

New perspectives on dimensionality and variability from large-scale cortical dynamics

Tatiana A Engel¹ and Nicholas A Steinmetz²



The neocortex is a multi-scale network, with intricate local circuitry interwoven into a global mesh of long-range connections. Neural activity propagates within this network on a wide range of temporal and spatial scales. At the micro scale, neurophysiological recordings reveal coordinated dynamics in local neural populations, which support behaviorally relevant computations. At the macro scale, neuroimaging modalities measure global activity fluctuations organized into spatiotemporal patterns across the entire brain. Here we review recent advances linking the local and global scales of cortical dynamics and their relationship to behavior. We argue that diverse experimental observations on the dimensionality and variability of neural activity can be reconciled by considering how activity propagates in space and time on multiple spatial scales.

Addresses

¹ Cold Spring Harbor Laboratory, Cold Spring Harbor, NY 11724, United States

² Department of Biological Structure, University of Washington, Seattle, WA 98195, United States

Corresponding authors: Engel, Tatiana A (engel@cshl.edu), Steinmetz, Nicholas A (nick.steinmetz@gmail.com)

Current Opinion in Neurobiology 2019, **58**:181–190

This review comes from a themed issue on **Computational neuroscience**

Edited by **Brent Doiron** and **Máté Lengyel**

For a complete overview see the [Issue](#) and the [Editorial](#)

Available online 1st October 2019

<https://doi.org/10.1016/j.conb.2019.09.003>

0959-4388/© 2019 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

Introduction

Behavior is driven by changes of neural activity throughout the brain. But even in the absence of behavior, the brain spontaneously generates massive tides of neural activity, which can be as large as the activity produced by external stimuli. Spontaneous activity (see [Box 1](#)) is present at all times and on all spatial scales — from single neurons to brain-wide networks — interacting with sensory inputs and affecting motor outputs. Due to varying patterns of ongoing activity, responses to identical stimuli are highly variable from trial to trial (see [Box 1](#)). Traditional computational theories dismiss ongoing activity as a nuisance, assuming that neural activity is composed of

task-related ‘signals’ mixed with a random ‘noise’ (see [Box 1](#)). Accordingly, the signals can be recovered by averaging activity data over trials, and the brain could likewise overcome the noise by appropriate averaging over neural populations [1,2]. However, evidence is accumulating that ongoing activity is not debilitating noise, but well-orchestrated dynamics with distinct, reproducible structure both within local microcircuits and across the whole brain. A theory of how these dynamics contribute, if at all, to perception, decisions, and actions, is only beginning to emerge.

The structure of spontaneous cortical activity is conventionally quantified by pairwise correlations [3] (see [Box 1](#), but other methods have been used as well [4,5], including time-dependent correlations [6]). At the local scale of a cortical microcircuit, correlations have been computed between fluctuations of spiking activity in pairs of simultaneously recorded neurons, termed noise correlations [7,8]. At the global scale of the entire cortex, correlations have been computed between spontaneous activity in pairs of distant cortical regions recorded with functional magnetic resonance imaging (fMRI), electroencephalography (EEG) or magnetoencephalography (MEG) in the absence of a stimulus or a task, a measure known as the resting-state functional connectivity [9,10]. Such zero-lag correlations only reveal precise synchrony, in which fluctuations of neurons or distant cortical regions are aligned in time.

With recent advances in recording technologies and data analysis methods, we are now beginning to increasingly appreciate that cortical activity is structured both in space and time, propagating through cortical networks as spatiotemporal waves [11,12,13*,14,15]. Propagating waves are commonly observed on multiple spatial and temporal scales and across different experimental setups. Spatiotemporal waves bind distant nodes within a network: activity of two nodes with a small zero-lag correlation can have a tight relationship at a non-zero lag, corresponding to wave propagation. Thus, understanding cortical activity will require a theory that accounts for the spatiotemporal nature of cortical dynamics. In this review, we survey new experimental techniques, data analysis methods, and models that enable us to understand the structure of spatiotemporal cortical activity on different scales, and moreover, to link this structure across scales. We argue that such links have allowed the field to make progress on classic problems understanding response dimensionality and trial-to-trial variability. We also

Box 1 Neural variability and dimensionality definitions.

Variability: seemingly stochastic differences in neural responses under precisely controlled experimental conditions, for example, differences in response to repeated presentations of the exact same sensory stimulus. Neural response variability can be in part attributed to activity driven by spontaneous, task unrelated behaviors (e.g. spontaneous movements), to uncontrolled changes of behavioral states (e.g. arousal, motivation), and to internal network dynamics.

Spontaneous activity: ongoing activity in the brain, which is not directly driven by sensory stimuli or overt behaviors. Spontaneous activity is most frequently studied in the absence of overt behavior (e.g. during rest, quiet wakefulness, or sleep). However, spontaneous activity, not driven by any apparent external causes, is also present during sensory stimulation and behavioral tasks and thus contributes to neural response variability.

Noise: a variable component of neural response that is considered to be task irrelevant, not carrying meaningful signals, but merely obscuring neural representations and deteriorating computations.

Pairwise (noise) correlation: Pearson correlation between activity fluctuations of two simultaneously recorded units (e.g. neurons, pixels, voxels) across repeated trials under the same experimental conditions. Noise correlation reflects not only random noise, but also other sources of variability, such as spontaneous movements, behavioral state changes, or cognitive factors such as learning and attention.

Principal components: a set of linearly uncorrelated variables obtained from the original, possibly correlated variables via an orthogonal transformation, such that the first principal component accounts for as much variability in the data as possible, and each succeeding component in turn accounts for as much variance as possible under the constraint that it is orthogonal to the preceding components. Principal components can be found by eigendecomposition of the data covariance matrix, where the eigenvalues indicate the proportion of variance accounted for by each principal component.

Linear dimensionality: the number of principal components required to account for a fixed proportion of variance in the data. For low-dimensional data, the variance is concentrated in the first few principal components with the largest eigenvalues. For high-dimensional data, the variance is distributed across many principal components with the eigenvalues of similar magnitude. Intuitively, the dimensionality corresponds to the extent of the linear subspace occupied by the data, or the number of separate patterns exhibited by the data.

speculate about computational functions of spatiotemporal activity.


Dynamics on the local scale

On the local scale of a cortical microcircuit, the activity of many neurons can be recorded simultaneously using multi-electrode arrays [16] or optical imaging techniques [17]. These methods sample activity from local neural populations within spatial regions that are typically smaller than a cortical area. The widespread availability of neural population recordings has recently facilitated our understanding of how spontaneous activity fluctuations, measured with pairwise correlations, are distributed over neural populations.

