NEUROSCIENCE of Consciousness

Neuroscience of Consciousness, 2023, **2023(1)**, 1–14

DOI: https://doi.org/10.1093/nc/niac018

Review Article

The functions of consciousness in visual processing

Dylan Ludwig*,†

Department of Philosophy, York University, 4700 Keele Street, Toronto, ON M3J 1P3, Canada

†Dylan Ludwig, http://orcid.org/0000-0002-7019-693X

*Correspondence address. Department of Philosophy, York University, 4700 Keele Street, Toronto, ON M3J 1P3, Canada. Tel: +1 (647) 668 4178;
E-mail: dylan.m.ludwig@gmail.com

Abstract

Conscious experiences form a relatively diverse class of psychological phenomena, supported by a range of distinct neurobiological mechanisms. This diversity suggests that consciousness occupies a variety of different functional roles across different task domains, individuals, and species; a position I call *functional pluralism*. In this paper, I begin to tease out some of the functional contributions that consciousness makes to (human) visual processing. Consolidating research from across the cognitive sciences, I discuss semantic and spatiotemporal processing as specific points of comparison between the functional capabilities of the visual system in the presence and absence of conscious awareness. I argue that consciousness contributes a cluster of functions to visual processing; facilitating, among other things, (i) increased capacities for semantically processing informationally complex visual stimuli, (ii) increased spatiotemporal precision, and (iii) increased capacities for representational integration over large spatiotemporal intervals. This sort of analysis should ultimately yield a plurality of functional markers that can be used to guide future research in the philosophy and science of consciousness, some of which are not captured by popular theoretical frameworks like global workspace theory and information integration theory.

Keywords: visual consciousness; functions; unconscious processing

Introduction

The theoretical and experimental investigation into consciousness' function(s) in information processing systems is still in its early stages. One potentially fruitful framework centres around the hypothesis that consciousness contributes a variety of different functions in different psychological task domains, such as vision, emotion, and social cognition. This position—which I call *functional pluralism*—broadens the possible scope of explanation in consciousness research and has the potential to help establish to a more accurate and nuanced picture of what experience *is* as a feature of certain complex systems. Identifying consciousness' function(s) in the psychologically complex domain of human vision is a natural starting point for the pluralistic project. Visual perception has long been the object of intensive analysis in Western philosophy (e.g. Locke 1689; Russell 1912; Burge 2010) and has historically been treated as paradigmatic of phenomenal content in philosophical discussions about consciousness (e.g. McDowell 1994; Tye 2000; Prinz 2011; Block and Phillips 2016). Similarly, the primate visual system is perhaps the most thoroughly studied sub-system in the psychological and neural sciences (e.g. Hubel and Wiesel 1968; van Essen *et al.* 1992; Hilgetag *et al.* 2016) and has become the primary experimental foothold into scientific issues surrounding consciousness (e.g. Dehaene and Changeux 2011; Tononi *et al.* 2016; Block 2019; Lamme 2020).

In this paper, I begin the process of carefully teasing out some of the specific functional contributions that consciousness makes to (human) visual processing. After briefly outlining some conceptual and methodological background (Section 1.1), I offer an empirically informed account of the functional capacities of the human visual system in the absence of conscious experience (Section 2). This account represents the basis of comparison for isolating some functional capacities that are unique to consciousness in the domain of visual processing (Section 3). Finally, I draw out some preliminary conclusions and discuss how my analysis challenges the explanatory scope of existing theories of consciousness like global workspace theory (GWT) and information integration theory (IIT) (Section 4).

Some conceptual and methodological preliminaries

A range of specific hypotheses about conscious experience can be tested with carefully designed visual processing tasks. Robust relationships between visual stimuli and patterns of neural, psychological, and behavioural response are relatively well established (e.g. Boly *et al.* 2017). In addition, many reliable psychophysical tools that can be used to objectively probe the functional–psychological components of visual processing (e.g. priming, adaptation paradigms) have been carefully developed over many decades of research (e.g. Kominsky and Scholl 2020). Finally, the ascription of conscious experience during visual processing tasks is becoming more systematic and experimentally reliable, thanks to increasingly stringent subjective report, no-report, and even

Received: 1 August 2022; **Revised:** 24 October 2022; **Accepted:** 6 December 2022

recently proposed no-cognition paradigms (see Brascamp *et al.* 2015; Block 2019). This has all been leveraged by consciousness researchers in the ongoing development of the field's primary methodology: task-based comparisons of conscious versus unconscious processing designed to isolate structural and functional features that are highly associated with conscious experience.

Like other psychological properties, there is good reason to think that consciousness contributes functionally to the overall causal dynamics of psychological systems. There is mounting evidence that the presence or absence of consciousness makes a difference in terms of psychological function, according to a causal interventionist approach to scientific explanation. Such an approach to consciousness can be seen most clearly in vision research. The most compelling experiments aim to hold all else fixed on some visual task (e.g. object categorization) except for whether or not experience is present (e.g. due to masking or suppression).[1] The functional consequences of this experimental manipulation provide strong theoretical reasons for assuming that there are task-relevant psychological capacities that require consciousness. This general methodological point is becoming orthodoxy in consciousness research (e.g. Dehaene *et al.* 2006; Cohen and Dennett 2011; Frith and Metzinger 2016; Boyle 2019; Birch 2020; Lamme 2020) and has already proven fruitful in our ability to advance claims about the psychological functions of conscious experience.

To understand precisely what it is that consciousness *does* in information processing systems, we have to look closely at particular functions employed for specific processing tasks in order to isolate the psychological capacities that truly are *sufficient for,*[2] or are unique to, conscious experience. As predicted by a functional pluralist approach, however, we should not expect these functions to be necessary for experience (i.e. present in every case), given how diverse conscious experiences are across the natural world (e.g. conscious vision versus conscious emotion versus conscious sonar). Much like the search for structural markers or neural correlates of consciousness, which similarly benefits from taking a pluralistic approach (e.g. Koch *et al.* 2016; Malach 2021), an investigation into the functional contributions of consciousness (FCCs) should ultimately yield a collection of functional markers that can be used as operationalizable proxies for experience in future research (Ludwig 2022).

Appreciating unconscious visual processing

Contrasting conscious and unconscious processing requires that we first get the best possible picture of the functional capacities of the visual system in the absence of awareness.[3] We want to see where unconscious processing, in the context of a specific task, individual, or species, can no longer handle the functional demands placed on the system. This means that even if there is redundant functionality among conscious and unconscious visual systems, establishing the upper functional limits of unconscious visual processing should yield a more informative functional comparison with conscious visual processing, as it helps us isolate precise functional shortcomings when visual information is processed in the absence of awareness.

Although it remains controversial in certain philosophical discussions (e.g. see the debate in Block and Phillips 2016), many cognitive scientists have endorsed the existence of unconscious visual perception (e.g. Kouider and Dehaene 2007; Tamietto and de Gelder 2010; Quilty-Dunn 2019). The idea is that a remarkable amount of visual processing seems to happen unconsciously and to such an extent that it seems plausible to ascribe paradigmatically perceptual capacities to human subjects that remain wholly outside of conscious experience. And while it is not entirely necessary for understanding conscious visual perception that there are unconscious processing capacities that meet some pre-established philosophical criteria for being genuinely 'perceptual'—and so I will not labour this specific point in what follows—employing this terminology can help us appreciate what the relevant research has revealed about the functional capabilities of unconscious processing in the domain of vision.

Several empirical paradigms bear on the functional capacities of unconscious visual processing. I will focus on just a couple of clusters of evidence in order to offer a thorough examination of each: one that emerges from laboratory settings and careful experimental manipulation and another that emerges from the study of particular pathological impairments to visual processing.

Continuous flash suppression

There are two dominant experimental paradigms for studying unconscious visual processing: visual masking and interocular suppression. Masking techniques allow a rapidly presented visual stimulus to be kept from conscious awareness due to the presentation of a second stimulus closely before and/or after the target. Subjects are subsequently asked to perform a range of different tasks that indirectly probe the extent to which this visual information has been registered by the visual system despite failing to reach consciousness (Breitmeyer and Öğmen 2006; Kouider and Dehaene 2007). Visual masking studies are also often performed while subjects are undergoing some form of neuroimaging (e.g. functional magnetic resonance imaging, magnetoencephalogram, electroencephalogram (EEG), single-cell recordings) aimed at identifying and contrasting the neural correlates of conscious and unconscious visual processing. Indeed, masking techniques played a significant role in the discovery that unconscious processes are typically marked by a feed-forward sweep of neural activity, whereas conscious vision appears to be marked by feedback or recurrent neural activity (e.g. Fahrenfort *et al.* 2007). Although masking has a long history in the vision science community, a new form of interocular suppression is quickly becoming the primary method for comparing the structural and functional elements of conscious and unconscious visual perception.

In general, interocular suppression techniques exploit a unique feature of binocular vision; namely, 'the reflexive suppression that occurs when different images are simultaneously presented to the two eyes' (Yang *et al.* 2014, 1). When different stimuli are shown to each eye separately but at the same retinal location, the visual system fails to fuse them into a single percept due to the natural constraint that dictates that two different objects cannot occupy the same location in space. Instead, the two stimuli 'compete' for subjective awareness; that is only one can be consciously

[1] One potential confound that deserves explicit mention is the degradation of internal signal strength (i.e. due to masking or suppression), which might account for some observed performance failures that are otherwise assumed to be driven solely by whether or not subjects were aware of the stimulus (e.g. Morales *et al.* 2015; Peters *et al.* 2017). Various proposals have been put forth in order to address and eliminate performance confounds based on degraded signal strength; including quantifying internal perceptual response and matching it to some pre-established set of (e.g. neural) criteria for awareness (Morales *et al.* 2015) or adjusting stimuli in order to compensate for differences in processing sensitivity (e.g. of contrast) on unconscious trials (Persaud *et al.* 2011).

