

Reading and writing the neural code

Garrett B Stanley

It has been more than 20 years since Bialek and colleagues published a landmark paper asking a seemingly innocuous question: what can we extract about the outside world from the spiking activity of sensory neurons? Can we read the neural code? Although this seemingly simple question has helped us shed light on the neural code, we still do not understand the anatomical and neurophysiological constraints that enable these codes to propagate across synapses and form the basis for computations that we need to interact with our environment. The sensitivity of neuronal activity to the timing of synaptic inputs naturally suggests that synchrony determines the form of the neural code, and, in turn, regulation of synchrony is a critical element in ‘writing’ the neural code through the artificial control of microcircuits to activate downstream structures. In this way, reading and writing the neural code are inextricably linked.

Neurons convey information through complex temporal patterns of action potentials across neuronal ensembles. Without sufficient prior knowledge, observers of this activity would be left to guess as to what a particular spatiotemporal pattern represents. In this way, patterns of neuronal activity relate to tangible features of our external environment, a relationship we often refer to as the neural code.

Reading the neural code

One clear litmus test as to whether we truly understand the neural code is whether we can tap into the activity of the neurons and make clear predictions about what is going on in the outside world or what is about to go on through the actions of the organism. We might refer broadly to this as our attempt to read the neural code. Upon observing a pattern of activity in the brain, we seek a dictionary of sorts, so that we may, for example, interpret a pattern in the visual pathway as representing a tree or a dog, or a pattern of activation in the motor pathway as representing a motor output such as an eye or hand movement. This is a perspective that neuroscientists have been taking for several decades¹, and indeed, some gross aspects of our sensory input and motor output can be decoded from the coarse firing rate activation of individual neurons with respect to the classical tuning curves that describe such things as visual orientation² and motion selectivity³ or target location in a center-out reaching motor task¹. Although this does not implicitly speak to how the brain may actually be using this information, and the decoding perspectives in sensory and motor

pathways are fundamentally different owing to their reversed roles in terms of causality, this general perspective is nonetheless a very useful tool in describing neural activity in the context of things that we can all relate to in the outside world. Although the primary focus of this Perspective is in coding in sensory pathways, and more specifically across cortical and subcortical structures, the general perspective is nonetheless generally applicable across pathways and brain regions.

Although the various pathways of the brain had long been described as communication lines, it was the work from Bialek and colleagues in the early 1990s that brought this idea into focus when working with the practical details of neuronal spike trains in detail⁴. Working with the fly visual pathway as a model system, they constructed a framework in which they could quantify how much information was conveyed about motion in the visual scene from neurons in the visual pathway, at a very detailed, spiking-neuron level. At the core of this perspective was the formalism of expressing the probability of observing a neuronal response R given the presentation of the visual stimulus S , or $p(R|S)$, which is generally referred to as an encoding model. Heuristically, the question they asked was, “How much does the observation of the neuronal response R reduce my uncertainty about the stimulus S ?” In this context, they introduced the perspective of decoding the visual input and demonstrated this explicitly by reconstructing the motion input using a matched pair of neurons with sensitivity to motion in opposite directions. Although previous studies had shown that static visual images could be reconstructed from the recordings of individual neurons in the visual pathway⁵, and this concept is implicit in the prediction of motor output¹ or visual orientation² from tuned cortical neurons, Bialek and colleagues introduced a formal approach that opened the problem up to complex, temporally varying signals and, perhaps more importantly, to quantification. Subsequent studies followed this lead, showing that spiking activity across neuronal populations could be used to reconstruct the artificial visual input in the periphery *in vitro*⁶, natural visual scenes from thalamic populations *in vivo*⁷ and locomotor trajectories from hippocampal ensembles^{8,9}. Many studies have since extended these ideas, both experimentally and computationally, and recent studies have even shown that rudimentary shapes of objects in the visual field can be reconstructed by applying these strategies to human functional magnetic resonance imaging data¹⁰. Among these studies, and many others, the techniques and approaches vary, but the general framework involves evaluating the posterior probability of a stimulus given the neural response, $p(S|R)$, derived from the encoding model $p(R|S)$ and prior probabilities of the stimuli, $p(S)$, or with direct estimates of $p(S|R)$, and it can use a variety of approaches¹¹ to estimate S . One common thread, however, is that all involve assumptions as to the very nature of the neural code itself.