In many studies, the activity fluctuations shared among neurons within a local population were found to be low-dimensional [18–22], meaning that only few activity patterns across neurons dominate population responses over time (Figure 1a). The low dimension of population activity is manifested in a rapidly decaying eigenvalue spectrum of the noise-correlation matrix, with only the first few eigenvalues set apart from zero (Figure 1b–d). In agreement with these observations, one-dimensional bulk measures such as local field potentials [23] or summed population firing rates [20,24] predict well a majority of pairwise noise correlations in local populations. Moreover, changes of behavioral and cognitive states, such as arousal, task engagement, and attention, appear to primarily modulate the shared, low-dimensional component of the population-wide fluctuations, without much effect on the uncorrelated fluctuations of single neurons [22,25•]. The picture of low-dimensional fluctuations was recently challenged by an experiment in which ~10,000 neurons from approximately 1 mm² of mouse visual cortex were recorded simultaneously using two-photon calcium imaging [26•]. Within this large population, activity fluctuations were found to be high-dimensional, comprising at least 100 linear dimensions.

The puzzling differences in reported dimensionality of local activity fluctuations may be reconciled by considering how it is related to spatiotemporal dynamics in the local network. In a network model with spatially structured connectivity — mimicking local connectivity in the cortex — spontaneous fluctuations form spatial clusters that propagate laterally as local irregular waves (Figure 2a) [27]. The typical cluster size depends on the spatial connectivity structure and neural excitability. If a neural population is sampled from a local region that is smaller than a typical cluster size, then most sampled neurons will equally participate in each wave passing through this region, simultaneously increasing and decreasing their activity. As a result, fluctuations of this population will be low-dimensional (Figure 2b,c). If, however, a neural population is sampled from a region that is larger than a typical cluster size, then different sampled neurons will participate in different local waves passing through remote parts of the network at different times, hence fluctuations will be high-dimensional (Figure 2b,c). This analysis predicts that fluctuations in a population of tightly interconnected neurons (such as within a cortical column) should be very low-dimensional. Indeed, in spiking activity recorded from all layers of single columns in the primate visual cortex, fluctuations are virtually one-dimensional: a column spontaneously transitions between phases of high and low spiking activity nearly synchronously across layers (Figure 2d,e) [25•,28]. Differences in spatial distribution of measured neurons and in network connectivity structure may therefore at least partially explain diverse observations about dimensionality of cortical activity.

Figure 1

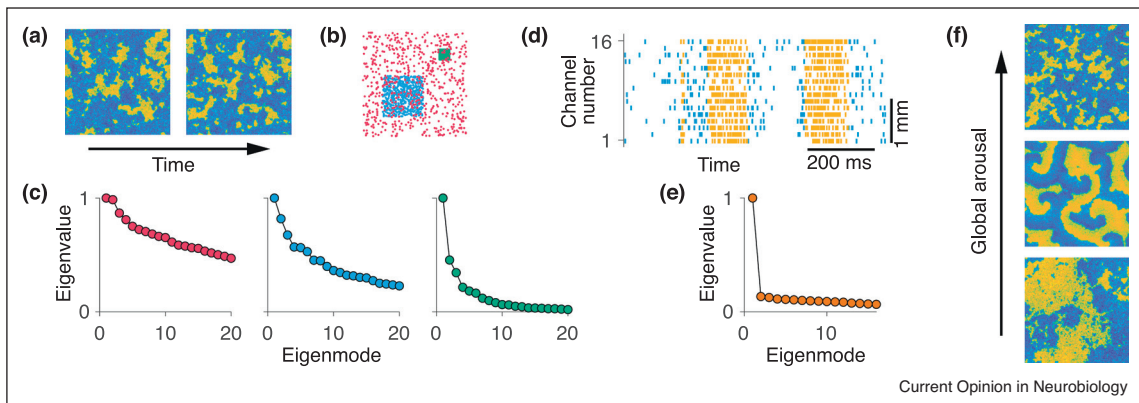


Dimensionality of population-wide fluctuations. **(a)** Responses of two simulated neural populations with low-dimensional (*upper row*) and high-dimensional (*lower row*) fluctuations on four example trials (different trials are offset vertically). Blue-to-green color code indicates projection of each trial on the first principal component (see *Box 1*) of the data correlation matrix. Neurons are sorted left-to-right by their activity on the first trial (blue). Across trials, low-dimensional population exhibits only a scaled version of the same activity pattern. High-dimensional population exhibits many diverse activity patterns. **(b)** Noise correlation (see *Box 1*) is a Pearson correlation coefficient between activities of a pair of neurons (*i* and *j*) in the population across trials, under identical stimulus conditions (each dot is the pair's activity on one trial). **(c)** Noise correlation matrix for all pairs in the population. Neurons are sorted according to their projection weight on the first principal component. **(d)** Eigenvalue spectrum of the noise-correlation matrix decays slowly for the high-dimensional population, but has only one eigenvalue set apart from zero for the low-dimensional population (see *Box 1*).

However, the spatial intermixing of topologically separate networks and the dependence of dynamics on behavioral state both complicate the picture. The spatial network structure supporting wave propagation is not always

aligned with anatomical dimensions of the cortex (depth and lateral distance). For example, in the mouse visual cortex, connectivity is stronger between cells with similar orientation preference, which are mixed in the cortical

Figure 2



The spatial scale of network dynamics determines the dimensionality of correlated fluctuations. **(a)** Spontaneous activity patterns in a two-dimensional network model of rate-units with spatially structured connectivity [27]. The network exhibits spatiotemporal activity with high (yellow) and low (blue) firing-rates clustered in space. The activity spreads through the network as local irregular waves. The typical cluster size depends on the spatial connectivity strength and local excitability. **(b)** Activities of 625 units in the network are sampled over time on spatial scales that are large (red), medium (blue) and small (green) relative to the typical cluster size. Dots indicate locations of the sampled units. **(c)** Eigenvalue spectra of the activity correlation matrix of units sampled on the large (red, *left*), medium (blue, *middle*) and small (green, *right*) scales. The activity is high-dimensional on scales larger than the typical cluster size, and low-dimensional on scales smaller than the typical cluster size. **(d)** Spontaneous transitions between periods of high (yellow) and low (blue) spiking activity recorded across layers of a single column in the primate visual cortex with a 16-channel linear multi-electrode array (data from Ref. [25*]). One example trial is shown. **(e)** Eigenvalue spectrum of the activity correlation matrix of units within a single column from panel (d). The single-column activity is dominated by one dimension. **(f)** Changes of behavioral states (e.g. global arousal), accompanied by changes in the level of neuromodulators that affect neural excitability (e.g. acetylcholine, ACh), can lead to qualitative changes of spatiotemporal activity patterns in the network (without any change in the network connectivity). The activity pattern in the network model shifts from global synchronous waves (*lower panel*), to spiral waves (*middle panel*), to local irregular waves (*upper panel*) when neural adaptation currents are enhanced, mimicking the effect of rising ACh level. Such shifts in spatiotemporal dynamics will produce changes in the dimensionality of correlated fluctuations in the same group of neurons.

tissue as ‘salt and pepper’ (imagine shuffling the pixels in Figure 2a). Therefore, a population sampled from a local cortical region may correspond to sampling from distant regions in the network rearranged according to the connectivity structure and thus exhibit high-dimensional fluctuations, such as in Ref. [26^{••}]. Intriguingly, recordings made with large-scale Neuropixels electrode arrays from ~3000 neurons distributed across brain regions showed similar dimensionality as the V1 populations, suggesting that coordinated spatiotemporal dynamics may extend far beyond local populations and even beyond cortex [26^{••}]. In addition, the spatial scale of spontaneous waves can be altered by changes in neural excitability due to variations in the level of neuromodulators accompanying changes of behavioral states (Figure 2f) [29]. Consequently, the dimensionality of correlated fluctuations in the same neural population may be modulated as a function of behavioral state. These considerations suggest that although correlations provide us a glimpse into organization of cortical activity, a full understanding will require elucidating how spatiotemporal activity propagation is related to network connectivity and how it is modulated during changes of behavioral states.

