
THEORIES OF CONSCIOUSNESS ARE SOLUTIONS IN NEED OF PROBLEMS

Stephen M. Fleming

Wellcome Centre for Human Neuroimaging
and Department of Experimental Psychology
University College London
stephen.fleming@ucl.ac.uk

November 12, 2020

ABSTRACT

Doerig et al. point out a number of shortcomings with existing theories of consciousness and argue they should be systematically constrained by empirical data. In this commentary I suggest a further constraint - the potential functions of (the contents of) consciousness. One such candidate function in humans is the social sharing of reportable mental states. The social function of consciousness provides a general framework within which to understand the evolution and neurobiology of conscious awareness.

A common approach in consciousness science is to start with subjective experience and then attempt to bridge an “explanatory gap” from experience to theoretical models or neural substrates. It is becoming increasingly clear that this approach to the problem of consciousness is in danger of losing its way, as nicely highlighted by Doerig’s target article. They point out a number of shortcomings with existing theories of consciousness, and in particular emphasise the need for “hard criteria” on what constitutes a useful theory. In this commentary I suggest a further constraint on these criteria – the functions of (the contents of) consciousness.

A useful framework for thinking about the function of computational systems is David Marr and Tomaso Poggio’s influential taxonomy of three distinct levels of description [1]. A computational level characterises the problem to be solved in terms of available inputs and desired outputs. An algorithmic level then specifies how this computation is performed – as a system of mathematical equations, lines of code or diagrams of interacting variables. Finally, an implementational level specifies how our chosen algorithm is implemented in hardware or wetware – for instance, how a computer program is implemented in the machine code of your laptop, or how dark adaptation is achieved by cells in the retina.

The definition of Marr and Poggio’s uppermost level as “computational” can be confusing to contemporary readers as the term computation is now often used to refer to different algorithms (such as reinforcement learning) rather than the problem to be solved (seeking food rewards). For the purposes of this commentary, let’s define this level as one identifying a putative (evolved) function, or the functional level for short. The paradigm cases of consciousness highlighted by Doerig et al. (such as visual masking) represent a significant step towards identifying such functions, as they aim to cleanly separate conscious and unconscious alternatives – such as being able to report we are experiencing contents X , versus not being able to do so. However, for the most part they leave the functional level unexplored – why might it be useful for a system to report that it is aware of X ?

The search for the functions of consciousness has often led to dead ends. For instance, perceptual responses, linguistic analysis, number processing and even executive functions may be triggered non-consciously [2], and the famous Libet experiments and phenomena such as choice blindness imply that consciousness is merely a post-hoc rationaliser of what we do or say [3, 4]. But just because this search has ruled out some promising candidates for the functions of consciousness does not mean there are none to be found. As Doerig et al. conclude: “Maybe consciousness is a ‘solution’, a by-product, or a core component of a computational challenge that information processing systems need to solve – and that we have not discovered yet.”

One candidate function is that consciousness is for sharing [5]. Unconscious mental states might be sufficient to get by on our own. But by definition, those mental states that can be shared with others need to be consciously reportable. Consider two hunters stalking a deer. They might wish to share information about what they are each seeing, both to pool their perceptual resources and to provide “common ground” for future coordinated action [6]. Experiments show that when two individuals share their confidence in their percepts, they can arrive at a joint decision that is better than that of the best individual working alone – a form of interpersonal Bayesian cue-combination [7].

Identifying candidate functions for consciousness can constrain theorising on its algorithms and neural substrates. For awareness reports to be socially useful they need to be flexible (we can communicate awareness of a deer, a lion, or some berries with equal fluency) and simple (occupying an abstract, low-dimensional space ranging from unaware to aware). Accordingly, linguistic analysis suggests that pairs of naïve observers rapidly hit upon a common confidence scale with which to communicate what they are seeing in psychophysical experiments [8]. Recent proposals including higher-order state space (HOSS) and predictive global neuronal workspace (PGNW) models have begun to flesh out how high-dimensional (but potentially unconscious) perceptual representations “connect up” to the capacity for such reports [9, 10, 11]. We can also make informed guesses about when the functionality for sharing mental states became useful during evolution. For instance, computer simulations show that the enhanced visual range of terrestrial compared to aquatic animals may have driven the need for planning based on the perception of distant objects [12]. Finally, functional criteria that constrain algorithms for awareness can help in interpreting candidate neural substrates (the implementational level). For instance, neural recordings in primates indicate that the dorsolateral prefrontal cortex (area 46) contains cells that show stimulus-independent representation of stimulus presence and absence – consistent with encoding a low-dimensional axis of awareness [13, 14].