What is the neural code and who reads it? In this probabilistic framework, precisely what neuronal response R do we mean? Although one

Coulter Department of Biomedical Engineering, Georgia Institute of Technology & Emory University, Atlanta, Georgia, USA. Correspondence should be addressed to G.B.S. (garrett.stanley@bme.gatech.edu).

Received 26 July 2012; accepted 11 January 2013; published online 25 February 2013; doi:10.1038/nn.3330

might imagine an infinite number of possible codes, a long-standing debate distills the argument¹², casting the problem as a code based on coarse measures of mean firing rate (a rate code) versus temporal codes based on precise patterns of ones and zeros (a timing code)^{13,14}. However, most now acknowledge that the issue is simply one of time-scale and that pathways likely operate over a range of timescales for which rate and temporal coding are the two extremes, and are likely not mutually exclusive.

So, how do we decide what the neural code really is? One approach that has been taken is to quantitatively evaluate candidate codes in terms of how much they reveal about the external world¹⁵. An equally important perspective, however, lies in the constraints placed on the coding by the anatomy and biophysics of the circuitry. Despite the visceral effect that recreating or reconstructing elements of sensory inputs or motor outputs from neural activity has on us, the results are ultimately less than satisfying unless we use this to tell us something about the physiological plausibility of the assumed codes. One way to attack this problem is to force the observer to take the role of the downstream neurons to which the population projects. In other words, who reads this information, and how do they do it? One can imagine all sorts of candidate codes that can be read out to provide useful information about the outside world¹⁵, but the challenge for us is to provide a plausible biological implementation of such a code and how it propagates. For example, it has been widely observed that the latency in the response to a sensory input is informative about the sensory stimulus, and this is often described as a ‘latency code’^{16,17}. That is, the identity of the sensory input can be inferred from the latency between the presentation of the sensory stimulus and the neuronal spiking. Although this is attractive from a theoretical perspective, the observer of course would require an internal clock that would somehow allow the computation of such a latency, as in general we do not know the precise time a sensory input occurs. Similarly, beyond the correlations in response probability (and spike count) due to common stimuli, neurons in several brain regions have been shown to exhibit trial-to-trial correlations that have been commonly referred to as ‘noise correlations’¹⁸, the potential utility of which has been speculated on vigorously^{19–21}. Although in some contexts the noise correlations appear to be harmful to information transmission¹⁸, these correlations can in other contexts theoretically provide more information about the sensory input¹⁹, providing an attractive addendum to a simpler code based just on response selectivity reflected in the tuning curves of individual neurons. However, to determine whether these are truly viable codes, we must identify a clear biophysical mechanism to propagate these codes across microcircuits and across regions.

The wiring suggests synchrony as an element of the code. Neuronal spiking in the CNS is generally not the result of a single synaptic input but instead reflects the spatiotemporal integration over a range of inputs, the number of which varies depending on the pathway and region. Beyond the obvious spatial and temporal integration properties associated with synaptic inputs, synchrony of presynaptic inputs has been shown to be an important element in effectively and reliably pushing information across synapses²². In a wide range of experimental studies, synchronous thalamic inputs (within 10 ms) have been shown to be more effective at driving cortical targets than asynchronous inputs^{23,24}, and this synchrony constitutes a mechanism for propagating signals from the sensory periphery to cortex²⁵. This timing argument is further compounded by the fact that short interspike intervals of single thalamocortical neurons are more likely than long interspike intervals to evoke spiking in the downstream cortical neuron²⁶. The relevant timescale for this phenomenon is on the order

of 10–20 ms. Taken together, these two effects are consistent with the notion of a ‘window of opportunity’ for integration on the part of the cortical cell that is in part dictated by disynaptic feedforward inhibition²⁷ that appears to be a recurring circuit motif at many stages.