Dynamics on the global scale

On the whole-brain scale, the activity from many regions can be recorded simultaneously with neuroimaging methods such as fMRI, EEG or MEG. fMRI provides complete volumetric images (at ~1 mm resolution), but measures slow (~1 s) hemodynamic responses associated with neural activity. EEG and MEG detect synchronized neural activity with high temporal resolution (~1 ms) but low spatial resolution (~1 cm). In these recordings, the structure of the whole-brain spontaneous activity is traditionally examined with pairwise correlations between brain regions, called functional connectivity. Functional connectivity analyses reveal that resting-state fluctuations are bilaterally symmetric and organized in multiple distributed networks of regions that activate and deactivate together, called resting-state networks [9,10,30,31]. Resting-state networks closely correspond with the large-scale anatomical connectivity of the human cortex [32]. Moreover, electrophysiological recordings in humans confirm synchronous patterns of neural activity in distant, often bilateral brain regions [33].

Recently, wide-field optical imaging methods enabled direct measurements of neural activity across broad regions of cortex at significantly higher spatial (<100 μ m) and temporal (~10 ms) resolution, and substantially enhanced signal-to-noise levels. Using large preparations with voltage sensitive dyes (VSDs) [34] or with transgenic expression of fluorescent voltage [35[•],36] or calcium sensors [37[•],38–40,35[•],41[•],42^{••}], a majority of both hemispheres of the mouse cortex can be imaged at once. The ability to observe the global cortical activity

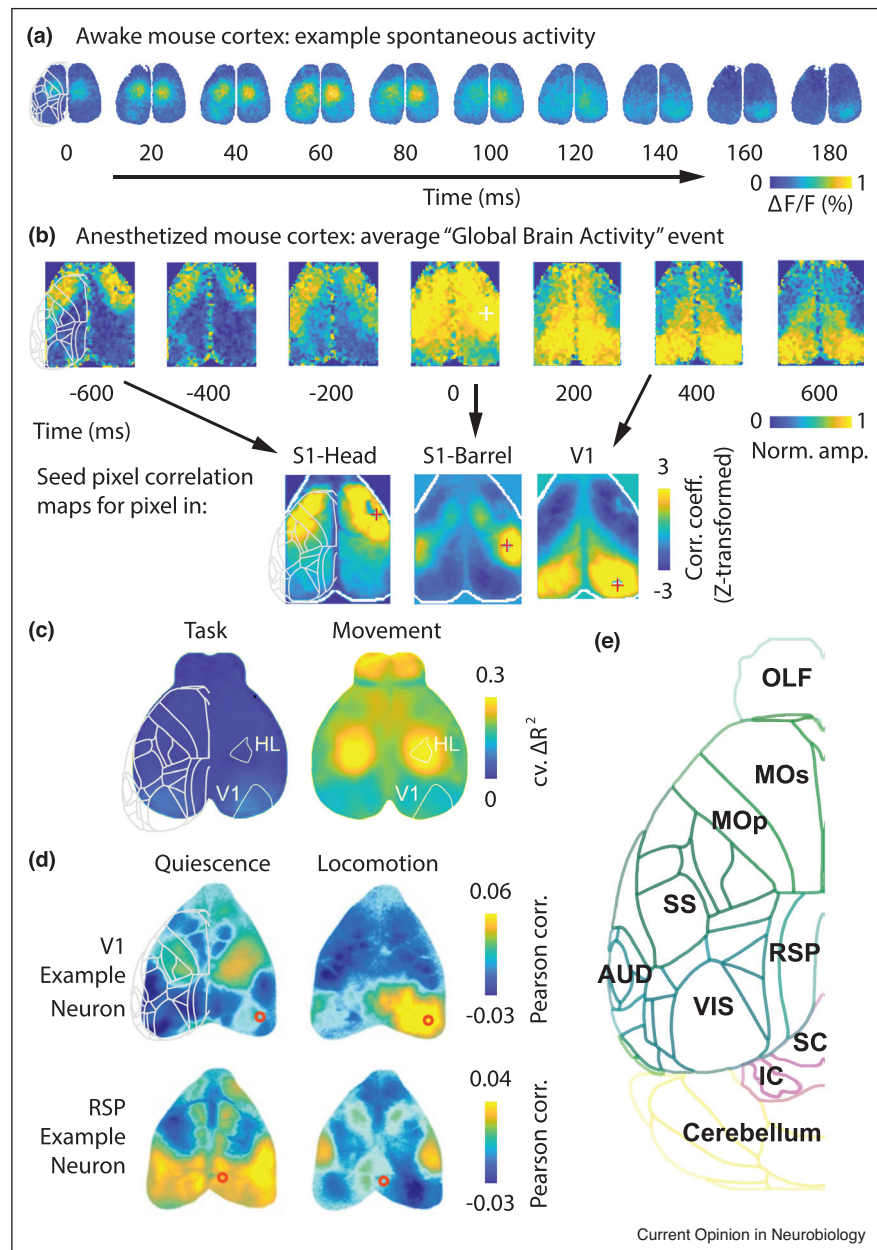
concurrently in space and time unveiled its striking spatiotemporal organization: spontaneous activity consists of bilaterally symmetric patterns that involve each part of cortex and occur both during behavior [39,41[•],42^{••}] and in the absence of overt behaviors (Figure 3a) [34]. Similar to the much slower resting state networks observed with fMRI, the spatial organization of these functional activity patterns are linked to known anatomical pathways [43[•]].

To what extent do distributed activity patterns in humans and rodents have spatiotemporal dynamics? A recent series of studies quantified the lead-lag relationships in the activity of human brain regions in resting-state EEG [45,46] or fMRI [47[•],48,49]. These analyses revealed that global activity travels slowly through the cerebral cortex along multiple, stereotypical spatiotemporal trajectories, and the directions of propagation differ across behavioral states, such as sleep and wakefulness [48,49]. The global activity propagation is largely unidirectional within conventional resting-state networks, which may therefore correspond to points along global spatiotemporal trajectories. In rodents, recordings during anesthesia reveal similar large, globally propagating waves, at multiple timescales [37[•],50,51]. As is the case in human slow-wave sleep [52], under some conditions these waves are observed to primarily travel anterior to posterior (Figure 3b, *upper panels*). Moreover, different sets of cortical areas, sharing high functional connectivity within each set, are coactivated at different moments during the propagating global waves, indicating that spatial information contained in functional-connectivity networks is embedded in the phase of the global waves (Figure 3b, *lower panels*) [37[•]]. The extent to which global spontaneous dynamics in the cortex of awake, un-anesthetized rodents can be described as waves [50,51] versus independent activation of patterns [43[•],38] remains unclear. Methods to describe and quantify spatiotemporally extended patterns such as waves when they propagate not just in one consistent direction, but through the same medium in multiple directions [53], will be valuable for better understanding the large-scale patterns observable on the global scale.