[2] This is not intended to represent sufficient *conditions* for consciousness, but rather a logical relation between psychological functions and consciousness (i.e. if sufficient, a function is never a feature of unconscious processing, and so its presence is enough to ascribe consciousness to the system).

[3] A similar point is often made about the search for neural correlates of consciousness (e.g. see Breitmeyer 2015).

seen at a time. There is typically a relatively spontaneous alternation between which stimulus reaches conscious awareness at a given moment and which remains supressed from awareness (Kovacs *et al.* 1996). Despite this, the primary advantage of interocular suppression over masking techniques is that stable input to the visual system can remain unconscious for longer periods of time. Modifications of this general experimental set-up have allowed researchers to more precisely identify when a stimulus is in fact consciously experienced and when it is not, which can ultimately provide compelling evidence of both the neurobiological correlates and the functional capabilities of unconscious visual processing.

One form of interocular suppression in particular has emerged as the primary psychophysical tool for comparing conscious and unconscious visual processing: continuous flash suppression (CFS). Mudrik *et al.* (2011) employ CFS and describe it as follows:

> In CFS, distinct color images (Mondrians) presented successively at approximately 10 Hz to one eye can reliably suppress the conscious awareness of an image presented to the other eye for a relatively long duration (Mudrik *et al.* 2011, 765).

Essentially, CFS is a unique modification of the interocular suppression paradigm that involves the controlled suppression of one monocular input from awareness by rapidly and repeatedly presenting a high-contrast colour-patterned image, or Mondrian, to the other eye, which dominates visual awareness. Stimuli like Mondrians have traditionally been presented to subjects on computer screens, although some researchers have begun to use CFS to supress real objects in the environment from awareness with the help of augmented reality goggles (Korisky *et al.* 2019). Masking and suppression techniques like CFS are also typically paired with other tools of experimental psychology like priming or the induction of adaptation after-effects (Yang *et al.* 2014). The idea is that the evidence of priming or adaptation involving visual stimuli that are suppressed from conscious awareness suggests that those stimuli are being processed unconsciously in ways that are paradigmatically perceptual. These tools are therefore used to probe the extent of unconscious visual processing.

A further modification of CFS, known as breaking continuous flash suppression (b-CFS), involves slowly reversing the contrast between the two images to determine the precise moment that the target image finally 'breaks through' suppression and becomes consciously seen (Stein and Sterzer 2014). Interestingly, reliable patterns emerge here; for example, some classes of stimuli reliably break through suppression faster than other classes of stimuli. The theoretical interpretation here is typically that the initial unconscious processing of the invisible stimulus can boost its input signal, somehow preferentially 'empowering' it to rise to the level of conscious experience (Yang *et al.* 2014, 5). Stein and Peelen (2021) suggest that the results of b-CFS experiments can be buttressed by including a second detection task that directly measures the visual dimension thought to be driving detection differences. The subject's failure to discriminate this dimension can be taken as evidence that stimuli were truly processed unconsciously. The underlying principle is that more 'meaningful' stimuli are brought into consciousness more quickly by way of unconscious processing. This accounts for some of the observed variability between subjects in b-CFS experiments, as 'meaningfulness'—understood here as the degree of perceptual or cognitive salience—depends on a range of contextual variables like individual learning history and experimental design (see Fig. 1).

Like any empirical paradigm, the results of CFS studies need to be carefully assessed before anything philosophically relevant can be extracted from them. Accordingly, several prominent philosophers and scientists have critically reviewed key aspects of the CFS paradigm (see Yang *et al.* 2014; Block and Phillips 2016). There are long-standing methodological concerns surrounding both the subjective and objective measures that are used in the laboratory to verify whether or not stimuli are consciously perceived, and these certainly also arise in CFS research. Despite acknowledging unconscious visual processing, Ian Philips (2018), e.g., has consistently raised the concern that subjective report, one of the primary methods used to determine whether a stimulus is consciously seen or not on CFS trials, faces the 'problem of criterion'. The concern here is that it is always possible to interpret subjects' reports as reflecting a conservative response bias, which opens up the possibility that subjects are perhaps at least 'partially aware' of stimuli that they report being unaware of on some trials.

There are also challenges in interpreting the behavioural measures that are used to objectively probe awareness, such as when subjects are asked either to detect the presence of a target or to discriminate between a target and a decoy. This is because (i) some argue that chance-level performance on these tasks does not necessarily entail that stimuli were unconscious, given that 'the absence of evidence is not the evidence of absence' (Altman and Bland 1995) and (ii) the aforementioned chance-level performance does not necessarily entail that stimuli were consciously seen, given that some of these behavioural tasks can be accomplished in the absence of awareness. Furthermore, subjective and objective measures of awareness can dissociate in laboratory settings, pointing to different conclusions about the conscious experiences of a subject (Yang *et al.* 2014). These all seem to make it very difficult to be certain whether or not awareness is truly absent during specific CFS trials.

Several responses to these methodological concerns have emerged in both theory and practice. In terms of subjective measures of awareness, report paradigms now commonly involve graded perceptual awareness scales in an attempt to nuance subjects' responses. Subjective reports are also increasingly combined with measures of confidence. Some studies employ 'post-decision wagering', e.g., in which a subject's confidence in their reports is indexed by the amount of money that they were willing to bet on their accuracy (e.g. Persaud *et al.* 2007). Other paradigms rely on subjects' metacognitive judgments without introducing the risk/loss aversion that comes with wagering (Maniscalo and Lau 2012). These supplemental methods can help researchers control for potential report biases and motivate subjects to respond without conservative response criteria.

It is important to note some further complications surrounding subjective report in consciousness research. For one thing, experience is often quantitatively richer than what subjects can report on at a given moment (e.g. Block 2011, 2014). This means that subjects may be conscious of a stimulus that they have not accessed by the mechanisms underlying verbal report. In other words, the ability to make a subjective report may not be a necessary condition for awareness. To be as clear as possible, the claim is typically that perceptual experiences can remain potentially reportable (say with the appropriate cue, as in the pioneering study by Sperling 1960) while not being actually accessed for report by subjects. Perceptual stimuli that are wholly unavailable for report, i.e. that are *unreportable* because they are presented too rapidly for instance, are not typically assumed to be part of this conscious 'overflow' in perception (Block 2011, 2014). On the other hand, this also means that the cognitive processing underlying subjective report is a potential confounding variable for debates about whether or not either access for global broadcasting or some meta-representational processing is necessary for perceptual consciousness above and beyond processing in perceptual systems, especially in the context of the search for neural correlates of consciousness.

As such, many researchers have developed 'no-report' paradigms, where a subjective report is used to eventually calibrate objective markers like eye movement patterns and pupil dilation that can be objectively linked to visual awareness (e.g. Frässle *et al.* 2014). There are still potential confounds here though, as even without explicit report, merely noting a change in the contents of perceptual awareness may produce cognitive effects that are unrelated to the conscious experience itself, but that still accompany it. Psychologists like Brascamp *et al.* (2015) and philosophers like Block (2019) have thus advocated for an even more rigorous 'no-cognition' paradigm. Using binocular rivalry, no-cognition trials involve the standard binocular alternation between stimuli, except that subjects cannot detect the change due to the nature of the stimuli used (e.g. randomly and sporadically moving patterns of the same colour dots), which eliminates any confounding cognitive processing. It remains to be seen how effectively CFS can accommodate a no-cognition paradigm. Regardless, much work has been done to supplement subjective report in consciousness research, which must ultimately be acknowledged as an extremely valuable foothold into the study of conscious experience.

In terms of other standard objective measures of awareness, scientists who employ CFS to study conscious and unconscious visual processing have recognized that different measures, like detection and categorization tasks, tap into different stages of processing. This means that apparent failures to detect the presence of a stimulus might be good evidence that it did not reach consciousness, whereas apparent failures to categorize a stimulus based on some feature might still occur either with minimal conscious awareness of the whole stimulus or in cases where certain (e.g. low level) stimulus features break through the suppression but not others (e.g. higher-level features). One solution, therefore, is to employ several independent and increasingly stringent measures of awareness—both subjective and objective—within a single study and to avoid generalization across these measures (e.g. Yokoyama *et al.* 2013; Gelbard-Sagiv *et al.* 2016). Another way to strengthen objective measures of awareness is to integrate them with the distinct measures used to probe the extent of visual processing (e.g. priming and adaptation) within a single experiment, so that both task design and levels of attention and motivation remain constant. These distinct tools that are used to assess the extent of visual processing on a given CFS trial, like priming or inducing adaptation effects, are much less controversial in the cognitive sciences and indeed have become popular methods in a wide variety of experimental paradigms. In summary, while there is always room for continued refinement of the methods employed, it is generally agreed that CFS is a viable and illuminating experimental design for rendering visual stimuli unconscious and assessing the extent of unconscious processing of visual information.

Given this arsenal of experimental tools and the massive body of relevant empirical literature, there is a robust enough body of research here to begin to get a sense of the kinds of visual processes that can occur unconsciously. Experiments using CFS have revealed a range of higher-level visual processes that do not require consciousness. Gelbard-Sagiv *et al.* (2016) provide a recent summary of some of the most compelling evidence (see also references therein):

> Remarkably, several recent studies demonstrated that many high-level processes can take place even when the stimuli are invisible: observers were found to read and process the meaning of words, process semantic incongruences in written sentences and visual scenes, perform arithmetic operations, categorize faces and other objects, process emotions, and exercise executive functions, in the absence of perceptual awareness.

