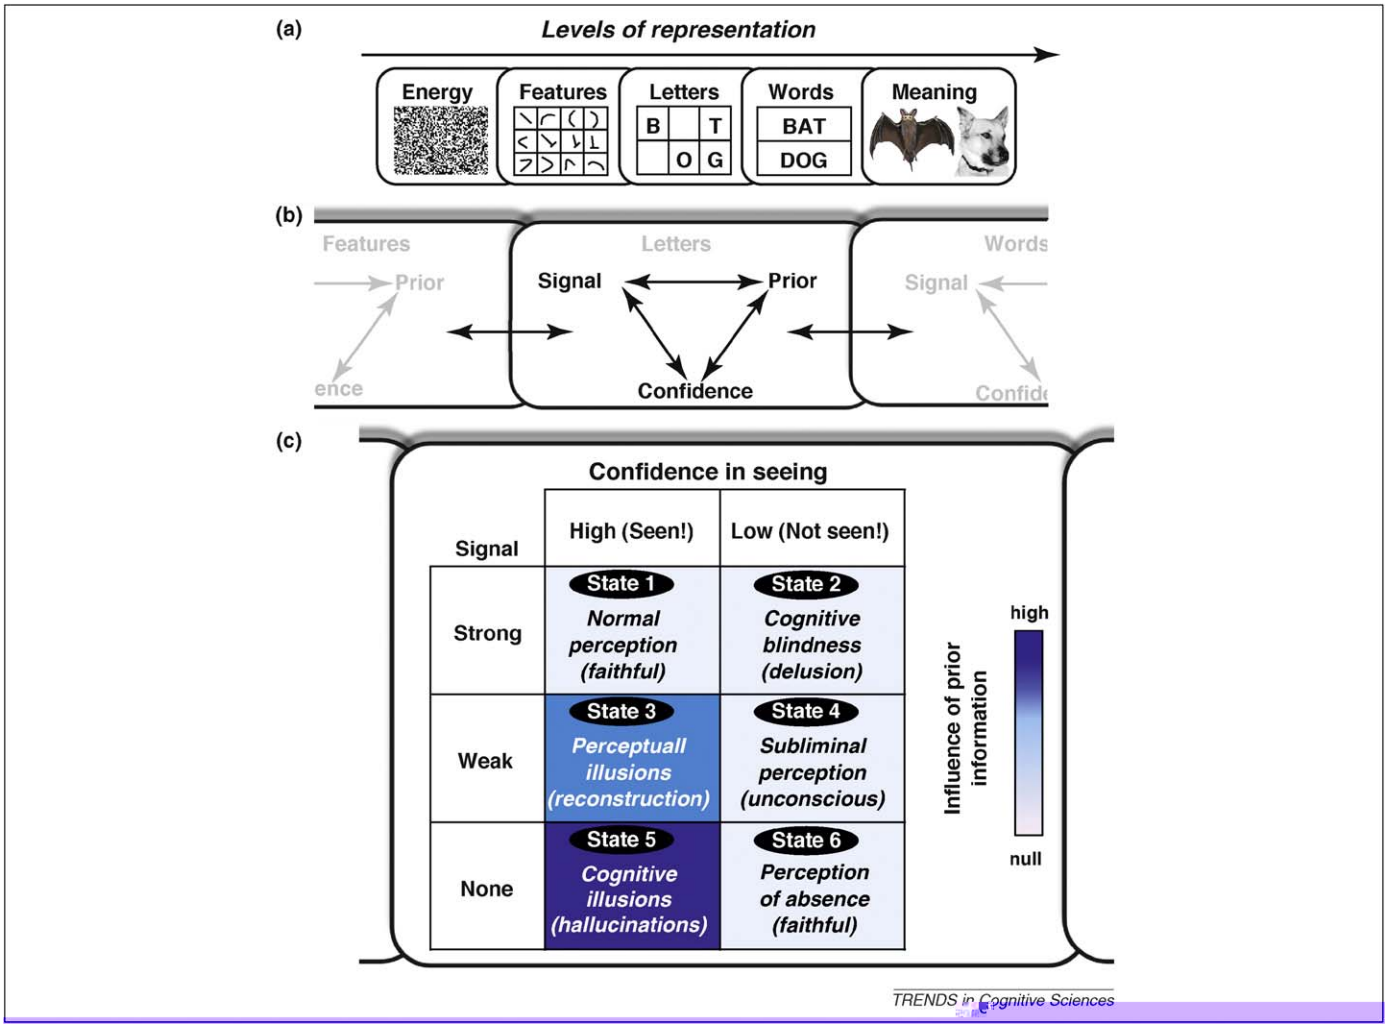

Figure 2. **(a)** Proposed hierarchical view of access in terms of levels of representation. Five example levels of representation are shown, ranging from lower to higher levels. **(b)** At each level the perceptual content reflects the interplay between the input signal, confidence in the signal and prior information. The levels are interconnected through the processing architecture. **(c)** Proposed typology of cognitive states (for each level of representation) resulting from the interplay between signal strength, confidence and prior information. A strong signal can lead to two possible states based on the degree of confidence in the signal. Normal faithful perception (state 1) corresponds to a strong signal and high confidence. Because the signal is strong, prior information can only have a slight impact. Cognitive blindness (state 2) is a state in which the same signal strength yields a very low confidence of seeing. This pathological state might be reflected in psychiatric reports of hysterical blindness, otherwise termed visual conversion disorders. A weak signal can lead to illusory perception (state 3) when the weak signal interacts with prior information in the presence of high confidence, or alternatively to subliminal perception (state 4) when confidence is low. Importantly, in subliminal processing, prior information is turned off by default (de Gardelle *et al.*, unpublished data). Finally, in the absence of a signal, hallucinations (state 5) arise when confidence is high. In this state, perceptual content is entirely driven by prior information. Normal non-perception (state 6) arises in the absence of a signal and with low confidence. In this state prior information is also turned off by default.

Figure 2, this corresponds to all levels of representation (features, letters, words, etc.) associated with state 1 (normal, faithful perception). (ii) Complete unawareness in which subjects do not have access to any of the levels and are confident that no stimulus has been presented. In Figure 2, all levels of representation are associated, with states 2, 4 or 6 reflecting cognitive blindness, subliminal perception or no perception at all, respectively. Thus, this situation does not preclude unconscious