

Opinion piece



Cite this article: Lamme VAF. 2018

Challenges for theories of consciousness:
seeing or knowing, the missing ingredient and
how to deal with panpsychism. *Phil.*

Trans. R. Soc. B **373**: 20170344.

<http://dx.doi.org/10.1098/rstb.2017.0344>

Accepted: 21 May 2018

One contribution of 17 to a theme issue
'Perceptual consciousness and cognitive access'.

Subject Areas:

neuroscience, behaviour, cognition

Keywords:

consciousness, access, phenomenal,
recurrent processing, panpsychism

Author for correspondence:

Victor A. F. Lamme

e-mail: victorlamme@gmail.com

Challenges for theories of consciousness: seeing or knowing, the missing ingredient and how to deal with panpsychism

Victor A. F. Lamme

Amsterdam Brain and Cognition (ABC), Department of Psychology, University of Amsterdam, 1001 NK
Amsterdam, The Netherlands

VAF, 0000-0003-4013-3687

Significant progress has been made in the study of consciousness. Promising theories have been developed and a wealth of experimental data has been generated, both guiding us towards a better understanding of this complex phenomenon. However, new challenges have surfaced. Is visual consciousness about the seeing or the knowing that you see? Controversy about whether the conscious experience is better explained by theories that focus on phenomenal (P-consciousness) or cognitive aspects (A-consciousness) remains, and the debate seems to reach a stalemate. Can we ever resolve this? A further challenge is that many theories of consciousness seem to endorse high degrees of panpsychism—the notion that all beings or even lifeless objects have conscious experience. Should we accept this, or does it imply that these theories require further ingredients that would put a lower bound on beings or devices that have conscious experience? If so, what could these 'missing ingredients' be? These challenges are discussed, and potential solutions are offered.

This article is part of the theme issue 'Perceptual consciousness and cognitive access'.

1. Introduction

Since the 1990s, behavioural science, neuroscience, biology, computational science and philosophy have joined forces to unravel one of the greatest mysteries, that of consciousness. How is it possible that with 1.5 kg of mainly fat and protein enclosed in a bony shell we have conscious experiences? Considerable progress has been made since Koch *et al.* [1], but we are clearly not there yet. Several promising theories and ideas have emerged, which have been backed up with lots of experimental data. This has led to consensus on some aspects of how, when and where brains do or do not produce consciousness. But strong controversy remains, in particular on what aspects of consciousness are the explanandum. Does the real mystery of consciousness lie in the fact that we experience the world that surrounds us, or in the ability to reflect on it and cognitively manipulate what we perceive; is consciousness about seeing or about knowing what we see? In the first part of this paper I will review the current state of that debate, starting with a brief overview of what we know about visual processing, and then delving into the question about where, when and how the transition from unconscious to conscious visual processing occurs.

In the second part, I will address a more recent issue about how theories of consciousness deal with the problem of panpsychism: the notion that all animals, or even all living beings and possibly even non-living items possess some form of consciousness. Major theories of consciousness can be argued to endorse high degrees of panpsychism. Does that falsify these theories? Should we accept panpsychism? Or is there a third way out? I will provide arguments for the latter, by showing that human and fly consciousness both exist, while at the same time are so different that swatting a fly should not worry you too much. It does appear, however, that many theories of

consciousness are too 'simple'. They seem to be missing a key ingredient, and more neurobiological avenues may point us in the right direction.

Note that the paper will focus almost entirely on visual consciousness. That is not because other sensory modalities of more executive faculties have no relevance. It is just a consequence of my ignorance on these other matters.

2. The intelligent reflex arc

There is little controversy on how visual processing proceeds. With every saccade—and we make about three per second—the eye lands on a new scene and the image is processed by the retina. In about 50 ms, this information has reached primary visual cortex, and from there on is distributed along a large number of other visual areas. Within 100–200 ms (depending on species and brain size) the whole brain is updated about the new image in front of us. During this rapid *feedforward sweep*, many features have been extracted from the image [2]. Neurons have detected shapes, colours, motion or depth. Higher level cells have signalled the presence of faces or animals [3] or other complex constellations that we are strongly familiar with (e.g. letters and words in humans [4]). We even have a rough idea about what the image is about: scene gist is extracted within that time-frame as well [5]. You might say that almost everything there is to know about the image has been signalled by the powerful machinery of visual cortex and its many low- and high-level feature selective cells (figure 1).

It does not end there. The feedforward sweep continues to proceed, now feeding into motor and frontal cortices. Here, visual information is transformed into execution—what can you do with what you see? If a ball is looming towards you, you instinctively dodge or catch it. The feats of such feedforward *sensorimotor transfer* are often quite remarkable. The eyes can saccade towards a visual stimulus within 120 ms [6], manual reaction times can be as short as 180 ms. Tennis players can return a ball serviced at 130 mph (world record is 163 mph), implying a reaction time of 400 ms, which includes deciding between forehand or backhand, doing a backswing and hitting the right way. You may have noticed yourself miraculously catching a glass of red wine that unexpectedly fell from a table, directly on its way to cause a horrible stain on the white carpet beneath. Before you knew what happened, the glass was in your hands, saving the day.


The feedforward sweep is the brain at its best. It shows how it can function as a reflex arc, not very different from the simple sensorimotor transfer that occurs in the knee-jerk reflex. Infinitely more complex and intelligent, of course, but a reflex arc nonetheless, where sensory input is as swiftly as possible translated into motor output. Note that this even involves some forms of decision-making: which of the many objects in a natural scene do we react to? How do we respond? In such reflex-like responses this strongly depends on the saliency or innate value of the external stimuli [6], but also internal variables (e.g. danger [7]) are taken into account. The feedforward sweep, therefore, is not a rigid process. What pathways are followed, and how deep each of these is penetrated depends on temporal and focused attention, task sets, expectation, brain state and many other factors [8–11].

3. Where does the magic happen?

From a behavioural or neuroscience point of view, this is all fairly straightforward. Of course some mysteries remain, among which the most prominent are how the brain detects objects for which there are no dedicated neurons, how it is possible to extract scene gist as fast as we can (the 'seeing the forest before the trees' problem) and how we select among the many motor plans that are activated by the multitude of stimuli in our surroundings. But when we take into consideration how we *experience* our brains doing this, an entirely different matter arises, that of consciousness. We are not mere automatons executing reflexes 'in the dark'. We see the ball that we hit during a tennis match, we feel the urgency when we grasp that glass. Somewhere along the way, a conscious experience of what we see arises. Where, when and how does that 'magic' happen? And why?

The short answer to the when and how is: after the fact. Feedforward processing itself is unconscious. The arguments for that have been laid out in many past reviews [12–15], showing considerable consensus on the matter. In short, when visual processing is artificially restricted to the feedforward sweep, it appears that features are extracted and that potential motor responses are activated yet that people remain completely unconscious of these events or of the visual stimulus that has set them in motion. The most widely applied method for doing so has been backward masking—i.e. showing a visual stimulus very briefly and then have it followed by another stimulus, the mask. This typically renders the visual stimulus completely invisible [16]. However, information about it travels pretty much throughout the brain, activating visual [17], motor [18] and even prefrontal areas [19,20]. The feedforward sweep, as intelligent as it may be, apparently does not suffice to give you a conscious sensation of the image that ignited it. So what is generating the conscious sensation that we automatically have when something hits our eyes?