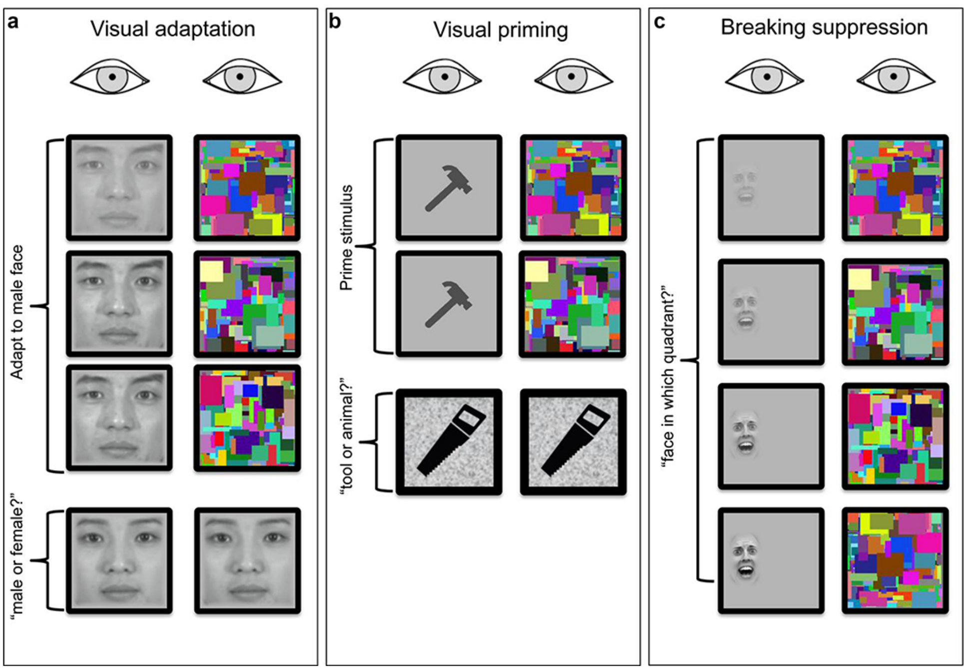

Figure 1. Schematic representation of the three main CFS paradigms. (a) CFS paired with the visual adaptation paradigm. (b) CFS paired with the visual priming paradigm. (c) Breaking-CFS paradigm. From Yang *et al.* (2014)

While the authors suggest that we be cautious about some of these results, as some experimental paradigms might not be sensitive to the possibility that some stimulus features are being suppressed from awareness while others are not, CFS has delivered some fairly robust evidence of the functional capacities of unconscious visual processing.

Unconscious visual integration

Interestingly, some laboratories have used this technique to challenge specific core assumptions of some of the most prominent theories of consciousness. IIT, e.g., assumes that conscious experience is required for the integration of disparate information into a unified whole (Tononi *et al.* 2016). In contrast, Liad Mudrik and colleagues have dedicated years of research to uncovering the kinds of integration of visual information that can be accomplished without subjective awareness, often employing CFS techniques (for review, see Mudrik *et al.* 2014). It has been shown, e.g., that unconscious visual processes can facilitate (i) the association of visually presented words, even with fairly large temporal integration windows of up to 78 seconds (Reber and Henke 2012), (ii) the integration of disparate features of visual stimuli, even with relatively high spatial integration windows (Oriet and Brand 2012), and (iii) high-level semantic and syntactic integration of visually presented words and numbers (Sklar *et al.* 2012).

Looking closely at a particular experiment can help illustrate this empirical finding regarding unconscious information integration. Plass *et al.* (2014) developed a CFS study that tested the extent of unconscious audio-visual integration, relying both on the logic of priming studies and the common finding that visual information about mouth movements can influence the auditory processing of words. They found that lip movements that were rendered invisible with CFS still facilitated performance on tasks that required subjects to categorize a spoken word (e.g. as either a tool word or non-tool word), when the target word was the same as that articulated by the suppressed lip movements. The study employed both subjective measures of awareness—namely, reports on whether the face was visible in addition to reports on the location of a circular probe near the mouth—and objective measures of awareness—namely, subjects were asked to indicate the colour of a translucent ellipse placed over the mouth region of the suppressed face. This study provides compelling evidence that unconsciously processed visual information was integrated with auditory and linguistic information when performing the word categorization task, as is standard in speech perception (see Venezia *et al.* 2015).

IIT theorists might deny that this unconscious visual processing counts as genuine information integration as defined by their research program, in the sense that involves combining distinct 'conceptual' structures into a single unified representation (Tononi *et al.* 2016). More specifically, this priming effect might either be considered too low level or informationally simplistic to capture the kinds of cause–effect repertoires underlying the IIT of consciousness, or it might instead be construed as the merely serial processing of distinct representational elements without integration. However, both of these objections are thwarted by the fact that the audio-visual processing in question fits perfectly well with IIT's definition of integration, when understood as the graded notion it was intended to be. That is the resulting causal and computational resources are more than the summed total of the resources provided by the component representations (Mudrik *et al.* 2014), and this is simply not the case with merely serial processing. Moreover, this sort of multi-modal perceptual integration specifically facilitates speech recognition and language comprehension, which are generally assumed to be 'higher' information-processing achievements. This is just a sample of the research aimed specifically at uncovering the extent to which unconscious processes can integrate visual information, but it brings us closer to identifying the upper limits of the functional capabilities of unconscious vision more generally. Unconscious visual processing seems to be capable of at least some genuine forms of information integration. Put differently, information integration turns out to be insufficient for conscious experience.

Accessing unconscious visual information

The same kinds of challenges have been made to assumptions held by GWT and specifically the claim that consciousness' functional role is to facilitate wide-ranging access to information (Dehaene 2014). A variety of studies have shown that unconsciously processed visual information is available to the same processing subsystems identified in the global workspace model (Dehaene and Changeux 2011); namely, evaluative systems, long-term memory systems, attentional systems, language systems, and motor systems. For instance, much work has been done to understand the role of unconsciously processed visual information in the guidance of action (e.g. Brogaard 2011; Goodale and Milner 2013). The visual system constructs representations in preparation for visually guided action that are available to motor and decision-making systems despite their failing to reach awareness (e.g. Bargh and Morsella 2008). Even the most basic visuomotor tasks might require that visual information processed outside of awareness be freely used by systems that predict, compare, and execute intended actions. CFS studies also repeatedly reveal that unconsciously processed words affect semantic processing networks and can even influence problem-solving strategies (e.g. Zabelina *et al.* 2013). Finally, a variety of research paradigms suggest that unconsciously processed visual information is available to systems responsible for emotional or evaluative assessment (e.g. Fang *et al.* 2016; Diano *et al.* 2017); in fact, visual stimuli are the primary experimental tool used to explore unconscious emotional processing (e.g. Morris *et al.* 1998; Williams *et al.* 2004; Mendez-Bertolo *et al.* 2016). In general, these experimental paradigms suggest that unconscious visual information is also accessed for use by a relatively wide range of downstream processing subsystems.

One study by Sklar *et al.* (2012) illustrates this sort of challenge to GWT. In one experimental set-up, relatively complex mathematical problems (e.g. three-digit subtraction equations) were suppressed from awareness using CFS. Both an objective forced-choice measure and a nuanced subjective measure consisting of direct questions about the trials were used to establish subjects' lack of awareness of the suppressed math problems. After the suppressed primes were presented, subjects were asked to pronounce out loud a target number that was either the correct or incorrect solution to the suppressed equations. The researchers found a significant priming effect in reaction times to correct responses, suggesting that 'the primed equation was mentally accessed (i.e. that the equation had been solved)' (Sklar *et al.* 2012, 19616), even though the visual information remained unconscious. Subliminal priming of this sort continues to be developed as a valuable tool for studying a wide variety of unconscious influences on thought and behaviour (see Elgendi *et al.* 2018).

GWT theorists have ruled out this interpretation of unconscious priming effects in vision, based on their definition of unconscious processing as 'a condition of information inaccessibility' (Dehaene *et al.* 2006, 3), according to which unconscious information cannot facilitate task performance in this way

(Bussche *et al.* 2008). But any facilitation in task performance indicates at least some form of access to unconscious visual information by decision- and action-guiding systems, whether or not the access is 'global' in the relevant sense. It should also be noted that the 'global' access proposed by GWT is not taken to be absolute, as information in the global workspace need not be accessible by every system—GWT is thought to be compatible with modularity and informational encapsulation (Dehaene *et al.* 1998). This ultimately suggests again that we adopt a graded notion of access, and there is no obvious reason to deny that this kind of cognitive operation crosses the conscious/unconscious divide. Evidence that unconscious visual information is used by systems that process mathematical equations, like other unconscious priming studies, provides compelling reason to doubt that access is the (sole) functional contribution that conscious experience makes to visual perception. Neither integration nor access is therefore unique to (or sufficient for) conscious experience, and so they are not likely candidates for picking out the functional contribution that consciousness makes in the domain of vision.

To summarize, in order to isolate what it is that consciousness contributes functionally to visual perception, it is necessary to get the best picture possible of the upper limits of unconscious visual processing capacities. It is becoming a widespread theoretical assumption that 'feature extraction, categorization, some interference, and inference occur regardless of whether one is conscious of the visual stimulus or not' (Lamme 2020). However, the CFS research programs outlined previously help to reveal a central point that has been emerging over the last few decades of consciousness research: unconscious visual processing is functionally impressive, but it is also functionally limited. Unconscious vision is itself functionally hierarchical and can take on increasingly demanding information processing tasks (Breitmeyer 2014), and yet it seems to max out at a certain level of functional complexity, at which point, the resources of consciousness are presumably recruited. The functions of access and integration exhibit this point in an interesting way: some kinds of access and integration can be carried out unconsciously while others cannot. This means that integration and access ought to be understood as graded notions that in some ways dissociate from conscious processing and therefore cannot be common denominators that exhibit the function of conscious experience in visual perception.

It is important to keep in mind the possibility that the specific limits on unconscious visual processing depends on a variety of factors, and so are likely to vary across different individuals, different task demands, and presumably different species. Nevertheless, while unconscious visual processes have an astonishing arsenal of functionality, perhaps even enough to establish genuine perception (e.g. say on the grounds of criteria like perceptual constancy, see Block in Peters *et al.* 2017), their limits suggest that conscious processes contribute functional resources that are otherwise lacking. This specific point about visual processing is significant both for (i) general theories that continue to deny that consciousness adds any functionality to the psychological system (e.g. Hassin's (2013) 'Yes It Can' principle, which states that unconscious processes have all the functional capability of conscious ones) and (ii) theories like IIT and GWT that mistakenly inflate the FCCs at the expense of appreciating the functional capacities of unconscious visual processing.

