State or content of consciousness?

A Reply to Valdas Noreika

Wolf Singer

An attempt is made to distinguish between brain states required to support consciousness and the neuronal underpinnings of conscious versus non-conscious processing in an awake, attentive brain, respectively. It is argued that brain states supporting consciousness are characterised by high dimensional dynamics exhibiting a high degree of complexity, implying that conscious states are graded. Different mechanisms determine whether signals are processed at the conscious or subconscious level. Thus, there is no unique neuronal correlate of consciousness.

Keywords
Brain dynamics | Complexity | Conscious processing | Conscious state | Content of consciousness | Dimensionality

Author

Wolf Singer
w.singer@brain.mpg.de
Max Planck Institute for Brain Research (MPI)
Frankfurt a. M., Germany

Commentator

Valdas Noreika
valdas.noreika@mrc-cbu.cam.ac.uk
Medical Research Council
Cambridge, United Kingdom

Editors

Thomas Metzinger
metzinger@uni-mainz.de
Johannes Gutenberg-Universität
Mainz, Germany

Jennifer M. Windt
jennifer.windt@monash.edu
Monash University
Melbourne, Australia

1 Introduction

My sincere thanks go to Valdas Noreika for having identified with succinct clarity the weaknesses in our current attempt to identify the neuronal correlates of consciousness (NCC). I would have sincerely appreciated these comments before finalising my manuscript, as they would have forced me to distinguish more clearly between the neuronal underpinnings of the conscious state and the neuronal correlates of conscious versus unconscious processing.

Noreika is absolutely right in pointing out that the search for the mechanisms permitting access to conscious processing falls short of identifying the NCC proper and, likewise, that the determination of variables required for the maintenance of a conscious state is insufficient if pursued without considering the contents of conscious processing. The mere fact that one can distinguish between the "conscious state" and the conditions required for "conscious processing", yet also consider both as targets in the search for the NCC, suggests that the explanandum is ill-defined. Presently, both studies devoted to the distinction between conscious and unconscious processing and those investigating the brain states required for conscious processing are considered as investigations of the NCC, although they clearly target different neuronal mechanisms. Thus, studies on consciousness are fraught with the problem of a lack of a clear definition of "the" consciousness for which we wish to find a neuronal correlate. Another problem is that we are still far from fully understanding the neuronal mechanisms underlying higher cognitive functions. Behavioural studies suggest, for example, that perception involves probabilistic Bayesian-matching operations in which sensory evidence is compared with stored knowledge about the probability of occurrence and the features of the respective perceptual objects. However, it is entirely unknown where and how the huge amount of priors are stored, how the specific priors can be retrieved on the fly within the few hundreds of milliseconds sufficient for recognition, and how the matching operations are realized in neuronal networks. Thus, at the present stage it is even impossible to precisely define the signatures of neuronal activity that could be considered the result of a perceptual process or as the neuronal representation of a percept.

In the light of these uncertainties, the distinctions between conscious and unconscious processing or between states compatible with conscious and unconscious processing, respectively, appear to be exploited primarily in order to learn more about mechanisms underlying pattern recognition, decision making, and intentionality, rather than serving the search for the neuronal underpinnings of the ill-defined phenomenon that we address as "consciousness". In contrast to NCC research, these more humble approaches have been quite successful, probably because the explananda are well-defined and can be operationalised.

2 The conscious state

The analysis of the neuronal prerequisites required for the maintenance of consciousness has a long history and has only recently been considered part of consciousness research. The reason for this is that the criteria used for distinctions between conscious and non-conscious states or altered states of consciousness can be tested in both humans and animals. Examples of these criteria are reactivity to noxious stimuli, the ability to move intentionally, and the ability to accomplish a number of well-defined cognitive tasks, involving attention, short and long term memory, recognition, and decision making. Thus, the plethora of studies performed both on animals and humans on the neuronal mechanisms underlying arousal, attention, wakefulness, sleep, anaesthesia, and coma all contribute to our understanding of the neuronal prerequisites of states permitting conscious processing. Accordingly, it is well-established that brain functions characteristic of the conscious state require that neuronal networks operate in a critical dynamical range. This range is regulated by half a dozen globally-acting modulatory systems that originate in deep and evolutionary ancient brain structures. The adjustable neuronal parameters are essentially the balance between excitatory and inhibitory effects and the time- and length constants of dendritic integration. These adjustments lead to marked modifications of the system's dynamics. These modulations are reflected by changes in the prevailing frequencies of oscillatory activity, the degree and spatial granularity of synchronisation (also addressed as correlation length), and the propagation of signals across the network.