information processing, potentially up to the highest level of representation [18]. (iii) Partial awareness corresponds to intermediate cases, combining awareness at some levels and unawareness at other levels of representation. This refers to situations in which subjects have high confidence about the informational content at some restricted levels (i.e. state 1 only for those levels), whereas for other levels subjects are typically unable to access information (states 2, 4 and 6). However, this can also refer to situations in which the other levels are actually accessed by being filled in with illusory contents (state 3 or 5). We return to the issue of illusory access below. Importantly, whereas complete awareness and complete unawareness correspond to extreme cases, partial awareness encompasses an enormous repertoire of possibilities, resulting from the combination of many levels of processing and several possible states at each level (accessed or not accessed, with or without subjective confidence that there is reliable information at this level). Any complex visual scene consists of many items, each of which belongs to one of the above categories. Conscious experience is only derived from these varieties of access and, crucially, none of these three situations requires that we dissociate phenomenal from access consciousness. Hence, we argue that phenomenal and access consciousness should be collapsed.

Reinterpretation of empirical arguments

In light of the assumptions described above, we propose that the empirical arguments for inaccessible consciousness

correspond either to situations of complete unawareness, in the case of neural purity, or to situations of partial awareness, in the case of the overflow argument. Complete unawareness occurs in paradigms such as attentional blink, inattentional blindness, or extinction in neglect patients, where subjects confidently report that they have not seen anything. In other words, subjects cannot successfully detect the critical stimulus, implying that they have no access to it at any level of representation. Of course, the absence of access does not imply the absence of processing, because the stimulus can induce unconscious neural activity and, eventually, have subliminal influences on behavior [10].

