Theoretical principles of multiscale spatiotemporal control of neuronal networks: a complex systems perspective

Nima Dehghani[1,2]
[1)] *Massachusetts Institute of Technology, Department of Physics, Cambridge, MA*
[2)] *New England Complex Systems Institute, Cambridge, MA*[a)]

(Dated: February 10, 2017)

Success in the fine control of the nervous system depends on a deeper understanding of how neural circuits control behavior. There is, however, a wide gap between the components of neural circuits and behavior. We advance the idea that a suitable approach for narrowing this gap has to be based on a multiscale information-theoretic description of the system. We evaluate the possibility that brain-wide complex neural computations can be dissected into a hierarchy of computational motifs that rely on smaller circuit modules interacting at multiple scales. In doing so, we draw attention to the importance of formalizing the goals of stimulation in terms of neural computations so that the possible implementations are matched in scale to the underlying circuit modules.

Keywords: Complexity, Information theory, Dynamical systems, Networks, Control theory

A. Overview

In this theoretical perspective article, we propose the need for a multiscale information theoretic framework on how to better control an adaptive complex system such as the mammalian nervous system. We start with a brief historical overview of neural stimulation from Galvani's pioneering work to modern opto-electrical methods. Pointing to control obstacles, we portray why an understanding of the information processing levels of neuronal networks is crucial to the proper design of a control paradigm. We then emphasize the importance of scale-independence and examine the shortcomings of network control without it. This interpretation lays out the need for attention to the dynamic nature of computation and the functional robustness in the light of structure variability thus illustrating the multiscale information processing nature of neuronal networks. We advance the idea that, within this framework, better control can be achieved through targeting the aggregate computational output rather than attempting to control the system at its finest scales.

B. Control obstacles and failures of neural stimulation

In 1780, Luigi Galvani discovered that an electrical spark causes the twitching of a dead frog's legs. This discovery was pivotal to the birth of bioelectromagnetism and led the to the idea of controlling action and behavior with electricity[13,97]. A century after Galvani's famous experiment, Gustav Fritsch and Eduard Hitzig showed that, in dogs, *in vivo* motor cortex stimulation causes limb movement[30]. Electrical neuromodulation as a therapeutic tool has since been recognized as a viable and concrete path to control neurological diseases[33]. Recognition of the implication of the Substantia Nigra's damaged cells in Parkinson's disease[46], better understanding of the operational principles of basal ganglia-thalamocortical loops[1] and the effective reduction of tremor by lesioning the subthalamic nucleus in monkey model of Parkinson's[8] provided the path to tackle Parkinson's. Following these discoveries and a century after Fritsch and Hitzig's experiments the first deep brain stimulation (DBS) operation was performed in Parkinsonian patients in 1987 and positive clinical effects were reported a few years later[44,51].

However, despite the technological advancements in electrical stimulation, control of the behavior has been plagued by improper precision and lack of understanding of the network response. In the case of macro-stimulation (such as TMS – transcranial magnetic stimulation, tDCS – transcranial direct stimulation, DBS), much of the effort is focused on improved targeting of a smaller area. The hope is that via fine-tuned targeting, one would eventually achieve the proper level of control. Nonetheless, electrical micro-stimulation has not shown much promise in precise control of the output of circuits. For example, microstimulation of MT (middle temporal visual area) directional-selective neurons shows variable behavioral efficacy across single trials, where this variability has been ascribed to attention gating[66] or spatial feature-selective gating[68,84]. A major issue with the electrical micro-stimulation is the imprecision in targeting a specific cell-group (excitatory vs inhibitory, or a given group of inhibitory cells) as well as the nebulous temporal control of the stimulation effect.

The discovery of optogenetics (use of light to control neuron) in cultured mammalian neurons[98] triggered the next wave of neuromodulatory attempts through temporally precise stimulation of individual excitatory and/or inhibitory neurons with millisecond resolution[12,99]. An ever-growing list of studies relying on optogenetics stim-

[a)] Electronic mail:
nima.dehghani@mit.edu
nima.dehghani@wyss.harvard.edu

ulation have since followed, including those targeting Parkinson's[43], aimed at behavioral conditioning[96] or fear-conditioning[36], as well as targeting deep structures (brainstem)[52], or for inducing fast (γ) rhythms[17], modulating sensory processing[87] and attempts to control network disorder (hippocampal seizure)[45].

