Hierarchy of Information Processing in the Brain: A Novel 'Intrinsic Ignition' Framework

Gustavo Deco^{1,2,3,4} and Morten L. Kringelbach^{5,6,7,*}

¹Center for Brain and Cognition, Computational Neuroscience Group, Department of Information and Communication Technologies, Universitat Pompeu Fabra, Roc Boronat 138, Barcelona 08018, Spain

²Institució Catalana de la Recerca i Estudis Avançats (ICREA), Passeig Lluís Companys 23, Barcelona 08010, Spain

³Department of Neuropsychology, Max Planck Institute for Human Cognitive and Brain Sciences, 04103 Leipzig, Germany

⁴School of Psychological Sciences, Monash University, Melbourne, Clayton VIC 3800, Australia

⁵Department of Psychiatry, University of Oxford, Oxford OX3 7JX, UK

⁶Center for Music in the Brain (MIB), Department of Clinical Medicine, Aarhus University, 8000 Aarhus C, Denmark

⁷Institut d'études avancées de Paris, 75004 Paris, France

*Correspondence: morten.kringelbach@psych.ox.ac.uk

http://dx.doi.org/10.1016/j.neuron.2017.03.028

A general theory of brain function has to be able to explain local and non-local network computations over space and time. We propose a new framework to capture the key principles of how local activity influences global computation, i.e., describing the propagation of information and thus the broadness of communication driven by local activity. More specifically, we consider the diversity in space (nodes or brain regions) over time using the concept of intrinsic ignition, which are naturally occurring intrinsic perturbations reflecting the capability of a given brain area to propagate neuronal activity to other regions in a given brain state. Characterizing the profile of intrinsic ignition for a given brain state provides insight into the precise nature of hierarchical information processing. Combining this data-driven method with a causal whole-brain computational model can provide novel insights into the imbalance of brain states found in neuropsychiatric disorders.

Introduction

Historically, within brain science, the most sophisticated and precise spatiotemporal information has come from single-neuron recordings of spiking activity (Hodgkin and Huxley, 1952; Hubel and Wiesel, 1959; Mountcastle, 1957). Measuring spiking activity could help derive the underlying mechanistic principles of brain function (notwithstanding the mounting evidence of the importance of glial processes). This local information has helped develop important computational models of attention, memory, and decision making (Amit and Brunel, 1997; Brunel and Wang, 2001; Deco and Rolls, 2005; Wang, 2002). Yet, non-local information is also very important given that neurons are connected to other neurons in sophisticated networks with specific properties (small world or not) that are crucial for efficient brain function (Markov et al., 2013, 2014; Sporns and Zwi, 2004; Watts and Strogatz, 1998). This local versus non-local processing can be characterized on many different levels, from the cortical column to brain region-specific networks to whole-brain connectivity (Sporns et al., 2005). Abundant evidence shows that computation happens on all these levels and that the field of brain science could benefit from a better understanding of the computational principles at the network level.

Here we propose a novel framework for measuring and understanding the intrinsic dynamics and communication principles of brain activity across all levels of networks (from local networks of neurons in a brain region to higher-order whole-brain networks). The starting point is how best to characterize the way that information propagates from a local node through the network. We are proposing a conceptual framework for studying the "intrinsic ignition" of brain activity across time and space, i.e., the diversity of computation in space and time. This framework allows the study not only of the propagation of brain activity, but also the underlying fluctuations and their functional network consequences. Informally, the concept of intrinsic ignition refers to the capability of a given local node (single neuron or brain area) in *space* to propagate feedforward and recurrent neuronal activity to other nodes in the network as measured by the whole-brain integration elicited. As such, intrinsic ignition is a novel concept that can be used to describe the specific profile of the ignition capabilities of brain regions in different brain states. Furthermore, defining the variability of the ignition-driven propagation of activity across *time*, we can classify each node (neuron or brain area) according to the local degree of functional variability, i.e., diversity and metastability.