Visual neglect

Another major source of evidence for unconscious visual processing comes from studying individuals who have experienced significant damage to the visual system. Neuropsychological analyses of pathological conditions have been a significant source of insight into the history of cognitive science, and again this is seen paradigmatically in the domain of vision. Disorders like blindsight, visual object agnosia, and visual neglect have continued to provide unique access into the workings of conscious and unconscious visual processing. Some cases of visual spatial neglect offer a particularly colourful illustration of the functional capacities and limitations of unconscious visual processing.

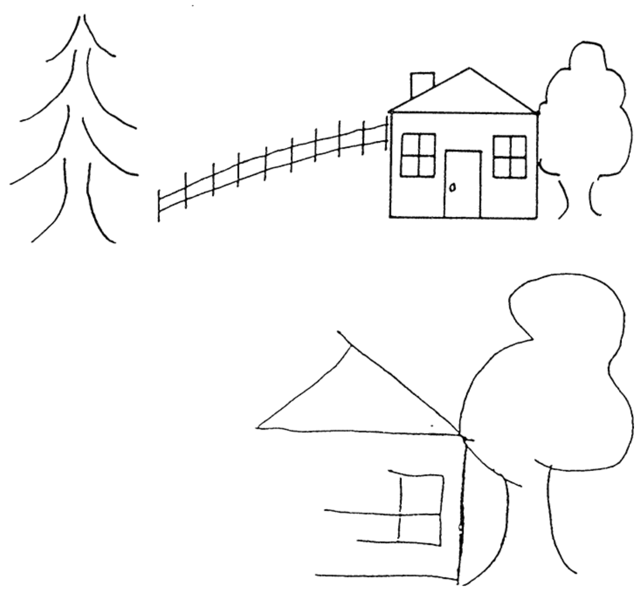

Figure 2. Top row: The Ogden Scene Test; A Test of Visual Hemineglect. Bottom Row: Copy of the drawing by a patient with a large glioma in the right posterior parietal lobe. From Ogden (2005)

Visual hemineglect is a particular neuropsychological disorder that is caused by cellular damage in areas of the posterior parietal cortex and surrounding structures, typically as a consequence of strokes or aggressive tumours (e.g. Ogden 2005). The resulting deficit involves a lack of visual awareness in the region of the visual field that is contralateral to the damaged neural tissue (Berti and Rizzolatti 1992). In many instances, patients even seem to neglect the corresponding space of visual imagery drawn from long-term memory, like one side of a familiar street they are asked to recall (e.g. Bisiach and Luzzatti 1978), suggesting that the lack of awareness cannot be explained away as a mere deficit in attentional mechanisms. Despite this lack of awareness in hemineglect patients, there is evidence that visual information in the damaged visual field is processed unconsciously to some extent. Striking dissociations between subjective measures of awareness and objective measures of performance are well documented in hemineglect patients.

The first body of evidence for unconscious visual processing in patients with hemineglect is more anecdotal in nature. Several researchers have captured unusual and theoretically fascinating behaviours that result from these specific deficits in conscious visual perception (see Fig. 2). In one famous case (Halligan and Marshall 1998), researchers presented two separate line drawings of a house to a hemineglect patient. In one of the images, the left side of the house (appearing in the affected region of the visual field) was on fire. The subject reported that there was no difference between the two images, suggesting that they were not consciously aware of the burning side of the house. Some might argue that the subject was at least partially aware of the neglected stimulus, although the extent of cortical damage caused by aneurysm and resulting haemorrhage, as well as broader patterns of behaviour (e.g. when asked to bisect horizontal lines, the patient was typically over 50% to the right of true centre), provides strong reasons to reject this interpretation (Halligan and Marshall 1998; Ogden 2005). Nevertheless, when subsequently asked which house they preferred to live in, the subject reliably indicated that they preferred the house that was not on fire. This has been taken by many cognitive scientists as compelling evidence of fairly sophisticated processing of the part of the visual image that the subject reported not being consciously aware of. Over the subsequent decades, a variety of these sorts of atypical behaviours and similarly unusual confabulations provide a prima facie reason to assume that although not consciously perceived, visual stimuli in neglected areas of visual space in hemineglect patients are still processed unconsciously (e.g. Verdon *et al.* 2010; Li and Malhorta 2015).

The second, and likely more convincing, body of evidence comes from additional experimental tools that are specifically designed to test the extent to which visual information is being processed in hemineglect patients. Like in CFS research, psychophysical tools such as priming can be used to probe unconscious visual processing under pathological conditions. Berti and Rizzolatti (1992), e.g., ran a standard priming study on patients with unilateral visual hemineglect. Their experiment showed that primes presented to the neglected portion of visual space facilitated task performance, even on trials where paired stimuli belonged to the same conceptual category despite being physically dissimilar. Once again, the extent of neural and behavioural pathology strongly indicates that visual consciousness was truly disrupted in the subjects. The authors concluded that 'patients with neglect are able to process stimuli presented to the neglected field to a categorical level of representation even when they deny the stimulus presence in the affected field' (Berti and Rizzolatti 1992, 345). Similarly, Nakamura *et al.* (2012) used a priming paradigm to assess the extent to which words are processed in hemineglect patients. They also found that primes facilitated task performance even when presented to the neglected part of the visual field.

As a result, Brogaard *et al.* (2020) have recently argued that research on unconscious processing in hemineglect patients casts further doubt on 'integrative' models like IIT and GWT as viable theories of consciousness and its function. They draw on evidence that certain visual illusions, like Kanizsa-style amodal completion illusions that require integration of visual information with 'amodal' assumptions about objects in the world, still occur in subjects with hemineglect who deny having an experience of half of the available visual cues (Vuilleumier and Landis 1998). Because this appears to be a genuine integration of visual information in the absence of awareness, integration is insufficient for conscious experience. The authors argue that the illusion only occurs when all the visual elements are integrated, and so the only viable interpretation is that despite failing to reach awareness, visual information is integrated in such a way as to play a constitutive role in establishing what becomes phenomenological content (Brogaard *et al.* 2020).

This sort of research further supports the picture that is emerging in experimental contexts: although functionally impressive, there are specific limits on unconscious visual processing. Sprenger *et al.* (2002), for instance, found that subjects with hemineglect had specific deficits in colour processing in the neglected parts of their visual field. Furthermore, much theoretical and empirical work has been done, e.g., to try to understand the extent to which visual information is processed in patients with similar pathological conditions like 'blindsight', which is caused by damage to primary visual cortices. Alexander and Cowey's (2010) research, e.g., suggests that when faces, colours, shapes, and patterns are presented to blindsight patients, only 'simple' stimulus features like luminance are processed unconsciously, as those features alone appear to be driving performance on perceptual discrimination tasks. Once again, several distinct research paradigms in neuropsychology are converging on the idea that there are functional limitations in unconscious processing that are plausibly related to the increasingly complex information processing demands that certain stimuli make on the visual system. Some specific performance failures under pathological conditions constitute compelling evidence that consciousness is recruited for tasks that are more functionally complex.

Comparing conscious visual perception

The same empirical tools that are used to investigate unconscious visual processes have been used to compare them with conscious ones. Once a minimal threshold for awareness is established, by leveraging report and/or appealing to objective indices of conscious experience, precise structural and functional comparisons can be made. The next step then is to look closely at some particular information-processing tasks in order to get the clearest picture possible of the similarities and differences in function between conscious and unconscious visual processing. This is necessary for isolating specific functional capacities that are unique to conscious visual processing. Building such a pluralistic functional profile will indeed be a conceptually nuanced project. There is wealth of research exploring very specific but at least conceptually unrelated visual functions like motion and colour processing; many individual differences in processing capability have been observed between subjects; and studies that use similar experimental paradigms can either produce different interpretations of the same results or produce different statistical results altogether. All of these will make drawing any substantial generalizations difficult at the moment. However, beginning to engage with this complexity by closely examining a few particular processes seems to be the only way forward for the philosophy and science of consciousness. In what follows, I will look at semantic and spatiotemporal processing carried out by the visual system, in order to more precisely compare functional capabilities on these different kinds of tasks both in the presence and absence of conscious experience.

Semantic processing in vision

One central set of functions carried out by the human visual system is the processing of semantic information. This occurs, e.g., both when we extract meaning from written language and when we recognize conceptual relations between bits of visual imagery. Semantic processing is a particularly good place to start building a functional profile for conscious experience in visual processing because the use of subliminally presented written words and images to prime performance on semantic tasks is extremely common in the cognitive sciences. It has certainly become one of the main battle grounds for debates about the extent and limits of unconscious visual processing (Kouider and Faivre 2017), which makes it particularly suitable for the kinds of contrastive analyses that dominate consciousness research.

A variety of masking and binocular rivalry studies, including experiments using CFS, have probed conscious versus unconscious semantic processing by using written words to prime behavioural responses. Robust unconscious priming effects with written words are generally well established both in psychological (e.g. Jiang *et al.* 2007; Costello *et al.* 2009; Reber and Henke 2012; Armstrong and Dienes 2013) and neuroimaging research (e.g. Nakamura *et al.* 2007; Axelrod *et al.* 2015). Zabelina *et al.* (2013), e.g., found that subjects performed better on compound association word problems, where three seemingly unrelated words (e.g. pine, crab, sauce) form familiar compounds with a solution word (e.g. apple), when the problem words were presented as subliminal primes during CFS. Crepaldi *et al.* (2010) similarly found that morphologically similar word pairs (e.g. 'fell' and 'fall') showed significant priming effects on a standard masking paradigm only when the words were conceptually related (e.g. so not 'bell' and 'ball'), suggesting unconscious processing of semantic relationships. And several studies have shown that more emotionally laden words reliably emerge from suppression faster than emotionally neutral words, which researchers have taken as compelling evidence of high-level semantic processing of written language in the absence of visual awareness (e.g. Yang and Yeh 2010; Sklar *et al.* 2012).