Interactions between local and global dynamics

Local and global neural dynamics co-exist within the same cortical networks and may interact across scales. Recent studies have begun to document such interactions. Understanding the relationship between activity across local and global scales may clarify the neural mechanisms of behavior, both goal-directed and spontaneous. For example, local one-dimensional dynamics within single cortical columns (Figure 2d) are modulated by allocation of spatial attention, the impact of which is confined to local regions within cortical maps [25^{••}]. However, they are also modulated by changes of arousal

Figure 3



Characteristics of global cortical activity in mice. **(a)** Spontaneous activity in the mouse cortex during quiet wakefulness recorded with wide-field VSD imaging (adapted from Ref. [34]). Observed activity is predominantly bilaterally synchronous. **(b)** Activity in anesthetized mouse cortex recorded with wide-field calcium imaging (adapted from Ref. [37]), revealing globally propagating slow waves, similar to those in humans. The average of many detected 'Global Brain Activity' events (*upper row*). Individual moments in this stereotyped wave resemble (arrows) the overall seed pixel correlation maps for different seed pixel locations (*lower row*). **(c)** Activity related to movement of the mouse's face or body dominates, over all cortical areas measured (adapted from Ref. [41]). Colormap represents the cross-validated unique variance explained by each group of predictors (those related to the task, or those related to other movements). **(d)** Measurements of individual neurons with simultaneous wide-field calcium imaging allows linking local and global dynamics. Each map represents a seed neuron correlation map: the correlation between spiking activity of an individual neuron recorded electrophysiologically (at the location of the red circle) and calcium fluorescence across cortex. These data reveal that individual neurons have distinct patterns, and that these correlations can vary strongly with behavioral state, from quiescence to locomotion (adapted from Ref. [44]). **(e)** For orientation, a map of the left hemisphere of the mouse brain, viewed from above. The outlines of these cortical areas are superimposed in gray over the left-most image of each preceding panel (superpositions done manually and approximately, for general guidance). *Abbreviations:* OLF, olfactory bulb; MOs, secondary motor cortex; MOp, primary motor cortex; SS, somatosensory areas; RSP, retrosplenial cortex; AUD, auditory areas; VIS, visual areas; SC, superior colliculus; IC, inferior colliculus. Atlas data from the Allen Institute for Brain Science.

[25**], which affect neural activity on the brain-wide scale and may relate to aspects of task performance such as engagement [54]. Modulation of global dynamics with arousal is also consistent with arousal-related local modulation of correlated fluctuations [55–58].

Widespread dynamics across cortex do not only reflect arousal changes, but are also related to behavior. Overt movements, with or without any goal-directed task, were recently shown to dominate rodent neural activity on the local [26**], cortex-wide (Figure 3c) [41*], and brain-wide scales [26**,59]. These studies recorded neural activity concurrently with multi-dimensional behavioral information in mice, including pupil-dilation [55,56], running [60], and whisking, as well as videos of facial and body movements. Surprisingly, they found that signals related to movements dominate neural activity across the entire cortex, dwarfing task-related variables and even sensory stimuli [41*,59]. Moreover, in visual cortex, sensory inputs do not interrupt this ongoing signal, but add onto it a representation of visual stimuli in orthogonal dimensions [26**].

Comparably, another study, using wide-field imaging, found that bilaterally coordinated activity dominates unilateral responses to local visual stimuli and adds to them, explaining much of the trial-to-trial response variability [35*]. This work builds on prior studies that examined the relationship between local ongoing activity and stimulus driven activity [61,28] by revealing that much of the local ongoing activity is in fact related to global cortical dynamics. The bilateral component of response variability must be accounted for when considering the relationship between neural activity and behavior. A model fit to these data predicted behavior accurately by taking a difference between activity in left and right visual cortex, thus removing the bilateral component, and the model therefore predicted that inactivation of visual cortex unilaterally should lead to an increase in ‘hallucinations’ — reports of ipsilateral stimuli not actually present — rather than ‘blindness’ — failure to report present stimuli [42**]. Strikingly, this prediction was borne out both qualitatively and quantitatively with no parameter adjustments. This result would not have been predictable from an examination of local visual cortical activity in one hemisphere alone. Together, these results argue that much local response variability can be accounted for by considering larger scales of neural activity, both large local populations and cross-hemispheric activity.

Experimental techniques are now emerging that enable simultaneous measurements of neural activity on the local and global scales, and thus can more directly probe their interactions. Specifically, cellular-resolution two-photon calcium imaging or electrophysiological recordings of a local microcircuit have been combined with simultaneous wide-field calcium imaging of the entire bilateral cortex

[44**,62**,63,64]. Such multi-scale recordings permit uncovering how different neurons within a local circuit participate in the brain-wide dynamics. In particular, functional connectivity maps between a single neuron and the entire cortex reveal that neighboring neurons in one cortical area are frequently coupled with distinct distal cortical regions (Figure 3d) [62**,44**]. One possibility is that heterogeneity of neurons’ functional connectivity may be aligned with diversity of their long-distance axonal projection patterns, which form multiple information processing streams [65–68]. In addition, some of the diversity of global correlations is accounted for by the genetic identity of the neurons (VIP-expressing interneurons versus non-VIP-expressing putative pyramidal neurons) [62**], consistent with differences of inputs and outputs of these types. However, a neuron’s functional connectivity depends at least in part on behavioral state [44**,69] (cf. Figure 2f), and may differ dramatically depending on the timescale under consideration [70], suggesting that anatomical pathways alone will not fully explain these features.

Another promising experimental direction is the use of novel large-scale electrophysiological techniques [71] to record populations of neurons distributed across the brain [26**,59]. This approach enables not only linking of local and global scales, but also incorporation of information about subcortical activity and dynamics, which are deeply intertwined with those of cortex [72,73]. With these new datasets available, computational analyses and models based on multi-scale measurements of neural dynamics with cell-type resolution can now begin to address the sources and mechanisms of local-global interactions.

Computational functions of local and global dynamics

What computational functions can be served by cortical dynamics on multiple spatiotemporal scales? Canonical models of information processing in the brain are oblivious to dynamics, portraying cortical neurons as static feature detectors that build increasingly complex representations through successive stages of cortical hierarchy [74–76]. However, state-of-the-art static feedforward architectures fall behind primates in recognizing ‘challenging’ images that require additional recurrent processing [77]. Since a shallow recurrent neural network is equivalent to a very deep feedforward network (e.g. ResNet) [78], it is possible that the brain’s recurrent circuitry efficiently implements what artificial systems achieve by stacking more feedforward layers, whereby the local dynamics of each layer act as additional nonlinear transformations. In addition, decoders of neural responses do not generalize across time [77,79], suggesting that the brain’s representation of images may be not static, but rather encoded in dynamic and transient neural trajectories. Moreover, static feedforward models are vulnerable to input perturbations, such as noise [80] or

partial occlusions [81^{*}], to which humans are robust. The robustness of feedforward architectures can be improved by augmenting the output layer with recurrent dynamics [81^{*}], which reinforces the idea that recurrent dynamics may serve to ‘fill in’ the missing information or apply priors for resolving ambiguity. One may wonder, why would the cortex need to resolve ambiguity at the level of lower sensory areas, rather than just at the highest level? Perhaps this relates to another key distinction between the cortex and modern artificial networks: whereas feedforward networks are ‘read out’ only at the last stage, nearly every cortical area — from primary sensory cortices onwards — projects to thalamus and striatum [82], and a majority project to the midbrain [83,84]. Thus nearly every area potentially influences behavior directly, hence local and global dynamics may be important for correcting representations at each level.