Ignition as Intrinsic Perturbation

Computation is the fundamental unit of a general theory of brain function, enabling a full characterization of the fundamental neuronal principles underlying the computations involved in cognitive, perceptual, and motor functions in health, as well as in disease. This is a complex problem given that there are billions of recurrently, synaptically coupled, non-homogeneous neurons in the human brain. More precisely, the non-linear character of these basic elements, primarily neurons and synapses, and the coupling through both feedforward and feedback connections makes the brain a complex, non-linear dynamical system (Deco et al., 2011; Sporns, 2014). Uncovering the fundamental mechanisms underlying the emerging properties of a complex

dynamical system will rely on solving the so-called "inverse problem." In neuroscience, the attempts of finding a solution has adopted three main strategies, namely: (1) correlating neural activity with task, (2) measuring spontaneous, non-task-related, on-going brain dynamics, and (3) lesioning or perturbing the internal dynamics of the brain.

Historically, single-neuron recordings have adopted the first task-evoked strategy, e.g., recording activity in primary visual cortices to simple visual cues such as orientation (Hubel and Wiesel, 1959). Once a correlation was established, using the third perturbation strategy, these neurons could then be lesioned (or perturbed) to show a causal functional role, e.g., lesioning neurons to show the effects on the brain of visual orientation cues (Calford et al., 2000).

The second strategy measuring spontaneous activity is difficult to apply to a local framework of such single-neuron recordings since the spontaneous activity within these recordings has traditionally been seen as merely noise. The advent of optical imaging (Grinvald et al., 1986; Kenet et al., 2003) and neuroimaging experiments of non-linear, baseline, resting-state activity (Biswal et al., 1995; Raichle et al., 2001) has led to a shift in focus to the second strategy of analyzing internal on-going brain dynamics (Smith et al., 2009; Zhang and Raichle, 2010). This has led to an emphasis on the functional significance of the non-local compared to the local computation. The first two strategies of measuring and analyzing brain dynamics with or without task are orthogonal and complementary yet only allow for understanding correlative, rather than causal, relationships.

The third perturbation strategy has been used extensively in experimental animals and measuring the behavioral outcomes. For ethical reasons, such systematic perturbations are not directly translatable to humans. Still, the field has been able to draw on perturbations from naturally occurring lesions linked to stroke and neurosurgical procedures. In addition, there is a growing field of electrical brain stimulation, often as part of neurosurgical evaluation of patients with epilepsy and resection of special cases of brain tumor (Parvizi et al., 2013; Selimbeyoglu and Parvizi, 2010; Winawer and Parvizi, 2016). Furthermore, there is even more widespread use of deep brain stimulation for symptom alleviation in movement disorders, such as Parkinson's disease (Kringelbach et al., 2007). Still, what has been missing is an understanding of how the propagation of information is influenced by perturbation (Borchers et al., 2011).

Ideally, measurements of activity across the whole brain would be useful to characterize the underlying mechanisms, similar to how perturbation of a physical system can be described in physics. Massimini and colleagues pioneered studies using electroencephalography (EEG) to characterize the degree and short-term latency of the dynamics (100–200 ms) elicited by the external perturbation by transcranial magnetic stimulation (TMS) (Massimini et al., 2005). This has been used successfully for separation of brain states in healthy subjects during wakefulness, dreaming, and sleeping and in different levels of anesthesia (Casali et al., 2013; Ferrarelli et al., 2010), as well as characterization of brain-injured patients emerging from coma (Rosanova et al., 2012). This approach can be extended in several ways by using whole-brain computational modeling of the underlying brain activity in a given brain state (e.g., wakefulness, sleep, or coma) measured, for example, with magnetoencephalography (MEG) or fMRI (Deco and Kringelbach, 2014). Such models can then be systematically perturbed in ways not possible experimentally (Cabral et al., 2014). Crucially, rather than just measuring the complexity of the elicited activity after perturbation, the whole-brain model allows for the possibility of generating a long-lasting, strong perturbation and measuring the latency of the recovery of brain dynamics, observable over much longer time spans (on the order of tens of seconds).

Similarly, rather than using extrinsic perturbation, one could study the effects of a naturally occurring intrinsic perturbation. We define intrinsic ignition as the capability of a given brain area to propagate neuronal activity to other regions in a given brain state, describing the whole-brain integration elicited from the propagation of both feedforward and recurrent activity. As such, intrinsic ignition is a novel concept that can be used to describe the specific profile of the ignition capabilities of regions across the network (from local networks of neurons in a brain region to higher-order whole-brain networks) in different brain states.