However, the attempts for control through optogenetics has faced many challenges. The observed variability of the evoked behavior following transient inactivation of the motor cortex in rats and nucleus interface (Nif) in songbirds[72] has seriously challenged the assumption that optogenetics can effectively overcome nonspecific cell targeting and temporal imprecision. As proposed, ignoring the indirect effects of downstream circuits[72] and disregarding the interconnectedness of the complex circuitry are among the key reasons that even a temporally-precise cell-specific control of individual elements (i.e. neurons) of the network can not yield precise control of the network function. The macroscopic behavior of the system (such as network balance of excitation/inhibition) is insensitive to the computational state of individual neurons[22]. This insensitivity is not because the functional symmetry of individual elements transcends to the total state[2], but because interconnectedness renders many details (at fine scale) to be irrelevant at the large-scale behavior of the system[34]. Thus, attempts for precise control of the system at its fine scale is precisely where it will fail. This failure roots in the breakdown of the constructionist view and its emphasis on the independence of behavior at microscopic and macroscopic scales[2]. In what follows, we emphasize on a the need for a paradigm shift to understand and control the system in a multiscale information theoretic framework.

C. Levels of information processing

Proper success in augmentation/alteration of the nervous system's behavior roots in the deep understanding of the complexity of this computational system. Here, by complexity we explicitly refer to a) existence of multiple scales of structure/dynamics, b) although the structural elements of coarser scales can be physically reduced to those of the finer scales, the macro-micro dynamical elements are neither fully independent nor completely coherent, and c) the number of possible states of the system (i.e the amount of information – in bits– needed to define the system) is scale-dependent. Early attempts to describe the nervous system at multiple scales, led to the formation of a tri-level hypothesis of information processing. This approach divides the description of the system into its computational, algorithmic, and implementational subsets after David Marr[62,63] or to semantic, syntactic, and physical categories according to[78]. While this view in neuroscience and philosophy/cognitive science has had a tremendous impact (specifically in the field of visual neuroscience), the missing links across the tri-levels have been the target of criticisms and a source of observed failures in the description of neural computation at the system level and its connection to behavior. Interestingly, even though Marr initiated this classification of levels, his own studies were always focused on only one level at a time. In fact, Marr's own quest was initially more focused on the fine scale level[58–60], but later he became skeptical[61] and switched to the top, i.e. computational, level[64,81]. The main issue in this approach is reliance on the separation of scales, a notion deeply in conflict with the nature of complex systems[6]. The behavior of the systems harboring hierarchical levels of sub-assemblies is defined by the interaction of the sub-assemblies at higher levels not by details in a given sub-assembly[88,89]. Surprisingly, a recent reversal of interest in the implementational level (mainly due to advances in microscopy and computational tools) has led to the revival of the constructionist view. This approach entails that by taking into account the adjacent biophysical details of cell type categories, their placement in micro-circuitry and their afferent and efferent projections, one could achieve proper control/alteration of neural information processing systems.

The recent surge of interest in the fine level details of both structure[14,39] and simulation[57] parallels Marr's initial emphasis on fine implementational details. This approach has drawn criticism from advocates for his later emphasis on the computational level. The underlying assumption is that through a combined study of in vivo physiology and network anatomy[9], one can build functional connectomics[85] and decipher the behavior of the system. In small systems lacking the hierarchical architecture of complexity, separation of scales is justifiable and such an approach may lead to a good understanding of the system's behavior (such as Dorsophila motion detection[94]), or yield relatively good control of simple behaviors (such as optogenetic control of simple motor in C elegans[49]). The opposing view – mainly driven by system neuroscience investigations of the mammalian neocortex – argues that the understanding the underlying details will not resolve the issues relevant to the neural computation of interest[15]. This vantage point suggests that a repetition of canonical computation such as linear filtering[69,82] or divisive normalization[37,90,101] serves as building blocks of computation and is the cornerstone of information processing[16] leading to progressively more complex representation in the hierarchy of sensory processing cortical areas[42]. It is suggested that such canonical computations themselves could be embedded in canonical circuits[24,26]. The canonical circuit is thus portrayed as the structural component of the computational unit[25] where slight modifications of either the hardware or software can shape a rich repertoire of network output[35,65]. Thus, an argument for straying away from the implementational level is rooted in computational identity of different circuits (across species) in the face of the multiplicity of their physical implementation[76,95].

We suggest that for fine scale information (connectome or biophysically-detailed simulations) to provide any sig-

nificant insight about the nature of computation, they should be further examined within a scale-dependent framework and considering:

The observation scale: Although one can ignore detailed connectivity profile as unnecessary while considering the overall statistical properties as important, the proper choice of the scale remains crucial.

The structure-function relationship: Physical details of the circuit affect the dynamics, suggesting that the structure constrains the computational function. However, constraining the space of possible computations does not provide an exact characterization of the performed computations because there exists no one-to-one mapping of structure-function.