Models of neural dynamics have been developed to account for specific computational mechanisms characteristic of frontal and association cortices, such as pattern completion [85], working memory and decision-making [86,87], selective attention [88], and executive control [89], as well as of motor cortex [90]. In these models, task-relevant computations are usually implemented by a single local circuit, for example, a recurrent neural network with connection weights optimized to transform a set of time-dependent inputs into desired outputs [91]. Such models, operating on a single spatial scale, lack the modular, multi-scale organization of the neocortex. Augmenting functional models of neural dynamics with a hierarchical structure and global dynamic interactions across hierarchical modules could potentially extend their cognitive capacity [92]. For example, a hierarchy of dynamical modules could represent a hierarchy of competing behavioral goals, and global dynamic interactions among the modules could resolve lower-level goals in service of high-level ones [93]. Parts of a hierarchical network could also engage in mental simulations of future actions and outcomes [94] or other metacognitive processes underlying flexible intelligent behavior. Another fascinating idea is that ongoing, spontaneous dynamics during slow-wave sleep may play a protective role against catastrophic forgetting (overwriting old memories with new ones) and enable brain networks to undergo continual learning [95], or to enhance learning from limited experience [96], both of which remain a challenge for artificial systems.

Taken together, with few exceptions, the existing hierarchical models of cortical processing lack dynamics, while the existing dynamical models lack the modular and hierarchical organization across spatial scales. Combining these computational elements in a unified architecture could lead to more powerful models that closer match the brain’s cognitive capacities. Mechanistic and descriptive models of multi-scale neural dynamics should

be developed to provide links with neurophysiological data and to reveal computational functions of the local and global dynamics in the brain. At the same time, normative models should be developed to predict what types of local and global dynamics could support efficient encoding of sensory stimuli, behavioral states and cognitive operations.

Conclusions

New experimental technologies continually push the boundaries for observing cortical activity across multiple spatial and temporal scales. Large-scale cortical recordings reveal massive spontaneous neural activity within local microcircuits and on the global scale. Novel analysis methods and models indicate that on each scale, neural activity is spatiotemporal, propagating along trajectories defined by the local and long-range connectivity structure and modulated by behavioral states. Interactions between dynamics on different scales reflect immediate behavioral goals. Theoretical frameworks now must account for the multi-scale, spatiotemporal nature of cortical activity to reveal its computational and behavioral consequences. Such multi-scale, spatiotemporal theories are beginning to clarify long-standing debates about dimensionality and variability of cortical responses, and they stand to unveil more about the enigmatic functions of brain’s spontaneous activity in the future.

Conflict of interest statement

Nothing declared.

Acknowledgements

Research supported by the NIH grant R01 EB026949 (TAE), the Swartz Foundation (TAE), and the Cold Spring Harbor Laboratory (TAE). The authors would like to thank E. Shea-Brown, D. Shimaoka, C. Stringer, M. Pachitariu, S. Musall, and M. Lengyel for helpful feedback.

References and recommended reading

Papers of particular interest, published within the period of review, have been highlighted as:

- of special interest
 - of outstanding interest
1. Averbeck BB, Latham PE, Pouget A: **Neural correlations, population coding and computation.** *Nat Rev Neurosci* 2006, **7**:358-366.
 2. Moreno-bote R, Beck J, Kanitscheider I, Pitkow X, Latham P, Pouget A: **Information-limiting correlations.** *Nat Neurosci* 2014, **17**:1410-1417.
 3. Cohen MR, Kohn A: **Measuring and interpreting neuronal correlations.** *Nat Neurosci* 2011, **14**:811-819.
 4. Kenet T, Bibitchkov D, Tsodyks M, Grinvald A, Arieli A: **Spontaneously emerging cortical representations of visual attributes.** *Nature* 2003, **425**:954-956.
 5. Berkes P, Orbán G, Lengyel M, Fiser J: **Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment.** *Science* 2011, **331**:83-87.
 6. Smith MA, Kohn A: **Spatial and temporal scales of neuronal correlation in primary visual cortex.** *J Neurosci* 2008, **28**:12591-12603.