As Doerig et al. highlight, a search for algorithms that are unmoored from function (such as causal structure theories) is liable to lead to an unconstrained exploration of often beautiful but potentially meaningless mathematics. A renewed focus on function instead aligns naturally with recent illusionist approaches to consciousness [15, 16]. Conscious awareness can sometimes seem like magic, appearing from nowhere out of a physical system. But a scientific explanation of a magic show cannot be given in terms of magic itself, but in terms of why the audience believe that magic is a plausible explanation of a trick, and why they go on to amazedly tell their friends about it. In the same way, rather than fret about the mystery of intrinsic subjectivity, the illusionist attempts to explain what people believe (and can communicate) about their experience using the lingua franca of cognitive science – a science of both the trick and the audience. Theories of consciousness constrained by social function can go one step further, and explain why putting on the show might have been useful in the first place.

Acknowledgements

This work was supported by a Philip Leverhulme Prize from the Leverhulme Trust and a Wellcome/Royal Society Sir Henry Dale Fellowship (206648/Z/17/Z).

References

- [1] David Marr and Tomaso Poggio. From understanding computation to understanding neural circuitry. 1976.
- [2] Simon van Gaal, Floris P De Lange, and Michael X Cohen. The role of consciousness in cognitive control and decision making. *Frontiers in human neuroscience*, 6:121, 2012.
- [3] Benjamin Libet, Elwood W Wright Jr, and Curtis A Gleason. Preparation-or intention-to-act, in relation to pre-event potentials recorded at the vertex. *Electroencephalography and clinical Neurophysiology*, 56(4):367–372, 1983.
- [4] Petter Johansson, Lars Hall, Sverker Sikström, and Andreas Olsson. Failure to detect mismatches between intention and outcome in a simple decision task. *Science*, 310(5745):116–119, 2005.
- [5] Chris D Frith. The social functions of consciousness. *Frontiers of Consciousness: Chichele lectures, L. Weiskrantz and M. Davies, eds.(Oxford: Oxford University Press)*, pages 225–244, 2008.
- [6] Andreas Roepstorff and Chris Frith. What’s at the top in the top-down control of action? script-sharing and ‘top-top’ control of action in cognitive experiments. *Psychological Research*, 68(2-3):189–198, 2004.
- [7] Bahador Bahrami, Karsten Olsen, Peter E Latham, Andreas Roepstorff, Geraint Rees, and Chris D Frith. Optimally interacting minds. *Science*, 329(5995):1081–1085, 2010.
- [8] Riccardo Fusaroli, Bahador Bahrami, Karsten Olsen, Andreas Roepstorff, Geraint Rees, Chris Frith, and Kristian Tylén. Coming to terms: Quantifying the benefits of linguistic coordination. *Psychological science*, 23(8):931–939, 2012.

- [9] Yoshua Bengio. The consciousness prior. *arXiv preprint arXiv:1709.08568*, 2017.
- [10] Stephen M Fleming. Awareness as inference in a higher-order state space. *Neuroscience of Consciousness*, 2020(1):niz020, 2020.
- [11] Christopher J Whyte and Ryan Smith. The predictive global neuronal workspace: A formal active inference model of visual consciousness. *bioRxiv*, 2020.
- [12] Ugurcan Mugan and Malcolm A MacIver. Spatial planning with long visual range benefits escape from visual predators in complex naturalistic environments. *Nature communications*, 11(1):1–14, 2020.
- [13] RE Passingham and HC Lau. Acting, seeing, and conscious awareness. *Neuropsychologia*, 128:241–248, 2019.
- [14] Katharina Merten and Andreas Nieder. Active encoding of decisions about stimulus absence in primate prefrontal cortex neurons. *Proceedings of the National Academy of Sciences*, 2012.
- [15] Keith Frankish. Illusionism as a theory of consciousness. *Journal of Consciousness Studies*, 23(11-12):11–39, 2016.
- [16] Michael S. A. Graziano. Consciousness engineered. *Journal of Consciousness Studies*, 23(11-12):98–115, 2016.