Plasticity: The ever-changing structure of the network implies a robustness in computational function despite changes in fine details of connectivity.

Subcellular control elements: These elements (whether they are gap junctions, ion channels, dendritic spines, etc) provide mechanisms for change of function without apparent change of structure at the scale of microcircuitry.

We have to recognize that since our system of interest harbors plasticity and homeostasis among other fundamental features outlined above (e.g. its hierarchy). These two aspects demand that our manipulations of the system be dynamic in nature in order to achieve that the intended effect is sustained over time despite the changing conditions (plasticity) and overall constraints (homeostasis).

D. What does network control entail?

In the search for a proper spatiotemporal dynamics of control, some may resort to modern interpretations of network control. It is essential to recognize the shortcomings of such an admixture of computational and implementational level. Although network structure determines certain properties of network dynamics such as limit-cycle oscillator synchrony[91] or the likelihood of reliable dynamical attractors[40], not all dynamics can be captured by the network structure alone. Some studies have considered a connectivity graph of complex networks equivalent to its dynamical nature, and have drawn the conclusion that recognizing the driver nodes is sufficient for understanding the strategy for controlling the network[54]. By reducing the dynamics to structural connectivity, they conclude that a large fraction of driver nodes (80%) are needed to control biological systems. Even though they point to the difficulty of controlling sparse inhomogeneous networks, in comparison to dense homogeneous ones[54], their assumptions of reduction of dynamics to network structure and the definition of control and numerical methods in defining driver nodes are criticized by a) the evidence that a few inputs can reprogram biological networks[70], b) trade-off between phase space nonlocality of the control trajectory and control input nonlocality[92] and c) that the node dynamics – not degree distributions – define the nature of controllability[19].

Another approach in advocating the structural control has been based on the assumption that active nodes can simultaneously activate all its connected neighbors but no further than that[71]. This assumption that is neither valid in the nervous system nor has any relevance to biological networks composed of elements with a variety of time constants and delayed communication. Additionally, not only the internal structure but also its connectivity pattern to the drivers and the depth of the network are essential in the emergent dynamics that follow stimulation. If the network is composed of a high external/internal node ratio, it will be largely influenced by the outside and its individual nodes will have a greater degree of independent behavior from the external stimulus. In contrast, if within system links are strong, then the system moves toward synchrony[18]. Likewise, the depth of the system is very crucial in its response to incoming stimuli. Systems with shallow depth are easier to force to behave in the desired way, as we can directly probe and influence the effects of external nodes on internal ones. In contrast, deep networks are harder to control, and along each step we have to tweak the internal vs external nodes to achieve the desired outcome. Overall, in high-dimensional systems that are governed by nonlinear dynamics, have dissipative properties (trajectories are confined only to a limited part of the permissible phase space), and where feedback imposes constraints on controllability, the mere identification of driver nodes and quantifying node variables are not sufficient to control the network[32,67]. Thus the control of complex networks requires both structure and dynamics across multiple scales of the system.

E. Dynamical nature of computation and structure/function variability

The alterations in network behavior can originate from constant structural reshaping of the network (addition or deletion of links), or through modified intrinsic properties of the network nodes (i.e. neurons) and via changes in possible states of the interaction among the network nodes (i.e. synaptic weights). These different modes of reconfiguration of the network behavior set the dynamics of the information flow and the active behavior of the organism. A functional translation of these attributes leads to few rules, where a) one neuron can be involved in more than one behavior, b) one behavior can involve several circuits and c) one neuromodulator can alter multiple circuits. Naturally the involvement of these different modes in reconfiguration of the network behavior changes according to the size of the system and the repertoire of its functional states. In non-vertebrates with a

very limited set of neurons, it is possible to pinpoint individual or several neurons as the key or even unique elements involved in highly reliable functions. For example, it is suggested that a single neuron in C elegans can be involved in highly reliable function[11]. In vertebrates (with larger brains), where routing and coordination of information flow becomes increasingly more important, local or global effects of neuromodulators and oscillatory rhythms play an essential role[41,48]. One can conclude that precision in control of the system depends on its complexity profile (i.e. the amount of information necessary to represent a system as a function of scale). As the number of neurons/circuits and diversity of network components increases, the reliability of functional dependence on individual components decreases. This increased complexity provides robustness at larger scale dynamics while it forgoes the details. Assessing the variability of structure/function thus becomes an essential aspect for directing stimulus at the right scale.

