

Consciousness Explored.

A review of Blackmore, S. & Troscianko, E. T. (2018). Consciousness: An Introduction (Third Edition). Abingdon, UK: Routledge.

Nick Chater,
Behavioural Science Group,
Warwick Business School,
University of Warwick,
Coventry, CV4 7AL, UK.

Acknowledgments. Nick Chater was supported by ERC grant 295917-RATIONALITY, the ESRC Network for Integrated Behavioural Science [grant numbers ES/K002201/1 and ES/P008976/1], the Leverhulme Trust [grant number RP2012-V-022], and RCUK Grant EP/K039830/1.

Consciousness is a puzzle. Indeed, as Blackmore and Troscianko's (B&T) stimulating and highly engaging review of the field amply illustrates, it is many puzzles.

As the very outset, B&T level with the prospective reader. Any introduction to the science of consciousness raise more perplexing and mystifying questions than it can possibly resolve. Indeed, one of the charms of the book is that it introduces such puzzling and delightful questions so vividly to a general audience, who need have no prior background in psychology, philosophy, and neuroscience. But it will also be a valuable resource for students of consciousness, ranging from undergraduates, to active researchers in the field—giving a thorough and dispassionate analysis of a huge quantity of evidence, arguments, and perspectives.

Consciousness is, of course, something that we are, in a sense, all too familiar with. On the one hand, it is this very familiarity that may lead us to place excessive trust in our intuitions. Yet, on the other hand, it is surely important to crosscheck any findings from psychology or neuroscience with our everyday interpretations of our conscious experience. It is not that these everyday interpretations should be taken at face value, still less as self-certifying truths (as was once the philosophical fashion); but our interpretations of conscious life are themselves data that we need somehow reconcile with the 'science.'

With this in mind, it is welcome that B&T do not focus exclusively on the scientific study of consciousness. They also encourage readers to reflect on their "inner lives," both through trying out a set of simple phenomenological exercises (e.g., "how much am I seeing?," "do I direct my attention?"), and to consider how consciousness is described in literature, popular culture, and religion (from Virginia Woolf's "stream of consciousness" writing through to Buddhist thought). The book is also full of intriguing entry points, drawing in the reader, including quotations from scholarly texts, pictures, and cartoons, as well as the conventional text, scientific diagrams, and the helpful lists of further reading that one would expect of a more conventional textbook. I suspect that most readers of this book will hop back and forth between one topic and another, rather than steadily progressing from the first page to the last. It is just too tempting!

The breadth of the book, in another sense, also makes it an excellent introduction to the field. Most technical discussions of consciousness start by trying to narrow the scope of the problem substantially. For example, many authors aim to persuade the reader that this aspect or that aspect of the problem is really the 'key' part to the problem of consciousness, whether it be the challenge of qualia, self-consciousness, the ability self-report, engage attention, whatever it may be (indeed, I'll go a little way in that direction myself later). And then the reader is led into detailed discussion of the problem, viewed from this narrow perspective. But that familiar strategy is entirely inappropriate for an introduction, where the reader, and perhaps future researcher, needs a broad overview, rather than a theoretically-driven slice through the field. B&T cover a terrific range of perspectives. They survey different viewpoints on the nature of the philosophical puzzles that consciousness raises; the apparent limits of perception vs richness of subjective experience; the unity or disunity of consciousness; the relevance of attention, agency, free will, the self, and the function of consciousness (if any!). They integrate perspectives from neuroscience, comparative psychology, mystical experiences, altered states of consciousness, dreams and much more.

Indeed, given this broad variety of topics that B&T survey, it is natural for the reviewer to ask: which are the really crucial areas of consciousness research, and where will most important future insights

be generated? My own suspicion is that the study of perception and attention may be particularly central (indeed, this is a key starting point for the argument of Chater, 2018, so I'm hardly impartial!). This theme chimes especially with B&T's third chapter, "The grand illusion," which discusses the mystifying disconnect between the startling limitations of perceptual processing revealed by psychological experiments, and our phenomenology of the color, detail, and general richness of the perceptual world.