How does this relate to conveying information about ethologically relevant signals? In response to sensory inputs, neurons in early sensory pathways have been shown to be temporally precise in their firing, but the degree of precision depends strongly on the statistical properties of the stimulus. A decoding approach reveals that the precision in the neural response matches the precision requisite for reconstructing the sensory input with high fidelity, which scales with the characteristic timescale of the sensory input and is approximately 10 ms for naturalistic visual stimuli²⁸. Finally, synchrony in thalamic inputs has been shown to provide the substrate on which orientation and direction selectivity in cortex may be constructed²⁹ and to maximize reliability and energetic efficiency in cortical firing²², where here the use of the term “synchrony” refers to precise timing across neuronal ensembles induced by an exogenous (sensory, motor, cognitive) event, as opposed to biophysical coupling across neurons. Taken together, the anatomical and biophysical constraints of the wiring suggest a synchrony code as a viable means for conveying information through neural pathways in the natural environment, where coding important sensory and motor features in the precise millisecond timing across neuronal subpopulations projecting to common post-synaptic targets would naturally be transmitted. Simply put, regardless of the spiking activity of an individual neuron and whether the classical tuning properties predict it, if this activity is not precisely coordinated with nearby neighbors on a timescale of a few milliseconds, the message is lost. The idea of a synchrony code would thus argue that we should be redefining classical tuning properties of individual neurons into tuning properties based on timing coordination across subpopulations²⁵.

Returning to the question of viability of candidate codes, in the context of the latency code example, it could be that the network is faced with the challenge of estimating the timing of the sensory input from the population response, thereby defining the relative latency of an individual neuron³⁰, but heuristically it is more reasonable to regard this as a population code that operates on the relative timing of neurons in the population, for which we might more easily envisage a way in which downstream neurons might encode this feature given their natural sensitivity to synchronization of synaptic input^{17,31}. It is likely that a subpopulation of neurons responding first would be more tightly synchronized in time and thus a more powerful input to downstream target neurons, a mechanism that should perhaps be more accurately described as a synchrony code³². This is more than just an issue of semantics: it has implications for the circuitry required for reading out the candidate code. Regarding the viability of the noise correlations as a candidate code, the answer must certainly lie in the patterns of connectivity and the integration properties of synaptic mechanisms. As experimental investigation requires simultaneous recordings from multiple neurons and the experiment duration required for analysis of joint correlations across increasing numbers of neurons grows quickly, investigations of noise correlations have been limited to pairs or small numbers of neurons. To realistically propagate such a code, we would have to assume that the correlations exist over subpopulations large enough to collectively influence downstream neurons to which they project, but this has not yet been established. Notably, in the visual pathway, where this issue has been explored most vigorously, such noise correlations have been observed at the level of the retina^{20,21} and cortex^{33,34} but not the thalamus³⁵, thus leaving us to speculate whether the presence of the noise correlations has something to do

with the degree of local connectivity and divergence or convergence of synaptic inputs that differ across these stages.

A dynamic code. Perhaps one of the more intriguing aspects of synchrony as an element of the neural code lies in its potential to dynamically gate information flow. Despite relatively static anatomical connectivity over short timescales, the input timing and the corresponding sensitivity to input timing can change, setting the stage for a dynamic code. For example, it has long been postulated that short inter-spike intervals within or across thalamic neurons could serve as a ‘wake-up call’ to cortex to signal the presence of a salient visual input, which in turn would result in depolarization of the thalamus through cortical feedback to faithfully transmit details of the visual scene³⁶. Indeed, thalamic bursting has been shown to effectively convey the arrival of salient objects in the natural visual environment³⁷, and the overall degree of depolarization of the thalamic neurons shapes the selectivity of thalamic firing to various luminance features of the visual scene³⁸. More recently, it has been shown that local inhibitory circuitry at the thalamic level directly regulates tonic versus burst firing of thalamocortical projections and generates state-dependent cortical spindles hypothesized to relate to behavioral state³⁹. How this is precisely coordinated to regulate information flow is unknown, but we can speculate that dynamic regulation of timing at the network level, either through bottom-up or top-down mechanisms, would provide a powerful means by which to control and shape information flow to cortex. From a bottom-up perspective, it has recently been shown that adaptation strongly shapes thalamic synchrony⁴⁰ and dictates the window of integration of the recipient cortical target²⁷. Put in terms of our simple mathematical framework, $p(R|S)$ changes, potentially upending our traditional perspective of static roles for specific pathways. Although there is certainly evidence for role-specific circuits in the nervous system (for example, lemniscal versus paralemniscal, magnocellular versus parvocellular or dorsal versus ventral streams, to name a few), it is more generally likely that pathways serve multiple roles in different contexts, making the interpretation of the activity in the pathway more challenging. In the thalamocortical circuit, for example, the degree of synchronous firing across thalamic subpopulations is being continuously modulated, regulating not only how much information is conveyed to cortex but the very nature of what information about the outside world is being conveyed to cortex, with adaptation switching the pathway from detector to discriminator⁴⁰. The adaptation effects on both synchrony and the corresponding ‘window of opportunity’ for downstream recipient neurons may more generally reflect a shift in excitatory/inhibitory balance, which has been shown to shape cortical sensitivity to input synchrony to ensure the relay of important events in the face of asynchronously arriving spikes²².