While the observed variability can be externally or intrinsically-driven[80], in order to achieve functional adaptability in the light of external variability, the nervous system benefits from adapting stochastic responses[27] where a given population activity can be construed as a likely sample from a posteriori distribution[38]. Intrinsic variability in the nervous system originates at different levels, from ion channels[29] to synapses[23], single cells[53,79] and population level[10]. As a result, in driving the system, instead of aiming for a pre-determined response, the target should be to induce population activity such that it could be interpreted as one stochastic instance of the likely probability distribution. This supposed probability distribution of set of responses should depend on a multidimensional parameter space, where robustness and flexibility are attained through overlapping redundant functions. This form of redundancy is in stark contrast to engineered systems where multiple instances of identical copies are used to guarantee the robustness[56]. Instead, in the biological systems (and specifically in the nervous system), the robustness is achieved through a) variability at fine scales, where single cell variability exists in parallel with a robust representation at the population level[22,53,79], b) networks with different configurations (of the underlying parameters) manifest robust response to neuromodulators[77] and c) structurally different circuits respond reliably to external perturbation[55]. The existence of distinct stable basins of attraction despite the permissible extensive variability at fine scales, is an essential characteristic of complex systems. Such systems' response to perturbation depends on the scale and the extent of the stimuli. In a sense, tolerance for small errors in complex systems comes at the price of the intolerance to large errors[47]. As a result, when the intrinsic variability is reduced due to excessively increased coupling, the abnormally greater ensemble correlation dramatically reduces the plethora of macroscopic states, a situation which is the hallmark of loss of complexity in organs with networks of excitable cells. Heart arrhythmia[74,75] or seizures[22] are examples of such loss of complexity.

F. Are complex systems with fine-scale variability and robust macroscopic features predictable?

In simple chaotic systems, the lack of predictability is due to sensitivity to initial conditions. In contrast, in complex adaptive systems, information exchange across the scales is the main obstacle in the proper control of the system and the lack of predictability is due to relevant interactions and novel information created by these interactions. Ignoring the multiscale levels of information processing, one may jump to the conclusion that for proper control, we need to have a detailed knowledge of the individual elements, the initial conditions and their interaction among elements. Since we know that this is impractical if not impossible, can we do better? To properly control a system through stimulation, one has to have a proper knowledge of the interaction in the subunits of the system. However, only a controlled trajectory of the macroscopic states is desirable and yet achievable. From a dynamical system's point of view, pushing the system at the right time and the right scale is the key to fine control. Given the scale-dependent interaction of the nervous system, we should ask: a) are complex systems with many scales inherently unpredictable? and b) if the answer is no, then how do we manipulate the system in a predictable way? Surprisingly, the answer to the first question is ''no, complex systems are not unpredictable". However the predictability requires a few specific conditions and particulars. The existence of nonlinear feedback creates either a fast or slow ''wait-time of divergence" and therefore, up to some point, prediction holds and it rapidly deteriorates afterwards. This is similar to the behavior of some bifurcation systems. In fact, many systems that are considered to be periodic are inherently chaotic. A prime example is the solar system where its very slow divergence time (4 million years) leads to the observation of chaotic behavior only when 100 million years of the entire solar system are examined[93]. Therefore to navigate the system in the aimed trajectory, we have to constantly adjust and push the system such that it does not deviate from the desired path after the implementation of the last stimulus. This strategy can be effective only if the system does not have low-dimensional chaotic behavior. Otherwise, any simple perturbation could lead to unrecoverable massive perturbation. In the case of the ensemble activity (such as cortical processing with the involvement of a much larger set of neurons in comparison to the simple nervous system of invertebrates such as C Elegans), this notion becomes highly relevant since low dimensional neural trajectories provide surprisingly accurate portrait of the circuit dynamics[79] and dimensionality reduction methods exhibit effectiveness in population decoding[20].