Classical brain theories have not attributed much attention to the significance of these dynamic variables for processing and assume that loss of consciousness in sleep and anaesthesia is essentially due to reduced excitability and signal transmission. However, in more recent theories, brain dynamics are thought to play a crucial role in information processing. This novel framework provides much more specific explanation of the breakdown of consciousness in sleep, anaesthesia, and coma. These theories posit that oscillations and the concomitant variables, such as synchronisation, phase locking, phase relations, and cross frequency coupling, are relevant for signal selection by attention, binding operations, and the representation of nested semantic relations (for review see Singer 1999; Buzsáki et al. 2013). In addition, these complex dynamics have been proposed as a substrate for the generation of the high-dimensional coding space required for the storage and superposition of priors, the matching of stored information with sensory evidence, and the segregation of patterns for classification (for review see Singer 2013). The basis of these operations is the transformation of low-dimensional input patterns into high-dimensional dynamic states, in order to perform the necessary computations in this space and to then retransform the results into low-dimensional output signals. The advantages of performing computations in high-dimensional dynamic space are currently explored in the conceptual framework of "reservoir computing" or "liquid state or echo state machines" (Bertschinger & Natschläger 2004; Buonomano & Maass 2009; Jaeger 2001).

Recent analysis of the properties of recurrent networks, such as those realized in neuronal systems and in particular the cerebral cortex, indicate that such high-dimensional dynamic states can indeed be generated in delay-coupled networks (Lazar et al. 2009; Buonomano & Maass 2009; Soriano et al. 2013; for review see Singer 2013). In the present context it is important to recall that the dynamics required for such computations can emerge only when the networks are in the appropriate state. The optimal state has been identified as the edge of chaos, slightly below self-organised criticality, the so-called SOC state, because in this state the dimensionality or the complexity of the system are very high. Computationally this range is optimal because it offers a maximum of possible bifurcation points and storage capacity. (Plenz & Thiagarajan 2007). In this conceptual framework, computational results should consist of substates with reduced dimensionality. Experimental evidence indicates that the high-dimensional resting states are actually reduced by sensory input, imagery, recall of memories, or focused attention. These processes are all associated with enhanced correlation between neuronal responses due to the induction of synchronized high-frequency oscillations—where enhancing correlations reduces dimensionality (for review see Singer 2013). The notion that SOC states are optimal prerequisites for processing also fits with the robust evidence that states compatible with consciousness are characterized by "desynchronized" brain activity, i.e., states characterized by uncorrelated activity, such as are typical for wakefulness and arousal. If, and evidence suggests this to be the case (for review see Singer 1999, 2013), establishment of lower-dimensional synchronous substates, e.g., the formation of transiently-synchronized assemblies of neurons, is an integral part of the computations, then dynamic states characterized by global, large scale synchrony would be inappropriate as background for computations underlying higher cognitive functions.

As outlined in the target paper and above, higher cognitive functions require fine-grained binding operations among semantically-related contents that need to be encoded in ad hoc-formed neuronal assemblies. Such concatenation of multiple assemblies by partial correlations and perhaps also cross-frequency coupling would be impossible in networks that are already highly synchronized to begin with and hence exhibit low complexity and dimensionality. The well-established notion that deep sleep, anaesthesia, and most forms of coma are associated with brain states that exhibit slow oscillations synchronized over considerable distances agrees with this interpretation. In agreement with the prediction that low-dimensional brain states are incompatible with sophisticated processing are also the recent stimulation experiments cited by Noreika. It is to be expected that stimulation of a dynamic system that is in a low-dimensional state and at an overall reduced level of excitability will elicit only a spatially-restricted responses of low complexity—in particular if the stimulus is itself very low-dimensional, as is the case for a TMS pulse.

Considering more recent theories on brain functions, it appears as if the prerequisite or the NCC of a conscious state is a dynamic state that assures a high degree of complexity and high-dimensionality of resting-state dynamics. It is only in this state that the higher cognitive functions can be realized that one expects from a conscious brain.

It should be noted, however, that this operational definition of consciousness makes no inferences about the subjective contents of consciousness or the awareness of particular qualia of experience. According to this definition, consciousness is simply a brain state that allows animals and humans to accomplish higher cognitive functions that include not only perception but also decision making, planning of actions, generation of procedural and episodic memories, and last but not least intentionality and reasoning. Thus, one would expect consciousness, defined in this way, to be a graded phenomenon. If the state of the brain changes towards reduced complexity and dimensionality, there should be a graded deterioration of functions. Those requiring integration of widely-distributed assemblies should become impeded first, while simple reactions to salient sensory stimuli would persist for much longer. This seems to be in perfect agreement with the gradual deterioration of cognitive functions as the brain state shifts from high levels of alertness to drowsiness and sleep.