What typically follows feedforward processing is recurrent or re-entrant processing. Via horizontal and feedback connections, neurons that initially responded to very different parts of the scene, or that had extracted different types of lower- or higher-level features, start to influence each other's activity patterns (figure 1) [2,21]. Neurons in the primary visual cortex (V1) typically have small receptive fields that are selective for particular features—say the orientation of line segments within a scene. Because of their small receptive fields, the neurons will respond identically, regardless of these line segments belonging to either a background or a figure. That is, up to about 100 ms. After that, the neurons' responses start to diverge for figure and background, as if they suddenly start to care about the larger perceptual context of the line segments they are responding to [22]. Quite strikingly for neurons in as 'low' a visual area as V1, this 'contextual modulation' of neural responses typically follows the perceptual interpretation that subjects have of the scene [23–25]. That is possible because it is mediated by horizontal connections and feedback connections from higher level areas [26], as lesions to these areas abolish the modulation [21,27]. Modelling has shown how neurons first do their independent job of detecting features and objects in a hierarchical cascade of feedforward processing, and then start to influence each other's activity via intricate inhibitory and excitatory interactions mediated by horizontal and feedback connections



Ample evidence emphasizes the importance of these recurrent interactions in conscious experience. They are abolished in anaesthesia, while feedforward processing is not [40,41]. Masking likewise selectively disrupts recurrent processing [17,42,43]. Dichoptic masking can render visual stimuli invisible for prolonged periods of time, which does not affect selective processing of faces in the fusiform cortex (FFA) yet disrupts recurrent interactions between the FFA and early visual cortex, as shown with psychophysiological interaction (PPI) analysis of fMRI [44]. Transcranial magnetic stimulation (TMS) can target visual areas with high temporal precision, which has shown that conscious percepts typically depend on early activity in high-level areas and on *later* activity in lower level ones, indicating that information has

to be sent back along the visual hierarchy to become conscious [45]. Various manipulations of conscious perception consistently show to influence EEG activity in the 200 ms latency range—known as the visual awareness negativity (VAN)—which is the typical latency of recurrent visual processing [46]. In monkeys trained to report the conscious perception of texture defined figures, contextual modulation of V1 neurons was selectively absent, while feedforward signals were unaffected [47]. In sum, recurrent interactions between visual neurons expressed as contextual modulation at the single neuron level or as 200 ms latency EEG signals (VAN) or PPI in humans, is crucial for the emergence of conscious experience.

Also, the disruption of longer latency recurrent interactions causes a failure to report conscious percepts. For example, when people fail to attend to stimuli, they will report not having seen them, as is the case in conditions like change blindness, inattention blindness or the attentional blink. These manipulations typically leave early recurrent signals relatively intact, while selectively disrupting later recurrent signals (e.g. P300 or N400 EEG responses [48] [49], but see [50]).

While the importance of recurrent processing for visual consciousness is widely accepted, the main controversy lies in the question of what extent of recurrent processing is sufficient for conscious experience to arise. The two main positions and their respective arguments (briefly) are

1. *Only when recurrent processing includes the fronto-parietal network do we experience the visual input.* At the basis of this line of reasoning is that for sensory information to become conscious it must become available to what is called the 'global (neural) workspace' (global ignition in figure 1) [15,51,52]. This global availability enables the cognitive manipulation of that information (e.g. storage in working memory), conscious access to the information, and eventually also the ability to report about it [53]. Hence, when subjects report not to have seen visual (or other modality) stimuli—as happens in cases like change blindness, attentional blink, neglect, and so on—this is taken at face value: it is then considered unconscious [54]. The neural correlates of such global availability are the involvement of fronto-parietal activation, and the later onset (P300) recurrent interactions [15]. Closely linked are theories that consider consciousness a higher-order thought (HOT) [55], implying that a higher-order representation of the visual experience must be present before it becomes conscious. A specific version of this idea is that consciousness requires (or is somehow linked to) metacognition, which can be briefly summarized as the 'knowing that you see' or 'the knowing that you know' [56,57].
2. *Recurrent processing between visual areas suffices for conscious visual experience to arise.* The hallmark of conscious vision, it is argued, is the integration of visual features into a coherent single scene. Once this is achieved, all the necessary and sufficient requirements for conscious vision have been fulfilled [14]. Unconscious vision, on the other hand, is characterized by unbound visual features, detected by neurons in distributed areas [58]. The boundary between unconscious and conscious vision is, therefore, put at the transition between feedforward and recurrent processing (figure 1b) [59]. Adding more

widespread recurrent interactions—including fronto-parietal activation—may give you the ability to cognitively manipulate and report the conscious visual percept, but this is not explaining the unconscious/conscious transition itself [60]. Conditions like change blindness, neglect, etc. are considered failures of attention and report, not of conscious vision [14,61]. Experimental support comes primarily from findings showing that report, cognitive access or attention do not change the quality of (neural) visual representations [62]; they just add report, access and attention [63].

Further arguments in favour of the one or the other have been laid out extensively elsewhere [14,15,53,60]. The discussion is reminiscent of some older philosophical discussions, for example about the distinction between A (access) and P (phenomenal) consciousness [64–66], or that of easy versus hard problems of consciousness [67]. The controversy between the two ideas seems to be in a stalemate, with support for one or the other swinging back and forth [54,60,61,68]. There is sufficient data for either point of view, it mainly boils down to the two theories having clearly different explananda: are you interested in explaining the seeing or the knowing that we see? Crucial is whether we should consider consciousness as closely intertwined with attention and cognition (as in the first view) or as a function or phenomenon with its own ontological status (as in the second view) [14,58,69,70].

Not that both types of theories would not argue to explain the 'seeing'. HOT, for example, argues that the higher-order representation of visual experience renders the experience conscious, hence truly 'seen', even when that higher-order representation has nothing to do with knowing or other higher levels of cognition. Similarly, in global workspace theory, it is argued that local recurrence leads to stimuli being 'pre-conscious' and that only after global ignition the 'seeing' arises.¹ Either way, however, according to these theories something beyond (recurrent) visual processing is necessary to explain the seeing. Similarly, neural correlates of these types of theories (when given) typically go beyond visual cortex and require the involvement of fronto-parietal or other non-visual regions. The second type of theories, however, confine all the necessary ingredients for seeing to the visual processes (or visual pathways) themselves. It may be fruitful to evaluate the two types of theories with respect to their *explanatory power* [13,14]. For example, does frontal involvement mechanistically explain something fundamental about the difference between conscious and unconscious vision, or rather about the difference between the ability to report or not [63]?

A more recent idea states that consciousness is all about the integration of information (information integration theory, IIT [72–74]). In this theory, it is taken as axiomatic that conscious experience is information that is integrated into a structured, united and intrinsic whole, at the exclusion of infinite other possible wholes. The properties of networks are linked to these axiomatic properties of consciousness in that representational systems, such as the brain, can be assigned a value (PHI) based on their connectivity, and to what extent it supports the integration of information. The higher this value, the more conscious the system is. Some aspects of the theory do quite well at explaining why highly coherent states (such as epilepsy) may yield

unconsciousness instead of consciousness, or why some connectivity types (such as those of the cerebellum) do not support consciousness, while others (such as in the cortex) do. IIT would predict low PHI for strictly feedforward systems, and higher PHI for recurrent ones. Strictly visual recurrency would have lower PHI than global recurrency, but the theory is somewhat ambiguous about the extent of information integration (and hence recurrency) that is necessary for the conscious experience to arise [75].

4. Escaping panpsychism, or what it is like to be a fly

Current models of consciousness all suffer from the same problem: at their core, they are fairly simple, too simple maybe. The distinction between feedforward and recurrent processing already exists between two reciprocally connected neurons. Add a third and we can distinguish between ‘local’ and ‘global’ recurrent processing. From a functional perspective, processes like integration, feature binding, global access, attention, report, working memory, metacognition and many others can be modelled with a limited set of mechanisms (or lines of MATLAB code). More importantly, it is getting increasingly clear that versions of these functions exist throughout the animal kingdom, and maybe even in plants.