However, there seem to be limits to the extent to which we can unconsciously process the semantic content of written words. One recent neuroimaging study (Nakamura *et al.* 2018), e.g., found unconscious semantic priming effects only if primes and targets were separated by no more than two words in a sequence. Other masked priming studies suggest that the extent of unconscious priming effects with unpracticed word pairs depend significantly on their associative strength and semantic similarity (e.g. Van den Bussche *et al.* 2012; Ortells *et al.* 2013). The point here is that although the visual system can extract meaning from written language in the absence of awareness, these capacities do in fact max out, either, e.g., if the stimuli are too informationally complex (e.g. there are too many represented elements to process) or if the conceptual relations among words are not salient enough from previous learning history.

Probing semantic processing in conscious and unconscious vision using images has produced a similar pattern of results, providing even more evidence of the functional resources that consciousness contributes to visual processing. A variety of independent research programs have revealed that processing certain semantic relationships in visual imagery requires the subject to have conscious experiences. By leveraging combinations of suppression techniques like CFS and behavioural measures like semantic priming, evidence is mounting that tasks that require the perceptual discrimination of images based on the basic-level category (Rosch 1978) that they belong to (e.g. snake or spider, tool or animal) cannot be accomplished when those images are unconsciously processed by the visual system (Hesselmann *et al.* 2016; Cox *et al.* 2018; Stein *et al.* 2020). Koivisto and Rientamo's (2016) study explicitly probed the kinds of semantic categorization that occur in the presence and absence of consciousness (see Fig. 3). They found unconscious priming in superordinate categorization tasks only (e.g. animal vs. non-animal), whereas no unconscious priming effect was observed when categorization tasks relied on basic-level categories (e.g. horse vs. non-horse). They take this as evidence that unconscious representations in the visual system are much coarser than conscious ones, thus limiting the kinds of discriminations they can support. Interestingly, unconscious priming effects with visual imagery are typically still observed in these studies when subsequent discrimination tasks rely on low-level information like shape rather than conceptual information, which is consistent with other research showing low-level feature-driven facilitation of performance in the absence of awareness (e.g. Koivisto and Grassini 2018).

In summary, there appear to be robust unconscious semantic priming effects when primes are written words that remain under a certain threshold of representational complexity, but only limited unconscious priming effects when primes are rich visual images and when tasks go beyond mere discrimination by superordinate category or shape. These sorts of empirical results need to be carefully interpreted in order to extract their underlying significance for philosophical theories of consciousness. How do we best describe the functional contributions that consciousness is making to visual processing of semantic information?

It is not a simple task to characterize the relevant functional roles here. Semantic processing is complex and plausibly involves a range of interacting computational capacities. The visual system's role, however, can be isolated in principle and understood independently. Briefly, semantically laden visual stimuli activate linguistic processing mechanisms (presumably supported in part by structures in the anterior temporal lobe region of the ventral visual stream) as a result of learned associations between those stimuli and some set of abstracted referential content. In the case of written language, associations between words and their referential content become deeply entrenched throughout development. Crucially, written language relies on extremely simple visual stimuli: a few straight lines appropriately arranged is all that is needed to convey semantically relevant information, once those semantic relationships are learned. If unconscious processes are limited in their representational capacities, i.e., they are only capable of trafficking in coarser representations with limited informational detail, then it makes sense that semantic processes can be activated unconsciously by words given the simplicity both of the stimulus itself and the internal processes required to represent it.

This in turn suggests that semantic tasks that rely on more informationally detailed representations require conscious experiences. Koivisto and Rientamo (2016) assume as much, arguing that discrimination tasks that require more fine-grained representational capacities depend on conscious experiences. There is a certain level of representational complexity needed to capture the appropriate stimulus features of a visual image in order to carry out certain semantically driven tasks. Here, complexity simply refers to the number of informational elements that comprise a given representation or the quantity of distinct qualitative dimensions of a stimulus that are captured in a representation. Whereas representations of words track simple features like line orientation, processing visual imagery requires the combination of a variety of potentially task-relevant visual features like colour, overall shape including depth, and motion (especially if the semantic task requires categorization in terms of animacy). In this sense, the emerging picture is that semantic processing tasks that require representations of stimulus features at a certain level of informational complexity can only be accomplished when subjects are consciously aware of that stimulus. That is, consciousness appears to contribute increased functional resources to the visual processing of semantically relevant stimuli by facilitating more informationally complex representations that are required for extracting meaning from more informationally complex written language and visual scenes.

Semantic processing by the visual system is also a good place to start to address the individual differences observed in CFS research and related experimental paradigms. While this remains speculative until future research addresses these questions more directly, it is plausible that individual differences in performance on conscious and unconscious semantic processing tasks in vision are the result of individual differences in the ways that individuals encode meaning. Differences in individual learning histories, e.g., will likely affect the stability of certain semantic relationships. In other words, more contact with certain semantic relations in the world (e.g. the meanings of words in a particular semantic domain) might increase the salience of these relations in subsequent episodes of visual processing. Much like how a skilled musician or athlete will eventually relegate much of their perceptual-motor processing to unconscious mechanisms, so too familiarity with certain semantic domains is likely to change the way that information is consolidated into memory systems and ultimately how it is drawn upon at later times (e.g. Lupyan *et al.* 2020). At the very least, there is a testable hypothesis in this vicinity that unconscious processing of semantic relationships between elements of a visual image might be more robust as a result of greater familiarity. Future research might also look for patterns of difference between different age groups or different linguistic communities. This would also support the assumption that consciousness is required for processing more novel stimuli, which might ultimately bottom out in representational complexity, if novel stimuli require richer representational resources. Despite making it harder to draw a clean line between conscious and unconscious functions across individual humans, let alone species, these individual differences add support to the pluralist claim that conscious experience occupies different functional roles in different systems based on individual processing requirements.

Spatial and temporal processing in vision

Another central set of functions carried out by the visual system involves processing information about space and time. This is especially relevant to visually guided action, where detailed spatial and temporal maps are constructed in order to plan, predict, and guide even the most elementary bodily movements (e.g. reaching for and grasping objects, see Crawford *et al.* 2011). Representing spatial and temporal information in vision also sometimes requires integrating spatially or temporally disparate elements in order to carry out a cognitive or behavioural task effectively (e.g. tracking motion). Once again, spatiotemporal visual processing is fairly common in CFS and related research and thus provides relatively stable grounds for comparison between the functions of conscious and unconscious visual processing.

Even though it is doubtful that the dorsal visual stream operates entirely unconsciously as was once assumed (e.g. see Wu 2020), research into visuomotor transformation does provide compelling evidence that much spatial and temporal information processed in vision remains unconscious and yet continues to

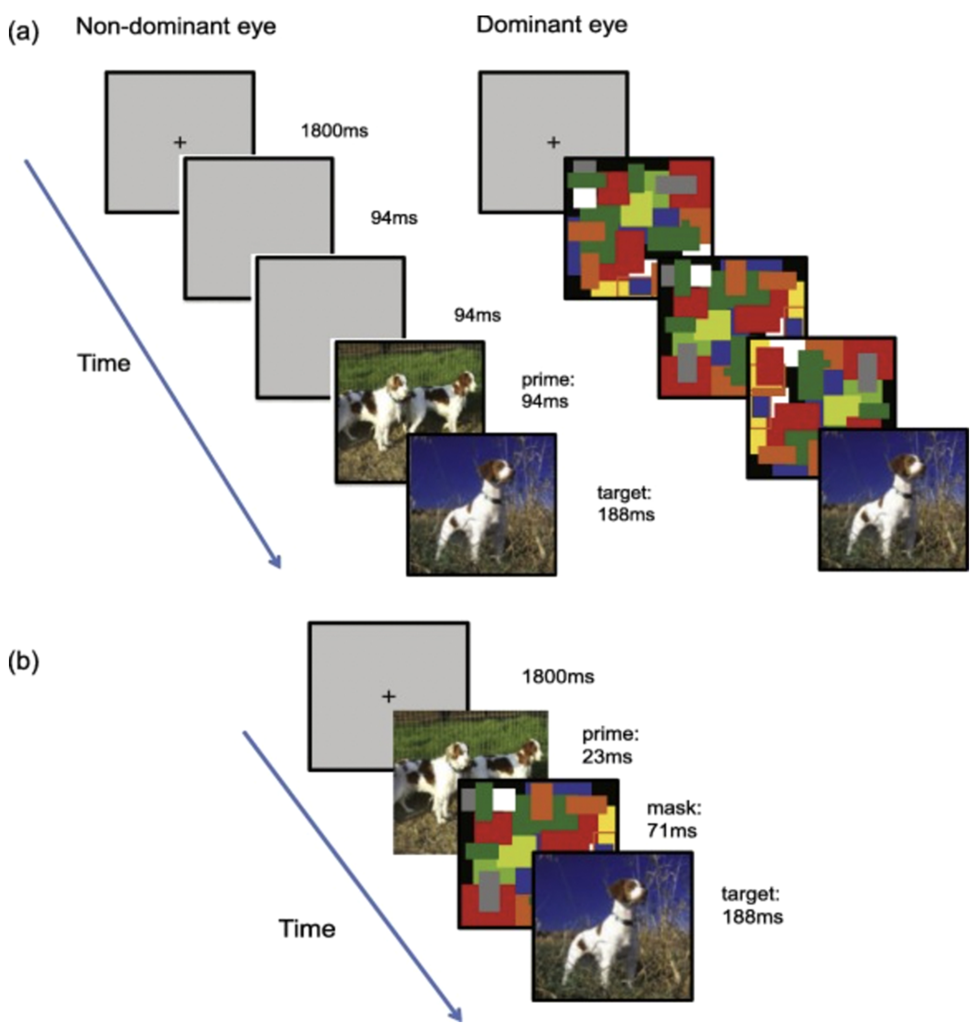

Figure 3. Schematic representation of the CFS/semantic priming paradigm. From Koivisto and Rientamo's (2016)

guide behaviour in sophisticated ways. The visual system creates, maintains, and updates representations of the positions of our bodies and other objects in 'egocentric' space, as well as representations of the absolute size of objects and relations in 'allocentric' space, that do not always enter into conscious experience (see Brogaard 2011 for review). These are impressive feats of unconscious visual processing. In order to isolate what it is that consciousness contributes functionally to the visual processing of spatial and temporal properties, we need to compare performance on tasks where the key variables can be manipulated along their relevant dimensions.

