

Out with folk psychology, in with what?

Review of *The Mind is Flat: The Remarkable Shallowness of the Improving Brain*, by Nick Chater. London: Allen Lane, 2018. 272 pp, hardcover.

Susan Blackmore and Emily Troscianko, 2019.

University of Plymouth and University of Oxford, susan.blackmore@plymouth.ac.uk, emily.troscianko@humanities.ox.ac.uk.

We loved this book, at least the first, and shorter, of its two parts. Here Chater boldly and skilfully shows how deluded we are about our perception, our emotional lives, and our decision-making faculties. Building on ideas usually restricted to studies of perception, he demolishes fondly held folk-psychological ideas about the depths of the mind. His thesis is simple if stark: we find it hard to plumb our mental depths not because they are so deep and murky, but because there are no mental depths to plumb. All those popular ideas about the subconscious and the unconscious, hidden motivations, suppressed fears and buried hopes, beliefs and desires have to go. We must abandon wholesale everything we think we know about the operation of our own minds, and in doing so we will discover that the surface of the mind does not lie above a rich swirling mass of unconscious thought and action, but that the surface – the ‘flat’ of the title – is all there is: ‘common-sense psychology *isn’t true*’ (p. 14).

Chater’s argument expands on the theory that vision is a ‘grand illusion’. This idea originated with the phenomenon of change blindness, in which even large changes to an image can be missed if they happen during a blink or when the image is moved. In one famous example, a military aircraft sitting on the tarmac is swapped for one without an engine, and people viewing alternating versions with a blank screen between them sometimes need up to 40 alternations to spot the difference. This implies not that the world we see is illusory, but that the richness of our perceptions is: we are wrong to believe that somewhere in our minds or brains there must be a rich and detailed picture of the outside world. These and many other experiments showed that no such inner picture is needed: most of the world’s detail remains where it always was, out in the world, and we get the illusion of having it inside our minds because we can always look again – just in time.

Chater doesn’t discuss change blindness, although he does talk about the closely related phenomenon of inattention blindness, and he gives other examples from areas like colour vision and text processing to show just how fragmentary our visual access to the world is. He also uses other classic perceptual phenomena like retinal stabilisation to show just how much vision depends on active interpretive creation. You might expect that if an image were kept completely still on your retina you’d just keep seeing the still image. In fact, within a few seconds the perceptual experience fades to uniform grey or black, and then parts of the image appear again, reorganise, and disintegrate again, in the effort to create meaning even out of arbitrary shapes.

'The brain is continually churning: despite the unaccustomed lack of new input, the brain is desperately attempting to disengage from the current organization, and to find another. When it cannot, the image entirely disappears' (pp. 50-51).

With examples like this, Chater proposes that the richness of our unconscious minds is just as illusory as the detailed mental pictures that are wrongly thought to be the basis of visual perception: 'the rich mental world we imagine that we are "looking in on" moment-by-moment, is actually a story that we are inventing moment-by-moment' (p. 14). Gathering evidence from such varied sources as *Gormenghast*, *Anna Karenina*, and *The Simpsons*; experiments on political opinions, crossing scary bridges, and inattentive deafness; and thought experiments on imagining a tiger's stripes, predicting what happens to coffee, ball bearings, and sugar when dropped on the kitchen floor, and how similar it feels to be frightened versus excited, Chater makes the point that our own beliefs, desires, hopes, and fears are just as much inventions as those of our favourite fictional characters. Thus we fall for the 'illusion of explanatory depth'. Dreams are likewise sketchy inventions made on the fly and not a royal road to the depths of the psychoanalyst's unconscious. Emotions, or at least the way we label and try to make sense of them, are fictions too: 'joy or anger do not well up from our inner depths' (p. 95), for there are no inner depths. Instead what we call our emotions are the results of in-the-moment interpretations based on the situation we are in and highly ambiguous evidence from our own bodily state.