Effects of anaesthesia are a good example of the problem we run into. At comparable doses of the volatile anaesthetics isoflurane or halothane, animals as different as worms, flies, goldfish, ducks, rats, horses, monkeys and man stop exhibiting responses to nociceptive or otherwise threatening stimuli [76]. Moreover, in all animals ‘higher’-level cognitive functions are more susceptible to anaesthesia than more primitive, reflex-like reactions. In the nematode, a creature with 302 neurons, the order of functions that ‘go’ under increasing levels of isoflurane anaesthesia are male mating, coordinated movement, chemotaxis and pharyngeal pumping until the worm stops moving at all [76]. Flies have larger brains, and neural activity shows similar spectra of local field potentials as can be found in mammals, with lower frequencies having more power than higher frequencies (roughly $1/f$ spectrum). In man and other mammals, it has been consistently shown that the higher frequencies (γ band) are indicative of feedforward or local processing, while lower frequencies (α and β bands) are linked to feedback or recurrent processing and top-down influences. Using Granger causality analysis of local field potentials, a similar distinction was found in the fly. Moreover, feedback and recurrent processing are more susceptible to anaesthesia than feedforward activation in mammals like man, monkey, rodents or ferrets. The same was recently found in the fly [77]. So in sum, both from functional and neural perspectives anaesthesia has the same effects in species as different as humans or flies. It may not even stop there. Plants such as the Venus flytrap stop responding to stimulation of their trigger hairs (normally excited by prey insects) when exposed to ether [78], a volatile anaesthetic also working in man.² The natural counterpart of anaesthesia—sleep—is equally ubiquitous in the animal and plant kingdoms. Animals that do not sleep have yet to be found [79], and many plants show clearly different metabolic activity patterns (and appearances) during day and night. We experience the difference between awake and asleep as one of the most fundamental transitions

of consciousness, and the fact that conscious sensations are lost during sleep is possibly the number one reason why we believe there *is* consciousness. If this transition between awake and asleep is present in all living beings, does that imply that to these other beings the transition ‘feels’ equally dramatic? And does it then follow that these beings should, therefore, be granted consciousness—when awake? We run into similar problems when we assign specific functions as central to consciousness. Cognitive function like attention, access or metacognition, often linked to consciousness in ‘type 1’ theories, seem similarly widespread throughout the animal kingdom. Attentional processes such as increased reaction times in the presence of distractors, or the suppression of behavioural responses to non-attended stimuli are present in bees and flies [80]. Ants can navigate home in the absence of a sensory stimulus (such as the scent of the nest) guiding their way. They seem to do so via an internally represented ‘home vector’, that is based on an integration of its own past movements relative to the polarization plane of the sky. This way, the ant, after having found food, can run straight back to the nest, regardless of the often haphazard way it left it. Importantly, however, the ant also can invert its home vector to then run straight back from the nest to the location where the food was found. Similarly, bees have the well-known capacity to convert similar vectors into dances that are signalling other bees where food can be found. Crabs combine several home vectors for other purposes than just homing, for example, to defend their burrow [81]. At the core, such capacities fulfil all requirements for what we call ‘cognitive manipulation’ or ‘access’ in higher animals like man, given that access is generally defined as the making flexibly available of sensed or memorized information for different types of behaviour. In line with this, bees can abstract at least two different concepts from a set of complex pictures and combine them in a rule for subsequent choices.

Metacognition is generally studied using paradigms where animals can choose to opt-out if they consider the task too difficult. The idea is that in doing so the animal expresses knowledge of its own performance, or of the probability that his stimulus–response coupling is the right one [82]. Also, opting out has been used to distinguish blindsight from the conscious vision in monkeys [83]. Honeybees were presented with visual targets that either indicated a positive (sucrose) or negative (quinine) reinforcement. The more difficult the visual discrimination became, the more often the bees opted to not respond at all, thereby increasing their overall performance. Moreover, they easily transferred this strategy to another task [80]. Whether it would qualify for metacognition is debatable, but even plants (peas) can adopt different root growing strategies depending on the variability of nutrients in their soils, showing risk appetite and aversion not unlike animals do [84]. What would follow from such findings? If we take the functional perspective, and state that attention, access or metacognitive faculties are the hallmarks of consciousness, we would have to conclude that these simple animals have conscious sensations of their surroundings, home vectors and stimulus response-couplings. If not willing to accept that, one would have to resort to meta- or higher-order variants of these functions and deny consciousness to animals that just execute the ‘simple versions’. But what higher order variants would that be: ‘super-attention’, ‘super-access’ or ‘super-metacognition’ of some sort? We

cannot argue that attention, access or metacognition are central to consciousness in humans, and then suddenly retreat if animals that we do not 'believe' to have consciousness execute those same functions. It is either that these functions are not constituting consciousness, or that the lower animals have consciousness too.

Functional and neural definitions of consciousness as currently formulated thus direct us towards a rather panpsychistic view,³ where all animals and possibly even plants would be conscious, or at least express the unconscious/conscious dichotomy. One may even start to worry over cultures of neurons in a Petri dish, or slices of cortex or hippocampus as often studied in electrophysiological experiments. Note that IIT, in particular, would probably assign higher PHI (and hence consciousness) to such preparations than to flies [74]. Another obvious extension would be consciousness in artificial intelligence (AI) systems. It has been argued that current AI generally lacks mechanisms for global access and availability of information or for metacognition (required for consciousness according to view 1 above), so that fears of conscious machines are premature [85]. Others, however, have opposed this view, arguing that many AI systems, in fact, do have these properties (either in rudimentary or more advanced forms). Moreover, if conscious experience is better explained by recurrent interactions or information integration (view 2 above), we may already have created conscious machines [86].

Do such considerations falsify current theories of consciousness? It seems entirely plausible that other animals than us have conscious sensations, fitting with the propensity of science to remove man from its pedestal, acknowledging that we are just animals too. On the other hand, it feels unsatisfactory to equate consciousness as we experience it to whatever happens in the mind of a fly. IIT as well as recurrent processing theories of consciousness offer some potential solutions to this problem. Both theories would argue that while the conscious/unconscious dichotomy should be present in many animals (or even devices), the 'what it is like's' may differ vastly between them. Of the senses, vision and olfaction are most important in the fly, and it has been shown that there is multisensory integration between the two. The visual system of the fly is dominated by neurons and pathways that detect motion (even second order and illusory motion), not surprising given its behaviour. Second are pathways related to colour vision. Much less prominent are systems that would enable the fly to discriminate between all sorts of objects and shapes [87]. So what 'is it like' to see like a fly? Obviously, conscious content can only be about information that the neural machinery detects. Recurrent processing theory (RPT) would argue that a fly sees an amalgam of motion and colour, quite 'objectless' in fact. Objectless vision is not unlike what patients with damage to visual object areas see, like the famous D.F., studied extensively by Goodale & Milner, who cannot identify objects and shapes, yet is seeing features like colour and motion [88]. Apperceptive agnosia patients have similar sensations of unbound basic features [89,90]. In blindsight, humans report to even have no conscious visual sensations at all (in the affected hemifield), apart from a vague sensation of objectless motion when things move fast enough [91,92]. So what the fly sees may be similar to the limited experiences of blindsight and other patients, who would qualify their vision as severely impaired or even 'unconscious'. Add to

that the obvious other differences between fly and human visual systems (resolution, numbers of neurons and areas, etc.), and it will be quite clear that we would probably qualify what the fly sees as not (or hardly) seeing at all.

IIT adds the prediction that conscious content is also determined by what is excluded by the current state of the system. For example, colour neurons in our brains contribute to our phenomenal experience even when silent, because they exclude the presence of colour, while patients with damage to colour regions do not even understand or 'miss' the notion of colour [72,75], much like we do not 'miss' perceiving the polarization of light. Because our brains support so much more concepts and sensory modalities than brains of flies, our visual experience would be richer than that of the fly even if the fly brain had the same visual machinery as we have.