Several specific studies have, e.g., shed light on consciousness' role in enhancing spatial and temporal resolution in vision. Koivisto *et al.* (2014), employ a go/no go animal/non-animal categorization task paired with EEG recording in order to investigate the categorization capacities of conscious versus unconscious vision. Their results represent compelling evidence that rapid categorization can occur when masking disrupts recurrent processing and stimuli are not consciously perceived. And, yet, when stimuli are unmasked and recurrent processing is established and maintained, (i) the 'clarity' or grain of visual categorization increases, (ii) categorization for unclear images gets faster and more accurate, and (iii) a greater electrophysiological difference is observed when categorizing animal versus non-animal stimuli. These functional advantages are explicitly construed by the authors as contributions that conscious awareness—facilitated by recurrent processing coalitions in the visual system—makes to the processing of spatial resolution. Along the same lines, Diano *et al.* (2017) argue that subcortical, feed-forward (and presumably unconscious) processing of visual stimuli results in low-frequency representations that trade detail for rapid categorization based on global properties. This contrasts with recurrent processing networks in the cortex that generate high-spatial-frequency representations in order to extract fine details from a visual scene for finer grained categorical distinctions. There seems to be growing consensus that unconscious processes simply traffic in coarser and therefore less spatially and temporally detailed representations of visual stimuli.

Another variable that can be manipulated in this regard is the extent to which spatial and temporal integration is needed in order to carry out specific visual functions. Faivre and Koch (2014), e.g., used CFS to compare performance on a task that requires the integration of motion-relevant visual information over increasing temporal periods. In one experimental set-up, they used 'apparent dot motion' stimuli, where multiple dots flashing in succession across a screen are reliably perceived as a single dot in motion. They employ an adaptation paradigm on both conscious and unconscious trials in order to assess and compare the extent of visual processing in each case. In both conscious and unconscious experimental conditions, they found adaptation effects in response to the presentation of the stimuli, suggesting paradigmatic perceptual capacities across the conditions. Crucially though, these adaptation effects only occur on unconscious trials when the 'temporal integration windows'—or the temporal 'distance' between the individual represented elements that need to be integrated—were sufficiently short (i.e. when the successive dot flashes were 100 ms apart). In contrast, integration over longer temporal windows (i.e. 400 ms, 800 ms, 1200 ms) only occurred when subjects were consciously aware of the stimuli. This suggests that although some spatial and temporal integration of motion information is possible in the absence of awareness, integrating over larger temporal distances seems to depend on conscious experiences.

Mudrik *et al.* (2014) review various such attempts to discover the capacities and limitations of unconscious spatiotemporal integration, which they take to be important for isolating the kinds of spatiotemporal processing tasks that require consciousness. Despite evidence that unconscious visual processing occurs on relatively spatiotemporally complex inputs (e.g. facial identity and natural scenes), the authors suggest that at a certain threshold, visual elements distributed across space and time likely cannot be processed together by the visual system without conscious awareness. One possible explanation they offer for this functional difference is that this threshold in unconscious processing represents a limit on ensemble encoding, understood as a subject's ability to extract summary statistics from arrays of simultaneously presented visual stimuli. The assumption is that there is an important relationship between this psychological capacity and conscious experience, and ultimately that conscious processing is indeed required for integration over a certain threshold of complexity.

Finally, some familiar visual illusions depend on the influence of particular contextual elements that are spatially or temporally distributed. This means that another way to manipulate key variables here is to vary the amount, distribution, and kind of contextual visual information that is drawn upon by the visual system in processing certain illusory stimuli. To this end, Harris *et al.* (2011) used CFS to selectively suppress contextual information from awareness while presenting subjects with different visual illusions. They found that simultaneous brightness illusions, in which identical stimuli look differently shaded due to differences in surrounding luminance, persist in the absence of awareness. In contrast, they found that Kanizsa-style contour illusions, in which the visual system represents illusory surface and edge information because of cues provided by shapes in the surrounding context, did not persist when *all* the surrounding cues—and not simply *half* of the cues as in the case of hemineglect patients discussed previously—were suppressed from awareness (see Fig. 4).

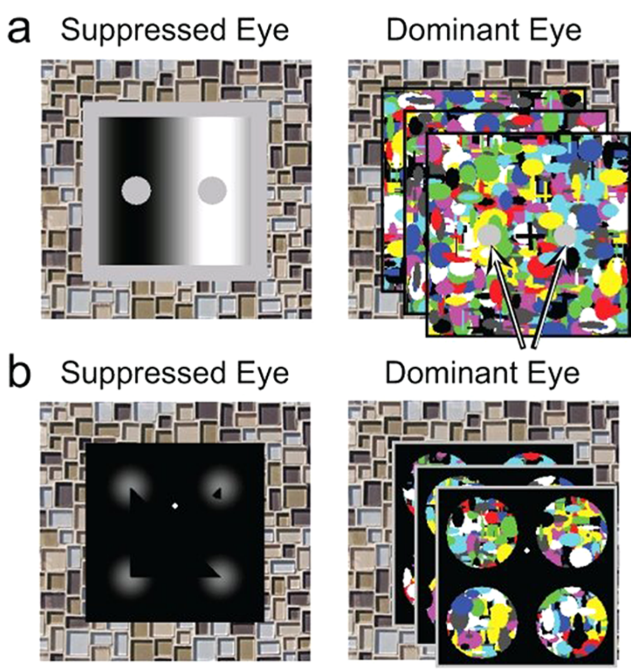

Figure 4. Schematic representation of the CFS/visual illusion paradigm. (a) Simultaneous brightness illusion under CFS. (b) Kanizsa-style illusion under CFS. From Harris *et al.* (2011)

A plausible interpretation of this observed functional limitation of unconscious visual processing is that simultaneous brightness requires the processing of less spatially distributed information (i.e. merely two points of comparison) than the Kanizsa-style illusions do (i.e. typically three or four spatially distributed shapes that provide contour cues). That the latter illusion persists when only half of the cues are unconsciously processed in hemineglect patients but does not persist when all of the cues fail to reach awareness due to suppression techniques gives us an even more precise idea of the FCCs to these visual processes. More recently, Chen *et al.* (2018) similarly found that the Ebbinghaus illusion persists even when the surrounding context is selectively suppressed from awareness using CFS, whereas the Ponzo illusion does not. Again, the interpretation given here is that Ebbinghaus illusions are the result of lower-level processes like contour interactions that require less spatially distributed information, whereas Ponzo illusions require holistic processing of much more spatially distributed contextual information. These all seems to support the hypothesis that one of consciousness' functions in vision is to facilitate more distributed, and hence in some ways more representationally demanding, spatial and temporal processing.

FCCs: visual perception

The goal here was to begin to isolate some of the functions that consciousness contributes to the domain of visual processing or at least formulate some testable hypothesis in this direction. After surveying the relevant empirical work and drawing out key theoretical implications, we are in a position to offer some preliminary remarks in this direction.

The most difficult aspect of the problem seems to be saying exactly *where* consciousness enters the visual processing hierarchy. Differences in both individual learning history and experimental design make it difficult, if not impossible at this point, to draw a hard line in the visual system with unconscious processing on one side and conscious processing on the other, at least in terms of their functional characteristics. But this result should be neither surprising nor troubling, given both the diversity and flexibility of the structures and functions that characterize human visual processing. In fact, exploring this shifting ontological boundary between unconscious and conscious vision will continue to strengthen our models of the FCCs. At this point, however, I think it is illuminating to acknowledge that the kinds of limited functional capacities that unconscious visual processes carry out are importantly related to the functions carried out by conscious visual processes; there is a sort of spectral continuity here. That is we typically find a relationship of degree between the functional capacities of unconscious versus conscious visual processing (e.g. spatiotemporal integration and resolution, semantic processing). According to the model emerging here, unconscious visual processing is in the business of constructing coarse representations that are useful within a certain limited range of informational complexity. This in turn implies that consciousness boosts processing capacities wherever a certain visual task requires informational richness that cannot be captured by coarse unconscious representations, which is clearly a multifaceted and highly context-dependent set of perceptual circumstances.

We are in a better position to answer *why* it is that consciousness enters the visual processing hierarchy; that is we can start to provide an account of the functional advantages conscious experience makes in the domain of vision. In particular, conscious processing appears to facilitate the following:

(i). Increased capacities for semantically interpreting visual stimuli (i.e. words and meaningful images)
(ii). Heightened spatiotemporal precision in visual representation
(iii). Integration of visual information over larger spatiotemporal windows (the latter of which also seems to enable integrating increasingly complex contextual information).

These candidate FCCs should be understood as working hypotheses awaiting more support from the psychological and neural sciences. The more comparative research that we can draw on in, the more confident we can be in our claims that these functions are sufficient for, or unique to, consciousness.