In non-chaotic system, our ability to modify the sys-

tem is limited[86]. If the system's behavior is periodic or multiply periodic (nested periodicity), we are stuck with the system's intrinsic periodicity dynamics. We can slightly change the orbit by small perturbations or we can only induce perturbations that would switch the dominant periodicity from one of the existent orbits to another existent one. This limitation also means that the achieved effect is robust because of the existent orbits, but it will not permit large alterations in the system. In contrast, the existence of chaos could be helpful in controlling a chaotic system. Since systems with a chaotic attractor have an infinite number of unstable periodic orbits, one can exploit this property and push the system towards one of the already existing unstable periodic orbits[73]. However, if the system is high dimensional evolving chaotically on a low dimensional attractor, this method no longer works. The reason for the ineffectiveness of this method in this case is that one would require an excessive amount of information extracted from the data and a very long history of the dynamics in order to be able to properly achieve control[4]. Though, interestingly, the situation is not hopeless, as one can exploit the low-dimensionality of the system dynamics and through a feedback control and repeated application of tiny perturbations, control the system[4]. Here we advance the idea that recurrent feedback in the neuronal networks might indeed be the evolutionary mechanism developed to specifically deal with the low dimensional dynamics of a high dimensional system for providing intrinsic control for response reliability. This method has been extended to control excitable biological systems in the past. It has been shown that by irregular (based on the chaotic time) delivery of the electrical stimulation, the cardiac arrhythmia can be pushed back to a low-order periodic regime[31]. In the hippocampal slice (CA3), the same method of entrained spontaneous burst discharges proved to be much more effective than periodic control[83]. Other methods, such as periodic[5,50] or stochastic[28] stimulation can affect chaotic systems too but it is hard to predict their effect for networks with many layers.

G. Concluding remarks

In linking the structure and dynamics of the neuronal network, mapping the space of possible interactions becomes important and a matter of challenge. Since all the possible interactions at a fine scale occupy a vast multidimensional space, in order to have robust behavior, the system is likely to rely on a more limited set of probable interactions between local densities (i.e. ensemble activity). As the scale grows, the set of possible interactions decreases yet the outcome of such interactions better matches with the behavioral (macroscopic) outputs of the system. Relevant parameters are those permissible sets of interactions that have increased probability of occurrence as the scale increases. Functional state transitions depend on the spatial variation and interaction of the ensemble activity densities. As a result, the complexity profile defined as the amount of information necessary to represent a system as a function of scale gives us the number of possible states of the system at a particular scale[6]. Therefore the finer the computational scale of the system, the more information is needed to describe it. It is through these mechanisms that the system can maintain its intrinsic dynamical balance yet manifest responsiveness across multiple time scales[22] and provide stereotypical macroscopic spatiotemporal patterns in the lights of microscopic variability[48]. This viewpoint defies the blind big-data approach of incorporating more details with the hope of yielding a better model and more accurate prediction[7]. Even with the assumption that at some point in the future we can map the connectivity and the activity of all the elements of the neuronal network, we are still in need of an information-theoretic formalism that shows how the output is sensitive to fine scale perturbations, and how the coarse-scale reflects redundancy and synergy of the aggregate activity of finer scales[21]. In order to enhance the effectiveness of targeting multiscale neural systems, alterations of biophysically-based features should match the level of the desired computation. In targeting a stream of computation, the link between the computation and the architecture is what defines the optimal solution for maximizing the efficiency of the stimulation. In summary, to better control the system, we have to focus on the information transfer across multiple scales. Only with this approach can engineering advancements in precise opto-electric stimulation open ways to alter the system in the desired way.

REFERENCES

[1] Alexander, G. E., Crutcher, M. D., and DeLong, M. R., in *Progress in Brain Research* (Elsevier BV, 1991) pp. 119–146.
[2] Anderson, P. W., Science **177**, 393 (1972).
[3] Ashby, W. R., Cybernetica **1**, 83 (1958).
[4] Auerbach, D., Grebogi, C., Ott, E., and Yorke, J. A., Phys. Rev. Lett. **69**, 3479 (1992).
[5] Azevedo, A. and Rezende, S. M., Phys. Rev. Lett. **66**, 1342 (1991).
[6] Bar-Yam, Y., *Dynamics Of Complex Systems (Studies in Nonlinearity)* (Westview Press, 2003).
[7] Bar-Yam, Y., Complexity (2016), 10.1002/cplx.21785.
[8] Bergman, H., Wichmann, T., and DeLong, M., Science **249**, 1436 (1990).
[9] Bock, D. D., Lee, W.-C. A., Kerlin, A. M., Andermann, M. L., Hood, G., Wetzel, A. W., Yurgenson, S., Soucy, E. R., Kim, H. S., and Reid, R. C., Nature **471**, 177 (2011).
[10] Boerlin, M., Machens, C. K., and Denève, S., PLoS Comput Biol **9**, e1003258 (2013).
[11] de Bono, M. and Maricq, A. V., Annu. Rev. Neurosci. **28**, 451 (2005).
[12] Boyden, E. S., Zhang, F., Bamberg, E., Nagel, G., and Deisseroth, K., Nature Neuroscience **8**, 1263 (2005).
[13] Bresadola, M., Brain Research Bulletin **46**, 367 (1998).
[14] Briggman, K. L., Helmstaedter, M., and Denk, W., Nature **471**, 183 (2011).
[15] Carandini, M., Nature Neuroscience **15**, 507 (2012).
[16] Carandini, M. and Heeger, D. J., Nature Reviews Neuroscience (2011), 10.1038/nrn3136.