Writing the neural code

As we have learned about how we might read the neural code, we are finally in a position to advance methodologies for controlling neuronal circuits in detail: to write the neural code. Beyond reading the neural code, an even better litmus test as to whether we understand the principles of neural coding is whether we can imprint the circuit with any code we desire and induce measurable effects on neural activity, perception or behavior. Can we artificially introduce patterns of activity in the brain that shape the way that information is propagated along the various pathways? Can we be made to see, hear or feel something that is not experienced naturally through our peripheral sensory organs? From the perspective of sensory pathways, can we induce a neuronal response R that our brain would interpret as arising from a particular sensory input S ? The very process of learning how

to control neurons and activate circuits in a way that is physiological and meaningful to the rest of the nervous system helps us learn more about the underlying system itself and issues of causality in a complex network. However, the implications for precise control of neural circuits are potentially far more profound. Beyond the immediate impact on neural pathologies (such as epilepsy, Parkinson’s disease and others), perhaps the most direct example is that of individuals with sensory deficits, for which surrogate inputs may provide a clinical ‘workaround’ for restoring function that has been lost to trauma or disease. The most salient example is that of cochlear implants, with more recent efforts at various stages of the visual pathway. The ability to finely control various circuits in the nervous system would theoretically open up the possibility of influencing or augmenting perception, behavior and perhaps even memories. At a practical level, however, our ability to precisely control neural circuits is ultimately limited by the spatial and temporal scales at which we can activate neural tissue and propagate signals by controlling the timing across neuronal ensembles.

The what and the when. When designing strategies for artificially activating neural circuits, we must ask the following questions: who is reading this information, what effect does the artificial activation have on the circuit receiving synaptic input from these neurons and to what extent can these downstream neurons distinguish between different surrogate signals? What is it that controls the ability to deliver a set of surrogate stimuli that result in downstream responses that are unique and distinguishable? This is directly a function of the spatial and temporal aspects of surrogate stimulation and the local anatomical projections to downstream targets. In other words, the ultimate effect of the activation of a subpopulation of neurons is expressed in the subsequent activation of neurons to which they project, and this is highly dependent on the wiring and the timing of the artificially induced presynaptic inputs. Thus, what matters is exactly which cells are activated, and when.

What neurons are activated? Although it is still poorly understood, there have been experimental and theoretical studies focused on the effects of microstimulation on the tissue local to the stimulation site. One huge problem facing us in the artificial activation of neural tissue is that injection of surrogate signals, whether electrical or optical in nature, tends to activate a large volume of neurons in the tissue. To the extent that these neurons are not just carbon copies of each other but represent different topographical regions or serve different functions, this creates a loss in specificity. For example, even injection of small ($\sim 10 \mu\text{A}$) currents into cortex can activate thousands of neurons proximal to the electrode tip, and optical stimulation at typical experimentally applied power ranges may activate tissue in a region on the order of a cubic millimeter⁴¹. Further, beyond the sensory periphery, the anatomy of the circuitry puts neuronal cell bodies and axonal fiber tracts in close proximity. As a result, microstimulation can activate neurons with cell bodies that are far from the site of stimulation⁴², beyond that predicted by conventional understanding of how electric fields grow with the magnitude of the current injection, having profound effects on the downstream neurons to which these neurons project⁴³.