For Chater, mystical and drug-induced states are likewise streams of invention and not voyages of inner discovery. As a long-term meditator and explorer of altered states, Sue would hate to think this means that any insights she might find in such states must be illusory. But this is not the only possible interpretation of Chater's view. Perhaps valid insights about the mind can arise from watching the surface – the ever-inventive flat mind – as long as we don't fall for the illusion that we have uncovered a rich inner store. Indeed, mystics and meditators often describe the conscious self disintegrating, inner and outer becoming one, or the realisation that there is 'only this moment' and the rest is invention – surely insights that Chater would appreciate.

Part Two concerns the 'improvised mind' and Chater's theory of the 'cycle of thought', and this, for us, was far less convincing. Contrasting human brains with fast digital computers, Chater notes that our brains depend on cooperation across a vast number of slow neural processing units, with the cooperation spanning whole networks or even entire regions of the brain. From here, he says, it is hard to see how so many interconnected neurons could coordinate on processing more than one thing at a time, and so he concludes that any given problem is split into tiny fragments to be dealt with across the entire, densely interconnected network. This implies that the networks make one giant, coordinated step at a time, running at several beats per second. This he calls 'the cycle of thought'.

The implication he draws is that any network can work on only one problem at a time so multitasking must be a rarity. Further, not only are we restricted to attending consciously to only one problem at a time but we cannot, even unconsciously, be thinking about another. There can be no background processing or clever unconscious thoughts going on: the brain circuits that would be needed for any such thoughts are 'blocked' by the conscious brain processes (p. 131). And the idea of the grand illusion is applied here too. It tricks us into believing that our focus of attention is far wider than it really is when in fact 'we can manage, roughly speaking, one thought at a time' (p. 143).

But is it true that we cannot have two thoughts at once? First of all, Sue tried to explore this question by paying attention to her own thinking. This is what she found:

At first the answer seemed to be 'no'; every time I caught a thought in progress it was overtaken by the observation 'I was thinking about x but now I'm thinking about the question'. However, after some practice the answer seemed to be 'yes'. For example, when cleaning my teeth and thinking about the structures of internet memes, I had the thought that two minutes must be nearly up, and I turned on the tap. The thought about memes continued, apparently without pause. But perhaps the two alternated too rapidly for me to detect, leaving Chater's claim intact.

Next, I tried deliberately creating two thoughts at once. For example, walking on a rough moorland track I paid attention to my feet walking while thinking about whether two thoughts are possible at once. This was easy, but wondering whether such simple observation does not count as thinking, I counted my steps, still thinking about thinking. This, too, was easy, but perhaps counting does not count as thinking either. Then suddenly I realised that the tune of 'The holly and the ivy' seemed to have been going on in my head for some time. Was this an unconscious thought continuing along with the counting and questioning? Presumably Chater would dismiss this possibility as a post-hoc invention, although it seemed reasonable to think that the song began when I passed a large holly tree some minutes before.

The most obvious question raised by Chater's rejection of unconscious thinking is: what counts as a thought? Happily, he gives us an answer: perceiving is a type of thinking and 'all other types of thought are really just extensions of perception' (p. 51). His reasoning is that because the sub-cortical brain areas he has concluded are the 'locations of conscious experience' (p. 141) are sensory relay stations, consciousness must be sensory experience. This requires him to argue that there is no such thing as a non-sensory thought, which also requires that language be brought in under the umbrella of the sensory. Given the scope of that task, he devotes oddly little effort to it. He claims that as soon as we're aware of the number 5, we're aware of a sensory representation of the shape of the symbol '5', or five

dots, or whatever – i.e., we can be conscious only of the sensory consequences of abstractions, not abstractions themselves.

Thinking through language, then, is just being conscious of ‘sensory impressions of the sound of things we might say’ (p. 183); and being conscious of a fact (about an abstract concept, say) is being conscious of ‘abbreviated snippets of English, running through our minds’ (p. 184). But this presupposes that the snippets of English we’re conscious of are – like numerical concepts – from the outset sensory in nature, and this is an open question. There’s evidence that processing of, for example, metaphorical and literal verb phrases (*grasp the idea* versus *grasp the scissors*) involves premotor activity only in the literal case, suggesting at least a metaphorical–literal or abstract–concrete continuum. And of course the neurophysiology of language processing is not the same as the experience of language processing: conclusions from one may contradict conclusions from the other. The evidence is complex, and deserves consideration.