[17]Cardin, J. A., Carlén, M., Meletis, K., Knoblich, U., Zhang, F., Deisseroth, K., Tsai, L.-H., and Moore, C. I., Nature **459**, 663 (2009).

[18]Chinellato, D. D., Epstein, I. R., Braha, D., Bar-Yam, Y., and de Aguiar, M. A. M., Journal of Statistical Physics **159**, 221 (2015).

[19]Cowan, N. J., Chastain, E. J., Vilhena, D. A., Freudenberg, J. S., and Bergstrom, C. T., PLoS ONE **7**, e38398 (2012).

[20]Cunningham, J. P. and Yu, B. M., Nature Neuroscience **17**, 1500 (2014).

[21]Daniels, B. C., Ellison, C. J., Krakauer, D. C., and Flack, J. C., Current Opinion in Neurobiology **37**, 106 (2016).

[22]Dehghani, N., Peyrache, A., Telenczuk, B., Quyen, M. L. V., Halgren, E., Cash, S. S., Hatsopoulos, N. G., and Destexhe, A., Sci. Rep. **6**, 23176 (2016).

[23]Destexhe, A. and Rudolph-Lilith, M., *Neuronal Noise* (Springer Science Business Media, 2012).

[24]Douglas, R. J. and Martin, K. A., Annu. Rev. Neurosci. **27**, 419 (2004).

[25]Douglas, R. J. and Martin, K. A., Neuron **56**, 226 (2007).

[26]Douglas, R. J., Martin, K. A., and Whitteridge, D., Neural Computation **1**, 480 (1989).

[27]Ernst, M. O. and Banks, M. S., Nature **415**, 429 (2002).

[28]Fahy, S. and Hamann, D. R., Phys. Rev. Lett. **69**, 761 (1992).

[29]Faisal, A. A., Selen, L. P. J., and Wolpert, D. M., Nature Reviews Neuroscience **9**, 292 (2008).

[30]Fritsch, G. and Hitzig, E., Epilepsy & Behavior **15**, 123 (2009).

[31]Garfinkel, A., Spano, M., Ditto, W., and Weiss, J., Science **257**, 1230 (1992).

[32]Gates, A. J. and Rocha, L. M., Sci. Rep. **6**, 24456 (2016).

[33]Gildenberg, P. L., Stereotactic and Functional Neurosurgery **83**, 71 (2005).

[34]Goldenfeld, N. and Kadanoff, L. P., Science **284**, 87 (1999).

[35]Harris, K. D. and Shepherd, G. M. G., Nature Neuroscience **18**, 170 (2015).

[36]Haubensak, W., Kunwar, P. S., Cai, H., Ciocchi, S., Wall, N. R., Ponnusamy, R., Biag, J., Dong, H.-W., Deisseroth, K., Callaway, E. M., Fanselow, M. S., Lüthi, A., and Anderson, D. J., Nature **468**, 270 (2010).

[37]Heeger, D. J., Vis Neurosci **9**, 181 (1992).

[38]Hoyer, P. O. and Hyvärinen, A., in *Advances in Neural Information Processing Systems 15*, edited by S. Becker, S. Thrun, and K. Obermayer (MIT Press, 2003) pp. 293–300.

[39]Kleinfeld, D., Bharioke, A., Blinder, P., Bock, D. D., Briggman, K. L., Chklovskii, D. B., Denk, W., Helmstaedter, M., Kaufhold, J. P., Lee, W.-C. A., Meyer, H. S., Micheva, K. D., Marcel Oberlaender, S. P., Reid, R. C., Smith, S. J., Takemura, S., Tsai, P. S., and Sakmann, B., Journal of Neuroscience **31**, 16125 (2011).

[40]Klemm, K. and Bornholdt, S., Proceedings of the National Academy of Sciences **102**, 18414 (2005).

[41]Kopell, N. J., Gritton, H. J., Whittington, M. A., and Kramer, M. A., Neuron **83**, 1319 (2014).

[42]Kouh, M. and Poggio, T., Neural Computation **20**, 1427 (2008).

[43]Kravitz, A. V., Freeze, B. S., Parker, P. R. L., Kay, K., Thwin, M. T., Deisseroth, K., and Kreitzer, A. C., Nature **466**, 622 (2010).

[44]Kringelbach, M. L., Jenkinson, N., Owen, S. L., and Aziz, T. Z., Nature Reviews Neuroscience **8**, 623 (2007).

