

Correction

NEUROSCIENCE, PSYCHOLOGICAL AND COGNITIVE SCIENCES

Correction for “Oscillatory recurrent gated neural integrator circuits (ORGaNICs), a unifying theoretical framework for neural dynamics,” by David J. Heeger and Wayne E. Mackey, which was first published October 21, 2019; 10.1073/pnas.1911633116 (*Proc. Natl. Acad. Sci. U.S.A.* **116**, 22783–22794).

The authors note that, on page 22784, right column, Eq. 1 appeared incorrectly. The equation should appear as follows:

$$\tau_y \frac{dy_j}{dt} = -y_j + \left(\frac{b_j^+}{1 + b_j^+} \right) z_j + \left(\frac{1}{1 + a_j^+} \right) \hat{y}_j,$$

$$\mathbf{z} = \mathbf{W}_{z\mathbf{x}} \mathbf{x},$$

$$\hat{\mathbf{y}} = \mathbf{W}_{\hat{y}\mathbf{y}} \mathbf{y},$$

$$\mathbf{r} = \mathbf{W}_r \mathbf{y},$$

$$a_j^+ \geq 0 \text{ and } b_j^+ \geq 0.$$

Published under the [PNAS license](#).

First published December 31, 2019.

www.pnas.org/cgi/doi/10.1073/pnas.1921632117



Oscillatory recurrent gated neural integrator circuits (ORGaNICs), a unifying theoretical framework for neural dynamics

David J. Heeger^{a,b,1} and Wayne E. Mackey^{a,b}

^aDepartment of Psychology, New York University, New York, NY 10003; and ^bCenter for Neural Science, New York University, New York, NY 10003

Contributed by David J. Heeger, August 27, 2019 (sent for review July 17, 2019; reviewed by Brent Doiron and Michael M. Halassa)

Working memory is an example of a cognitive and neural process that is not static but evolves dynamically with changing sensory inputs; another example is motor preparation and execution. We introduce a theoretical framework for neural dynamics, based on oscillatory recurrent gated neural integrator circuits (ORGaNICs), and apply it to simulate key phenomena of working memory and motor control. The model circuits simulate neural activity with complex dynamics, including sequential activity and traveling waves of activity, that manipulate (as well as maintain) information during working memory. The same circuits convert spatial patterns of premotor activity to temporal profiles of motor control activity and manipulate (e.g., time warp) the dynamics. Derivative-like recurrent connectivity, in particular, serves to manipulate and update internal models, an essential feature of working memory and motor execution. In addition, these circuits incorporate recurrent normalization, to ensure stability over time and robustness with respect to perturbations of synaptic weights.

computational neuroscience | recurrent neural network | normalization | working memory | motor control

Neuroscience research on working memory has largely focused on sustained delay-period activity (1–4). A large body of experimental research has measured sustained activity in prefrontal cortex (PFC) and/or parietal cortex during delay periods of memory-guided saccade tasks (5–9) and delayed-discrimination and delayed match-to-sample tasks (10–13). Most of the models of working memory, based on neural integrators (see *SI Appendix, Figs. S1–S3* for a primer on neural integrators), are aimed to explain sustained delay-period activity or to explain behavioral phenomena associated with sustained activity (14, 15).

Working memory, however, involves much more than simply holding a piece of information online. In cognitive psychology, the idea of working memory includes manipulating online information dynamically in the context of new sensory input. For example, understanding a complex utterance (with multiple phrases) often involves disambiguating the syntax and/or semantics of the beginning of the utterance based on information at the end of the sentence. Doing so necessitates representing and manipulating long-term dependencies, that is, maintaining a representation of the ambiguous information, and then changing that representation when the ambiguity is resolved. In addition, there are a variety of experimental results that are difficult to reconcile with sustained activity and neural integrator models. Some (if not the majority of) neurons either exhibit sequential activity such that activity is handed off from one neuron to the next during a delay period with each individual neuron being active only transiently (16–21) or they exhibit complex dynamics during delay periods (21–27). Complex dynamics (including oscillations) are evident also in the combined activity (e.g., as measured with local field potentials) of populations of neurons (28, 29). We hypothesize that these complex dynamics serve a purpose, to manipulate working memory representations.

Models of perceptual decision making, like working memory models, are also based on simple neural integrators. Specifically, perceptual decision making has been proposed to involve

integration of noisy sensory information (30–34), a simple form of manipulation, in which neurons literally sum sensory-evoked activity over a period of time. However, a more general theoretical framework for representing and manipulating long-term dependencies is lacking.