Note that these functions do not seem to be captured by the leading monolithic theories like GWT and IIT. GWT, e.g., which assumes that *the* function of consciousness is to access and broadcast information via the global workspace, is ill equipped to explain exactly how consciousness contributes functionally to aspects of visual perception that are not accessed for global broadcasting. Much of the conscious spatiotemporal visual integration, for instance, that contributes to the complexity of a visual scene remains unaccessed by cognitive or attentional mechanisms (Block 2014). In fact, it seems that none of the functions specified above are significantly associated with the activity of the global workspace; that is global broadcasting is neither necessary nor sufficient for realizing these psychological functions. The same is true of IIT. A candidate FCC like spatiotemporal precision, e.g., is not necessarily related to the psychological function of informational integration but rather represents other functions with other dimensions of complexity that are associated with consciousness. Again, the mere integration of information, as formulated by IIT, is neither necessary nor sufficient for carrying out these functions of conscious visual processing. Both theories employ graded psychological constructs, namely integration and access, that we have good reason to think occur unconsciously at least to some theoretically significant degree. In this way, integration and access, understood as monolithic constructs, are psychological functions that conceptually dissociate from consciousness, at least until we can discover the different qualifying conditions that might link them to different conscious processes (e.g. perhaps integration at reliable spatiotemporal thresholds). It will likely turn out that at least some kinds or degrees of access and integration are sufficient for consciousness and therefore can act as viable markers of experience, among the many FCCs we might amass by pursuing the pluralist project. Future research ought to explore exactly what these functions are.

Thus, the functional pluralism hypothesis seems to be corroborated even by an investigation into the domain of visual processing. On one hand, the candidate functions outlined above, while likely sufficient for conscious experience, are likely not necessary. That is while these functions require consciousness, many conscious experiences are not marked by these domain-specific functions. The way that experience enhances visual processing is likely not to be the same as the way it contributes to emotional processing, or to the processing of social bias, e.g., where the same dimensions of representational complexity are not applicable. The pluralist intuition is strengthened when one considers the vast range of different kinds of specialized perceptual and cognitive mechanisms that emerge across the animal kingdom. On the other hand, there is already enough diversity among the FCCs identified here to assume that we are dealing with importantly distinct functions. Some of these functions associated with consciousness might be understood as falling under the broader functional categories of global access (e.g. semantic capacities) and integration (e.g. apparent dot motion integration), while others cannot (e.g. spatiotemporal precision). This further exhibits the functional diversity underlying the pluralist framework.

Functional pluralism, therefore, points to rich avenues for future inquiry. Researchers ought to shift their focus to domain-specific analyses of the functions of consciousness by performing similar comparisons between specific functional tasks carried out in other psychological domains. There is also further need to taxonomize the different functions that emerge from such an analysis under scientifically useful categories. The end result of this inquiry should allow us to compile a range of different local, domain-specific, non-dissociable markers that can guide ongoing research in the philosophy and science of consciousness.

Conclusion

Several lines of research support the idea that although there are some fairly sophisticated unconscious visual processes, they are importantly limited in functional capacity. A careful analysis of the functional differences when awareness alone is manipulated in visual processing tasks reveals a cluster of capacities that consciousness likely facilitates. These include increased capacities for semantically interpreting visual stimuli, heightened spatiotemporal precision in visual representation, and integration of visual information over larger spatiotemporal windows. I urge that this sort of domain-specific functional analysis should precede overarching theories of the nature of consciousness. Moreover, the functions that consciousness contributes to vision are likely to be different from the functions it contributes to other psychological processes across the natural world, and perhaps in artificial systems. Further domain-specific analyses are required in order to determine whether 'consciousness' remains as a unified construct or whether functional and structural differences prompt conceptual distinctions among phenomena previously unified under this conceptual category.

Conflict of interest statement

None declared.

References

Alexander I, Cowey A. Edges, colour and awareness in blindsight. *Conscious Cogn* 2010;**19**:520–33.

Altman DG, Bland JM. Absence of evidence is not evidence of absence. *BMJ* 1995;**311**:485.

Armstrong A-M, Dienes Z. Subliminal understanding of negation: unconscious control by subliminal processing of word pairs. *Conscious Cogn* 2013;**22**:1022–40.

Axelrod V, Bar M, Rees G *et al.* Neural correlates of subliminal language processing. *Cereb Cortex* 2015;**25**:2160–9.

Bargh JA, Morsella E. The unconscious mind. *Perspect Psychol Sci* 2008;**3**:73–9.

Berti A, Rizzolatti G. Visual processing without awareness: evidence from unilateral neglect. *J Cogn Neurosci* 1992;**4**:345–51.

Birch J. The search for invertebrate consciousness. *Noûs* 2020;**56**:133–53.

Bisiach E, Luzzatti C. Unilateral neglect of representational space. *Cortex* 1978;**14**:129–33.

Block N. Perceptual consciousness overflows cognitive access. *Trends Cogn Sci* 2011;**15**:567–75.

Block N. Rich conscious perception outside focal attention. *Trends Cogn Sci* 2014;**18**:445–7.

Block N. What is wrong with the no-report paradigm and how to fix it. *Trends Cogn Sci* 2019;**23**:1003–13.

Block N, Phillips I. Debate on unconscious perception. In: Nanay B (ed.), *Current Controversies in Philosophy of Perception*. New York: Routledge, Taylor & Francis Group, 2016, 165–92.

Boly M, Massimini M, Tsuchiya N *et al.* Are the neural correlates of consciousness in the front or in the back of the cerebral cortex? Clinical and neuroimaging evidence. *J Neurosci* 2017;**37**:9603–13.

Boyle A. The impure phenomenology of episodic memory. *Mind Lang* 2019;**35**:1–20.

Brascamp J, Blake R, Knapen T . Negligible fronto-parietal BOLD activity accompanying unreportable switches in bistable perception. *Nat Neurosci* 2015;**18**:1672–8.

Breitmeyer B. Functional hierarchy of unconscious object processing. In: Breitmeyer B (ed.), *The Visual (Un)conscious and Its (Dis)contents*. Oxford: Oxford University Press, 2014, 89–102.

Breitmeyer BG. Psychophysical "blinding" methods reveal a functional hierarchy of unconscious visual processing. *Conscious Cognit* 2015;**35**:234–50.

Breitmeyer BG, Öğmen H. *Oxford Psychology Series. Visual Masking: Time Slices through Conscious and Unconscious Vision*, 2nd edn. New York: Oxford University Press, 2006.

Brogaard B. Are there unconscious perceptual processes? *Conscious Cogn* 2011;**20**:449–63.

Brogaard B, Chomanski B, Gatzia DE. Consciousness and information integration. *Synthese* 2020;**198**:763–92.

Burge T. *Origins of Objectivity*. New York: Oxford University Press, 2010.

Bussche EV, Segers G, Reynvoet B. Conscious and unconscious proportion effects in masked priming. *Conscious Cogn* 2008;**17**:1345–58.

Chen L, Qiao C, Wang Y *et al.* Subconscious processing reveals dissociable contextual modulations of visual size perception. *Cognition* 2018;**180**:259–67.

Cohen MA, Dennett DC. Consciousness cannot be separated from function. *Trends Cogn Sci* 2011;**15**:358–64.

Costello P, Jiang Y, Baartman B *et al.* Semantic and subword priming during binocular suppression. *Conscious Cogn* 2009;**18**:375–82.

Cox EJ, Sperandio I, Laycock R *et al.* Conscious awareness is required for the perceptual discrimination of the threatening animal stimuli: a visual masking and continuous flash suppression study. *Conscious Cogn* 2018;**65**:280–92.

Crawford JD, Henriques DYP, Medendorp WP. Three-dimensional transformations for goal-directed action. *Annu Rev Neurosci* 2011;**34**:309–31.

Crepaldi D, Rastle K, Coltheart M *et al.* "Fell" primes "fall", but does "bell" prime "ball"? Masked priming with irregularly-inflected primes. *J Mem Lang* 2010;**63**:83–99.

Dehaene S. *Consciousness and the Brain: Deciphering How the Brain Codes Our Thoughts*. New York: Viking, 2014.

Dehaene S, Changeux JP. Experimental and theoretical approaches to conscious processing. *Neuron* 2011;**70**:200–27.

Dehaene S, Changeux J-P, Naccache L *et al.* Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends Cogn Sci* 2006;**10**:204–11.

Dehaene S, Kerszberg M, Changeux J. A neuronal model of a global workspace in effortful cognitive tasks. *Proc Natl Acad Sci U S A* 1998;**95**:14529–34.

Diano M, Celeghin A, Bagnis A *et al.* Amygdala response to emotional stimuli without awareness: facts and interpretations. *Front Psychol* 2017;**7**:2029.

Elgendi M, Kumar P, Barbic S *et al.* Subliminal priming—state of the art and future perspectives. *Behav Sci (Basel)* 2018;**8**:54.

Fahrenfort JJ, Scholte HS, Lamme VAF. Masking disrupts reentrant processing in human visual cortex. *J Cogn Neurosci* 2007;**19**:1488–97.

Faivre N, Koch C. Temporal structure coding with and without awareness. *Cognition* 2014;**131**:404–14.

Fang Z, Han L, Chen G *et al.* Unconscious processing of negative animals and objects: role of the amygdala revealed by fMRI. *Front Hum Neurosci* 2016;**10**:146.

Frässle S, Sommer J, Jansen A *et al.* Binocular rivalry: frontal activity relates to introspection and action but not to perception. *J Neurosci* 2014;**34**:1738–47.

Frith CD, Metzinger T. What's the use of consciousness? How the stab of conscience made us really conscious. In: Engel AK, Friston KJ, Kragic D (eds), *The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science*. Strüngmann Forum Reports, vol. 18, J. Lupp, series editor. Cambridge, MA: MIT Press, 2016, 193–214.

Gelbard-Sagiv H, Faivre N, Mudrik L *et al.* Low-level awareness accompanies "unconscious" high-level processing during continuous flash suppression. *J Vis* 2016;**16**:3, 1–16.

Goodale M, Milner D. *Sight Unseen: An Exploration of Conscious and Unconscious Vision*, 2nd edn. Oxford: Oxford University Press, 2013.

Halligan PW, Marshall JC. Neglect of awareness. *Conscious Cogn* 1998;**7**:356–80.