7. Cohen MR, Maunsell JHR: **Attention improves performance primarily by reducing interneuronal correlations.** *Nat Neurosci* 2009, **12**:1594-1600.
 8. Mitchell JF, Sundberg KA, Reynolds JH: **Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4.** *Neuron* 2009, **63**:879-888.
 9. Fox MD, Raichle ME: **Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging.** *Nat Rev Neurosci* 2007, **8**:700-711.
 10. Buckner RL, Vincent JL: **Unrest at rest: default activity and spontaneous network correlations.** *Neuroimage* 2007, **37**:1091-1096.
 11. Muller L, Chavane F, Reynolds J, Sejnowski TJ: **Cortical travelling waves: mechanisms and computational principles.** *Nat Rev Neurosci* 2018, **19**:255-268.
 12. Muller L, Piantoni G, Koller D, Cash SS, Halgren E: **Rotating waves during human sleep spindles organize global patterns of activity that repeat precisely through the night.** *Elife* 2016, **5**: e17267.
 13. Muller L, Reynaud A, Chavane F, Destexhe A: **The stimulus-evoked population response in visual cortex of awake monkey is a propagating wave.** *Nat Commun* 2014, **5**:3675.
- A method is developed for detecting propagating waves in noisy recordings on single trials. VSD recordings in visual cortex of the awake monkey show that stimulus-evoked population response propagates as a traveling wave, with consistent dynamics across trials. Evoked waves propagating in several visual areas have precise phase relations.
14. Sato TK, Nauhaus I, Carandini M: **Traveling waves in visual cortex.** *Neuron* 2012, **75**:218-229.
 15. Huang X, Xu W, Liang J, Takagaki K, Gao X, Wu Jy: **Spiral wave dynamics in neocortex.** *Neuron* 2010, **68**:978-990.
 16. Buzsáki G: **Large-scale recording of neuronal ensembles.** *Nat Neurosci* 2004, **7**:446-451.
 17. Grewe BF, Helmchen F: **Optical probing of neuronal ensemble activity.** *Curr Opin Neurobiol* 2009, **19**:520-529.
 18. Rabinowitz NC, Goris RLT, Cohen MR: **Attention stabilizes the shared gain of V4 populations.** *Elife* 2015.
 19. Lin IC, Okun M, Carandini M, Harris KD: **The nature of shared cortical variability.** *Neuron* 2015, **87**:644-656.
 20. Scholvinck ML, Saleem AB, Benucci A, Harris KD, Carandini M: **Cortical state determines global variability and correlations in visual cortex.** *J Neurosci* 2015, **35**:170-178.
 21. Williamson RC, Cowley BR, Litwin-Kumar A, Doiron B, Kohn A, Smith MA, Yu BM: **Scaling properties of dimensionality reduction for neural populations and network models.** *PLoS Comp Biol* 2016, **12**.
 22. Huang C, Ruff DA, Pyle R, Rosenbaum R, Cohen MR, Doiron B: **Circuit models of low-dimensional shared variability in cortical networks.** *Neuron* 2019, **101** 337-348.e4.
 23. Kelly RC, Smith MA, Kass RE, Lee TS: **Local field potentials indicate network state and account for neuronal response variability.** *J Comput Neurosci* 2010, **29**:567-579.
 24. Okun M, Steinmetz NA, Cossell L, Iacaruso MF, Ko H, Barthó P, Moore T, Hofer SB, Mrsic-Flogel TD, Carandini M, Harris KD: **Diverse coupling of neurons to populations in sensory cortex.** *Nature* 2015, **521**:511-515.
 25. Engel TA, Steinmetz NA, Gieselmann MA, Thiele A, Moore T, Boahen K: **Selective modulation of cortical state during spatial attention.** *Science* 2016, **354**:1140-1144.
- In the awake primate visual cortex, the ensemble neural activity within single columns fluctuates between phases of vigorous and faint spiking synchronously across layers. These ongoing dynamics, reminiscent of slow-wave sleep oscillations, are modulated globally by arousal and locally within the retinotopic map during selective attention.
26. Stringer C, Pachitariu M, Steinmetz NA, Reddy CB, Carandini M, Harris KD: **Spontaneous behaviors drive multidimensional, brainwide activity.** *Science* 2018, **364** pg. eaav7893.
- Simultaneous recordings of 10 000 neurons in the mouse visual cortex exhibited high-dimensional fluctuations, which were partially related to the mouse's ongoing behavior. Sensory inputs did not interrupt this ongoing signal, but added onto it a representation of visual stimuli in orthogonal dimensions. Recordings of 3000 neurons distributed brain-wide revealed similarly high dimensionality and dependence on ongoing behavior.
27. Shi Y, Steinmetz NA, Moore T, Boahen K, Engel TA: **Linking noise correlations to spatiotemporal population dynamics and network structure.** *28th Annual Computational Neuroscience Meeting (CNS*2019)* 2019.
 28. Luczak A, Barthó P, Harris KD: **Spontaneous events outline the realm of possible sensory responses in neocortical populations.** *Neuron* 2009, **62**:413-425.
 29. Roach JP, Ben-Jacob E, Sander LM, Zochowski MR: **Formation and dynamics of waves in a cortical model of cholinergic modulation.** *PLoS Comp Biol* 2015, **11**:e1004449.
 30. Fox MD, Snyder AZ, Vincent JL, Corbetta M, Van Essen DC, Raichle ME: **The human brain is intrinsically organized into dynamic, anticorrelated functional networks.** *Proc Natl Acad Sci U S A* 2005, **102**:9673-9678.
 31. Vincent JL, Patel GH, Fox MD, Snyder AZ, Baker JT, Van Essen DC, Zempel JM, Snyder LH, Corbetta M, Raichle ME: **Intrinsic functional architecture in the anaesthetized monkey brain.** *Nature* 2007, **447**:83-86.
 32. Honey CJ, Sporns O, Cammoun L, Gigandet X, Thiran JP, Meuli R, Hagmann P: **Predicting human resting-state functional connectivity from structural connectivity.** *Proc Natl Acad Sci U S A* 2009, **106**:2035-2040.
 33. Nir Y, Mukamel R, Dinstein I, Privman E, Harel M, Fisch L, Gelbard-Sagiv H, Kipervasser S, Andelman F, Neufeld MY, Kramer U, Arieli A, Fried I, Malach R: **Interhemispheric correlations of slow spontaneous neuronal fluctuations revealed in human sensory cortex.** *Nat Neurosci* 2008, **11**:1100-1108.
 34. Mohajerani MH, McVea DA, Fingas M, Murphy TH: **Mirrored bilateral slow-wave cortical activity within local circuits revealed by fast bihemispheric voltage-sensitive dye imaging in anesthetized and awake mice.** *J Neurosci* 2010, **30**:3745-3751.
 35. Shimaoka D, Steinmetz NA, Harris KD, Carandini M: **The impact of bilateral ongoing activity on evoked responses in mouse cortex.** *Elife* 2019, **8**:e43533.
- Wide-field imaging, both with voltage-sensitive and calcium-sensitive fluorescent proteins, during decision-making reveals that bilateral fluctuations in mouse cortical activity account for a large proportion of the trial-to-trial variability in visual responses. These components of bilateral activity add to the visual responses linearly, and notably do not affect behavioral reports.
36. Madisen L, Garner AR, Shimaoka D, Chuong AS, Klapoetke NC, Li L, van der Bourg A, Niino Y, Egoif L, Monetti C, Gu H, Mills M, Cheng A, Tasic B, Nguyen TN, Sunkin SM, Benucci A, Nagy A, Miyawaki A, Helmchen F, Empson RM, Knöpfel T, Boyden ES, Reid RC, Carandini M, Zeng H: **Transgenic mice for intersectional targeting of neural sensors and effectors with high specificity and performance.** *Neuron* 2015, **85**:942-958.
 37. Matsui T, Murakami T, Ohki K: **Transient neuronal coactivations embedded in globally propagating waves underlie resting-state functional connectivity.** *Proc Natl Acad Sci U S A* 2016, **113**:6556-6561.
- Simultaneous wide-field imaging of calcium and hemoglobin signals across most of the mouse cortex reveals bilaterally symmetric spontaneous waves propagating across the entire neocortex, under anesthesia. Cortical areas sharing high functional connectivity are co-activated at different times during the propagated global waves.
38. Ma Y, Shaik MA, Kozberg MG, Kim SH, Portes JP, Timerman D, Hillman EMC: **Resting-state hemodynamics are spatiotemporally coupled to synchronized and symmetric neural activity in excitatory neurons.** *Proc Natl Acad Sci U S A* 2016, **113**:E8463-E8471.
 39. Allen WE, Kauvar IV, Chen MZ, Richman EB, Yang SJ, Chan K, Gradinaru V, Deverman BE, Luo L, Deisseroth K: **Global**

- representations of goal-directed behavior in distinct cell types of mouse neocortex.** *Neuron* 2017, **94**:891–907.e6.
40. Makino H, Ren C, Liu H, Kim AN, Kondapaneni N, Liu X, Kuzum D, Komiyama T: **Transformation of cortex-wide emergent properties during motor learning.** *Neuron* 2017, **94**:880–890.e8.
 41. Musall S, Kaufman MT, Juavinett AL, Gluf S, Churchland AK:
 - **Single-trial neural dynamics are dominated by richly varied movements.** *bioRxiv* 2019:308288.
 Wide-field calcium imaging of the mouse cortex during decision-making shows that activity related to ongoing task-unrelated behavior dominates across the entire cortex, dwarfing representations of task-related variables and sensory stimuli. Multi-dimensional representation of spontaneous movements and behavioral states predicts large portion of neural responses across the cortex.
 42. Zatzka-Haas P, Steinmetz NA, Carandini M, Harris KD: **Distinct contributions of mouse cortical areas to visual discrimination.** *bioRxiv* 2019:501627.