Motor preparation and execution, analogous to working memory, involves maintaining a neural representation of a motor plan while manipulating that representation to generate the desired movement dynamics. Neural circuits and systems subserving motor preparation and execution exhibit analogous sustained and sequential activity phenomena (35–39), and there are analogous challenges reconciling these phenomena with neural integrator models.

Long short-term memory units (LSTMs) are machine-learning (ML) algorithms that represent and manipulate long-term dependencies (40). LSTMs are a class of recurrent neural networks. A number of variants of the basic LSTM architecture have been developed and tested for ML applications including language modeling, translation, and speech recognition (41–45). In these and other tasks, the input stimuli contain information across multiple timescales, but the ongoing presentation of stimuli makes it difficult to correctly combine that information over time. This is analogous to the problem of representing and manipulating long-term dependencies mentioned above in working memory, decision making, and motor control. An LSTM handles this problem by updating its internal state over time with a pair of gates: The update gate selects which part(s) of the current input to process, and

Significance

Oscillatory recurrent gated neural integrator circuits (ORGaNICs) are a family of recurrent neural circuits that can simulate a wide range of neurobiological phenomena, various of which have each been explained by different models. This theoretical framework can be used to simulate neural activity with complex dynamics, including sequential and traveling waves of activity. When used to model cognitive processing in the brain, these circuits can both maintain and manipulate information during a working memory task. When used to model motor control, these circuits convert spatial patterns of premotor activity to temporal profiles of motor control activity. ORGaNICs offer a conceptual framework; rethinking cortical computation in these terms should have widespread implications, motivating a variety of experiments.

Author contributions: D.J.H. and W.E.M. designed research; D.J.H. performed research; and D.J.H. wrote the paper.

Reviewers: B.D., University of Pittsburgh; and M.M.H., Massachusetts Institute of Technology.

The authors declare no competing interest.

Published under the PNAS license.

Data deposition: The MATLAB code for recreating the simulation results is available at <https://archive.nyu.edu/handle/2451/60439>.

¹To whom correspondence may be addressed. Email: david.heeger@nyu.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1911633116/-DCSupplemental.

First published October 21, 2019.

NEUROSCIENCE

PSYCHOLOGICAL AND COGNITIVE SCIENCES

The modulators are themselves dynamical systems that depend on the inputs and outputs:

$$\tau_a \frac{da}{dt} = -a + W_{ax}x + f(y), \quad [2]$$

$$\tau_b \frac{db}{dt} = -b + W_{bx}x.$$

The values of τ_a and τ_b are the intrinsic time constant of the modulator neurons. The recurrent modulator \mathbf{a} depends on a function of the output responses $f(\mathbf{y})$, to incorporate recurrent normalization (*Robustness via Normalization and SI Appendix*). ORGaNICs (by analogy with LSTMs) use the modulators to encode a dynamically changing state. The modulators depend on the current inputs and the current outputs, which in turn depend on past inputs and outputs, so the state depends on the current inputs and past context. The modulators can be controlled separately for each neuron so that each neuron can have a different state (different values for a_i and b_i) at each instant in time. In the example that follows, however, all of the neurons in the circuit shared the same state, but that state changed over time.

ORGaNICs are inherently a nonlinear dynamical system because the input drive and the recurrent drive are each multiplied by nonlinear functions of the modulators (Eq. 1) and because the recurrent modulator depends nonlinearly on the output responses (Eq. 2). However, there are circumstances when these equations can be analyzed as a linear system, specifically when the modulators are constant over time, because the only remaining nonlinearity is due to the normalization which simply acts to rescale the responses.

There is considerable flexibility in the formulation of ORGaNICs, with different variants corresponding to different hypothesized neural circuits (*SI Appendix*). In one such variant, each of the modulators can depend on both the inputs and the outputs, unlike Eq. 2 in which only a depends on the output responses. In another variant, the 2 modulators have analogous effects such that larger values of a increase the gain of the recurrent drive, unlike Eq. 1 in which larger values of a decrease the gain of the recurrent drive. In yet another variant, the 2 modulators are coordinated to govern balance between input drive and recurrent drive.

The following subsections describe some examples of ORGaNICs. We begin with a simplified example of a sustained activity circuit, then modify the recurrent weights to simulate sequential activity and traveling waves, and then add multiple recurrent terms for manipulation. Simulated neural responses shown in the figures are intended to exhibit qualitative aspects of neurophysiological phenomena, that is, the models have not (yet) been optimized to replicate published data by tuning or fitting the model parameters. The weights in the various weight matrices were prespecified (not learned) for each of the simulations in this paper (although ORGaNICs are compatible with modified versions of ML algorithms; see *SI Appendix*).