There has been a recent resurgence of interest in artificially stimulating neural circuits, propelled by the development of optical methods of activating neurons⁴¹. In general, neural tissue is a heterogeneous combination of cell types, perhaps most fundamentally divided between neurons that release excitatory and those that release inhibitory neurotransmitter at their synaptic targets. Electrical

stimulation is generally nonselective for excitatory and inhibitory subpopulations, eliciting a complex, likely nonphysiological response of the local circuit. Given the often close proximity of excitatory and inhibitory neurons in the circuit, microstimulation is analogous to driving a car in which the gas and brake pedals have been welded together. Fortunately, optogenetically based techniques may provide the much-needed selectivity among cell types that will permit differential activation of subpopulations of neurons that are not segregated spatially. Through genetic expression or viral transfection, different cell types can be targeted to express specific channels that can be optically activated. This advance is not only appealing in trying to disentangle the functions of various subcircuits in the network but may also, when coupled with the extra ability to differentially activate (depolarize) and silence (hyperpolarize) neurons in the circuit, provide a needed degree of control of the circuit that is simply not afforded by electrical microstimulation.

When are neurons activated? Electrical microstimulation induces stimulus artifacts that preclude precise measurement of local neural activity, neurons can be driven at rates exceeding 100 Hz through microstimulation. Unlike electrical microstimulation, the first generation of optogenetic tools were unable to drive neurons above 50 Hz and often produced artifactual extra spikes, a hurdle recently overcome⁴⁴ through mutations of channelrhodopsin-2 (ChR2) to allow modulations of up to 200 Hz. How fast is fast enough for control? Neurons have been shown to be temporally precise on very fine time-scales²⁸. One could argue, therefore, that, owing to the natural time-scales of synaptic integration, control on the timescale of 5–10 ms may be necessary to truly control the circuit beyond local activation. This timescale would theoretically be attainable with the 100–200 Hz resolution of stimulation. However, stimulation of neural tissue, through either electrical microstimulation or optical methods, is further thought to synchronize neurons in its volume of influence. Given the sensitivity of downstream neuronal targets to the relative timing of synaptic inputs^{23,24}, and given that the regulation of neuronal synchrony has been proposed as a means by which information flow in sensory pathways is gated⁴⁰, this potentially leads to an all-or-none effect of surrogate stimulation that may be related to the seemingly invariant qualities of phosphene generation in the stimulation of the visual pathway. Synchrony is thus a double-edged sword. Too little of it and signals do not propagate across synapses, and too much of it results in loss of specificity.

Thus, given the inherent geometric and timing constraints on the surrogate activation, a principal challenge we face in writing the neural code lies in the ability to match the spatial and temporal scales of activation that ultimately are functionally relevant. Implicit is the need for a model, even if conceptual, of the coding at the particular stage of processing to which the stimulation is applied and how the synchronization of neuronal ensembles controls the propagation of information across brain structures. Hence, the reading and writing of the neural code are inextricably linked.

Challenges for the field

Embracing the dynamic code. Much of what we know about the neural code comes from experiments in which we work very hard to maintain relatively stationary conditions and assume that there is one and only one neural code. But if we know anything about the brain, it is that it is fluidly changing over a range of timescales because of a variety of factors.

Neural activity is highly dependent upon brain state that can be modulated through a combination of bottom-up (external) and

top-down (internal) influences, a salient example of which occurs in the thalamocortical circuit. Increased thalamic activity associated with arousal causes a depression of the thalamocortical synapse and a regulation of thalamic bursting and synchronization across the subpopulation⁴⁵, leading to an 'adapted' state in cortex characterized by low background firing, a higher signal-to-noise ratio and sharpened receptive fields. These kinds of internal influences, when combined with similar modulations from bottom-up external influences, are likely to profoundly affect the neural code through changes in state, regulating not only how much information is relayed to cortex but, more importantly, what kind of information is relayed to cortex through the regulation of timing⁴⁰, all possibly detailed elements of the larger umbrella of the efficient coding hypothesis originally posed a half century ago⁴⁶. Thus, incorporating the state dependence is critical both for understanding the neural code and for controlling neural circuitry in different contexts.