[45]Krook-Magnuson, E., Armstrong, C., Oijala, M., and Soltesz, I., Nature Communications **4**, 1376 (2013).

[46]Langston, J., Ballard, P., Tetrud, J., and Irwin, I., Science **219**, 979 (1983).

[47]Laughlin, R. B., Neuron **83**, 1253 (2014).

[48]Le Van Quyen, M., Muller, L. E., Telenczuk, B., Halgren, E., Cash, S., Hatsopoulos, N. G., Dehghani, N., and Destexhe, A., Proceedings of the National Academy of Sciences **113**, 9363 (2016).

[49]Leifer, A. M., Fang-Yen, C., Gershow, M., Alkema, M. J., and Samuel, A. D. T., Nature Methods **8**, 147 (2011).

[50]Lima, R. and Pettini, M., Phys. Rev. A **41**, 726 (1990).

[51]Limousin, P., Pollak, P., Benazzouz, A., Hoffmann, D., Bas, J.-F. L., Perret, J., Benabid, A.-L., and Broussolle, E., The Lancet **345**, 91 (1995).

[52]Lin, J. Y., Knutsen, P. M., Muller, A., Kleinfeld, D., and Tsien, R. Y., Nature Neuroscience **16**, 1499 (2013).

[53]Litwin-Kumar, A. and Doiron, B., Nature Neuroscience **15**, 1498 (2012).

[54]Liu, Y.-Y., Slotine, J.-J., and Barabási, A.-L., Nature **473**, 167 (2011).

[55]Marder, E., Proceedings of the National Academy of Sciences **108**, 15542 (2011).

[56]Marder, E. and Goaillard, J.-M., Nature Reviews Neuroscience **7**, 563 (2006).

[57]Markram, H., Muller, E., Ramaswamy, S., Reimann, M. W., Abdellah, M., Sanchez, C. A., Ailamaki, A., Alonso-Nanclares, L., Antille, N., Arsever, S., Kahou, G. A. A., Berger, T. K., Bilgili, A., Buncic, N., Chalimourda, A., Chindemi, G., Courcol, J.-D., Delalondre, F., Delattre, V., Druckmann, S., Dumusc, R., Dynes, J., Eilemann, S., Gal, E., Gevaert, M. E., Ghobril, J.-P., Gidon, A., Graham, J. W., Gupta, A., Haenel, V., Hay, E., Heinis, T., Hernando, J. B., Hines, M., Kanari, L., Keller, D., Kenyon, J., Khazen, G., Kim, Y., King, J. G., Kisvarday, Z., Kumbhar, P., Lasserre, S., Bé, J.-V. L., Magalhães, B. R., Merchán-Pérez, A., Meystre, J., Morrice, B. R., Muller, J., Muñoz-Céspedes, A., Muralidhar, S., Muthurasa, K., Nachbaur, D., Newton, T. H., Nolte, M., Ovcharenko, A., Palacios, J., Pastor, L., Perin, R., Ranjan, R., Riachi, I., Rodríguez, J.-R., Riquelme, J. L., Rössert, C., Sfyrakis, K., Shi, Y., Shillcock, J. C., Silberberg, G., Silva, R., Tauheed, F., Telefont, M., Toledo-Rodriguez, M., Tränkler, T., Geit, W. V., Díaz, J. V., Walker, R., Wang, Y., Zaninetta, S. M., DeFelipe, J., Hill, S. L., Segev, I., and Schürmann, F., Cell **163**, 456 (2015).

[58]Marr, D., The Journal of Physiology **202**, 437 (1969).

[59]Marr, D., Proceedings of the Royal Society B: Biological Sciences **176**, 161 (1970).

[60]Marr, D., Philosophical Transactions of the Royal Society B: Biological Sciences **262**, 23 (1971).

[61]Marr, D., Science **190**, 875 (1975).

[62]Marr, D., *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information.* (W.H. Freeman, San Francisco, 1982).

[63]Marr, D. and Poggio, T., "From Understanding Computation to Understanding Neural Circuitry (A.I. Memo 357)," Tech. Rep. (Massachusetts Institute of Technology, 1976).

[64]McClamrock, R., Minds and Machines **1**, 185 (1991).

[65]Miller, K. D., Current Opinion in Neurobiology **37**, 75 (2016).

[66]Moran, J. and Desimone, R., Science **229**, 782 (1985).

[67]Motter, A. E., Chaos **25**, 097621 (2015).

[68]Motter, B. C., J Neurosci **14**, 2190 (1994).