This definition of intrinsic ignition makes it possible to create a full characterization of each "activity event" of the nodes (neuron or brain areas) for a given brain state. Such intrinsic activity event can be seen as an internal intrinsic perturbation that could eventually elicit (or not) the propagation of activity across the whole network. As an example, one simple way to characterize an "ignition event" is to binarize each measurement resulting from the thresholding of activity at different possible levels in multimodal neuroimaging data (see Figure 1). This method lends itself naturally to spiking neuronal activity but can easily be extended to local field potential (LFP), MEG, and fMRI signals. However, please note that our framework could equally well use other more sophisticated mathematical methods for extracting point processes (Caballero Gaudes et al., 2013; Karahanoğlu et al., 2013; Petridou et al., 2013). Such methods have been shown to be able to describe many important aspects of dynamics, such as, for example, resting-state networks and complexity (Karahanoğlu et al., 2013; Karahanoğlu and Van De Ville, 2015).

The proposed new paradigm allows for a description of how the ignition capability of each node, i.e., the integration elicited, varies for different brain states by averaging the event-related changes (across many occurrences of the same event) induced for each time point aligned after each event. The differences between ignition profiles of the different nodes (i.e., neurons or brain areas) provide a marker of different brain states.

Note that our novel concept of intrinsic ignition is different from the ignition defined by Dehaene and colleagues (Moutard et al., 2015), as the rapid and sustained activity elicited after stimulation in contrast to the ultra-slow (<0.1 Hz) fluctuations of relatively low amplitude in resting state. Both modes can emerge from the same underlying connectome as "two dynamic faces" of the strong recurrent loops built by the brain networks. Indeed, the dense lateral intra- and inter-areal connections that characterize brain networks make possible the emergence of a



INTEGRATION

Figure 1. Schematic of the Process of Measuring Intrinsic Ignition We measure the activity of a node in the network and the changes in subsequent activity across the whole network.

(Å) Here we show the spiking activity in a neuron (green area superimposed on the brain and with activity in the stippled green square). For each driving event, we measure the activity in the rest of the network (in stippled red area) in the gray time window.

(B) This corresponds to a binary connectivity matrix, where spikes are cooccurring.

(C) In this matrix, we find the largest subcomponent as a measure of the global integration, i.e., the broadness of communication across the network. This is repeated for each of the driving events, producing a mean and standard deviation of the intrinsic ignition for each network node.

reverberatory dynamic when the level of excitation exceeds the level of inhibition, which can be propagated globally across the brain. This imbalance between excitation and inhibition could appear spontaneously (resting state) or rapidly induced (ignition) by the action of an external stimulation, explaining in this way both modes. Nevertheless, such concept of ignition is not covering how an intrinsic (i.e., in the absence of any kind of external stimulation) local activity event in a given brain state (e.g., wakefulness and sleep) is eliciting a propagation of activity across the network.

The novel concept of intrinsic ignition allows us to measure the spatial and temporal diversity of propagation of information in the whole network using our existing integration measure (Deco et al., 2015). Briefly, for each intrinsic ignition event, a value can be computed that measures the integration, i.e., the degree of broadcasting of the information arising from both feedforward and recurrent processing (Figure 1). This integration value is defined as the length of the largest connected component in a binarized functional connectivity matrix of the whole network at a given window of time following the intrinsic ignition triggering event. Note that the binarization procedure of the activity in each node allows us to construct the functional connectivity matrix in this window of time, where the largest subcomponent is defined as the length of the connected component of the undirected graph defined by the binarized matrix considered as an adjacency matrix. This is the largest sub-graph in which any two vertices are connected to each other by paths and which connects to no additional vertices in the super-graph. Similar to event-related potential (ERP) analysis, the intrinsic ignition for each node in a given network is fully characterized by the mean and standard deviation across events.