This hints towards an important conclusion. What we consider highly impaired vision or even blindness, is what the fly may see consciously. Both species have an unconscious/conscious dichotomy. When a fly escapes from a swat it executes a very intelligent reflex arc, where it positions its middle legs to jump in a direction orthogonal to the approaching danger [93], not unlike what we do when a large object is looming. After the fact, when recurrent processing (or information integration) has kicked in, the fly may experience the fear associated with it (flies show fear conditioning [94]). But the experience will never transcend beyond 'something nasty that moved'. In our minds, however, we will come to recognize that the looming object was a pick-up truck, running off the road, steered by a bearded man. We heard the engine roaring, felt the gush of wind when the car almost hit us. And in doing so we exclude millions of other possible percepts. So although our reflex may be similar to that of the fly, the sensory experience that follows is infinitely richer. In fact, what the fly experiences may be very similar to what we experience in a masking experiment,⁴ or when we choose 'not seen' on a perceptual awareness scale. What we call unconscious may be the richest experience a fly or worm ever experiences.

Should we then conclude that flies do not have consciousness after all? That is missing the point. The notion of conscious versus unconscious is still useful—or even necessary—to grasp what goes on in the mind of a fly. The unconscious to conscious transition occurs in flies just like it does in humans, in the sense that it is a radical—and probably nonlinear—phase transition, supported by the transition between feedforward to recurrent processing. Accordingly, the amount of integrated information (PHI) suddenly increases strongly, which is accompanied by a sudden transition from processing 'in the dark' to conscious phenomenology. Therefore, to the fly, being awake will 'feel' clearly different from being anesthetized. Hence, there should be a clear and rather dichotomous distinction between the two states from the perspective of the fly, just as there is this distinction in our phenomenology when we go from awake to asleep. But conscious experience in the fly is just not comparable in richness, quality and extent to ours.

Such considerations call for another distinction, that between unconscious and non-conscious, or death. No matter how deep the sleep or anaesthesia, there remain clear differences between brains that live and those that do not, both in terms of remaining functionality and neural processing. Anesthetized brains process information in a feedforward

fashion, and potentially also exhibit limited recurrent information exchange; PHI will be low but not zero. Clinically, we know there are differences between brain death, coma, vegetative state or minimally conscious state. In sleep, we recognize different depths. In other words, there are *gradations* of unconscious. These gradations are all accompanied by differences in the extent of neural processing, for example, expressed in the spread of neural signals evoked by TMS of localized regions of cortex [96,97]. But when dead, all that stops; clearly another nonlinear transition in the level of consciousness. Does that imply that being asleep or anesthetized ‘feels’ different from being dead? Nobody knows, of course, but it is an interesting speculation. It is getting increasingly clear that sleep is not always as unconscious as we may think it is. Sensory processing continues, allowing a child sobbing in the distance to wake up a worried mother, or outside events to get integrated into dreams. People with neurophysiological signs of sleep may feel awake [98]. This idea of ‘multiple gradations’ of unconsciousness is also suggested by studies in awake subjects, where masked and unattended items may feel equally ‘unconscious’ (as in ‘not accessible by the self’), yet for entirely different reasons, and supporting taxonomically different types of unconsciousness (in that case unconscious and pre-conscious) [99].

Together, we arrive at a view where we can objectively define strong and nonlinear transitions in states of consciousness that most likely occur in many animals. Because of differences in cognitive machinery that these animals are endowed with (different sensory modalities and capabilities, emotion, social cognition, language, self, etc.), these different states of consciousness support highly different phenomenal experiences. These will feel so entirely different that what we ‘experience’ during deep sleep is similar in richness to what a fly experiences during wakefulness (figure 2). It would be human-chauvinist to deny the fly its state transitions, however, and claim that it would not feel different to the fly to be awake, anesthetized or dead. More importantly, however, it would deny consciousness its own ontological status, separate from the cognitive machinery it supports [14].

5. The missing ingredient

Another, and highly related problem is equally present in theories of consciousness—that of the missing ingredient. Recurrent processing will not grind to a complete halt in anaesthesia. Although monkey experiments have shown the abolishment of interareal (e.g. between V1 and extrastriate areas) or maybe even local (horizontal) interactions [40], it is conceivable that at more local levels (between nearby neurons) recurrent interactions remain. In brain death, recurrent processing is probably completely absent, as even feedforward signals such as the N20 SSEP are absent. In coma, however, some recurrent signals may be present, expressed in burst suppression EEG, where the complete absence of coherent activity alternates with bursts of synchronous cortical activity. In vegetative state,⁵ connectivity is similar to that of deep sleep or anaesthesia, all showing coherent EEG signals [96]. If some levels of recurrent processing are compatible with unconsciousness, ‘something else’ beyond mere recurrency apparently is needed for consciousness.

IIT offers a solution by arguing that not all recurrency equally supports the integration of information. With too little connectivity, the information stays too ‘localized’, but when units are connected too strongly neurons lose their ability to convey independent information, also lowering PHI.⁶ Indeed, the high amplitude and widespread slow waves seen in sleep, anaesthesia or vegetative state EEG are signals of too strong coherency [97]. This may also be the reason why in monkey experiments anaesthesia completely disrupts neural signals that are signalling perceptual organization [40] (i.e. meaningful information integration), while at the same time more ‘meaningless’ interactions remain [100].

Another additional ingredient may be that some types of recurrency are more likely to induce changes to the connectivity of the neural network, and do so at different time scales, and using different molecular mechanisms. A long-standing idea is that recurrent interactions between neurons, particularly if they involve burst firing of neurons, are likely to induce synaptic plasticity, involving NMDA receptor activation and subsequent Ca^{2+} influx, which induces a cascade of molecular and genetic events, changing the efficacy of the synapses involved [101]. Indeed, NMDA receptor activation has been tightly linked to cortico-cortical feedback [37], and NMDA receptor blockers impair feature integration [102]. Some have proposed that the final common pathway of many anaesthetics is NMDA receptor inactivation [103]. Synaptic plasticity typically occurs at time scales too slow for a direct involvement in generating conscious experience, yet NMDA receptor activation may be one of the required preconditions for effective large-scale recurrent interactions to occur [104]. Also, more rapid changes in synaptic efficacy occur, such as paired pulse facilitation or suppression, each having various potential underlying mechanisms [105]. Another important role may be played by γ -aminobutyric acid, as it is actively involved in the dynamics of the neuronal competition underlying binocular rivalry [106]. In sum, it may be that recurrent interactions are just the first step in a more complex cascade of molecular events that induce changes to the network at various time scales. These changes to the network—and not the recurrency *per se*—may then be the ingredient that generates the conscious experience, and these changes may only happen when recurrency fulfils certain temporal or spatial conditions (figure 3).

IIT has formulated some axioms about conscious experience, for which it seeks a physical substrate (PSC), for example, in the brain. This would pose a problem, however, when these axioms are not unique to conscious experience yet also hold for unconscious processing. That seems to be the case for two of these axioms, those of information and of integration. The axiom of information states that ‘experience is specific, being composed of a particular set of phenomenal distinctions (qualia), which make it what it is and different from other experiences. ... the content of my current experience might be composed of seeing a book (rather than seeing no book), which is blue (rather than not blue), and so on ...’ [75, p. 451]. This is hardly specific to conscious processing, as all unconscious processing equally involves specific information, for example, in that neurons encode all sorts of features and high-level objects (at the exclusion of others) even when anesthetized, under conditions of masking, etc. [58]. The axiom of integration states that ‘experience is unitary, meaning that it is composed of a




Figure 2. Phenomenal content is conscious in the fly, yet comparable to unconscious content in humans. The figure depicts the richness or quality of phenomenal content (calculated as PHI in IIT) on the vertical axis. In both fly and man there is a nonlinear transition in this richness, marking the transition between unconscious to conscious processing, corresponding to the transition from feedforward to recurrent processing at the neural level. At the lower end, there is a second nonlinear transition, that between dead or alive. In both species, it therefore makes sense to talk about ‘conscious’ versus ‘unconscious’. However, the content of conscious processing in the fly is comparable to unconscious processing (e.g. deep sleep) in humans. (Online version in colour.)

set of phenomenal distinctions, bound together in various ways, that is irreducible to non-interdependent subsets’ [75, p. 452]. This smacks of what others call feature binding, Gestalt grouping, perceptual organization, etc. While previously thought to be exclusively linked to conscious perception, feature binding does occur to quite some extent also in unconscious visual processing [107]. This holds in particular for the binding of features that are detected simultaneously by a single neuron, such as orientation and movement or colour (so-called base-groupings), as these are already extracted during the feedforward sweep [108]. More problematic is the integration of distant features that constitute many perceptual illusions or perceptual inferences, which often also occur unconsciously [58]. Only some levels of perceptual organization or Gestalt grouping seem to require consciousness. In that sense, IIT suffers from the same ‘missing ingredient’ as the other theories

mentioned, in that consciousness is more than the mere integration of information.