With respect to the overflow argument, we contend that phenomenal consciousness without access is actually supported by the very same cognitive mechanisms of access applied to different levels along the hierarchy of representations. Situations in which subjects report the impression of a richer perceptual experience than they can describe can be operationally defined as detection in the absence of discrimination, or discrimination at certain (usually low) levels of representation in the absence of discrimination at other (usually higher) levels. In sum, rather than invoking an indemonstrable phenomenal form of consciousness without access, situations of overflow, such as in Sperling's study (see above), are best described in terms of partial awareness (Figure 1).

Additional assumptions: exploring the illusion of phenomenal richness

Although the basic assumptions presented above constitute the core of our hypothesis, they primarily account for how the signal is accessed at each level of representation. However, when subjects fail to accurately report items at one level (e.g. letters), they might still claim that they see all the letters. A number of additional assumptions, although tentative, are necessary to account for the illusion of phenomenal richness in functional terms. First, access at each level comes with a certain confidence that reflects the subjective likelihood that the signal is caused by an external event, as opposed to random internal noise. It is influenced both by the strength of the signal and by internal states (vigilance, attention, etc.). The fact that confidence plays an important role in consciousness has recently been supported by results from both theoretical and empirical studies [24–26]. However, our second assumption contends that accessible contents also result from the integration of bottom-up signals with contextual prior information. Classically, under Bayesian integration rules [20,21], the weaker the signal, the stronger the impact of prior information will be, yielding perceptual illusions. Finally, the integration of bottom-up signals and prior information at each level of representation is modulated by the degree of confidence. In particular, when the degree of confidence is low, prior information is turned off. This modulation is essential to prevent the system from experiencing constant hallucinations in the absence of stimuli, or in the presence of strong prior information (Figure 2).

These additional assumptions enable us to capture the phenomenological aspects in dissociative theories by recasting the feeling of a rich phenomenology without complete access as high confidence with inappropriate access. For instance, in two modified Stroop and Sperling experiments (see Figure I in Box 2), subjects were strongly confident of seeing real color words and real letters even though the stimuli were false color words and false letters, respectively. These are both instances of high confidence in accessing one piece of information even though the information reflects the erroneous reconstruction of the stimuli due to a conjunction of degraded inputs and high priors.

Continuous vs. all-or-none conscious access

Before concluding, we would like to point out that the partial awareness hypothesis might also offer a parsimonious solution to the controversy about whether consciousness is an all-or-none or graded phenomenon [24,27–29]. Indeed, by incorporating the notion of independently accessible levels of representation, we can attain a graded perspective on conscious experience even with the all-or-none mechanisms for access proposed in workspace theories [3,27] (Box 1). Indeed, although representations within each level might be accessed in an all-or-none manner, in accordance with the winner-takes-all rule of workspace models, this does not imply that the full set of representations associated with the stimulus must be conscious. Access consciousness can be graded in the sense that representations of an object can be more or less complete, whereas the mechanisms of access can still be all-or-none.

Recent studies have argued that conscious access is evidenced as gradual, relying on continuous rather than discrete, all-or-none, measures of awareness. One example is the perceptual awareness scale, in which participants rate their conscious experience from 1 (no experience) to 4 (clear experience), including intermediate possibilities (e.g. 2, brief glimpse; a feeling that something was present, even though a content cannot be specified any further) [28–31]. However, intermediate situations can also be explained as partial awareness situations and thus reflect access to different levels of representation (e.g. access to low-level features without access to object identity). Thus, it is possible that intermediate cases on these continuous scales reflect, in fact, all-or-none conscious access to limited levels of representation.

Concluding remarks and future directions

Several aspects of the preliminary proposal presented here remain speculative and require further specification, particularly regarding the exact interplay between prior information, confidence and signal processing, as well as the interaction between levels of representation. Although extended discussion is beyond the scope of the present article, many of these interactions (i.e. between signal and prior information, between levels of representations) can be construed in terms of Bayesian inference, as described in recent hierarchical Bayesian accounts of vision [32,33]. In addition, although we proposed here that each level can be accessed independently, several important issues remain. For instance, it remains an open question as to whether conscious access can be performed for several levels in parallel, with a read out of appropriate

representations solely at the decision level. In this case, how many conscious contents can we track simultaneously and how does this relate to limits in working memory resources? Alternatively, conscious contents might actually be accessed only one level at a time by switching back and forth between task-relevant representations.