[69]Movshon, J., Thompson, I., and Tolhurst, D., J Physiol **283**, 53 (1978).

[70]Müller, F.-J. and Schuppert, A., Nature **478**, E4 (2011).

[71]Nacher, J. C. and Akutsu, T., Sci. Rep. **3** (2013), 10.1038/srep01647.

[72]Otchy, T. M., Wolff, S. B. E., Rhee, J. Y., Pehlevan, C., Kawai, R., Kempf, A., Gobes, S. M. H., and Ölveczky, B. P., Nature **528**, 358 (2015).

[73]Ott, E., Grebogi, C., and Yorke, J. A., Phys. Rev. Lett. **64**, 1196 (1990).

[74]Pincus, S. M. and Goldberger, A. L., Am J Physiol **266**, H1643 (1994).

[75]Pincus, S. M. and Singer, B. H., Proceedings of the National Academy of Sciences **93**, 2083 (1996).

[76]Priebe, N. J. and Ferster, D., Neuron **75**, 194 (2012).

[77]Prinz, A. A., Bucher, D., and Marder, E., Nature Neuroscience **7**, 1345 (2004).

[78]Pylyshyn, Z. W., *Computation and cognition: toward a foundation for cognitive science* (MIT Press, 1984).

[79]Rabinovich, M., Huerta, R., and Laurent, G., Science **321**, 48 (2008).
[80]Renart, A. and Machens, C. K., Current Opinion in Neurobiology **25**, 211 (2014).
[81]Rolls, E. T., Brain **134**, 913 (2011).
[82]Rust, N. C., Mante, V., Simoncelli, E. P., and Movshon, J. A., Nature Neuroscience **9**, 1421 (2006).
[83]Schiff, S. J., Jerger, K., Duong, D. H., Chang, T., Spano, M. L., and Ditto, W. L., Nature **370**, 615 (1994).
[84]Seidemann, E., Zohary, E., and Newsome, W. T., Nature **394**, 72 (1998).
[85]Seung, H. S., Nature **471**, 170 (2011).
[86]Shinbrot, T., Grebogi, C., Yorke, J. A., and Ott, E., Nature **363**, 411 (1993).
[87]Shusterman, R., Smear, M. C., Koulakov, A. A., and Rinberg, D., Nature Neuroscience **14**, 1039 (2011).
[88]Simon, H. A., Proceedings of the American Philosophical Society **106**, 467 (1962).
[89]Simon, H. A., "The Sciences of the Artificial," (M.I.T. Press, 1969) Chap. The Architecture of Complexity.
[90]Solomon, S. G., Lee, B. B., and Sun, H., Journal of Neuroscience **26**, 8715 (2006).
[91]Strogatz, S. H., Nature **410**, 268 (2001).
[92]Sun, J. and Motter, A. E., Phys. Rev. Lett. **110** (2013), 10.1103/physrevlett.110.208701.
[93]Sussman, G. J. and Wisdom, J., Science **257**, 56 (1992).
[94]Takemura, S.-y., Bharioke, A., Lu, Z., Nern, A., Vitaladevuni, S., Rivlin, P. K., Katz, W. T., Olbris, D. J., Plaza, S. M., Winston, P., Zhao, T., Horne, J. A., Fetter, R. D., Takemura, S., Blazek, K., Chang, L.-A., Ogundeyi, O., Saunders, M. A., Shapiro, V., Sigmund, C., Rubin, G. M., Scheffer, L. K., Meinertzhagen, I. A., and Chklovskii, D. B., Nature **500**, 175 (2013).
[95]Tank, D., seminars in The Neurosciences **1**, 67 (1989).
[96]Tsai, H.-C., Zhang, F., Adamantidis, A., Stuber, G. D., Bonci, A., de Lecea, L., and Deisseroth, K., Science **324**, 1080 (2009).
[97]Whittaker, E. T., *A history of the theories of aether and electricity from the age of Descartes to the close of the nineteenth century* (Smithsonian Institution, 1910).
[98]Zemelman, B. V., Lee, G. A., Ng, M., and Miesenböck, G., Neuron **33**, 15 (2002).
[99]Zhang, F., Wang, L.-P., Brauner, M., Liewald, J. F., Kay, K., Watzke, N., Wood, P. G., Bamberg, E., Nagel, G., Gottschalk, A., and Deisseroth, K., Nature **446**, 633 (2007).
[100]Zikopoulos, B. and Barbas, H., PLoS ONE **2**, e848 (2007).
[101]Zoccolan, D., Cox, D. D., and DiCarlo, J. J., Journal of Neuroscience **25**, 8150 (2005).