At the network level, the mean of the intrinsic ignition in a network allows us to show the spatial diversity as the differences in average intrinsic ignition profiles across the different nodes for different brain states. Furthermore, at the network level, we can characterize the temporal diversity by measuring the variability of the intrinsic ignition-driven increase of integration across events by classifying nodes according to the local degree of functional variability. This functional variability is a measure of diversity across time, which is related to the fundamental functionality of a given node, linked to high or low levels of local metastability, i.e., the variability of a state that falls outside the natural equilibrium state of the system but persists for an extended period of time. In other words, different levels of temporal diversity in a given node can be thought of as a measure of local functional variability or metastability and thus describe the versatility of a given network node.

Hierarchies of Computation

There is evidence from the field of connectomics using graph theory to reveal a topological hierarchy in structural connectivity of the brain (Deco et al., 2015; Sporns et al., 2000). But it is presently unclear whether this structural hierarchical connectivity is reflected in the dynamical processing hierarchy. Measuring the intrinsic ignition of nodes (mean and variability) in a network makes it possible to investigate the important question of whether there is a hierarchy of information processing.

Figure 2 shows the spectrum of possibilities for hierarchical processing. At one end (shown in Figure 2A), the intrinsic ignition is equal for all nodes in the network. This would be the equivalent of a weak, flat non-hierarchy, demonstrated by a single circle, where each node has equal computational importance. At the other end (shown in Figure 2D), intrinsic ignition is uniform, linear gradation between nodes, which corresponds to a strong hierarchy in which a node at the top of the linear hierarchy has the



Figure 2. Spectrum of Possible Dynamical Processing Hierarchies in the Brain

(A) Weak, flat non-hierarchy is derived when the intrinsic ignition is equal for all nodes in the network, which is shown by a single level of nodes.
 (B) On the other hand, a "staircase" hierarchy is suggestive of distinct circles (or orbits) of groups of nodes with equal computational importance. Such a scheme would correspond to the ideas of a global workspace with a clear computational quantum jump between sensory regions and regions in the global neuronal workspace.

(C) There are many other possibilities of graded, non-uniform hierarchies with non-uniform, clustered circular orbits for the network nodes.

(D) Strong, graded, uniform hierarchy can occur when a node at the top of the linear hierarchy has the highest intrinsic ignition, which is demonstrated by a uniform distribution of levels (rings) for each node in the network.

highest intrinsic ignition, demonstrated by a uniform distribution of rings for each node in the network.

Between these poles of weak and strong hierarchies, there are many significantly different functional possibilities. One possibility is where the intrinsic ignition resembles a staircase, suggestive of distinct circles (or orbits) of groups of nodes with equal computational importance. These groups are, however, clearly stratified, shown in Figure 2B by two orbits with nodes of similar intrinsic ignition. This scheme would correspond to the idea of the global workspace, where there is a clear, computational, quantum jump between sensory regions and regions in the global neuronal workspace (Baars, 1989; Dehaene and Changeux, 2011; Dehaene et al., 1998).

Another possibility is that the profile of network nodes shows a non-uniform gradation in intrinsic ignition, suggestive of a hierarchical organization of a non-uniform graded variety. This is schematically demonstrated in Figure 2C by the non-uniformity circular orbits for each node.

Whole-Brain Evidence for Hierarchy of Intrinsic Ignition

The key question of the dynamical processing hierarchy across human brain regions can be addressed using the concept of spatiotemporal diversity as defined by intrinsic ignition. To this end, we here investigated this issue by using the spontaneous brain activity from a group of 16 healthy individuals measured with fMRI (van Hartevelt et al., 2015).

The procedure of defining events used the simple method of Tagliazucchi and colleagues to binarize the fMRI series in the following way (Tagliazucchi et al., 2012). An event is defined as a binary signal resulting from the transformed functional time series (BOLD fMRI) into Z scores, z_i(t), and imposing a threshold, θ , such that the binary sequence $\sigma_i(t) = 1$ if $z_i(t) > \theta$ and is crossing the threshold from below and $\sigma_i(t) = 0$ otherwise (shown schematically in Figure 3). If the top signal (in red) refers to the analysis of the ignition capability of a given brain region, an event is the point where the signal crosses the threshold from below. The x axis represents the time whereas the y ordinates represent the different brain areas. The brain was segmented into 758 regions (based on running a K-means algorithm run on the standard AAL atlas, Automated Anatomical Labeling, including all cortical and subcortical areas) (Tzourio-Mazoyer et al., 2002). Each single black vertical bar refers to an event for the corresponding brain region.