Note that theories that propose higher cognitive functioning as the key ingredient for consciousness also suffer from the ‘missing ingredient’ problem, as many high-level cognitive functions occur unconsciously. For example, masked no-go stimuli activate prefrontal inhibition networks in such a way that they induce a slowing of reactions [20], indicating that cognitive control can be executed unconsciously. Also, task switching, attention shifting, conflict monitoring, error detection and evidence accumulation can occur unconsciously [109,110]. Do such results suggest the existence of unconscious access or metacognition? Using unconscious cues to adapt or control behaviour, to change or make decisions or to switch strategies certainly qualifies for unconscious access, as these are often ways in which conscious access is operationalized. The case for unconscious metacognition is more difficult.⁷ However, post error slowing is very similar to opting out, in that the system seems to know something about its own performance and adapts its behaviour accordingly. Post error slowing does occur for unconscious errors [111], yet is typically weaker than after conscious errors. When comparing conscious to unconscious cognition it is often found that the former is more flexibly deployed and has stronger effects on behaviour as well as stronger associated brain signals. This does not solve the ‘missing ingredient’ issue, however, unless that ingredient is ‘more’ or ‘stronger’. Some have, therefore, proposed that cognitive functioning *per se* is not the key ingredient of consciousness but learning to continuously redescribe the cause–effect relations of these functions, whether simple or complex, and internal or external. Once these relations are well established, they leave the conscious domain (as in well-learned skills) [112]. Naccache [113] proposes that a multitude of functions (or ingredients) may be simultaneously required to produce consciousness. In his dialectic approach, the starting point would be the set of cognitive functions that in a healthy awake adult support self-reportability (in its widest sense, so also non-verbal), as this is the aspect of consciousness that is central to the phenomenon, at least according to longstanding philosophical and introspective intuitions. The next step would be to find the neural mechanisms that support such self-reportability, which in the healthy human seem to be recurrent interactions between highly educated and sophisticated cortical regions [52]. Only when all the necessary ingredients coincide, that is when both sufficiently educated cortical areas are active and engage in recurrent interactions, does consciousness as experienced by a normal adult emerge. It should be possible to infer the psychological states going along with recurrent interactions between less sophisticated cognitive modules (say in animals, babies, etc.), such that at some point one would have to conclude that ‘consciousness as we experience it’ is no longer present. The same logic could be applied to difficult to categorize patients suffering from disorders of consciousness [114].

6. Concluding remarks: let us talk

I have tried to highlight some pregnant issues in our current knowledge about consciousness and how it may relate to what happens in brains. The focus was on three theories, in




Figure 3. Recurrent processing-induced network plasticity as a potential ‘missing ingredient’. A pre-stimulus network state is depicted, with varying strengths of connections. A stimulus evokes a feedforward sweep, followed by recurrent processing. Recurrent processing may then either induce changes to the strengths of connections (via long or short term, or pre- or postsynaptic plasticity; right) or not (left). Only in the case of network changes will conscious experience arise. (Online version in colour.)

particular, being recurrent processing theory (RPT), global neuronal workspace theory (GNWT) and integrated information theory (IIT). I noticed the stalemate between RPT and GNWT, primarily caused by them having different explananda. I discussed how the three theories endorse panpsychism and how this could potentially be solved, and that each of them still miss ‘key ingredients’ to satisfactorily distinguish unconscious from conscious processing. Many other theories exist, of course. Readers are encouraged to comment on how some of these other theories may solve the issues raised. Because the science of consciousness will clearly benefit from an open exchange of ideas, based on the wealth of experimental data that exists and is produced every day, I hope I have inspired such an exchange.

Data accessibility. This article has no additional data.

Competing interests. I declare I have no competing interests.

Funding. I received no funding for this study.

Endnotes

¹Although often GWS is rather ambiguous about whether it explains seeing or access to seeing. For example, the somewhat cryptical phrase ‘access to consciousness’ is often used to describe what the GWS is providing [71].

²The effect is caused by a suppression of action potentials generated by touch-sensitive ion channels, which is feedforward as far as we know, so vegetarians need not worry yet. . .

³Real panpsychism goes somewhat further in that it would also grant consciousness to any item that processes information (such as thermostats), or even those that do not (such as stones).

⁴There is some recent discussion as to whether masking indeed takes away all visual experience, or just renders the experience highly degraded. In other words, a masked visual stimulus is never the same as no visual stimulus at all. Peters and Lau, for example, had subjects discriminate oriented Gabor patches and bet on their performance to gauge metacognitive awareness. They found above chance betting only for stimuli that also can be discriminated, and that subjects behave as ideal observers with different levels of noise for the two tasks. They assert that masking may in fact never completely remove awareness when other functions remain, implying that previous experiments showing unconscious cognition in fact study cognition under limited rather than eliminated awareness [95].

⁵This refers to VS patients in which conscious sensations seem absent. In some (approx. 20%) of VS patients clear signs of conscious experience are present, as indicated by fMRI experiments that even showed some patients being able to communicate using their brain signals.

⁶This is also purported to be the reason that we are not conscious of what happens in the cerebellum, where connectivity is too dense compared to the cortex.

⁷Also because in the ‘knowing that you know’ the first ‘knowing’ to some researchers or philosophers only qualifies when it is ‘conscious knowing’. In that case, unconscious metacognition cannot exist by definition.