Giving ‘conscious thought’ the broad definition of ‘extensions of perception’ also leaves Chater in the odd position of having to conclude that if we are perceiving the road we’re walking along while thinking about thinking, or while talking to a companion, it’s the thinking or the talking that counts as the thought, even though the perception of the road is the more obviously sensory activity. But if perception in this case is not conscious thought, doesn’t it revert to being the standard unconscious ‘background processing’ he said did not exist? Alongside the first-person inquiry and the problems with language being exclusively sensory, this makes for another reason to question whether it’s really impossible to have two thoughts running along at once.

We recognise the severe limits of introspection, but there is no technology (yet) which could let us observe individual thoughts coming and going. Our best guess is that thinking of many kinds does go on at once, but that introspection can easily mislead us into thinking otherwise. Just as whenever we ask ourselves ‘Am I conscious now?’ the answer is always ‘yes’, so when we ask, ‘How many thoughts am I having now?’ the answer is always ‘one’ because asking that question brings together any strong enough threads of thought into one new combined thought and ignores the rest. This is a salutary reminder that attending to the realities of our experience takes practice. We await any possible objective answer to the question with impatience!

Fortunately, the strong ‘one thought at a time’ claim is not critical to Chater’s wider thesis about the flatness of the mind, and his wholesale rejection of folk beliefs stands. Even if it turns out that several neural networks regularly operate at once, this need not imply the existence of an unconscious mind populated with hidden motives and intentions – it is just neurons doing what neurons do. This is the great contribution of Chater’s book: to see how much of folk psychology we can do

without, and how much better our understanding of the human mind could be without it.

Meanwhile, however, the underlying question of what thought is raises other questions about the new theory Chater replaces the folk intuitions with. For Chater, thinking is interpreting. Interpretation is one of the key concepts on which his thesis hangs, but as Part II of the book kicks off, there begins a series of extremely broad and often mutually contradictory statements about interpretation, along with two other crucial concepts: attention and consciousness. First, Chater says that '*attention is the process of interpretation*' (p. 140) and 'our only conscious experience is our interpretation of sensory information' (p. 141). This seems nicely clear: attention = interpretation = consciousness. But later we learn that 'without attention, there is no interpretation, analysis or understanding' (p. 156), so now attention seems to be a prerequisite for interpretation, not identical with it. And later, when explaining how we make sense of sensory inputs, he says that 'conscious experience is merely the output of this remarkable process' (p. 143). So now maybe we have a causal chain: attention → interpretation → consciousness. But he also says that 'all conscious thought concerns the meaningful interpretation of sensory information' (p. 141). By saying 'concerns' he seems to mean that consciousness is *about* interpretation (or attention) rather than caused by it. And later, 'the meaning that we impose in each step in our cycle of thought corresponds with the contents of our stream of consciousness' (p. 153). So now the *contents* of consciousness are the results of interpretation: consciousness has been split up into container and contained.

Further confusion ensues when Chater adopts a bottleneck theory of attention without mentioning the many drawbacks of such theories, and even claims to locate the bottleneck in sub-cortical structures: 'Here, perhaps, somewhere in these deep brain structures, is the crucial bottleneck of attention; and whatever passes through the bottleneck is consciously experienced' (p. 134). This is a rephrasing of his idea that 'conscious' brain circuits 'block' any potential sophisticated 'unconscious' thoughts. It suggests that consciousness (or maybe its contents) is actually the information that makes it to a certain point in a particular neural pathway – a point which is later also described as a gateway structure and a crucible. These sub-cortical structures 'may determine both the contents of the "flow" of consciousness and, indeed, whether we are conscious at all' (p. 136).