For each brain region, we computed the mean and standard deviation of the integration associated with events. Specifically, we used a window size restricted by the intra-event duration for observing the evolution of the elicited integration and used the maximum value of integration within that given time window.



Figure 3. Intrinsic Ignition Hierarchical Organization Demonstrated with Empirical Data

(A) Neuroimaging signals, such as BOLD, can be treated in the same manner as spiking data by applying a threshold method to define events. This again gives rise
to a connectivity matrix for each intrinsic ignition event for which the integration, i.e., the largest subcomponent, can be measured.
 (B) When applied to fMRI signals of spontaneous activity in normal healthy participants using a very fine-grained parcellation, sorted from largest the intrinsic

(b) when applied to twint signals of spontaneous activity in normal healthy participants using a very inte-grained parcentation, sorted from largest the intrinsic ignition of nodes, we found, for both mean and standard deviation, an inverted sigmoid profile. This result indicates that there is a hierarchy of function across brain regions, compatible with the global workspace theory. Yet, this hierarchy is not stratified in a staircase manner. Still, there are clearly regions with higher intrinsic ignition variability, which would be more computationally relevant and could play a central role in broadcasting information than the regions with low intrinsic ignition.

Figure 3 shows the profiles of the mean and variability of the ignition-driven integration. This shows an inverse sigmoidal curve, suggesting that the hierarchical organization is of the non-uniform, graded variety and thus closer to the scenario described in Figure 2C than to the other scenarios. Although, it should be noted that, with respect to the global workspace scenario, the results confirm that there is a hierarchy of function across brain regions. However, this hierarchy is not stratified in a staircase manner as suggested by the strong version of the global workspace theory, which is a surprising and non-trivial result. Still, there are clearly regions with higher intrinsic ignition variability, suggesting that they are more computationally relevant and could play a central role in broadcasting information than the regions with low intrinsic ignition, which are more likely to be related to sensory processing.

The Role of Whole-Brain Computational Modeling for Binding Information

The diversity of computation has to be understood at the network level, but while intrinsic ignition is an important way to uncover the computational role of a given node, this does not provide information about what happens when this node is eliminated. This complementary question cannot be answered by simple data-driven methods but requires causal whole-brain models. Given such a model, carefully adjusted to empirical data, it is possible to lesion and perturb the model offline and study the consequences (van Hartevelt et al., 2015). For example, it is possible to study how the spatiotemporal diversity will change following the elimination of any given node.

In order to investigate these challenging questions more carefully, we propose future studies testing the robustness of computational diversity using a whole-brain model fitting resting-state empirical data. Briefly, whole-brain models link anatomical structure with functional dynamics. Structural connectivity data can be obtained by diffusion weighted/tensor imaging (DWI/DTI) combined with probabilistic tractography, which is a measure of the density of fibers between brain regions. The global dynamics of the whole-brain model results from the mutual interactions of local node dynamics coupled through the underlying empirical anatomical structural connectivity matrix. Typically, the temporal dynamics of local brain areas in these models is taken to be either asynchronous (spiking models or their respective mean-field reduction) or oscillatory (Deco and Kringelbach, 2014). Recently, evidence has emerged for a promising version using, for each



Figure 4. Using Whole-Brain Models to Determine Causal Functional Roles of Nodes (A) Causal whole-brain methods can be used to investigate the effects on the computational diversity fitting the model to the empirical data. (B) The removal or perturbation of the brain regions with the highest level of intrinsic ignition can provide causal information on the functional importance of individual brain regions in terms of how the profiles of intrinsic ignition might change, similar to how cutting off the head of a hydra may sprout many more heads. As such, this makes it possible to investigate the change of intrinsic ignition of individual nodes across different brain states and in neuropsychiatric disorders.

brain area, a local dynamical model given by a normal form of bifurcations (e.g., a supercritical Hopf bifurcation) (Deco et al., 2016; Kringelbach et al., 2015). The normal form of a Hopf bifurcation can describe the transition from asynchronous noisy behavior to full oscillations and thus unify previous asynchronous and full oscillatory scenarios. The main parameter that is manipulated for fitting the empirical data and for analyzing the model is the so-called global coupling parameter G. The global coupling parameter G corresponds to the conductivity of the synaptic connections, which is considered here, for simplicity, uniform across the brain. Our preliminary results indicate that the working point of a whole-brain computational model of the resting state corresponds to where there is maximal hierarchical processing and thus maximal entropy. Further preliminary results suggest that this is not the case for other dynamic brain states, such as deep sleep or drug-altered states. As such, this could provide novel insights into the dynamic processing of hierarchy in different brain states.