Despite its tentative nature, the partial awareness hypothesis offers several advantages in addition to parsimony. First, it provides an operational framework for predicting and manipulating the subjective impression of richness (e.g. by modulating reliance on prior information). Second, contrary to previous accounts that either deny the problem [4] or refer to it as an unspecified cognitive illusion (Box 2), characterization of the illusion of richness is essential in our framework and might make it possible to overcome the divide between phenomenology and psychology promoted by dissociative theories of consciousness. Nevertheless, although the impression of richness might be based on illusory perceptual contents, understanding the function of this illusion will be critical for future research on consciousness.

References

1 Koch, C. (2004) *T Q o Co c : A N ob o o ca A oac* Roberts & Co.
2 Kouider, S. (2009) Neurobiological theories of consciousness. In *E c c o a o Co c* (Banks, W., ed.), pp. 87–100, Elsevier
3 Baars, B.J. (1989) *A Co T o o Co c* , Cambridge University Press
4 Dennett, D.C. (1991) *Co c E a* , The Penguin Press
5 Crick, F. and Koch, C. (1995) Are we aware of neural activity in primary visual cortex? *Na* 375, 121–123
6 Dehaene, S. and Naccache, L. (2001) Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Co o* 79, 1–37
7 de Gardelle, V. and Kouider, S. (2009) Cognitive theories of consciousness. In *E c c o a o Co c* (Banks, W., ed.), pp. 135–146, Elsevier
8 Chalmers, D. (1996) *T Co c M* , Oxford University Press
9 Block, N. (1995) On a confusion about a function of consciousness. *B . B a Sc* . 18, 227–287
10 Kouider, S. and Dehaene, S. (2007) Levels of processing during non-conscious perception: a critical review of visual masking. *P o . T a . R. Soc. Lo . B. B o . Sc* . 362, 857–875
11 Sperling, G. (1960) The information available in brief visual presentation. *P c . Mo o* . 74, 1–29
12 de Gardelle, V. *a* . (2009) Perceptual illusions in brief visual presentations. *Co c Co* . 18, 569–577
13 Dehaene, S. *a* . (2006) Conscious, preconscious, and subliminal processing: a testable taxonomy. *T Co* . *Sc* . 10, 204–211
14 Block, N. (2007) Consciousness, accessibility, and the mesh between psychology and neuroscience. *B . B a Sc* . 30, 481–499 (discussion 499-548)
15 Lamme, V.A. (2006) Towards a true neural stance on consciousness. *T Co* . *Sc* . 10, 494–501
16 Rosenthal, D.M. (2002) How many kinds of consciousness? *Co c* . *Co* . 11, 653–665
17 Baars, B.J. and Laureys, S. (2005) One, not two, neural correlates of consciousness. *T Co* . *Sc* . 9, 269 (author reply 270)
18 Gescheider, G. (1997) *P c c : T* *a a* , Lawrence Erlbaum Associates
19 O'Regan, J.K. and Noe, A. (2001) A sensorimotor account of vision and visual consciousness. *B . B a Sc* . 24, 939–973 (discussion 973-1031)
20 Kouider, S. *a* . (2007) Cerebral bases of subliminal and supraliminal priming during reading. *C b. Co* 17, 2019–2029
21 Craik, F. and Lockhart, R. (1972) Levels of processing: a framework for memory research. *J. V ba L a . V ba B* . 11, 671–684
22 Greene, M.R. and Oliva, A. (2009) The briefest of glances: the time course of natural scene understanding. *P c . Sc* . 20, 464–472
23 Greene, M.R. and Oliva, A. (2009) Recognition of natural scenes from global properties: seeing the forest without representing the trees. *Co . P c* . 58, 137–176
24 Cleeremans, A. (2008) Consciousness: the radical plasticity thesis. *P o . B a R* . 168, 19–33