3 Conscious versus subconscious processing

As Noreika points out, "consciousness" defined by the status of phenomenal content is something very different from a conscious state, as this connotation of consciousness can only be investigated in human subjects. The reason for this is that the distinguishing criterion is the degree of subjective awareness of a cognitive content, and this variable can only be assessed through verbal report. It is simply not possible to know whether a monkey trained to press a lever to signal that it has recognized a particular pattern has the subjective experience that we equate with conscious perception. The monkey brain has the same mechanisms as humans for the allocation of attention, the selection of objects for perception, and the routing of experiences to the different storage modes (working memory, procedural and episodic memory). Thus it is very likely that monkeys are aware of their perceptions in a similar way to us, and that the distinction between conscious and non-conscious processing holds for them as well—but we have no way of knowing. Conditioned lever presses in response to stimuli do not require conscious perception of the stimuli, just as stopping at a red light while being engaged in a conversation does not require conscious recollection of having perceived the light. It is for this reason that the criterion for conscious processing is the reportability of the perceived stimulus, and hence this aspect of consciousness can only be studied in humans.

Attempts to identify the differences between the neuronal processes that accompany non-conscious and conscious processing, respectively, are of course interesting in their own right. The expectation is that they will provide answers to the question of why certain processes are reportable and have access to working and episodic memory while others are excluded, or the question of why certain forms of reasoning and decision-making require conscious deliberations while others do not. However, as pointed out so stringently by Noreika, these attempts fall short of identifying the NCC proper, and at best cover some aspects of conscious processing while being fraught with problems. The most difficult problems are related to the distinction between the processes that are essential for subjective awareness and reportability and those that are the consequence of having become aware of something or that simply provide favourable conditions for becoming aware, such as the allocation of attention or the saliency of stimuli. So far the only neuronal signatures distinguishing between reportable and non-reportable processes have been found to be transitory, lasting at most a few hundred milliseconds. Noreika argues rightly that this disqualifies these events as NCCs because the stream of consciousness is continuous and the awareness of contents can persist for quite some time.

4 Conclusion and outlook

We need to be more cautious when using the term NCC and to define precisely, each time we perform a search for underlying neuronal mechanisms, which of the many aspects of "consciousness" we actually intend to investigate. We need to differentiate between processes assuring access to conscious processing, which are expected to be transient, and processes necessary for sustaining the stream of consciousness that has longer time-constants. And finally, we need to distinguish processes assuring sustained awareness of contents that are most likely related to the transfer of material to short- and long-term memories. If we proceed in this way, subdividing “consciousness” into subfunctions including reportability and defining these as explananda, some of the present problems may dissolve. However, the consequence is that we shall have to give up the search for “the” overarching NCC.

If we pursue this agenda, it is to be expected that correlates will be found for all aspects of consciousness except those associated with the “hard” problem, which appears to be a specific human problem. As I argued in the target paper, searching for the neuronal correlates of qualia in individual brains is unlikely to be successful because the immaterial and therefore somewhat mysterious connotations of qualia are likely to have the status of social realities. What we can achieve, however, is an identification of brain processes that underlie those cognitive functions required for generating social realities. These would be the ability to engage in social interaction, to develop a theory of mind, to find symbolic descriptions of internal states, and to reach consensus on the “reality” of these through communication with others.

To conclude this brief reply to the extremely inspiring commentary on my target paper, I want to express my sincere gratitude to Noreika for having pointed out the critical issues in our research on the NCC. The reply forced me to engage with this research again and helped me substantially in clarifying my own position in the debate.

References

Bertschinger, N. & Natschläger, T. (2004). Real-time computation at the edge of chaos in recurrent neural networks. *Neural Computation*, *16* (7), 1413-1436. 10.1162/089976604323057443

Buonomano, D. V. & Maass, W. (2009). State-dependent computations: Spatiotemporal processing in cortical networks. *Nature Reviews Neuroscience*, *10*, 113-125. 10.1038/nrn2558

Buzsáki, G., Logothetis, N. & Singer, W. (2013). Scaling brain size, keeping timing: Evolutionary preservation of brain rhythms. *Neuron*, *80* (3), 751-764. 10.1016/j.neuron.2013.10.002

Jaeger, H. (2001). The “echo state’’ approach to analysing and training recurrent neural networks - with an Erratum note. *German National Research Center for Information Technology, GMD Report*, *148*

Lazar, A., Pipa, G. & Triesch, J. (2009). SORN: A self-organizing recurrent neural network. *Frontiers in Computational Neuroscience*, *3* (23), 1-9. 10.3389/neuro.10.023.2009

Plenz, D. & Thiagarajan, T. C. (2007). The organizing principles of neuronal avalanches: Cell assemblies in the cortex? *Trends in Neurosciences*, *30* (3), 99-110. 10.1016/j.tins.2007.01.005

Singer, W. (1999). Neuronal synchrony: A versatile code for the definition of relations? *Neuron*, *24* (1), 49-65. 10.1016/S0896-6273(00)80821-1

——— (2013). Cortical dynamics revisited. *Trends in Cognitive Sciences*, *17* (12), 616-626. 10.1016/j.tics.2013.09.006

Soriano, M. C., Garcia-Ojalvo, J., Mirasso, C. R. & Fischer, I. (2013). Complex photonics: Dynamics and applications of delay-coupled semiconductors lasers. *Review of Modern Physics*, *85* (1), 421-470. 10.1103/RevModPhys.85.421