Similar to our recent work, to the effect of removing so-called "binding nodes" (Deco et al., 2017), future studies could use causal whole-brain methods to investigate the effects on the computational diversity of the removal or perturbation of the brain

regions with the highest level of intrinsic ignition (see Figure 4). This would provide information on the functional importance of individual brain regions. As such, it would also be possible to investigate the change of intrinsic ignition of individual nodes across different brain states (Figure 4B). It would also be of high interest to use direct brain stimulation of such binding nodes in patients undergoing neurosurgical investigations to predict and investigate their causal role in cognition (Parvizi et al., 2013).

Conclusion

In this Perspective, we have provided evidence for a powerful novel framework for understanding the underlying basis of neural communication and network organization. The concept of intrinsic ignition provides a simple way to investigate the computational spatiotemporal diversity of network nodes and provides direct evidence on the hierarchical structuring of information processing in the network. This method can be applied to multimodal neuronal data; from the single neuron to local field potentials, MEG, and even fMRI. Combining this data-driven method with a causal whole-brain model can provide strong evidence on the functional role of a network node by measuring the

consequences of its elimination. As such, this is a powerful new method that can provide strong evidence on brain computation on many levels from the single neuron to groups of neurons to the whole-brain level.

There are thus many important implications for using this framework to understand hitherto poorly understood problems such as altered states of consciousness. But perhaps most importantly, it provides new tools for understanding the imbalances in functional organization found in neuropsychiatric disorders, and, when combined with whole-brain models, may even provide insights into novel ways of rebalancing the whole-brain networks in health and disease.

ACKNOWLEDGMENTS

G.D. is supported by the ERC Advanced Grant DYSTRUCTURE (n. 295129), by the Spanish Research Project PSI2016-75688-P, and by the European Union's Horizon 2020 research and innovation programme under grant agreement n. 720270 (HBP SGA1). M.L.K. is supported by the ERC Consolidator Grant: CAREGIVING (n. 615539) and Center for Music in the Brain, funded by the Danish National Research Foundation (DNRF117).

REFERENCES

Amit, D.J., and Brunel, N. (1997). Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. Cereb. Cortex 7, 237–252.

Baars, B.J. (1989). A Cognitive Theory of Consciousness (Cambridge University Press).

Biswal, B., Yetkin, F.Z., Haughton, V.M., and Hyde, J.S. (1995). Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. Magn. Reson. Med. 34, 537–541.

Borchers, S., Himmelbach, M., Logothetis, N., and Karnath, H.O. (2011). Direct electrical stimulation of human cortex - the gold standard for mapping brain functions? Nat. Rev. Neurosci. *13*, 63–70.

Brunel, N., and Wang, X.J. (2001). Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. J. Comput. Neurosci. *11*, 63–85.

Caballero Gaudes, C., Petridou, N., Francis, S.T., Dryden, I.L., and Gowland, P.A. (2013). Paradigm-free mapping with sparse regression automatically detects single-trial functional magnetic resonance imaging blood oxygenation level dependent responses. Hum. Brain Mapp. 34, 501–518.

Cabral, J., Kringelbach, M.L., and Deco, G. (2014). Exploring the network dynamics underlying brain activity during rest. Prog. Neurobiol. 114, 102–131.

Calford, M.B., Wang, C., Taglianetti, V., Waleszczyk, W.J., Burke, W., and Dreher, B. (2000). Plasticity in adult cat visual cortex (area 17) following circumscribed monocular lesions of all retinal layers. J. Physiol. *524*, 587–602.