Some aspects of the illusion are incontrovertible. The sparsity of cone cells outside the fovea leaves us roughly color blind in most of the visual field; yet we have the sense of experiencing a fully colored environment. Similarly, the rapid diminution of visual acuity away from the fovea tells us, unequivocally, that most of the visual world can yield no more than a blurry image---we simply don't pick up high spatial frequencies in the visual periphery. Yet, looking up from my laptop now, I find myself confronted with the familiar subjective impression of an entire room, with papers, books, journals, computers, pens and mugs, all apparently present in considerable detail. Visual crowding---the phenomenon by which recognizing objects, faces, and words in the periphery can be severely impaired by interference from the clutter of other nearby objects, faces, and words---just makes matters worse (e.g., Whitney & Levi, 2011). Or consider the wonderfully illuminating tradition of work in visual search, especially associated with the path-breaking findings of the late Anne Treisman (e.g., Treisman & Gelade, 1980; Treisman & Gormican, 1988). Whatever one's precise theoretical model (and these days there are so many!), it is clear that finding a red square in a field of green squares and red triangles is a slow and effortful process. Indeed, it seems plausible that outside one or a few items that are in some sense in attentional focus, most items in the array are simply not encoded as having a specific color and shape. Indeed, as B&T point out, this viewpoint is further strengthened by results on change blindness: large changes to an image are frequently entirely unnoticed when, for example, made while the eye is in motion (Henderson & Hollingworth, 1999), or where pre-attentive cues to the change are masked when brief 'mud-splashes,' or similar patterns, are interleaved with successive presentations of an image.

According to believers that we are all subject to a "grand illusion" (Dennett, 1991, but also, for example, Blackmore, Breilstaff, Nelson & Troscianko, 1995; Chater, 2018; O'Regan, 1992) then our intuitions about our conscious experience are extremely unreliable: indeed, we are deluded about the most elementary facts of our immediate perceptual experience. We think we see a detailed, colorful world full of recognisable words objects and people; but almost none of this is loaded into a perceptual system, or represented anywhere in the brain.

Our sense of a rich visual world is, rather, the sense of *potential*: the sense that, whatever question we ask about the colour, detail, or identity of some aspect of our surroundings, that question will have, after a quick flick of the eye and re-engagement of attention, a ready answer. Indeed, the answer can be generated so fluently, that it seems instantaneous: rich information about the visual world may not be encoded in our perceptual systems (and hence, not in immediate conscious experience), but it is, as it were, at our visual fingertips.

The same may be true for the feeling of conscious awareness more broadly. For example, suppose that I ask myself why I chose a particular route to work this morning. Perhaps my sense of conscious awareness of the reason (e.g., "this is the route I always take") is nothing more than the ability to verbally formulate a plausible answer on the spot. Indeed, whatever answer I gave can, itself, be

interrogated, and so on, indefinitely. Why, I might wonder, do I always take that particular route? Perhaps because it is the quickest? So why do I prioritize speed? Perhaps because my main constraint in life is time. But shouldn't I be worried instead about minimising costs, or my carbon footprint, or reducing the risk of accidents, or some other criterion? And, in any case, how do I know that this route *is* the fastest? Of course, I can come up with lengthy responses to these and innumerable other points, on request.

Thinking about such everyday issues, then, we find that the interplay of questions and answers can continue indefinitely, in relatively quick-fire fashion. But it can't be, of course, that I have simultaneously loaded into my consciousness all possible questions of all possible answers. This implies that my sense of *being conscious of* my motivations, beliefs, desires, cannot reflect their *being present in consciousness*. As with the case of perception, perhaps my sense of being conscious of the source of my own behaviour is instead the sense that I can verbally justify my behaviour, fluently and easily, whenever required. This general viewpoint aligns with the philosopher Daniel Dennett's important analyses of conscious experience, of which B&W give an excellent treatment (e.g., Dennett, 1991, 2005).

From this point of view, to say that I am conscious of my beliefs, desires, or motivations, is to say nothing more than that I can generate, at will, a stream of conscious inner speech; and that this stream of speech provides an explanation (of sorts) of my behaviour in terms of such beliefs desires and motivations. And, on reflection, at any particular moment, surely I am really only conscious of is the *perceptual* qualities of that inner speech, rather than the beliefs, desires and motives the inner speech expresses. I experience the auditory imagery of the sounds of the words as they flow through my mind; and, conceivably, I may have some visual images to accompany them (perhaps quick flashes of motorway, country lanes, fragments of traffic jams, if I'm thinking about my journey to work). Indeed, is not the conscious flow of experience in itself always entirely perceptual?

