CellPress
REVIEWS

Forum

Robust, Transient Neural Dynamics during Conscious Perception

Biyu J. He[1,*]

While neuroscientific research on perceptual awareness has traditionally focused on the spatial and temporal localizations of neural activity underlying conscious processing, recent development suggests that the dynamic characteristics of spatiotemporally distributed neural activity contain important clues about the neural computational mechanisms underlying conscious processing. Here, we summarize recent progress.

Introduction

Much of neuroscientific research on perceptual awareness has centered around two questions: first, which brain regions are crucial for generating conscious perception? Second, is the neural underpinning of conscious perception fast or slow? Here, we argue that in order to develop a theory that explains neural computational mechanisms underlying conscious processing, simply asking the questions of where and when is insufficient. Understanding the dynamic characteristics of spatiotemporally distributed neural activity within a region, across regions and networks, and at the whole-brain level, could provide key insights. In particular, across-trial variability and stability/transience measures applied to multivariate data can provide a handle for dissecting the dynamical mechanisms embedded in distributed neural activity; and state–space representations allow compact descriptions of spatiotemporally evolving neural activity patterns.

Clarifying Concepts

A recent review proposed that the effects of nonconscious stimuli on the brain 'tend to be small, variable, and short-lived' and, by contrast, conscious processing is associated with 'stable, reproducible representation . . . by a distributed activity pattern in higher cortical areas' [1]. This assertion involves multiple concepts: response magnitude (small), response duration (short-lived), across-trial variability (variable vs reproducible), and stability (stable) (Figure 1). Below, we provide definitions and clarifications for these concepts.

Stability refers to the steadiness of a neural activity pattern over time. It can be measured by variability of neural activity patterns across time within a trial, or the rate at which neural activity patterns change over time. It is a different, empirically dissociable concept from across-trial variability. For instance, neural activity could be fast changing, yet the specific sequence of activity patterns it follows can be reproducible across trials (hence, reproducible but not stable). Alternatively, neural activity could exhibit a sustained, stable pattern within a trial, yet settle on different patterns in different trials (hence, stable but not reproducible).

Response magnitude and duration typically refer to trial-averaged responses, after subtracting the prestimulus baseline. Both concepts are orthogonal to across-trial variability and stability. A well-replicated finding is that consciously perceived stimuli tend to elicit larger-magnitude and longer-lasting responses than those that fail to reach awareness [2,3]. However, only recently have studies begun to examine the stability and across-trial variability dimensions of neural activity during conscious as compared to unconscious processing.

Neural Activity Underlying Conscious Processing Is More Reproducible

Two recent magnetoencephalography (MEG) studies, using dichopic color masking [4] and threshold-level visual perception [5], found that across-trial variability of neural activity is significantly lower in seen than unseen trials, and that this effect far outlasts the duration of the stimulus. Lower across-trial variability in seen trials was observed not only at the single-sensor level [5], but also in the multivariate activity pattern across the whole brain [4,5]. Furthermore, both studies found that this multivariate effect was driven by the angle of the population activity vector, which captures the relative activity pattern across sensors (i.e., after the effect of the mean was removed). A functional magnetic resonance imaging study also reported that the angle of population activity vector, as measured by voxel-wise activity in fusiform gyrus, was more reproducible across seen trials than unseen trials [6]. Thus, convergent results from different imaging modalities and using different analytical approaches show that neural activity is more reproducible across trials during conscious processing.

Neural Activity Underlying Conscious Processing: Stable or Transient?

Thus far, empirical studies on whether neural activity underlying conscious processing is stable over time have reported conflicting findings. Using angle of population activity vector obtained from whole-head MEG, one group found that within-trial, across-time variability was lower in seen than unseen trials, leading them to conclude that activity pattern in seen trials is more stable over time [4]. However, another study from the same