From the perspective of writing the neural code, longer term plasticity obviously has implications for how we might think about strategies for surrogate signaling. Given that the brain is known to be highly plastic, can we not just plug into the brain whatever signal we want and rely on reorganization at the cellular level and learning at the more global level to take care of things? Despite dramatic demonstrations of the functional consequences of plasticity⁴⁷, its potential is limited by the dependence on developmental stage. At this point we simply do not know what the limits of plasticity are and whether these findings generalize to problems beyond a simple substitution, in which the individual is asked to remap neural activation to represent the intended sensory input. In any case, any set of stimuli that evoke indistinguishable neural responses will theoretically be indistinguishable to the individual, which is likely to be particularly relevant in the widespread synchronization of local populations through surrogate stimulation. One can imagine that at least we need to be able to provide a set of distinct surrogate signals that would be discriminable to downstream neurons receiving projections from the area of stimulation or would drive the plasticity for functional remapping, in essence giving the pathway enough degrees of freedom to work with and to determine the real limits of plasticity.

Natural sensorimotor environment. Finally, much of what we know about sensory and motor function is through highly simplified, artificial scenarios. Although for practical reasons these simplifications are often well justified, both reading and writing the neural code need to be addressed in the context of more naturalistic scenarios for applicability in more ethological contexts. In the case of natural sensory stimuli, it is likely that the correlation structure that is the hallmark of the natural world induces precise patterns of synchronized activity across neuronal ensembles, which in turn are critical for propagating codes across brain structures^{28,35,48}. Further, the spatiotemporal correlations present in natural visual stimuli preferentially activate bursting of individual thalamic neurons³⁷, which is likely part of a stimulus-driven synchronization across the local neuronal ensemble and has profound implications for the activation of downstream brain structures. Similarly, sensory signals are likely profoundly affected by the relative influences of our own passive versus active movement, and vice versa, which has only begun to be studied in detail⁴⁹. New technologies that permit us to measure and stimulate neural activity during more unconstrained, naturalistic behaviors will thus undoubtedly be a key to directly addressing these kinds of questions⁵⁰. Understanding how neuronal populations collectively represent our natural environment and how this shapes feature selectivity in a way that takes into account the relationship

between the external world and our internal brain state is the next big step in understanding and controlling the neural code.

Conclusions

Because of the rapid tool development in both measuring and activating neuronal circuits, we are at a crossroads in both understanding and controlling the neural code. The anatomical and biophysical constraints across the various circuits in the brain naturally suggest that the synchronization of spiking across neurons is an important element in propagating signals across brain structures, pushing us to question whether a 'synchrony code' based on this premise forms a basis of the neural code. Although there is some compelling evidence in its favor, there is still much work to be done to determine how ubiquitous this phenomenon is in different contexts and across brain regions. And although synchrony may sound like a good thing, as it is a prerequisite for propagating across brain structures, the profound synchrony induced through artificial stimulation poses a real challenge for us in creating a wide range of neural representations. Neuronal synchrony is something that is modulated in normal function, serving as an exquisite gating mechanism, and is thus an important element in the control of neuronal circuitry.

ACKNOWLEDGMENTS

I would like to thank D.A. Butts, J.-M. Alonso, C. Schwarz and D.C. Millard for comments on the manuscript. G.B.S. was supported by US National Science Foundation Collaborative Research in Computational Neuroscience grants IIS-0904630 and IOS-1131948, and US National Institutes of Health National Institute of Neurological Disorders and Stroke grant 2R01NS048285.

COMPETING FINANCIAL INTERESTS

The author declares no competing financial interests.

Published online at <http://www.nature.com/doi/10.1038/nn.3330>.

Reprints and permissions information is available online at <http://www.nature.com/reprints/index.html>.