Another obvious problem with the 'one at a time' claim is that it would seem to force Chater to claim that we can't, in his example, walk and chew gum while listening to a conversation – or, in Sue's example, walk and think about thinking. He creates a loophole to wriggle through by distinguishing between interpretation proper and the kind of processing of sensory information that's more automatic and less 'meaningful'. Walking doesn't involve *interpretation* because we're not trying to make sense of what's in our mental 'focus' (p. 146). But this seems circular: the sense-

making involved in navigating a route is 'mindless' and 'automatic' and not real interpretation because it's not what we're attending to, so can't be real interpretation. Again, 'no background processing in the brain' actually just seems to reduce to: background processing yes, but no clever stuff, and (thanks to a circular argument) the not-clever stuff can't be conscious. And this is more or less where every previous attempt to distinguish conscious from unconscious neural processing has ended up.

So in the end, we are left not knowing what Chater means by any of these three core terms, *consciousness*, *attention*, and *interpretation*, and not knowing whether they're all meant to be basically the same thing or in some causal relationship to each other or defined with reference to brain activity or brain structures or something else. This is already a classic problem in discussions of the relation of consciousness to attention, and introducing interpretation – another great monolith of cultural and cognitive analysis – into the mix doesn't help matters.

Part of the wider problem is perhaps that Chater assumes that the thinking (and the attending, the interpreting, the being-conscious) is all done by the brain. So the mereological fallacy (of imputing to part of an organism a function which is performed by the whole) is all over the place: our brains are interpreting, attending, playing tricks on 'us' (who and where 'we' are he doesn't explore). And so the question he is trying to answer, in Part II, becomes a rewrite of Chalmers' classic brain-based Hard Problem: 'how electrical and chemical activity in our neural circuits can somehow generate our stream of thought' (p. 35). Trying to solve this problem has arguably created far more new problems than solutions, and it's an oddly traditional problem for someone who wants to overturn so much else. It also leads to traditional attempts at solutions: the common search for the neural correlates of consciousness or the neural location of consciousness are other fashions Chater falls in line with.

In this vein come three further intuitions that Chater adopts without challenge: that there is obviously something it means for a thought or sensation or emotion to be conscious as opposed to unconscious, and that consciousness is a stream, and has contents. The arguments about all three are complex, but Chater doesn't engage with them, taking these 'facts' for granted. So he urges us to 'revise our intuitions about "unconscious thoughts"' (p. 131), but not to the extent of questioning the underlying assumption that all brain processes and thoughts must be either conscious or unconscious. Likewise, the contents of the stream of consciousness are the meanings we derive from each step of our cycles of thought. And in the end he doesn't reject the idea of the unconscious, either, but redefines it in strangely conventional terms: 'This is the real nature of the unconscious: the vastly complex patterns of nervous activity that create and support our slow, conscious experience' (p. 175).

Naturally, this comes back round to asking 'What, then, are the underlying (and unconscious) calculations that our brain networks are carrying out in order to

generate the conscious experience'? (p. 177). And in these mentions of streams and contents and *underlying* calculations, it's strange to observe, as the second part of the book unfolds, how the depths creep right back in despite the all first half's valiant attempts to flatten the mind. Human beings end up characterised as 'like corals layered, polyp by polyp, into infinitely diverse forms', with our thoughts like water droplets flowing from the high ground to the sea, 'each droplet cut[ting] those channels just a little more deeply' (p. 202).