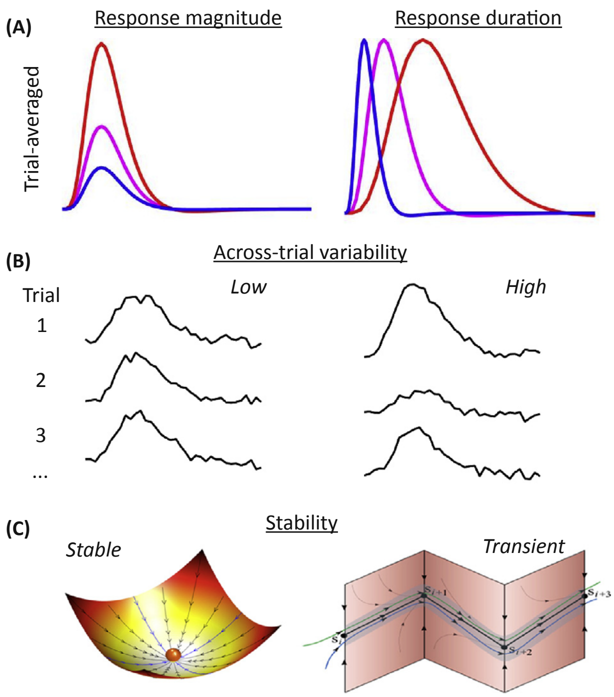

Figure 1. Illustration of Concepts. (A) Response magnitude and duration. (B) Across-trial variability. (C) Left: state–space representation of a stable fixed-point attractor; activity pattern associated with the minimum in the energy landscape is sustained over time unless terminated by external perturbation or shift in the energy landscape. Right: state–space representation of robust transient dynamics. Each black dot represents a point in the state space that is stable in many directions but unstable in certain direction (e.g., a saddle node). Together, they form a robust trajectory, such that the system's activity evolves from one point to the next. Adapted from [12].

group trained a decoder to classify stimulus location in seen and unseen trials separately (also using whole-head MEG data), and found that the decoder generalized across time more successfully in unseen trials, suggesting that neural activity pattern is in fact more stable over time in unseen trials [7].

A recent study [5] offers a potential reconciliation for this discrepancy. The key insight is that different analysis methods are sensitive to different frequency contents (i.e., slower vs faster activity) when fed with similarly broadband data. Here, the authors found that multivariate decoding is most sensitive to the low-frequency component of the signal; by contrast, analyses assessing changes over time appear more sensitive to the high-frequency component. Thus, although both these previous studies used broadband data, the analysis in [4] is likely more sensitive to β- and γ-band activity, while that in [7] is likely contributed most by δ and sub-δ activity (also called slow cortical potentials, SCPs [8]).

This recent study [5] further showed that for the population activity vector constructed using whole-head MEG data, its rate of change over time (velocity, with higher velocity indicating lower stability) is higher in seen than unseen trials in the SCP band, but exhibits an opposite pattern in higher-frequency ranges. Thus, for high-frequency activity, their finding was consistent with [4]. However, single-trial decoding of seen versus unseen perceptual outcome could be robustly obtained only in the SCP band. Thus, at least for large-scale brain activity, in the frequency band that is most relevant for encoding the status of conscious perception, activity pattern changes faster over time in seen trials than unseen trials – a finding opposite to the stability in conscious processing proposal [1].

Robust, Transient Neural Dynamics and State–Space Representation

The picture emerging from these recent studies [4–7] suggests the following: first, neural activity during conscious processing is more reproducible across trials as compared to unconscious processing of the same stimulus; and second, large-scale neural activity (in the SCP band) follows a sequence of activity patterns that distinguishes seen from unseen trials, with seen activity evolving faster over time. One way to compactly visualize spatiotemporally evolving neural activity patterns and gain intuitions for these results is through dimensionality reduction and state–space representation. At each time point, large-scale neural activity pattern can be represented as a point in a high-dimensional space, where each dimension describes the activity of one sensor. Due to correlated activity across sensors, the actual dimensionality of the data is much lower. For example, using standard principle component (PC) analysis, the top five PCs already account for >70% of variance in the whole-head MEG data [5]. Thus, projecting the original high-dimensional data to a low-dimensional space defined by the top few PCs allows a compact representation of the temporal evolution of neural activity patterns (e.g., Figure 2). Abstract quantitative measures such as across-trial variability and velocity can also be visualized in this space as the spread of single-trial trajectories and the distance between activity states at consecutive time points, respectively.

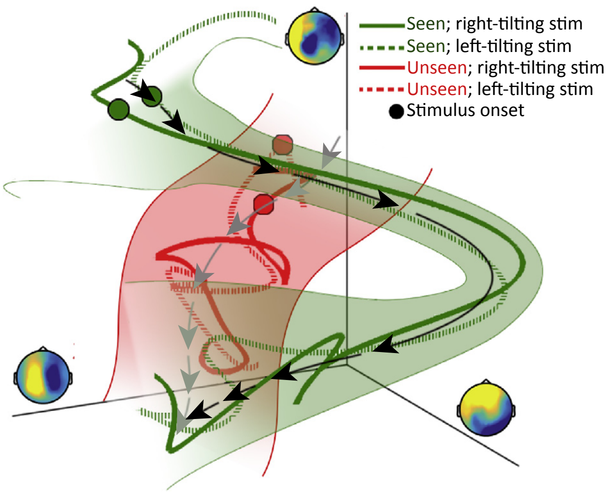

Figure 2. Example of Robust, Transient Neural Dynamics in Conscious Perception [5]. Whole-brain MEG activity in the SCP band is projected to a low-dimensional state–space using the first three PCs. Scalp topography corresponding to each PC is shown next to the axis. Trial-averaged activity for seen and unseen conditions are plotted from 1 s before to 3 s after stimulus onset. The stimulus was a threshold-level Gabor patch with one of two possible orientations lasting ~50 ms. Seen and unseen trajectories are well separated throughout the trial. In seen trials, following stimulus onset, activity patterns evolve quickly (indicated by the length of arrows) and across-trial variability decreases substantially (shown as shading). Abbreviations: MEG, magnetoencephalography; PC, principle component; SCP, slow cortical potential.