Sustained Activity. We used ORGaNICs to simulate sustained activity during a memory-guided saccade task (Fig. 2), using the circuit depicted in Fig. 1A. In this task, a target is flashed briefly while a subject is fixating the center of a screen (Fig. 1B). After a delay period of several seconds, the fixation point disappears, cueing the subject to make an eye movement to the remembered location of the target.

The modulators in the simulation were constant during each successive phase of the behavioral task. Many experimental protocols in behavioral neuroscience comprise a sequence of distinct phases (including the oculomotor delayed response task; see below figures for more examples). The behavioral cues built into the experimental protocol set the state of the modulators via W_{ax} and W_{bx} in Eq. 2, and the state changed from one phase to the next. During each phase, the modulators were constant and

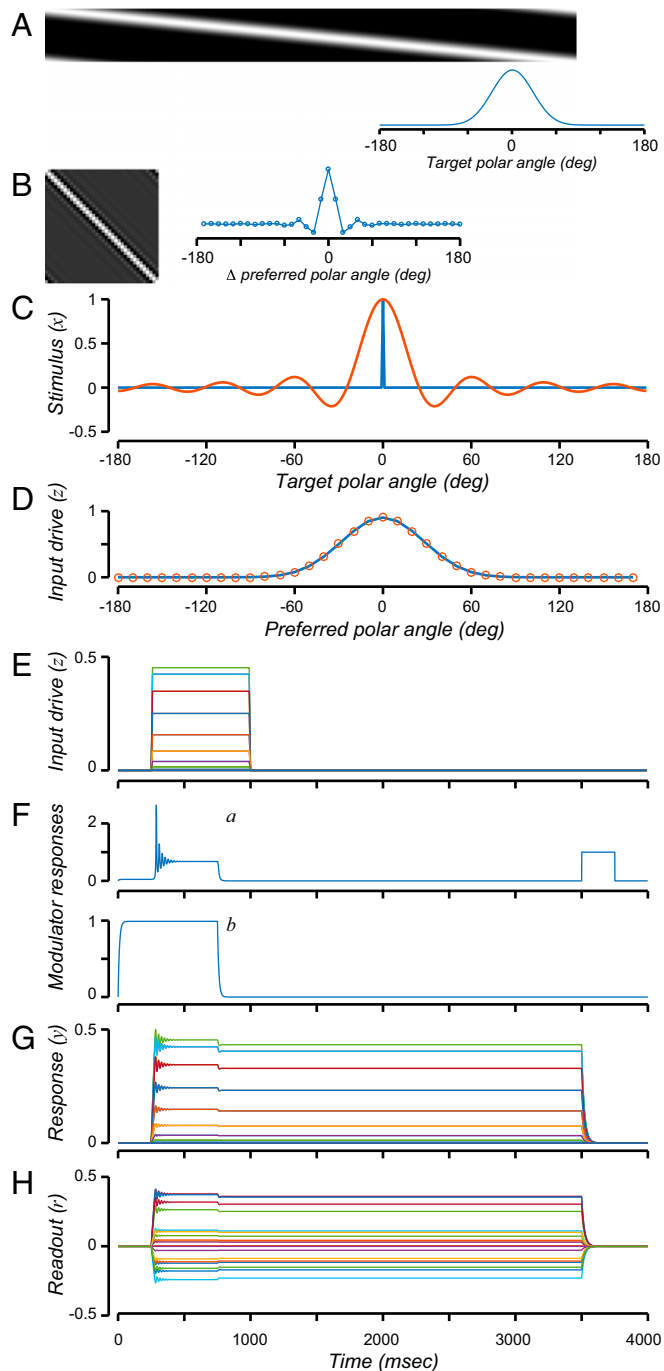


Fig. 2. Sustained activity. (A) Encoding matrix (W_{zx}), each row of which corresponds to a neuron's response field. Graph, response field corresponding to the middle row of the matrix. (B) Recurrent weight matrix (W_{yy}), each row of which corresponds to the recurrent synaptic weights from other neurons in the population. Graph, recurrent weights corresponding to the middle row of the matrix. (C) Input stimulus and reconstructed stimulus. Blue, input stimulus (x) corresponding to target location. Orange, reconstructed stimulus, computed as a weighted sum of the reconstructed input drive (D). (D) Input drive and reconstructed input drive. Blue, input drive (z) to each neuron as a function of that neuron's preferred target location. Orange, reconstructed input drive, computed as a weighted sum of the readout (H). (E) Input drive (z) over time. Each color corresponds to a different neuron. (F) Modulator responses. Top row, a . Bottom row, b . (G) Output responses (y). Each color corresponds to a different neuron. (H) Readout (r). Each color corresponds to a different component of the readout.