This study employed a combination of wide-field calcium imaging and systematic cortex-wide optogenetic manipulations to build and validate a model of visually guided behavior accounting for global cortical dynamics. The model incorporated measured activity from left and right visual and frontal cortex, predicted behavior on individual trials, and predicted the effects of inactivating these regions with no further parameter adjustments. Accounting for global dynamics therefore enabled a clear determination of the relative roles of each brain region in visually guided behavior.
 43. Mohajerani MH, Chan AW, Mohsenvand M, LeDue J, Liu R, McVea DA, Boyd JD, Wang YT, Reimers M, Murphy TH: **Spontaneous cortical activity alternates between motifs defined by regional axonal projections.** *Nat Neurosci* 2013, **16**:1426–1435.

Wide-field VSD imaging of the mouse cortex reveals multiple patterns of hemisphere-wide motifs in spontaneous activity. Similar activity motifs are evoked with sensory stimulation or direct cortical activation by optogenetics. Maps of intracortical monosynaptic structural connections predict hemisphere-wide patterns of spontaneous and sensory-evoked depolarization.
 44. Clancy KB, Orsolic I, Mrcsic-Flogel TD: **Locomotion-dependent remapping of distributed cortical networks.** *Nat Neurosci* 2019, **22**(5):778–786.

Pairing electrophysiology with wide-field calcium imaging enables simultaneous cellular-resolution recording of a local microcircuit along with whole-cortex activity measurement in awake, behaving mice. Individual neurons can have idiosyncratic cortical correlation maps, and these maps frequently differ to some extent between quiescent and locomoting behavioral conditions.
 45. Ito J, Nikolaev AR, Leeuwen Cv: **Spatial and temporal structure of phase synchronization of spontaneous alpha EEG activity.** *Biol Cybern* 2004, **92**:54–60.
 46. Ito J, Nikolaev AR, van Leeuwen C: **Dynamics of spontaneous transitions between global brain states.** *Hum Brain Mapp* 2007, **28**:904–913.
 47. Mitra A, Snyder AZ, Blazey T, Raichle ME: **Lag threads organize the brain's intrinsic activity.** *Proc Natl Acad Sci U S A* 2015, **112**:E2235–E2244.

Resting-state human fMRI fluctuations comprise multiple, highly reproducible, temporal sequences of propagated activity. This propagated activity is largely unidirectional within conventionally understood resting-state networks. Resting-state networks naturally emerge as a consequence of shared patterns of propagation.
 48. Mitra A, Snyder AZ, Tagliazucchi E, Laufs H, Raichle ME: **Propagated infra-slow intrinsic brain activity reorganizes across wake and slow wave sleep.** *Elife* 2015, **4**.
 49. Mitra A, Snyder AZ, Hacker CD, Pahwa M, Tagliazucchi E, Laufs H, Leuthardt EC, Raichle ME: **Human cortical-hippocampal dialogue in wake and slow-wave sleep.** *Proc Natl Acad Sci U S A* 2016, **113**:E6868–E6876.
 50. Mitra A, Kraft A, Wright P, Acland B, Snyder AZ, Rosenthal Z, Czerniewski L, Bauer A, Snyder L, Culver J, Lee JM, Raichle ME: **Spontaneous infra-slow brain activity has unique spatiotemporal dynamics and laminar structure.** *Neuron* 2018, **98**:297–305.e6.
 51. Shimaoka Knöpfel D, Song C, Knöpfel T: **State-dependent modulation of slow wave motifs towards awakening.** *Front Cell Neurosci* 2017, **11**:108.
 52. Massimini M: **The sleep slow oscillation as a traveling wave.** *J Neurosci* 2004, **24**:6862–6870.
 53. Townsend RG, Gong P: **Detection and analysis of spatiotemporal patterns in brain activity.** *PLoS Comp Biol* 2018, **14**:e1006643.
 54. Jacobs EAK, Steinmetz NA, Carandini M, Harris KD: **Cortical state fluctuations during sensory decision making.** *bioRxiv* 2019:348193.
 55. Vinck M, Batista-Brito R, Knoblich U, Cardin JA: **Arousal and locomotion make distinct contributions to cortical activity patterns and visual encoding.** *Neuron* 2015, **86**:740–754.
 56. Reimer J, Froudarakis E, Cadwell CR, Yatsenko D, Denfield GH, Tolias AS: **Pupil fluctuations track fast switching of cortical states during quiet wakefulness.** *Neuron* 2014, **84**:355–362.
 57. Pinto L, Goard MJ, Estandian D, Xu M, Kwan AC, Lee SH, Harrison TC, Feng G, Dan Y: **Fast modulation of visual perception by basal forebrain cholinergic neurons.** *Nat Neurosci* 2013, **16**:1857–1863.
 58. McGinley MJ, Vinck M, Reimer J, Batista-Brito R, Zagha E, Cadwell CR, Tolias AS, Cardin JA, McCormick DA: **Waking state: rapid variations modulate neural and behavioral responses.** *Neuron* 2015, **87**:1143–1161.
 59. Steinmetz NA, Zatzka-Haas P, Carandini M, Harris K: **Distributed coding of choice, action and engagement across the mouse brain.** *Nature* 2019. in press.
 60. Niell CM, Stryker MP: **Modulation of visual responses by behavioral state in mouse visual cortex.** *Neuron* 2010, **65**:472–479.
 61. Arieli A, Sterkin A, Grinvald A, Aertsen A: **Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses.** *Science (New York, NY)* 1996, **273**:1868–1871.
 62. Barson D, Hamodi AS, Shen X, Lur G, Constable RT, Cardin J,
 - Crair M, Higley M: **Simultaneous mesoscopic and two-photon imaging of neuronal activity in cortical circuits.** *bioRxiv* 2018:468348.
 The paper demonstrates a method for simultaneous cellular-resolution two-photon calcium imaging of a local microcircuit and wide-field calcium imaging of the entire cortical mantle in awake, behaving mice. Functional connectivity maps between individual, genetically defined neurons and the entire cortex uncover diverse subnetworks within a local microcircuit.
 63. Xiao D, Vanni MP, Mitelut CC, Chan AW, LeDue JM, Xie Y, Chen AC, Swindale NV, Murphy TH: **Mapping cortical mesoscopic networks of single spiking cortical or sub-cortical neurons.** *Elife* 2017, **6**:437.
 64. Peters AJ, Steinmetz NA, Harris KD, Carandini M: **Striatal activity reflects cortical activity patterns.** *bioRxiv* 2019:703710.
 65. Harris KD, Shepherd GMG: **The neocortical circuit: themes and variations.** *Nat Neurosci* 2015, **18**:170–181.
 66. Oberlaender M, de Kock CPJ, Bruno RM, Ramirez A, Meyer HS, Dercksen VJ, Helmstaedter M, Sakmann B: **Cell type-specific three-dimensional structure of thalamocortical circuits in a column of rat vibrissal cortex.** *Cereb Cortex (New York, NY: 1991)* 2011, **22**:2375–2391.
 67. Lu J, Tucciarone J, Padilla-Coreano N, He M, Gordon JA, Huang ZJ: **Selective inhibitory control of pyramidal neuron ensembles and cortical subnetworks by chandelier cells.** *Nat Neurosci* 2017, **20**:1377–1383.
 68. He M, Huang ZJ: **Genetic approaches to access cell types in mammalian nervous systems.** *Curr Opin Neurobiol* 2018, **50**:109–118.
 69. Deco G, Hagmann P, Hudetz AG, Tononi G: **Modeling resting-state functional networks when the cortex falls asleep: local and global changes.** *Cereb Cortex (New York, NY: 1991)* 2014, **24**:3180–3194.
 70. Okun M, Steinmetz NA, Lak A, Dervinis M, Harris KD: **Distinct structure of cortical population activity on fast and infraslow timescales.** *Cereb Cortex* 2019, **29**(5):2196–2210.