1. Koch C, Massimini M, Boly M, Tononi G. 2016 Neural correlates of consciousness: progress and problems. *Nat. Rev. Neurosci.* **17**, 307–321. (doi:10.1038/nrn.2016.22)
2. Lamme VA, Roelfsema PR. 2000 The distinct modes of vision offered by feedforward and recurrent processing. *Trends Neurosci.* **23**, 571–579. (doi:10.1016/S0166-2236(00)01657-X)
3. Li FF, VanRullen R, Koch C, Perona P. 2002 Rapid natural scene categorization in the near absence of attention. *Proc. Natl Acad. Sci. USA* **99**, 9596–9601. (doi:10.1073/pnas.092277599)
4. Grainger J, Rey A, Dufau S. 2008 Letter perception: from pixels to pandemonium. *Trends Cogn. Sci.* **12**, 381–387. (doi:10.1016/j.tics.2008.06.006)
5. Oppermann F, Hassler U, Jescheniak JD, Gruber T. 2011 The rapid extraction of gist—early neural correlates of high-level visual processing. *J. Cogn. Neurosci.* **24**, 521–529. (doi:10.1162/jocn_a_00100)
6. Kirchner H, Thorpe SJ. 2006 Ultra-rapid object detection with saccadic eye movements: visual processing speed revisited. *Vision Res.* **46**, 1762–1776. (doi:10.1016/j.visres.2005.10.002)
7. Buetti S, Juan E, Rinck M, Kerzel D. 2012 Affective states leak into movement execution: automatic avoidance of threatening stimuli in fear of spider is visible in reach trajectories. *Cogn. Emotion* **26**, 1176–1188. (doi:10.1080/02699931.2011.640662)
8. Naccache L, Blandin E, Dehaene S. 2002 Unconscious masked priming depends on temporal attention. *Psychol. Sci.* **13**, 416–424. (doi:10.1111/1467-9280.00474)
9. Kunde W, Kiesel A, Hoffmann J. 2003 Conscious control over the content of unconscious cognition. *Cognition* **88**, 223–242. (doi:10.1016/S0010-0277(03)00023-4)
10. Naccache L. 2008 Conscious influences on subliminal cognition exist and are asymmetrical: validation of a double prediction. *Conscious Cogn.* **17**, 1359–1360. (doi:10.1016/j.concog.2008.01.002)
11. Super H, van der Togt C, Spekreijse H, Lamme VA. 2003 Internal state of monkey primary visual cortex (V1) predicts figure–ground perception. *J. Neurosci.* **23**, 3407–3414. (doi:10.1523/JNEUROSCI.23-08-03407.2003)
12. Lamme VA. 2000 Neural mechanisms of visual awareness: a linking proposition. *Brain Mind* **1**, 385–406. (doi:10.1023/A:1011569019782)
13. Lamme VA. 2006 Towards a true neural stance on consciousness. *Trends Cogn. Sci.* **10**, 494–501. (doi:10.1016/j.tics.2006.09.001)
14. Lamme VA. 2010 How neuroscience will change our view on consciousness. *Cogn. Neurosci.* **1**, 204–220. (doi:10.1080/17588921003731586)
15. Dehaene S, Changeux JP, Naccache L, Sackur J, Sergent C. 2006 Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends Cogn. Sci.* **10**, 204–211. (doi:10.1016/j.tics.2006.03.007)
16. Breitmeyer BG, Ogmen H. 2000 Recent models and findings in visual backward masking: a comparison, review, and update. *Percept. Psychophys.* **62**, 1572–1595. (doi:10.3758/BF03212157)
17. Lamme VAF, Zipser K, Spekreijse H. 2001 Masking interrupts figure-ground signals in V1. *J. Vis.* **1**, 32.
18. Dehaene S, Naccache L, Le Clec'h G, Koehlin E, Mueller M, Dehaene-Lambertz G, van de Moortele PF, Le Bihan D. 1998 Imaging unconscious semantic priming. *Nature* **395**, 597–600. (doi:10.1038/26967)
19. van Gaal S, Ridderinkhof KR, Fahrenfort JJ, Lamme VAF. 2008 Electrophysiological correlates of unconsciously triggered inhibitory control. *Perception* **37**, 139.
20. van Gaal S, Ridderinkhof KR, Fahrenfort JJ, Scholte HS, Lamme VAF. 2008 Frontal cortex mediates unconsciously triggered inhibitory control. *J. Neurosci.* **28**, 8053–8062. (doi:10.1523/jneurosci.1278-08.2008)
21. Lamme VA, Super H, Spekreijse H. 1998 Feedforward, horizontal, and feedback processing in the visual cortex. *Curr. Opin Neurobiol.* **8**, 529–535. (doi:10.1016/S0959-4388(98)80042-1)
22. Lamme VA. 1995 The neurophysiology of figure-ground segregation in primary visual cortex. *J. Neurosci.* **15**, 1605–1615. (doi:10.1523/JNEUROSCI.15-02-01605.1995)
23. Lamme VA, Van Dijk BW, Spekreijse H. 1993 Contour from motion processing occurs in primary visual cortex. *Nature* **363**, 541–543. (doi:10.1038/363541a0)
24. Lamme VA, Rodriguez-Rodriguez V, Spekreijse H. 1999 Separate processing dynamics for texture elements, boundaries and surfaces in primary visual cortex of the macaque monkey. *Cereb. Cortex* **9**, 406–413. (doi:10.1093/cercor/9.4.406)
25. Zipser K, Lamme VA, Schiller PH. 1996 Contextual modulation in primary visual cortex. *J. Neurosci.* **16**, 7376–7389. (doi:10.1523/JNEUROSCI.16-22-07376.1996)
26. Hupé JM, James AC, Payne BR, Lomber SG, Girard P, Bullier J. 1998 Cortical feedback improves discrimination between figure and background by V1, V2 and V3 neurons. *Nature* **394**, 784–787. (doi:10.1038/29537)
27. Supér H, Lamme VAF. 2007 Altered figure-ground perception in monkeys with an extra-striate lesion. *Neuropsychologia* **45**, 3329–3334. (doi:10.1016/j.neuropsychologia.2007.07.001)
28. Roelfsema PR, Lamme VA, Spekreijse H, Bosch H. 2002 Figure-ground segregation in a recurrent network architecture. *J. Cogn. Neurosci.* **14**, 525–537. (doi:10.1162/08989290260045756)
29. Caputo G, Casco C. 1999 A visual evoked potential correlate of global figure-ground segmentation. *Vision Res.* **39**, 1597–1610. (doi:10.1016/S0042-6989(98)00270-3)
30. Scholte HS, Spekreijse H, Lamme VAF. 2001 Neural correlates of global scene segmentation are present during inattention blindness. *J. Vis.* **1**, 346. (doi:10.1167/1.3.346)
31. Lee TS, Yang CF, Romero RD, Mumford D. 2002 Neural activity in early visual cortex reflects behavioral experience and higher-order perceptual saliency. *Nat. Neurosci.* **5**, 589–597. (doi:10.1038/nn860)
32. Squire RF, Noudoost B, Schafer RJ, Moore T. 2013 Prefrontal contributions to visual selective attention. *Annu. Rev. Neurosci.* **36**, 451–466. (doi:10.1146/annurev-neuro-062111-150439)
33. Roelfsema PR, Lamme VA, Spekreijse H. 1998 Object-based attention in the primary visual cortex of the macaque monkey. *Nature* **395**, 376–381. (doi:10.1038/26475)
34. Schmitz TW, De Rosa E, Anderson AK. 2009 Opposing influences of affective state valence on visual cortical encoding. *J. Neurosci.* **29**, 7199–7207. (doi:10.1523/JNEUROSCI.5387-08.2009)
35. Huang W-J, Chen W-W, Zhang X. 2015 The neurophysiology of P 300—an integrated review. *Eur. Rev. Med. Pharmacol. Sci.* **19**, 1480–1488.
36. Ekstrom LB, Roelfsema PR, Arsenault JT, Bonmassar G, Vanduffel W. 2008 Bottom-up dependent gating of frontal signals in early visual cortex. *Science* **321**, 414–417. (doi:10.1126/science.1153276)
37. Self MW, Kooijmans RN, Supér H, Lamme VA, Roelfsema PR. 2012 Different glutamate receptors convey feedforward and recurrent processing in macaque V1. *Proc. Natl Acad. Sci. USA* **109**, 11 031–11 036. (doi:10.1073/pnas.1119527109)
38. Shao Z, Burkhalter A. 1996 Different balance of excitation and inhibition in forward and feedback circuits of rat visual cortex. *J. Neurosci.* **16**, 7353–7365. (doi:10.1523/JNEUROSCI.16-22-07353.1996)
39. Sillito AM, Cudeiro J, Jones HE. 2006 Always returning: feedback and sensory processing in visual cortex and thalamus. *Trends Neurosci.* **29**, 307–316. (doi:10.1016/j.tics.2006.05.001)
40. Lamme VA, Zipser K, Spekreijse H. 1998 Figure-ground activity in primary visual cortex is suppressed by anesthesia. *Proc. Natl Acad. Sci. USA* **95**, 3263–3268. (doi:10.1073/pnas.95.6.3263)
41. Alkire MT, Hudetz AG, Tononi G. 2008 Consciousness and anesthesia. *Science* **322**, 876–880. (doi:10.1126/science.1149213)
42. Fahrenfort JJ, Scholte HS, Lamme VA. 2007 Masking disrupts reentrant processing in human visual cortex. *J. Cogn. Neurosci.* **19**, 1488–1497. (doi:10.1162/jocn.2007.19.9.1488)
43. Ro T, Breitmeyer B, Burton P, Singhal NS, Lane D. 2003 Feedback contributions to visual awareness in human occipital cortex. *Curr. Biol.* **13**, 1038–1041. (doi:10.1016/S0960-9822(03)00337-3)
44. Fahrenfort JJ, Snijders TM, Heinen K, van Gaal S, Scholte HS, Lamme VA. 2012 Neuronal integration in visual cortex elevates face category tuning to conscious face perception. *Proc. Natl Acad. Sci. USA* **109**, 21 504–21 509. (doi:10.1073/pnas.1207414110)