The content of the explanations produced by our internal verbal commentary should be viewed with extreme scepticism. B&W discuss Gazzaniga's famous work on people whose left and right cerebral cortices have been divided surgically. Gazzaniga found, remarkably, that the language areas of the left hemisphere were able to fluidly "explain" actions taken by the right hemisphere, of which they could have no knowledge. Indeed, these explanations were demonstrably nothing more than plausible confabulation. Yet, as B&W point out, citing Pinker (2002), once we realise that the left hemisphere language generating system confabulates in people with split brains, we should begin to suspect that it may be confabulating just a wildly in normally functioning brains. Indeed, there is a long tradition in social psychology suggesting that our interpretations of our own mental states are, indeed, post hoc: that we seek to explain our own behaviour in terms of beliefs, desires, and motives, in just the same way that we seek to explain the behavior of other people (e.g., Albarracin & Wyer Jr, 2000; Bem, 1972).

Perhaps, then, one way of attacking the puzzle of consciousness, in the light of the huge breadth of material in B&W's survey, is to focus on the narrow problem of the conscious awareness of the output of the perceptual (and imagistic) system. Indeed, as we have seen, it may be that the scope of this problem is much narrower than we think--or at least that contents of consciousness are much sparser than we typically imagine.

One remarkable window into the astonishing narrowness of consciousness is the technique of eye tracking (e.g., Rayner, 1998). The very fact that, when we read, we roughly hop one word at a time left to right is, I think something of revelation. Similarly, the fact that when we inspect, say a family photograph, our eyes jump from face to face; and, when inspecting a face, our eyes jump from eyes to nose to mouth, and so on, is by no means introspectively obvious. We have the sense of taking in whole groups of people, interfaces, and great swathes of text, a single visual gulp; and have little to no intuitive awareness that our eyes are moving at all, let alone that they moves in discrete saccades. Gaze contingent eye tracking methods (pioneered by McConkie & Rayner, 1975), where the image can be systematically changed as the eyes move (for example, to replace or blur text that we are not directly looking at), provide further, and very direct evidence that our contact with visual world is much more tenuous than we imagine. Indeed, this, and other evidence, suggest to me that, roughly speaking, the brain is only able to “pick up” one object at a time; and we may even be limited to paying attention to one property at a time (Huang & Pashler, 2007).

Why is the flow of conscious experience so limited? One starting point for an explanation may be the distributed nature of neural computation. If solving a particular perceptual or cognitive problem is spread across one or many networks in the brain (for example, perhaps a single cognitive step corresponds to something analogous to a neural network solving a constraint satisfaction problem), then we should expect interference between problems that draw on overlapping networks. And, indeed, many complex cognitive problems, involving high-level perception, decision-making, problem-solving, or language processing, may draw on a sufficient variety of brain networks that each test task interferes with any of others.

If this is right, then perhaps among the many puzzles of consciousness that B&W survey, there are three which are particularly central underlying:

1. What corresponds to the immediate, and presumably very limited, contents of the flow of conscious experience? Here, research on attention, including eye-tracking, provides one line of attack. Another is the neural basis of conscious experience---for example, does the conscious flow of experience correspond to a bottleneck of neural activity in subcortical structures (Merker, 2007)?
2. What cognitive processes support, or undermines, the grand illusion, that so misleads us about the richness of conscious experience? Evidence may come, for example, from people with visuo-spatial neglect, who may fail to attend to fully half of the visual field often have no subjective sense of any abnormality in conscious experience (Marshall & Robertson, 1993). This might, perhaps, be analogous to having a sense that we can answer questions on a particular topic, even if the relevant knowledge has been lost, because we are never able to ask ourselves relevant questions--- that is, in the case of visual neglect, it seems impossible to shift attention to the relevant location of the visual field. Conversely, people with simultagnosia (Coslett & Saffran, 1991) seem to lose the grand illusion, and have conscious awareness of just one object at a time, perhaps because they are unable to “freely question” the rest of the visual field, perhaps because they cannot readily disengage attention.
3. Finally, there remains the “hard” problem of consciousness (Chalmers, 1995): why do the contents of consciousness have a subjective “feel” at all, and where does the quality of that feel come from.

Regarding this last, and particularly intractable, category of questions, the psychology of perception does seem to tell us something useful. For example, subjective sensations appear to correspond, in general, to the inferred properties of the external world rather than sensory input (e.g., color and lightness signal inferred surface reflectance functions of external objects, rather than the flux of light entering the eye). Relatedly, subjective sensations appear largely to be coded relatively rather than in absolute terms (e.g., turning up the lights in a room has little on the perceived lightness of the objects within it). But why are there any subjective sensations of anything?