Casali, A.G., Gosseries, O., Rosanova, M., Boly, M., Sarasso, S., Casali, K.R., Casarotto, S., Bruno, M.A., Laureys, S., Tononi, G., and Massimini, M. (2013). A theoretically based index of consciousness independent of sensory processing and behavior. Sci. Transl. Med. *5*, 198ra105.

Deco, G., and Kringelbach, M.L. (2014). Great expectations: using whole-brain computational connectomics for understanding neuropsychiatric disorders. Neuron *84*, 892–905.

Deco, G., and Rolls, E.T. (2005). Attention, short-term memory, and action selection: a unifying theory. Prog. Neurobiol. 76, 236–256.

Deco, G., Jirsa, V.K., and McIntosh, A.R. (2011). Emerging concepts for the dynamical organization of resting-state activity in the brain. Nat. Rev. Neurosci. *12*, 43–56.

Deco, G., Tononi, G., Boly, M., and Kringelbach, M.L. (2015). Rethinking segregation and integration: contributions of whole-brain modelling. Nat. Rev. Neurosci. *16*, 430–439. Deco, G., Kringelbach, M.L., Jirsa, V., and Ritter, P. (2016). The dynamics of resting fluctuations in the brain: metastability and its dynamical core. bioRxiv. http://dx.doi.org/10.1101/065284.

Deco, G., Van Hartevelt, T.J., Fernandes, H.M., Stevner, A., and Kringelbach, M.L. (2017). The most relevant human brain regions for functional connectivity: Evidence for a dynamical workspace of binding nodes from whole-brain computational modelling. Neuroimage *146*, 197–210.

Dehaene, S., and Changeux, J.P. (2011). Experimental and theoretical approaches to conscious processing. Neuron *70*, 200–227.

Dehaene, S., Kerszberg, M., and Changeux, J.P. (1998). A neuronal model of a global workspace in effortful cognitive tasks. Proc. Natl. Acad. Sci. USA *95*, 14529–14534.

Ferrarelli, F., Massimini, M., Sarasso, S., Casali, A., Riedner, B.A., Angelini, G., Tononi, G., and Pearce, R.A. (2010). Breakdown in cortical effective connectivity during midazolam-induced loss of consciousness. Proc. Natl. Acad. Sci. USA *107*, 2681–2686.

Grinvald, A., Lieke, E., Frostig, R.D., Gilbert, C.D., and Wiesel, T.N. (1986). Functional architecture of cortex revealed by optical imaging of intrinsic signals. Nature *324*, 361–364.

Hodgkin, A.L., and Huxley, A.F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. J. Physiol. *117*, 500–544.

Hubel, D.H., and Wiesel, T.N. (1959). Receptive fields of single neurones in the cat's striate cortex. J. Physiol. *148*, 574–591.

Karahanoğlu, F.I., and Van De Ville, D. (2015). Transient brain activity disentangles fMRI resting-state dynamics in terms of spatially and temporally overlapping networks. Nat. Commun. 6, 7751.

Karahanoğlu, F.I., Caballero-Gaudes, C., Lazeyras, F., and Van de Ville, D. (2013). Total activation: fMRI deconvolution through spatio-temporal regularization. Neuroimage *73*, 121–134.

Kenet, T., Bibitchkov, D., Tsodyks, M., Grinvald, A., and Arieli, A. (2003). Spontaneously emerging cortical representations of visual attributes. Nature 425, 954–956.

Kringelbach, M.L., Jenkinson, N., Owen, S.L.F., and Aziz, T.Z. (2007). Translational principles of deep brain stimulation. Nat. Rev. Neurosci. 8, 623–635.

Kringelbach, M.L., McIntosh, A.R., Ritter, P., Jirsa, V.K., and Deco, G. (2015). The rediscovery of slowness: exploring the timing of cognition. Trends Cogn. Sci. *19*, 616–628.

Markov, N.T., Ercsey-Ravasz, M., Van Essen, D.C., Knoblauch, K., Toroczkai, Z., and Kennedy, H. (2013). Cortical high-density counterstream architectures. Science *342*, 1238406.

Markov, N.T., Vezoli, J., Chameau, P., Falchier, A., Quilodran, R., Huissoud, C., Lamy, C., Misery, P., Giroud, P., Ullman, S., et al. (2014). Anatomy of hierarchy: feedforward and feedback pathways in macaque visual cortex. J. Comp. Neurol. *522*, 225–259.