45. Ro T. 2010 What can TMS tell us about visual awareness? *Cortex* **46**, 110–113. (doi:10.1016/j.cortex.2009.03.005)
46. Koivisto M, Revonsuo A. 2010 Event-related brain potential correlates of visual awareness. *Neurosci. Biobehav. Rev.* **34**, 922–934. (doi:10.1016/j.neubiorev.2009.12.002)
47. Supér H, Spekreijse H, Lamme VAF. 2001 Two distinct modes of sensory processing observed in monkey primary visual cortex (V1). *Nat. Neurosci.* **4**, 304–310. (doi:10.1038/85170)
48. Sergent C, Baillet S, Dehaene S. 2005 Timing of the brain events underlying access to consciousness during the attentional blink. *Nat. Neurosci.* **8**, 1391–1400. (doi:10.1038/nn1549)
49. Koivisto M, Revonsuo A. 2003 An ERP study of change detection, change blindness, and visual awareness. *Psychophysiology* **40**, 423–429. (doi:10.1111/1469-8986.00044)
50. Shafto JP, Pitts MA. 2015 Neural signatures of conscious face perception in an inattentive blindness paradigm. *J. Neurosci.* **35**, 10 940–10 948. (doi:10.1523/JNEUROSCI.0145-15.2015)
51. Baars BJ. 2005 Global workspace theory of consciousness: toward a cognitive neuroscience of human experience. *Prog. Brain Res.* **150**, 45–53. (doi:10.1016/S0079-6123(05)50004-9)
52. Dehaene S, Naccache L. 2001 Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition* **79**, 1–37. (doi:10.1016/S0010-0277(00)00123-2)
53. Kouider S, de Gardelle V, Sackur J, Dupoux E. 2010 How rich is consciousness? The partial awareness hypothesis. *Trends Cogn. Sci. (Regul. Ed.)* **14**, 301–307. (doi:10.1016/j.tics.2010.04.006)
54. Cohen MA, Dennett DC. 2011 Consciousness cannot be separated from function. *Trends Cogn. Sci. (Regul. Ed.)* **15**, 358–364. (doi:10.1016/j.tics.2011.06.008)
55. Cleeremans A, Timmermans B, Pasquali A. 2007 Consciousness and metarepresentation: a computational sketch. *Neural Netw.* **20**, 1032–1039. (doi:10.1016/j.neunet.2007.09.011)
56. Seth AK. 2010 The grand challenge of consciousness. *Front. Psychol.* **1**, 5. (doi:10.3389/fpsyg.2010.00005)
57. Fleming SM, Lau HC. 2014 How to measure metacognition. *Front. Hum. Neurosci.* **8**, 443. (doi:10.3389/fnhum.2014.00443)
58. Lamme VAF. 2015 The crack of dawn: perceptual functions and neural mechanisms that mark the transition from unconscious processing to conscious vision. In *Open MIND* (eds TK Metzinger, JM Windt). Frankfurt am Main, Germany: MIND Group.
59. Lamme VA. 2003 Why visual attention and awareness are different. *Trends Cogn. Sci.* **7**, 12–18. (doi:10.1016/S1364-6613(02)00013-X)
60. Boly M, Massimini M, Tsuchiya N, Postle BR, Koch C, Tononi G. 2017 Are the neural correlates of consciousness in the front or in the back of the cerebral cortex? Clinical and neuroimaging evidence. *J. Neurosci.* **37**, 9603–9613. (doi:10.1523/JNEUROSCI.3218-16.2017)
61. Fahrenfort JJ, Lamme VA. 2012 A true science of consciousness explains phenomenology: comment on Cohen and Dennett. *Trends Cogn. Sci.* **16**, 138–139; author reply 139–40. (doi:10.1016/j.tics.2012.01.004)
62. Vandenbroucke AR, Fahrenfort JJ, Sligte IG, Lamme VA. 2014 Seeing without knowing: neural signatures of perceptual inference in the absence of report. *J. Cogn. Neurosci.* **26**, 955–969. (doi:10.1162/jocn_a_00530)
63. Tsuchiya N, Wilke M, Frässle S, Lamme VA. 2015 No-report paradigms: extracting the true neural correlates of consciousness. *Trends Cogn. Sci.* **19**, 757–770. (doi:10.1016/j.tics.2015.10.002)
64. Block N. 2007 Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behav. Brain Sci.* **30**, 481–499. (doi:10.1017/S0140525X07002786)
65. Block N. 2005 Two neural correlates of consciousness. *Trends Cogn. Sci.* **9**, 46–52. (doi:10.1016/j.tics.2004.12.006)
66. Block N. 1996 How can we find the neural correlate of consciousness? *Trends Neurosci.* **19**, 456–459. (doi:10.1016/S0166-2236(96)20049-9)
67. Chalmers DJ. 1995 Facing up to the problem of consciousness. *J. Conscious. Stud.* **2**, 200–219.
68. Odegaard B, Knight RT, Lau H. 2017 Should a few null findings falsify prefrontal theories of conscious perception? *J. Neurosci.* **37**, 9593–9602. (doi:10.1523/JNEUROSCI.3217-16.2017)
69. Lamme VAF. 2010 What introspection has to offer, and where its limits lie. *Cogn. Neurosci.* **1**, 232–235. (doi:10.1080/17588928.2010.502224)
70. Lamme VA. 2005 Independent neural definitions of visual awareness and attention. In *Cognitive penetrability of perception: attention, action, strategies, and bottom-up constraints* (ed. A Raftopoulos), pp. 171–191. Hauppauge, NY: Nova Science Publishers.
71. Dehaene S, Changeux J-P. 2005 Ongoing spontaneous activity controls access to consciousness: a neuronal model for inattentive blindness. *PLoS Biol.* **3**, e141. (doi:10.1371/journal.pbio.0030141)
72. Tononi G. 2004 An information integration theory of consciousness. *BMC Neurosci.* **5**, 42. (doi:10.1186/1471-2202-5-42)
73. Tononi G. 2008 Consciousness as integrated information: a provisional manifesto. *Biol. Bull.* **215**, 216–242. (doi:10.2307/25470707)
74. Tononi G. 2012 Integrated information theory of consciousness: an updated account. *Arch. Ital. Biol.* **150**, 293–329.
75. Tononi G, Boly M, Massimini M, Koch C. 2016 Integrated information theory: from consciousness to its physical substrate. *Nat. Rev. Neurosci.* **17**, 450–461. (doi:10.1038/nrn.2016.44)
76. Zalucki O, van Swinderen B. 2016 What is unconsciousness in a fly or a worm? A review of general anesthesia in different animal models. *Conscious Cogn.* **44**, 72–88. (doi:10.1016/j.concog.2016.06.017)
77. Cohen D, van Swinderen B, Tsuchiya N. 2018 Isoflurane impairs low frequency feedback but leaves high frequency feedforward connectivity intact in the fly brain. *eNeuro* **5**, 0329. (doi:10.1523/eneuro.0329-17.2018)
78. Yokawa K, Kagenishi T, Pavlovic A, Gall S, Weiland M, Mancuso S, Baluška F. 2017 Anaesthetics stop diverse plant organ movements, affect endocytic vesicle recycling and ROS homeostasis, and block action potentials in Venus flytraps. *Ann. Bot.* **mx155**. (doi:10.1093/aob/mcx155)
79. Cirelli C, Tononi G. 2017 The sleeping brain. *Cerebrum* 2017. See http://www.dana.org/Cerebrum/2017/The_Sleeping_Brain/.
80. Giurfa M. 2015 Learning and cognition in insects. *WIREs Cogn. Sci.* **6**, 383–395. (doi:10.1002/wcs.1348)
81. Webb B. 2012 Cognition in insects. *Phil. Trans. R. Soc. B* **367**, 2715–2722. (doi:10.1098/rstb.2012.0218)
82. Terrace HS, Son LK. 2009 Comparative metacognition. *Curr. Opin Neurobiol.* **19**, 67–74. (doi:10.1016/j.conb.2009.06.004)
83. Moore T, Rodman HR, Repp AB, Gross CG. 1995 Localization of visual stimuli after striate cortex damage in monkeys: parallels with human blindsight. *Proc. Natl Acad. Sci. USA* **92**, 8215–8218. (doi:10.1073/pnas.92.18.8215)
84. Dener E, Kacelnik A, Shemesh H. 2016 Pea plants show risk sensitivity. *Curr. Biol.* **26**, 1763–1767. (doi:10.1016/j.cub.2016.05.008)
85. Dehaene S, Lau H, Kouider S. 2017 What is consciousness, and could machines have it? *Science* **358**, 486–492. (doi:10.1126/science.aan8871)
86. Carter O, Hohwy J, van Bostel J, Lamme V, Block N, Koch C, Tsuchiya N. 2018 Conscious machines: defining questions. *Science* **359**, 400. (doi:10.1126/science.aar4163)
87. Zhu Y. 2013 The *Drosophila* visual system. *Cell Adh. Migr.* **7**, 333–344. (doi:10.4161/cam.25521)
88. Goodale MA, Milner AD. 1992 Separate visual pathways for perception and action. *Trends Neurosci.* **15**, 20–25. (doi:10.1016/0166-2236(92)90344-8)
89. Shelton PA, Bowers D, Duara R, Heilman KM. 1994 Apperceptive visual agnosia: a case study. *Brain Cogn.* **25**, 1–23. (doi:10.1006/brcg.1994.1019)
90. Grossman M, Galetta S, D'Esposito M. 1997 Object recognition difficulty in visual apperceptive agnosia. *Brain Cogn.* **33**, 306–342. (doi:10.1006/brcg.1997.0876)
91. Weiskrantz L. 1996 Blindsight revisited. *Curr. Opin Neurobiol.* **6**, 215–220. (doi:10.1016/S0959-4388(96)80075-4)
92. Zeki S, Fyfe DH. 1998 The Riddoch syndrome: insights into the neurobiology of conscious vision. *Brain* **121**, 25–45. (doi:10.1093/brain/121.1.25)
93. Card G, Dickinson MH. 2008 Visually mediated motor planning in the escape response of *Drosophila*. *Curr. Biol.* **18**, 1300–1307. (doi:10.1016/j.cub.2008.07.094)
94. Kottler B, van Swinderen B. 2014 Taking a new look at how flies learn. *eLife* **3**, e03978. (doi:10.7554/eLife.03978)