Fig. 3. Normalization. (A) Output responses (y), corresponding to the sustained activity circuit depicted in Figs. 1 and 2, but with the recurrent weight matrix scaled by a factor of 1.02. Each color corresponds to a different neuron. (Inset) Full range of responses on an expanded (240 \times) ordinate. (B) Output responses with normalization. Dashed oval, high frequency, coherent, synchronized oscillations following target onset.

drive amplitude) even though some neurons responded strongly to the input whereas others did not.

Recurrent normalization made the circuit robust with respect to imperfections in the recurrent weight matrix (Fig. 3). Without normalization, responses depended critically on fine tuning. For example, we used the sustained activity circuit (Figs. 1 and 2), but with $f(y) = 0$ so that normalization was disabled, and we scaled the recurrent weight matrix by a factor of 1.02. The responses were unstable, growing without bound (Fig. 3A). Including normalization automatically stabilized the activity of the circuit (Fig. 3B). The increases in activity evoked by the recurrent weight matrix (with largest eigenvalues = 1.02) were countered by normalization such that the total activity in the circuit was roughly constant over time ($\|y\|^2 \sim 1$). The ratios of the responses were maintained (Eq. 4), enabling an accurate readout, throughout the delay period. Analogous results were obtained with the other example circuits described below, including those that exhibited oscillatory and sequential dynamics, because the normalization depends on the squared norm of the responses, which was constant over time during the delay period for each of these example circuits. The stability of the normalized responses did not depend on fine-tuning any of the other synaptic weights in the circuit; perturbing those synaptic weights by random values within $\pm 5\%$ yielded virtually identical simulated responses and the responses were stable even when those synaptic weights were perturbed by random values ranging from $0.5\times$ to $2\times$ (see *SI Appendix* for details). We have also implemented a generalization of this recurrent normalization circuit in which each neuron's response can be normalized by an arbitrary (nonnegative) weighted sum of the other neurons in the circuit.

The normalized responses exhibited high-frequency oscillations following target onset that were synchronized across all of the neurons in the circuit (Fig. 3B, dashed oval). There are 2 nested recurrent circuits in ORGaNICs: 1) the recurrent drive and 2) the multiplicative modulators. The high-frequency oscillations emerged because of the inherent delay in the second of these recurrent circuits (i.e., because of the multiplicative modulator underlying normalization). The oscillation frequency depended on the membrane time constants. For the time constants used for Fig. 3, the responses exhibited oscillations in the gamma frequency range. Different intrinsic time constants yielded different oscillation frequencies. The oscillation frequency would have depended also on axon length if we were to include conduction delays.

The responses exhibited lower-frequency oscillations during the delay period (Fig. 3B). These lower-frequency oscillations

emerged because of the recurrent drive in combination with normalization; the recurrent weight matrix was scaled to have eigenvalues greater than 1, which drove the responses to increase over time, but this increase was countered by normalization. These oscillations were synchronized so the ratios of the responses were maintained (Eq. 4), enabling an accurate readout, despite the oscillations.

Sequential Activity. ORGaNICs can be used to generate delay-period activity with complex dynamics, including sequential activity and traveling waves of activity, in addition to sustained activity, and the same theoretical framework was used to analyze them. The key idea is that the recurrent weight matrix can have complex-valued eigenvectors and eigenvalues. One way for this to happen is when the recurrent weights and output responses are complex-valued (*SI Appendix*, Fig. S4). The complex-number notation is just a notational convenience (*SI Appendix*). Another way to generate complex dynamics is for the recurrent weight matrix to be real-valued but asymmetric, such that the responses are real-valued but the eigenvectors are eigenvalues are complex-valued.

One such example circuit was designed to generate sequential activity (Fig. 4). In this example circuit, there were again 36 neurons with the same response fields as in the preceding example (Fig. 2A). The modulators were also the same as in the preceding example, including recurrent normalization. The recurrent weight matrix was real-valued but asymmetric (Fig. 4A).