Massimini, M., Ferrarelli, F., Huber, R., Esser, S.K., Singh, H., and Tononi, G. (2005). Breakdown of cortical effective connectivity during sleep. Science *309*, 2228–2232.

Mountcastle, V.B. (1957). Modality and topographic properties of single neurons of cat's somatic sensory cortex. J. Neurophysiol. *20*, 408–434.

Moutard, C., Dehaene, S., and Malach, R. (2015). Spontaneous fluctuations and non-linear ignitions: two dynamic faces of cortical recurrent loops. Neuron 88, 194–206.

Parvizi, J., Rangarajan, V., Shirer, W.R., Desai, N., and Greicius, M.D. (2013). The will to persevere induced by electrical stimulation of the human cingulate gyrus. Neuron *80*, 1359–1367.

Petridou, N., Gaudes, C.C., Dryden, I.L., Francis, S.T., and Gowland, P.A. (2013). Periods of rest in fMRI contain individual spontaneous events which are related to slowly fluctuating spontaneous activity. Hum. Brain Mapp. *34*, 1319–1329.

Raichle, M.E., MacLeod, A.M., Snyder, A.Z., Powers, W.J., Gusnard, D.A., and Shulman, G.L. (2001). A default mode of brain function. Proc. Natl. Acad. Sci. USA 98, 676–682.

Rosanova, M., Gosseries, O., Casarotto, S., Boly, M., Casali, A.G., Bruno, M.A., Mariotti, M., Boveroux, P., Tononi, G., Laureys, S., and Massimini, M. (2012). Recovery of cortical effective connectivity and recovery of consciousness in vegetative patients. Brain *135*, 1308–1320.

Selimbeyoglu, A., and Parvizi, J. (2010). Electrical stimulation of the human brain: perceptual and behavioral phenomena reported in the old and new literature. Front. Hum. Neurosci. *4*, 46.

Smith, S.M., Fox, P.T., Miller, K.L., Glahn, D.C., Fox, P.M., Mackay, C.E., Filippini, N., Watkins, K.E., Toro, R., Laird, A.R., and Beckmann, C.F. (2009). Correspondence of the brain's functional architecture during activation and rest. Proc. Natl. Acad. Sci. USA *106*, 13040–13045.

Sporns, O. (2014). Contributions and challenges for network models in cognitive neuroscience. Nat. Neurosci. 17, 652–660.

Sporns, O., and Zwi, J.D. (2004). The small world of the cerebral cortex. Neuroinformatics 2, 145–162.

Sporns, O., Tononi, G., and Edelman, G.M. (2000). Theoretical neuroanatomy: relating anatomical and functional connectivity in graphs and cortical connection matrices. Cereb. Cortex *10*, 127–141.

Sporns, O., Tononi, G., and Kötter, R. (2005). The human connectome: a structural description of the human brain. PLoS Comput. Biol. 1, e42.

Tagliazucchi, E., Balenzuela, P., Fraiman, D., and Chialvo, D.R. (2012). Criticality in large-scale brain FMRI dynamics unveiled by a novel point process analysis. Front. Physiol. 3, 15.

Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B., and Joliot, M. (2002). Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. Neuroimage 15, 273–289.

van Hartevelt, T.J., Cabral, J., Møller, A., FitzGerald, J.J., Green, A.L., Aziz, T.Z., Deco, G., and Kringelbach, M.L. (2015). Evidence from a rare case study for Hebbian-like changes in structural connectivity induced by long-term deep brain stimulation. Front. Behav. Neurosci. 9, 167.

Wang, X.J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. Neuron *36*, 955–968.

Watts, D.J., and Strogatz, S.H. (1998). Collective dynamics of 'small-world' networks. Nature 393, 440–442.

Winawer, J., and Parvizi, J. (2016). Linking electrical stimulation of human primary visual cortex, size of affected cortical area, neuronal responses, and subjective experience. Neuron *92*, 1213–1219.

Zhang, D., and Raichle, M.E. (2010). Disease and the brain's dark energy. Nat. Rev. Neurol. 6, 15–28.