This linguistic equivocation matters when so much of the book's argument not only depends on the concept of interpretation, but uses that concept specifically to deconstruct an age-old *metaphor* of depth. Chater's argument, that language is part of the wide swathe of sensory processing that is 'thought', fits well with George Lakoff and Mark Johnson's *conceptual metaphor theory*, according to which many of our most common linguistic forms derive directly from our sensorimotor realities, and linguistic forms in turn help determine how we live and think about our lives: *love as journey* versus *love as heat* create quite different possibilities for experience and understanding. In the first half, Chater seems to agree that the sneaky metaphors of depth matter, and near the end of the book he includes a section on 'The Ubiquity of Metaphor', citing Lakoff and Johnson and reminding us that 'metaphors are just as apt to *mislead* as they are to inform' (p. 210). But by returning nonetheless to depth as a quality of thoughts, concepts, and processing, and using *inner-outer* constructions repeatedly too, he nudges our readerly processing back into conceptual channels that preserve all the associations so actively broken down earlier. If the purpose of the book is to overturn the status quo, and the status quo is crystallised in the turns of phrase that shape, and are shaped by, how we think about our own minds, then we need more help in resisting the old ways, rather than being allowed to fall back into them.

Our take on this is influenced by Emily's background in literary studies, a field that Chater dismisses as uselessly relativist. His sketch makes it sound as though everyone who works in the humanities is a late 20th-century postmodernist nihilist. Yet lots of humanities scholars rejected that fashion too, and the *cognitive humanities* is one of the results. The cognitive humanities reject limitless relativism, without rejecting the idea that the methods and insights of the humanities can be relevant to understanding the mind, nor indeed that disciplines devoted to the mind can help us understand texts better, including how they're constructed and processed. Chater argues that mental activity is interpretive activity, and his opening move is to provoke us as readers into comparing what we do with literary characters to what we do with flesh-and-blood people. So texts that are created to be interpreted and the rest of a world that is endlessly interpretable are all the same kind of stuff. Cognitive literary studies, the formal study of interpretive activity grounded in complex verbal texts, would therefore seem like a helpful contributor to Chater's project.

But in the end, there's a second failure of nerve alongside the metaphor-manifested one: Chater ultimately seems worried that if we explain too much, the wonder of the human mind will be ruined. He rejects the search for an ultimate purpose that lies 'beyond our everyday understanding, or perhaps deep inside our innermost core', but also rejects its opposite: 'that no such transcendental meaning exists and that human life is no more than a brief, purposeless biochemical agitation at one edge of a vast and lifeless cosmos' (p. 191). Chater considers this a mistake that leads often to despair, but to us it sounds rather true and beautiful, and makes us think of the long line of thinkers and explorers of the human condition who have come to more or less this conclusion: from ancient Chinese daoists such as Zhuangzi, to modern philosophers like Julien de La Mettrie and his *L'Homme machine*, to contemporary thinkers like Sam Harris. Their consensus (which we agree with) seems to be that even if you despair at first, you can get through the despair and lead a full, happy, and moral life, living more calmly once you let go of the search for non-existent purposes.

For Chater, however, life is not purposeless, because it's all about the attempt to create meaning. In precisely what sense this makes human life purposeful is open to debate. But Chater treads a tricky line between (in most of the book) treating interpretation as something that is so fundamental to every single cognitive act that it's as automatic as breathing, and (very late on) trying to make this interpretive capacity into a strike for human freedom. So now the compulsive pareidolia that means we *can't help* seeing faces in pieces of toast is transformed into something we *can* do, and so into an emancipation from 'blind repetition' (p. 202), or into a bastion for 'purposeful' reshaping of our futures. The final sentence of the last chapter reads: 'If imagination and metaphor is the secret of our intelligence, then that secret may, perhaps, be safely locked away in the human brain for centuries and perhaps for ever' (p. 219). This is strangely reminiscent of the obscurantist humanities he so viciously lampooned back in Part I. It also runs counter to a book whose purpose has been to expose so many misconceptions in 'the project of understanding ourselves through the application of science' (p. 34).

This book makes an excellent stab at something really difficult – just how difficult is made clear by the knots the book ties itself in while trying. Our sense is that the first half alone, published as a short, potent, self-sufficient demolition of untenable intuitions, would have been a more forcefully unignorable contribution to psychology and consciousness studies. But the problems with the second half don't change the fact that this will be a valuable and provocative read for anyone interested in inquiring into 'common-sense psychology', which everyone who studies consciousness should be.