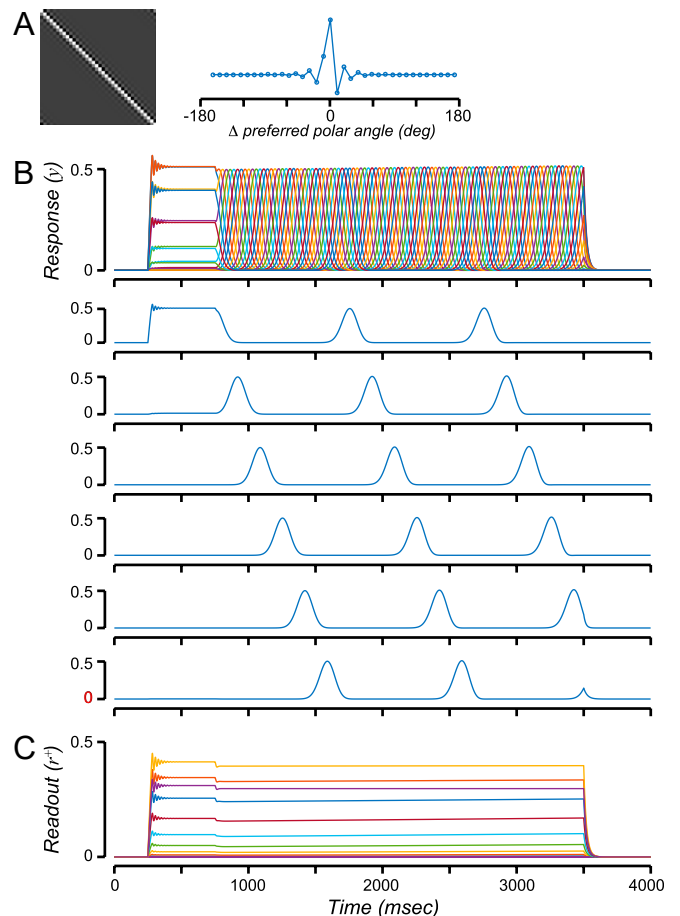


Fig. 4. Sequential activity. (A) Recurrent weight matrix (W_{yy}). Graph, recurrent weights corresponding to the middle row of the matrix. (B) Output responses (y). Each color corresponds to a different neuron. Successive rows, responses of a few example neurons. (C) Readout (r^+). Each color corresponds to a different component of the readout.

Because of the asymmetry, the eigenvectors and eigenvalues of the recurrent weight matrix were complex-valued, and the output responses exhibited oscillatory dynamics (Fig. 4B). The recurrent weight matrix was designed so that the recurrent connectivity depended on the spatial derivative of the neural activity (55), that is, the difference in activity between nearby neurons (SI Appendix). Consequently, the activity was handed off from one neuron to the next during the delay period, analogous to a synfire chain (56–59), but with activity that continuously tiled time (60).

Despite the complex dynamics, the readout was constant over time (Fig. 4C). The readout matrix was again, as for the preceding sustained activity circuit (Fig. 2), computed as a unitary basis for the subspace spanned by the eigenvectors of W_{yy} with corresponding eigenvalues that had real parts = 1. However, the readout was computed as $r^+ = |W_{ry} y|$, that is, the modulus (square root of the sum of squares of real and imaginary parts) of a weighted sum of the responses. Consequently, this circuit was capable of maintaining some (but not all) information about the input during the delay period. Unlike the preceding example, it was not possible to reconstruct the input drive from the readout at arbitrary points in time during the delay period. A linear reconstruction (like that used for the preceding example) generated a copy of the input drive that shifted over time like a traveling wave (SI Appendix, Fig. S5). That is, the information maintained during the delay period was sufficient for discriminating some inputs (e.g., 2 targets with different contrasts or 2 pairs of targets with different spacings) but incapable of discriminating between other inputs (e.g., a single target of the same contrast presented at 2 different locations).

Motor Preparation and Motor Control. ORGANICs are also capable of generating signals, like those needed to execute a complex sequence of movements (e.g., speech, bird song, backside double McTwist 1260 on a snowboard out of the halfpipe). Some actions are ballistic (open loop), meaning that they are executed with no sensory feedback during the movement. Others are closed loop, meaning that the movements are adjusted on the fly based on sensory feedback. ORGANICs evoke patterns of activity over

time that may underlie the execution of both open- and closed-loop movements.