Harris JJ, Schwarzkopf DS, Song C *et al.* Contextual illusions reveal the limit of unconscious visual processing. *Psychol Sci* 2011;**22**:399–405.

Hassin RR. Yes it can: on the functional abilities of the human unconscious. *Perspect Psychol Sci* 2013;**8**:195–207.

Hesselmann G, Darcy N, Ludwig K *et al.* Priming in a shape task but not in a category task under continuous flash suppression. *J Vis* 2016;**16**:17.

Hilgetag CC, Medalla M, Beul SF *et al.* The primate connectome in context: principles of connections of the cortical visual system. *Neuroimage* 2016;**134**:685–702.

Hubel DH, Wiesel TN. Receptive fields and functional architecture of monkey striate cortex. *J Physiol* 1968;**195**:215–43.

Jiang Y, Costello P, He S. Processing of invisible stimuli: advantage of upright faces and recognizable words in overcoming interocular suppression. *Psychol Sci* 2007;**18**:349–55.

Koch C, Massimini M, Boly M *et al.* Neural correlates of consciousness: progress and problems. *Nat Rev Neurosci* 2016;**17**:307–21.

Koivisto M, Grassini S. Unconscious response priming during continuous flash suppression. *PLoS One* 2018;**13**:e0192201.

Koivisto M, Kastrati G, Revonsuo A. Recurrent processing enhances visual awareness but is not necessary for fast categorization of natural scenes. *J Cogn Neurosci* 2014;**26**:223–31.

Koivisto M, Rientamo E. Unconscious vision spots the animal but not the dog: masked priming of natural scenes. *Conscious Cognit* 2016;**41**:10–23.

Kominsky JF, Scholl BJ. Retinotopic adaptation reveals distinct categories of causal perception. *Cognition* 2020;**203**:104339.

Korisky U, Hirschhorn R, Mudrik L. "Real-life" continuous flash suppression (CFS)-CFS with real-world objects using augmented reality goggles. *Behav Res* 2019;**51**:2827–39.

Kouider S, Dehaene S. Levels of processing during non-conscious perception: a critical review of visual masking. *Philos Trans R Soc Lond B Biol Sci* 2007;**362**:857–75.

Kouider S, Faivre N. Conscious and unconscious perception. In: Schneider S, Velmans M (eds), *The Blackwell Companion to Consciousness*, 2nd edn. Hoboken, NJ: John Wiley & Sons, 2017, 551–61.

Kovács IP, Papathomas TV, Yang M *et al.* When the brain changes its mind: interocular grouping during binocular rivalry. *Proc Natl Acad Sci U S A* 1996;**93**:15508–11.

Lamme V. Visual functions generating conscious seeing. *Front Psychol* 2020;**11**:83.

Li K, Malhotra PA. Spatial neglect. *Pract Neurol* 2015;**15**:333–9.

Locke J. *An Essay Concerning Human Understanding*. London: Oxford University Press, 1689.

Ludwig D. The functional contributions of consciousness. *Conscious Cogn* 2022;**104**:103383.

Lupyan G, Rahman RA, Boroditsky L *et al.* Effects of language on visual perception. *Trends Cogn Sci* 2020;**24**:930–44.

Malach R. Local neuronal relational structures underlying the contents of human conscious experience. *Neurosci Conscious* 2021;**2021**:niab028.

Maniscalco B, Lau H. A signal detection theoretic approach for estimating metacognitive sensitivity from confidence ratings. *Conscious Cogn* 2012;**21**:422–30.

McDowell J. The content of perceptual experience. *Philos Q* 1994;**44**:190–205.

Mendez-Bertolo C, Moratti S, Toledano R *et al.* A fast pathway for fear in human amygdala. *Nat Neurosci* 2016;**19**:1041–9.

Morales J, Chiang J, Lau H. Controlling for performance capacity confounds in neuroimaging studies of conscious awareness. *Neurosci Conscious* 2015;**2015**:niv008.

Morris JS, Ohman A, Dolan RJ. Conscious and unconscious emotional learning in the human amygdala. *Nature* 1998;**393- 4**:467–70.

Mudrik L, Breska A, Lamy D *et al.* Integration without awareness: expanding the limits of unconscious processing. *Psychol Sci* 2011;**22**:764–70.

Mudrik L, Faivre N, Koch C. Information integration without awareness. *Trends Cogn Sci* 2014;**18**:488–96.

Nakamura K, Dehaene S, Jobert A *et al.* Task-specific change of unconscious neural priming in the cerebral language network. *Proc Natl Acad Sci U S A* 2007;**104**:19643–8.

Nakamura K, Makuuchi M, Oga T *et al.* Neural capacity limits during unconscious semantic processing. *Eur J Neurosci* 2018;**47**:929–37.

Nakamura K, Oga T, Fukuyama H. Task-sensitivity of unconscious word processing in spatial neglect. *Neuropsychologia* 2012;**50**:1570–7.

Ogden JA editor. *Fractured Minds: A Case Study Approach to Clinical Neuroscience*. New York: Oxford University Press, 2005.

Oriet C, Brand J. Size averaging of irrelevant stimuli cannot be prevented. *Vision Res* 2012;**79**:8–16.

Ortells JJ, Marí-Beffa P, Plaza-Ayllón V. Unconscious congruency priming from unpracticed words is modulated by prime–target semantic relatedness. *J Exp Psychol Learn Mem Cogn* 2013;**39**:394–413.

Persaud N, Davidson M, Maniscalco B *et al.* Awareness-related activity in prefrontal and parietal cortices in blindsight reflects more than superior visual performance. *NeuroImage* 2011;**58**:605–11.

Persaud N, McLeod P, Cowey A. Post-decision wagering objectively measures awareness. *Nat Neurosci* 2007;**10**:257–61.

Peters MAK, Kentridge RW, Phillips I *et al.* Does unconscious perception really exist? Continuing the ASSC20 debate. *Neurosci Conscious* 2017;**2017**:1–11.

Phillips I. Unconscious perception reconsidered. *Anal Philos* 2018;**4**:471–514.

Plass J, Guzman-Martinez E, Ortega L *et al.* Lip reading without awareness. *Psychol Sci* 2014;**25**:1835–7.

Prinz J. The sensory basis of cognitive phenomenology. In: Bayne T, Montague M (eds), *Cognitive Phenomenology*. New York: Oxford University Press, 2011, 174–96.

Quilty-Dunn J. Unconscious perception and phenomenal coherence. *Analysis* 2019;**79**:461–9.

Reber TP, Henke K. Integrating unseen events over time. *Conscious Cogn* 2012;**21**:953–60.

Rosch E. Principles of categorization. In: B.E. Rosch, Lloyd BB (eds.), *Cognition and Categorization*. Hillsdale: Erlbaum, 1978, 28–49.

Russell B. *The Problems of Philosophy*. London: Williams and Norgate. New York: Henry Holt and Company, 1912.

Sklar AY, Levy N, Goldstein A, Levy N Goldstein A, Goldstein A *et al.* Reading and doing arithmetic nonconsciously. *Proc Natl Acad Sci U S A* 2012;**109**:19614–9.

Sperling G. The information available in brief visual presentations. *Psychol Monogr* 1960;**74**:1–29.

Sprenger A, Kömpf D, Heide W. Visual search in patients with left visual hemineglect. *Prog Brain Res* 2002;**140**:395–416.

Stein T, Peelen MV. Dissociating conscious and unconscious influences on visual detection effects. *Nat Hum Behav* 2021;**5**:612–24.

Stein T, Sterzer P. Unconscious processing under interocular suppression: getting the right measure. *Front Psychol* 2014;**5**:387.

Stein T, Utz V, van Opstal F. Unconscious semantic priming from pictures under backward masking and continuous flash suppression. *Conscious Cogn* 2020;**78**:102864.

Tamietto M, de Gelder B. Neural bases of the non-conscious perception of emotional signals. *Nat Rev Neurosci* 2010;**11**:697–709.

Tononi G, Boly M, Massimini M *et al.* Integrated information theory: from consciousness to its physical substrate. *Nat Rev Neurosci* 2016;**17**:450–61.

Tye M. *Consciousness, Color and Content*. Cambridge, MA: MIT Press, 2000.

Van den Bussche E, Smets K, Sasanguie D *et al.* The power of unconscious semantic processing: the effect of semantic relatedness between prime and target on subliminal priming. *Psychol Belg* 2012;**52**:59–70.

van Essen DC, Anderson CH, Felleman DJ. Information processing in the primate visual system: an integrated systems perspective. *Science* 1992;**255**:419–23.

Venezia JH, Matchin W, Hickok G. Multisensory integration and audiovisual speech perception. In: *Brain Mapping: An Encyclopedic Reference*, Vol. 2, 2015, 565–72.

Verdon V, Schwartz S, Lovblad K-O *et al.* Neuroanatomy of hemispatial neglect and its functional components: a study using voxel-based lesion-symptom mapping. *Brain: J Neurol* 2010;**133**:880–94.

Vuilleumier P, Landis T. Illusory contours and spatial neglect. *NeuroReport* 1998;**9**:2481–4.

Williams MA, Morris AP, McGlone F *et al.* Amygdala responses to fearful and happy facial expressions under conditions of binocular suppression. *J Neurosci* 2004;**24**:2898–904.

Wu W. Is vision for action unconscious? *J Philos* 2020;**117**:413–33.

Yang E, Brascamp J, Kang M-S *et al.* On the use of continuous flash suppression for the study of visual processing outside of awareness. *Front Psychol* 2014;**5**:724.

Yang YH, Yeh SL. Accessing the meaning of invisible words. *Conscious Cogn* 2010;**20**:223–33.

Yokoyama T, Noguchi Y, Kita S. Unconscious processing of direct gaze: evidence from an ERP study. *Neuropsychologia* 2013;**51**:1161–8.

Zabelina DL, Guzman-Martinez E, Ortega L *et al.* Suppressed semantic information accelerates analytic problem solving. *Psychon Bull Rev* 2013;**20**:581–5.