95. Peters MAK, Lau H. 2015 Human observers have optimal introspective access to perceptual processes even for visually masked stimuli. *eLife Sci.* **4**, e09651. (doi:10.7554/eLife.09651)
96. Gosseries O, Di H, Laureys S, Boly M. 2014 Measuring consciousness in severely damaged brains. *Annu. Rev. Neurosci.* **37**, 457–478. (doi:10.1146/annurev-neuro-062012-170339)
97. Tononi G, Massimini M. 2008 Why does consciousness fade in early sleep? *Ann. NY Acad. Sci.* **1129**, 330–334. (doi:10.1196/annals.1417.024)
98. Hsiao F-C *et al.* 2018 The neurophysiological basis of the discrepancy between objective and subjective sleep during the sleep onset period: an EEG-fMRI study. *Sleep* **41**, zsy056. (doi:10.1093/sleep/zsy056)
99. Kouider S, Dehaene S. 2007 Levels of processing during non-conscious perception: a critical review of visual masking. *Phil. Trans. R. Soc. B* **362**, 857–875. (doi:10.1098/rstb.2007.2093)
100. Allman J, Miezin F, McGuinness E. 1985 Stimulus specific responses from beyond the classical receptive field: neurophysiological mechanisms for local-global comparisons in visual neurons. *Annu. Rev. Neurosci.* **8**, 407–430. (doi:10.1146/annurev.ne.08.030185.002203)
101. Bi G, Poo M. 2001 Synaptic modification by correlated activity: Hebb's postulate revisited. *Annu. Rev. Neurosci.* **24**, 139–166. (doi:10.1146/annurev.neuro.24.1.139)
102. Meuwese JD, van Loon AM, Scholte HS, Lirk PB, Vulink NC, Hollmann MW, Lamme VA. 2013 NMDA receptor antagonist ketamine impairs feature integration in visual perception. *PLoS ONE* **8**, e79326. (doi:10.1371/journal.pone.0079326)
103. Flohr H, Glade U, Motzko D. 1998 The role of the NMDA synapse in general anesthesia. *Toxicol. Lett.* **100–101**, 23–29. (doi:10.1016/S0378-4274(98)00161-1)
104. Flohr H. 2006 Unconsciousness. *Best Pract. Res. Clin. Anaesthesiol.* **20**, 11–22. (doi:10.1016/j.bpa.2005.08.009)
105. Zucker RS, Regehr WG. 2002 Short-term synaptic plasticity. *Annu. Rev. Physiol.* **64**, 355–405. (doi:10.1146/annurev.physiol.64.092501.114547)
106. van Loon AM, Knapen T, Scholte HS, St John-Saaltink E, Donner TH, Lamme VA. 2013 GABA shapes the dynamics of bistable perception. *Curr. Biol.* **23**, 823–827. (doi:10.1016/j.cub.2013.03.067)
107. Keizer AW, Hommel B, Lamme VA. 2015 Consciousness is not necessary for visual feature binding. *Psychon. Bull. Rev.* **22**, 453–460. (doi:10.3758/s13423-014-0706-2)
108. Roelfsema PR. 2006 Cortical algorithms for perceptual grouping. *Annu. Rev. Neurosci.* **29**, 203–227. (doi:10.1146/annurev.neuro.29.051605.112939)
109. van Gaal S, Lamme VA. 2012 Unconscious high-level information processing: implication for neurobiological theories of consciousness. *Neuroscientist* **18**, 287–301. (doi:10.1177/1073858411404079)
110. de Lange FP, van Gaal S, Lamme VA, Dehaene S. 2011 How awareness changes the relative weights of evidence during human decision-making. *PLoS Biol.* **9**, e1001203. (doi:10.1371/journal.pbio.1001203)
111. Cohen MX, van Gaal S, Ridderinkhof KR, Lamme VAF. 2009 Unconscious errors enhance prefrontal-occipital oscillatory synchrony. *Front. Hum. Neurosci.* **3**, 54. (doi:10.3389/fpsyg.2009.054.2009)
112. Cleeremans A. 2011 The radical plasticity thesis: how the brain learns to be conscious. *Front. Psychol.* **2**, 86. (doi:10.3389/fpsyg.2011.00086)
113. Naccache L. 2018 Reply: response to 'Minimally conscious state or cortically mediated state?' *Brain* **141**, e27. (doi:10.1093/brain/awy026)
114. Naccache L. 2018 Minimally conscious state or cortically mediated state? *Brain* **141**, 949–960. (doi:10.1093/brain/awx324)