- Georgopoulos, A.P., Kettner, R.E. & Schwartz, A.B. Primate motor cortex and free arm movements to visual targets in three-dimensional space. II. Coding of the direction of movement by a neuronal population. *J. Neurosci.* **8**, 2928–2937 (1988).
- Paradiso, M.A. A theory for the use of visual orientation information which exploits the columnar structure of striate cortex. *Biol. Cybern.* **58**, 35–49 (1988).
- Britten, K.H., Shadlen, M.N., Newsome, W.T. & Movshon, J.A. The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J. Neurosci.* **12**, 4745–4765 (1992).
- Bialek, W., Rieke, F., de Ruyter van Steveninck, R.R. & Warland, D. Reading a neural code. *Science* **252**, 1854–1857 (1991).
- Creutzfeldt, O.D. & Nothdurft, H.C. Representation of complex visual stimuli in the brain. *Naturwissenschaften* **65**, 307–318 (1978).
- Warland, D.K., Reinagel, P. & Meister, M. Decoding visual information from a population of retinal ganglion cells. *J. Neurophysiol.* **78**, 2336–2350 (1997).
- Stanley, G.B., Li, F.F. & Dan, Y. Reconstruction of natural scenes from ensemble responses in the lateral geniculate nucleus. *J. Neurosci.* **19**, 8036–8042 (1999).
- Brown, E.N., Frank, L.M., Tang, D., Quirk, M.C. & Wilson, M.A. A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. *J. Neurosci.* **18**, 7411–7425 (1998).
- Zhang, K., Ginzburg, I., McNaughton, B.L. & Sejnowski, T.J. Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. *J. Neurophysiol.* **79**, 1017–1044 (1998).
- Miyawaki, Y. *et al.* Visual image reconstruction from human brain activity using a combination of multiscale local image decoders. *Neuron* **60**, 915–929 (2008).
- Quiñero, R. & Panzeri, S. Extracting information from neuronal populations: information theory and decoding approaches. *Nat. Rev. Neurosci.* **10**, 173–185 (2009).
- Shadlen, M.N. & Newsome, W.T. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J. Neurosci.* **18**, 3870–3896 (1998).
- Fellous, J.-M., Tiesinga, P.H.E., Thomas, P.J. & Sejnowski, T.J. Discovering spike patterns in neuronal responses. *J. Neurosci.* **24**, 2989–3001 (2004).
- Bazhenov, M., Rulkov, N.F., Fellous, J.-M. & Timofeev, I. Role of network dynamics in shaping spike timing reliability. *Phys. Rev. E* **72**, 041903 (2005).
- Jacobs, A.L. *et al.* Ruling out and ruling in neural codes. *Proc. Natl. Acad. Sci. USA* **106**, 5936–5941 (2009).
- Van Rullen, R. & Thorpe, S.J. Rate coding versus temporal order coding: what the retinal ganglion cells tell the visual cortex. *Neural Comput.* **13**, 1255–1283 (2001).
- Gollisch, T. & Meister, M. Rapid neural coding in the retina with relative spike latencies. *Science* **319**, 1108–1111 (2008).
- Zohary, E., Shadlen, M.N. & Newsome, W.T. Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* **370**, 140–143 (1994).
- Abbott, L.F. & Dayan, P. The effect of correlated variability on the accuracy of a population code. *Neural Comput.* **11**, 91–101 (1999).
- Nirenberg, S. & Latham, P.E. Decoding neuronal spike trains: how important are correlations? *Proc. Natl. Acad. Sci. USA* **100**, 7348–7353 (2003).
- Schneidman, E., Bialek, W. & Berry, M.J. Synergy, redundancy, and independence in population codes. *J. Neurosci.* **23**, 11539–11553 (2003).
- Wang, H.-P., Spencer, D., Fellous, J.-M. & Sejnowski, T.J. Synchrony of thalamocortical inputs maximizes cortical reliability. *Science* **328**, 106–109 (2010).
- Alonso, J.M., Usrey, W.M. & Reid, R.C. Precisely correlated firing in cells of the lateral geniculate nucleus. *Nature* **383**, 815–819 (1996).
- Bruno, R.M. & Sakmann, B. Cortex is driven by weak but synchronously active thalamocortical synapses. *Science* **312**, 1622–1627 (2006).
- Brette, R. Computing with neural synchrony. *PLoS Comput. Biol.* **8**, e1002561 (2012).
- Swadlow, H.A. & Gusev, A.G. Receptive-field construction in cortical inhibitory interneurons. *Nat. Neurosci.* **5**, 403–404 (2002).