Together, these results indicate that large-scale neural dynamics during conscious processing is well described by a class of theoretical models called robust, transient dynamics [9]. Robust refers to robustness to noise, such that small variations in the initial state do not prevent the system from following a robust trajectory in the state space (i.e., a robust sequence of activity pattern evolution) upon sensory input, thereby conferring reproducibility across trials. Transient refers to activity states that are constantly evolving, in which no stable equilibrium is reached.

State–space representation also revealed another intriguing finding in this study: up to 1 s before stimulus onset, large-scale SCP activity is well separable between seen and unseen trials, suggesting that the initial brain state influences which trajectory of activity pattern evolution will be adopted upon sensory input, which in turn lead to different perceptual outcomes (Figure 2) [5]. This indicates that neural activity is both robust to small variations in the initial state and sensitive to large variations – a finding that can be easily accommodated within the robust, transient dynamics framework. Supposing that there are two robust trajectories coexisting in the state space, if brain activity is near the attractor basin of either trajectory, it would be attracted onto that one. Importantly, this perspective points to the futility of focusing on the exact temporal latency of the neural correlate of conscious perception. Depending on the initial brain state, the desired location in the state space (corresponding to conscious perception) may be reached sooner or later, or not at all. While many previous studies have reported the effect of contextual factors, such as spontaneous brain activity and cognitive states including attention and expectation, on conscious perception (e.g., [10,11]), the framework of robust, transient neural dynamics offers a concrete, testable mechanism for how this may happen through the influence of initial brain state on the ensuing neural activity evolvement.

Acknowledgments

BJH was supported by New York University Langone Medical Center, the Intramural Research Program of the National Institutes of Health/National Institute of Neurological Disorders and Stroke, Leon Levy Foundation, and Klingenstein–Simons Neuroscience Fellowship. Gratitude goes to Brian Maniscalco and Alex Baria for astute observations and insightful discussions during the course of developing these ideas.

Disclaimer Statement

The author declares no competing interest.

[1]Departments of Neurology, Neuroscience and Physiology, and Radiology, Neuroscience Institute, New York University Langone Medical Center, NY 10016, USA

*Correspondence: biyu.he@nyumc.org (B.J. He).
https://doi.org/10.1016/j.tics.2018.04.005

References

1. Dehaene, S. *et al.* (2017) What is consciousness, and could machines have it? *Science* 358, 486–492
2. Dehaene, S. and Changeux, J.P. (2011) Experimental and theoretical approaches to conscious processing. *Neuron* 70, 200–227
3. Koch, C. *et al.* (2016) Neural correlates of consciousness: progress and problems. *Nat. Rev. Neurosci.* 17, 307–321
4. Schurger, A. *et al.* (2015) Cortical activity is more stable when sensory stimuli are consciously perceived. *Proc. Natl. Acad. Sci. U. S. A.* 112, E2083–E2092
5. Baria, A.T. *et al.* (2017) Initial-state-dependent, robust, transient neural dynamics encode conscious visual perception. *PLoS Comput. Biol.* 13, e1005806
6. Schurger, A. *et al.* (2010) Reproducibility distinguishes conscious from nonconscious neural representations. *Science* 327, 97–99
7. Salti, M. *et al.* (2015) Distinct cortical codes and temporal dynamics for conscious and unconscious percepts. *Elife* 4, e05652
8. He, B.J. and Raichle, M.E. (2009) The fMRI signal, slow cortical potential and consciousness. *Trends Cogn. Sci.* 13, 302–309
9. Rabinovich, M. *et al.* (2008) Neuroscience: Transient dynamics for neural processing. *Science* 321, 48–50
10. Melloni, L. *et al.* (2011) Expectations change the signatures and timing of electrophysiological correlates of perceptual awareness. *J. Neurosci.* 31, 1386–1396
11. Wyart, V. and Tallon-Baudry, C. (2009) How ongoing fluctuations in human visual cortex predict perceptual awareness: baseline shift versus decision bias. *J. Neurosci.* 29, 8715–8725
12. Rabinovich, M.I. and Varona, P. (2011) Robust transient dynamics and brain functions. *Front. Comput. Neurosci.* 5, 24