B&W note that there are a number of sophisticated sceptics about the very existence of the “hard” problem--- suspecting that its intractability arises from confusion, rather than depth (e.g., Churchland, 1996; Dennett, 2001; O’Hara & Scott, 1996). Perhaps strangely, one of the most powerful intuitions of the problem just won’t go away comes from ethics: it seems that there must be a crucial question concerning which creatures suffer pain (and, perhaps, one day, which robots suffer pain, although as something of an AI sceptic, I suspect this question will not be pressing for many centuries). The question of which creatures suffer pain seems to crucially direct how we treat our fellow humans and nonhuman animals; and does not seem to be a matter we can set aside, by blithely claiming that the very question of subjective experience is a pseudo-problem. Hardly a rigorous argument perhaps; but it is the consideration that makes me unable to set aside the very existence phenomenal experience as just one more layer of the grand illusion.

The 600 pages of B&W’s delightful, balanced, and wide-ranging introduction to consciousness cover far more than has been touched on in this review, of course. I hope it is clear that, for anyone interested in plunging into the murky but entrancing depths of the psychology, philosophy, and neuroscience of consciousness, B&W’s book is the perfect place to start. It is, moreover, a stimulating and challenging resource for specialists in the field.

References

- Albarracín, D., & Wyer Jr, R. S. (2000). The cognitive impact of past behavior: influences on beliefs, attitudes, and future behavioral decisions. *Journal of Personality and Social Psychology, 79*(1), 5-22.
- Bem, D. J. (1972). Self-perception theory. In L. Berkowitz (Ed.), *Advances in Experimental Social Psychology*, Vol. 6 (pp. 1-62). New York: Academic Press.
- Blackmore, S.J., Breilstaff, G., Nelson, K., Troscianko, T. (1995). Is the richness of our visual world an illusion? Transsaccadic memory for complex scenes, *Perception, 24*, 1075–81.
- Chalmers, D. J. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies, 2*(3), 200-219.
- Chater, N. (2018). *The Mind is Flat*. London, UK: Penguin Allen Lane/New Haven, CT: Yale University Press.
- Churchland, P. S. (1996). The Hornswoggle problem. *Journal of Consciousness Studies, 3*, 402-408.
- Coslett, H. B., & Saffran, E. (1991). Simultanagnosia: To see but not two see. *Brain, 114*(4), 1523-1545.

- Dennett, D. C. (1991). *Consciousness Explained*. Boston, MA: Little, Brown & Co.
- Dennett, D. C. (2001). Are we explaining consciousness yet? *Cognition*, 79, 221-237.
- Dennett, D. C. (2005). *Sweet dreams: Philosophical obstacles to a science of consciousness*. Cambridge, MA: MIT Press.
- Henderson, J. M., & Hollingworth, A. (1999). The role of fixation position in detecting scene changes across saccades. *Psychological Science*, 10(5), 438-443.
- Huang, L., & Pashler, H. (2007). A Boolean map theory of visual attention. *Psychological review*, 114(3), 599-631.
- Marshall, J., & Robertson, I. (Eds.). (1993). *Unilateral Neglect: Clinical and Experimental Studies*. Hove, UK: Erlbaum.
- McConkie, G. W., & Rayner, K. (1975). The span of the effective stimulus during a fixation in reading. *Perception & Psychophysics*, 17(6), 578-586.
- Merker, B. (2007). Consciousness without a cerebral cortex: A challenge for neuroscience and medicine. *Behavioral and Brain Sciences*, 30(1), 63-81.
- O'Hara, K., & Scutt, T. (1996). There is no hard problem of consciousness. *Journal of Consciousness Studies*, 3, 290-302.
- O'Regan, J. K. (1992). Solving the "real" mysteries of visual perception: the world as an outside memory. *Canadian Journal of Psychology*, 46, 461-88.
- Pinker, S. (2002). *The Blank Slate: The Modern Denial of Human Nature*. New York: Viking.
- Rayner, K. (1998). Eye movements in reading and information processing: 20 years of research. *Psychological Bulletin*, 124(3), 372-422.
- Simons, D. J., & Levin, D. T. (1997). Change blindness. *Trends in Cognitive Sciences*, 1(7), 261-267.
- Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12(1), 97-136.
- Treisman, A. M., & Gormican, S. (1988). Feature analysis in early vision: evidence from search asymmetries. *Psychological Review*, 95(1), 15-48.
- Whitney, D., & Levi, D. M. (2011). Visual crowding: A fundamental limit on conscious perception and object recognition. *Trends in Cognitive Sciences*, 15(4), 160-168.