- Gabernet, L., Jadhav, S.P., Feldman, D.E., Carandini, M. & Scanziani, M. Somatosensory integration controlled by dynamic thalamocortical feed-forward inhibition. *Neuron* **48**, 315–327 (2005).
- Butts, D.A. *et al.* Temporal precision in the neural code and the timescales of natural vision. *Nature* **449**, 92–95 (2007).
- Stanley, G.B. *et al.* Visual orientation and directional selectivity through thalamic synchrony. *J. Neurosci.* **32**, 9073–9088 (2012).
- Chase, S.M. & Young, E.D. First-spike latency information in single neurons increases when referenced to population onset. *Proc. Natl. Acad. Sci. USA* **104**, 5175–5180 (2007).
- Johansson, R.S. & Birznieks, I. First spikes in ensembles of human tactile afferents code complex spatial fingertip events. *Nat. Neurosci.* **7**, 170–177 (2004).
- Masquelier, T. Relative spike time coding and STDP-based orientation selectivity in the early visual system in natural continuous and saccadic vision: a computational model. *J. Comput. Neurosci.* **32**, 425–441 (2012).
- Smith, M.A. & Kohn, A. Spatial and temporal scales of neuronal correlation in primary visual cortex. *J. Neurosci.* **28**, 12591–12603 (2008).
- Ohiorhenuan, I.E. *et al.* Sparse coding and high-order correlations in fine-scale cortical networks. *Nature* **466**, 617–621 (2010).
- Desbordes, G. *et al.* Timing precision in population coding of natural scenes in the early visual system. *PLoS Biol.* **6**, e324 (2008).
- Sherman, S.M. A wake-up call from the thalamus. *Nat. Neurosci.* **4**, 344–346 (2001).
- Lesica, N.A. & Stanley, G.B. Encoding of natural scene movies by tonic and burst spikes in the lateral geniculate nucleus. *J. Neurosci.* **24**, 10731–10740 (2004).
- Lesica, N.A. *et al.* Dynamic encoding of natural luminance sequences by LGN bursts. *PLoS Biol.* **4**, e209 (2006).
- Halassa, M.M. *et al.* Selective optical drive of thalamic reticular nucleus generates thalamic bursts and cortical spindles. *Nat. Neurosci.* **14**, 1118–1120 (2011).
- Wang, Q., Webber, R.M. & Stanley, G.B. Thalamic synchrony and the adaptive gating of information flow to cortex. *Nat. Neurosci.* **13**, 1534–1541 (2010).
- Bernstein, J.G. *et al.* Prosthetic systems for therapeutic optical activation and silencing of genetically-targeted neurons. *Proc. Soc. Photo Opt. Instrum. Eng.* **6854**, 68540H (2008).
- Butovas, S. & Schwarz, C. Spatiotemporal effects of microstimulation in rat neocortex: a parametric study using multi-electrode recordings. *J. Neurophysiol.* **90**, 3024–3039 (2003).
- Wang, Q., Millard, D.C., Zheng, H.J.V. & Stanley, G.B. Voltage-sensitive dye imaging reveals improved topographic activation of cortex in response to manipulation of thalamic microstimulation parameters. *J. Neural Eng.* **9**, 026008 (2012).
- Gunaydin, L.A. *et al.* Ultrafast optogenetic control. *Nat. Neurosci.* **13**, 387–392 (2010).
- Stoelzel, C.R., Bereshpolova, Y. & Swadlow, H.A. Stability of thalamocortical synaptic transmission across awake brain states. *J. Neurosci.* **29**, 6851–6859 (2009).
- Barlow, H. Possible principles underlying the transformation of sensory messages. In *Sensory Communication* (ed. Rosenbluth, W.A.) 217–234 (MIT Press, Cambridge, Massachusetts, USA, 1961).
- Sharma, J., Angelucci, A. & Sur, M. Induction of visual orientation modules in auditory cortex. *Nature* **404**, 841–847 (2000).
- Lesica, N.A. *et al.* Adaptation to stimulus contrast and correlations during natural visual stimulation. *Neuron* **55**, 479–491 (2007).
- Seki, K. & Fetz, E.E. Gating of sensory input at spinal and cortical levels during preparation and execution of voluntary movement. *J. Neurosci.* **32**, 890–902 (2012).
- Szuts, T.A. *et al.* A wireless multi-channel neural amplifier for freely moving animals. *Nat. Neurosci.* **14**, 263–269 (2011).