71. Steinmetz NA, Koch C, Harris KD, Carandini M: **Challenges and opportunities for large-scale electrophysiology with Neuropixels probes.** *Curr Opin Neurobiol* 2018, **50**:92-100.
72. Guo ZV, Inagaki HK, Daie K, Druckmann S, Gerfen CR, Svoboda K: **Maintenance of persistent activity in a frontal thalamocortical loop.** *Nature* 2017, **545**:181-186.
73. Gao Z, Davis C, Thomas AM, Economo MN, Abrego AM, Svoboda K, De Zeeuw CI, Li N: **A cortico-cerebellar loop for motor planning.** *Nature* 2018, **563**:113-116.
74. Hubel DH, Wiesel TN: **Receptive fields of single neurones in the cats striate cortex.** *J Physiol (Lond)* 1959, **148**:574-591.
75. Riesenhuber M, Poggio T: **Hierarchical models of object recognition in cortex.** *Nat Neurosci* 1999, **2**:1019-1025.
76. Yamins DLK, DiCarlo JJ: **Using goal-driven deep learning models to understand sensory cortex.** *Nat Neurosci* 2016, **19**:356-365.
77. Kar K, Kubilius J, Schmidt K, Issa EB, DiCarlo JJ: **Evidence that recurrent circuits are critical to the ventral stream's execution of core object recognition behavior.** *Nat Neurosci* 2019, **22**:974-983.
78. Liao Q, Poggio T: **Bridging the Gaps Between Residual Learning, Recurrent Neural Networks and Visual Cortex.** 2016 <https://arxiv.org/abs/1604.03640v1>.
79. Meyers EM, Freedman DJ, Kreiman G, Miller EK, Poggio T: **Dynamic population coding of category information in inferior temporal and prefrontal cortex.** *J Neurophysiol* 2008, **100**:1407-1419.
80. Szegedy C, Zaremba W, Sutskever I, Bruna J, Erhan D, Goodfellow I, Fergus R: **Intriguing Properties of Neural Networks.** 2013 <https://arxiv.org/abs/1312.6199v4>.
81. Tang H, Schrimpf M, Lotter W, Moerman C, Paredes A, Ortega Caro J, Hardesty W, Cox D, Kreiman G: **Recurrent computations for visual pattern completion.** *Proc Natl Acad Sci U S A* 2018, **115**:8835-8840.
- The study shows that human object recognition abilities remain robust under heavy occlusion, but the performance of feedforward deep convolutional networks are impaired. Recognition performance was recovered when the model was augmented with attractor dynamics in the output layer.
82. Oh SW, Harris JA, Ng L, Winslow B, Cain N, Mihalas S, Wang Q, Lau C, Kuan L, Henry AM, Mortrud MT, Ouellette B, Nguyen TN, Sorensen SA, Slaughterbeck CR, Wakeman W, Li Y, Feng D, Ho A, Nicholas E, Hirokawa KE, Bohn P, Joines KM, Peng H, Hawrylycz MJ, Phillips JW, Hohmann JG, Wohnoutka P, Gerfen CR, Koch C, Bernard A, Dang C, Jones AR, Zeng H: **A mesoscale connectome of the mouse brain.** *Nature* 2014, **508**:207-214.
83. Fries W: **Cortical projections to the superior colliculus in the macaque monkey: a retrograde study using horseradish peroxidase.** *J Comp Neurobiol* 1984, **230**:55-76.
84. Wang Q, Harris JA, Oh SW, Ng L, Bernard A, Mortrud MT, Ouellette B, Hohmann JG, Koch C, Zeng H: **Cortical projections to the superior colliculus in mouse.** *Society for Neuroscience Annual Meeting* 2013:1-2.
85. Hopfield JJ: **Neural networks and physical systems with emergent collective computational abilities.** *Proc Natl Acad Sci U S A* 1982, **79**:2554-2558.
86. Wang XJ: *The Prefrontal Cortex as a Quintessential Cognitive-Type Neural Circuit.* Oxford University Press; 2013.
87. Mante V, Sussillo D, Shenoy KV, Newsome WT: **Context-dependent computation by recurrent dynamics in prefrontal cortex.** *Nature* 2013, **503**:78-84.
88. Ardid S, Wang XJ, Compte A: **An integrated microcircuit model of attentional processing in the neocortex.** *J Neurosci* 2007, **27**:8486-8495.
89. Ardid S, Wang XJ: **A tweaking principle for executive control: neuronal circuit mechanism for rule-based task switching and conflict resolution.** *J Neurosci* 2013, **33**:19504-19517.
90. Sussillo D, Churchland MM, Kaufman MT, Shenoy KV: **A neural network that finds a naturalistic solution for the production of muscle activity.** *Nat Neurosci* 2015, **18**:1025-1033.
91. Sussillo D: **Neural circuits as computational dynamical systems.** *Curr Opin Neurobiol* 2014, **25**:156-163.
92. Eliasmith C, Stewart TC, Choo X, Bekolay T, DeWolf T, Tang C, Rasmussen D: **A large-scale model of the functioning brain.** *Science* 2012, **338**:1202-1205.
93. Johnson A, van der Meer MA, Redish AD: **Integrating hippocampus and striatum in decision-making.** *Curr Opin Neurobiol* 2007, **17**:692-697.
94. Johnson A, Redish AD: **Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point.** *J Neurosci* 2007, **27**:12176-12189.
95. González OC, Sokolov Y, Krishnan GP, Bazhenov M: **Can sleep protect memories from catastrophic forgetting?** *BioRxiv* 2019:569038.
96. Caze R, Khamassi M, Aubin L, Girard B: **Hippocampal replays under the scrutiny of reinforcement learning models.** *J Neurophysiol* 2018, **120**:2877-2896.