25 Lau, H.C. (2008) A higher order Bayesian decision theory of consciousness. *P o . B a R* . 168, 35–48
26 Lau, H.C. and Passingham, R.E. (2006) Relative blindsight in normal observers and the neural correlate of visual consciousness. *P oc. Na . Aca . Sc . U. S. A.* 103, 18763–18768
27 Sergent, C. and Dehaene, S. (2004) Is consciousness a gradual phenomenon? Evidence for an all-or-none bifurcation during the attentional blink. *P c . Sc* . 15, 720–728
28 Seth, A.K. *a* . (2008) Measuring consciousness: relating behavioural and neurophysiological approaches. *T Co* . *Sc* . 12, 314–321
29 Overgaard, M. *a* . (2006) Is conscious perception gradual or dichotomous? A comparison of report methodologies during a visual task. *Co c Co* . 15, 700–708
30 Dienes, Z. and Seth, A. (2010) Gambling on the unconscious: a comparison of wagering and confidence ratings as measures of awareness in an artificial grammar task. *Co c . Co* . 19, 674–681
31 Sandberg, K. *a* . (2010) Measuring consciousness: is one measure better than the other? *Co c . Co* DOI: 10.1016/j.concog.2009.12.013
32 Kording, K.P. and Wolpert, D.M. (2006) Bayesian decision theory in sensorimotor control. *T Co* . *Sc* . 10, 319–326
33 Lee, T.S. and Mumford, D. (2003) Hierarchical Bayesian inference in the visual cortex. *J. O . Soc. A . A O . I a Sc . V* . 20, 1434–1448
34 Chalmers, D. (1995) Facing up to the problem of consciousness. *J. Co c . S* . 2, 200–219
35 Block, N. (2005) Two neural correlates of consciousness. *T Co* . *Sc* . 9, 46–52
36 Milner, A.D. and Goodale, M.A. (1995) *T V a B a Ac o* , Oxford University Press
37 Goodale, M. (2007) Duplex vision: separate cortical pathways for conscious perception and the control of action. In *T B ac Co a o o Co c* (Velmans, M. and Schneider, S., eds), pp. 616–627, Blackwell
38 Zeki, S. (2007) A theory of micro-consciousness. In *T B ac Co a o o Co c* (Velmans, M. and Schneider, S., eds), pp. 580–588, Blackwell
39 Zeki, S. and Bartels, A. (1999) Toward a theory of visual consciousness. *Co c . Co* . 8, 225–259
40 Lamme, V.A. (2003) Why visual attention and awareness are different. *T Co* . *Sc* . 7, 12–18
41 Bar, M. *a* . (2006) Top-down facilitation of visual recognition. *P oc. Na . Aca . Sc . U. S. A.* 103, 449–454
42 Summerfield, C. and Egner, T. (2009) Expectation (and attention) in visual cognition. *T Co* . *Sc* . 13, 403–409
43 Kleinschmidt, A. *a* . (2002) The neural structures expressing perceptual hysteresis in visual letter recognition. *N o* 34, 659–666
44 Kouider, S. and Dupoux, E. (2004) Partial awareness creates the "illusion" of subliminal semantic priming. *P c . Sc* . 15, 75–81
45 Koch, C. and Tsuchiya, N. (2007) Attention and consciousness: two distinct brain processes. *T Co* . *Sc* . 11, 16–22
46 Koch, C. and Tsuchiya, N. (2007) Phenomenology without conscious access is a form of consciousness without top-down attention. *B . B a Sc* . 30, 509–510
47 Prinz, J. (2000) A neurofunctional theory of visual consciousness. *Co c . Co* . 9, 243–259
48 Reddy, L. and Koch, C. (2006) Face identification in the near-absence of focal attention. *V o R* . 46, 2336–2343
49 Li, F.F. *a* . (2002) Rapid natural scene categorization in the near absence of attention. *P oc. Na . Aca . Sc . U. S. A.* 99, 9596–9601
50 Treisman, A. and Gelade, G. (1980) A feature-integration theory of attention. *Co . P c* . 12, 97–136
51 Huang, L. *a* . (2007) Characterizing the limits of human visual awareness. *Sc c* 317, 823–825