An example of open-loop control (Fig. 5) was implemented using the sequential activity circuit described above, but with a different readout. The encoding matrix and the recurrent matrix were identical to those in the sequential activity circuit. The modulators were also the same as in the preceding examples, including recurrent normalization. The readout was different, simply summing the components, $r^\Sigma = \Sigma \text{Re}(W_{ry} y)$. Different spatial patterns of inputs led to different temporal dynamics of the responses. When the input was chosen to drive a particular eigenvector (i.e., because the input drive was orthogonal to the other eigenvectors), then the readout during the period of motor execution (same as the delay period in the preceding example circuits) was a 1-Hz sinusoid (Fig. 5A). When the input was chosen to drive another eigenvector, then the readout was an 8-Hz sinusoid (Fig. 5C). A linear sum of these inputs evoked a readout that was proportional (because of normalization) to the linear sum of the readouts (Fig. 5D).

How are these temporal profiles of activity generated? Each eigenvector of the recurrent weight matrix is associated with a basis function, a pattern of activity across the population of neurons and over time. Each basis function is a complex exponential (i.e., comprising sine and cosine), the frequency of which is specified by the imaginary part of the corresponding eigenvalue:

$$\omega_i = \left(\frac{1,000}{2\pi\tau_y} \right) \text{Im}(\lambda_i). \quad [5]$$

The value of λ_i is the imaginary part of the i^{th} eigenvalue of the recurrent weight matrix, and ω_i is the corresponding oscillation frequency (in hertz). The factor of 1,000 is needed because the time constant τ_y is presumed to be specified in milliseconds but the oscillation frequency is specified in hertz (cycles per second). The responses exhibit an oscillating traveling wave (Fig. 5B); the response of any individual neuron oscillates over time and the entire pattern of activity across the population of neurons shifts over time (Fig. 5B, orange – yellow – purple – green – cyan – red). For inputs corresponding to different eigenvectors, the responses

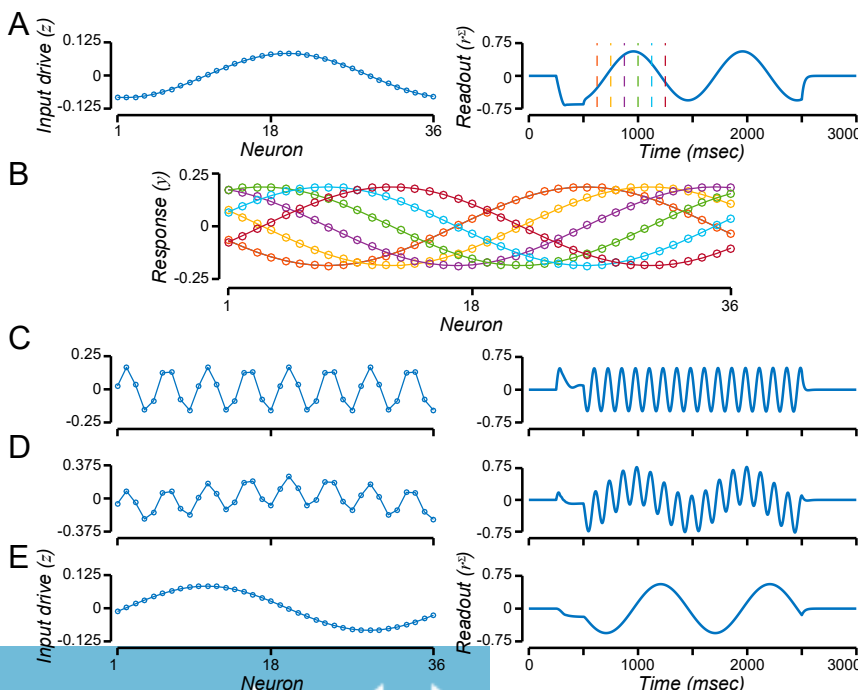


Fig. 5. Motor preparation and motor control. (A) Input drive and readout corresponding to input that drives only the 1-Hz component of the recurrent weight matrix. (A, Left) Input drive (z), spatial pattern activity across the 36 neurons during the premotor time period (250 to 500 ms). (A, Right) Readout (r^Σ) over time. Vertical dashed lines, times corresponding to curves in B. (B) Responses exhibit an oscillating traveling wave of activity. Different colors correspond to different time points, indicated in A. (C) Input drive and readout corresponding to the 8-Hz component of the recurrent weight matrix. Same format as A. (D) Summing the inputs from A and C evokes the sum of the responses. (E) Input drive from A is shifted in space, generating a readout that is shifted in time.

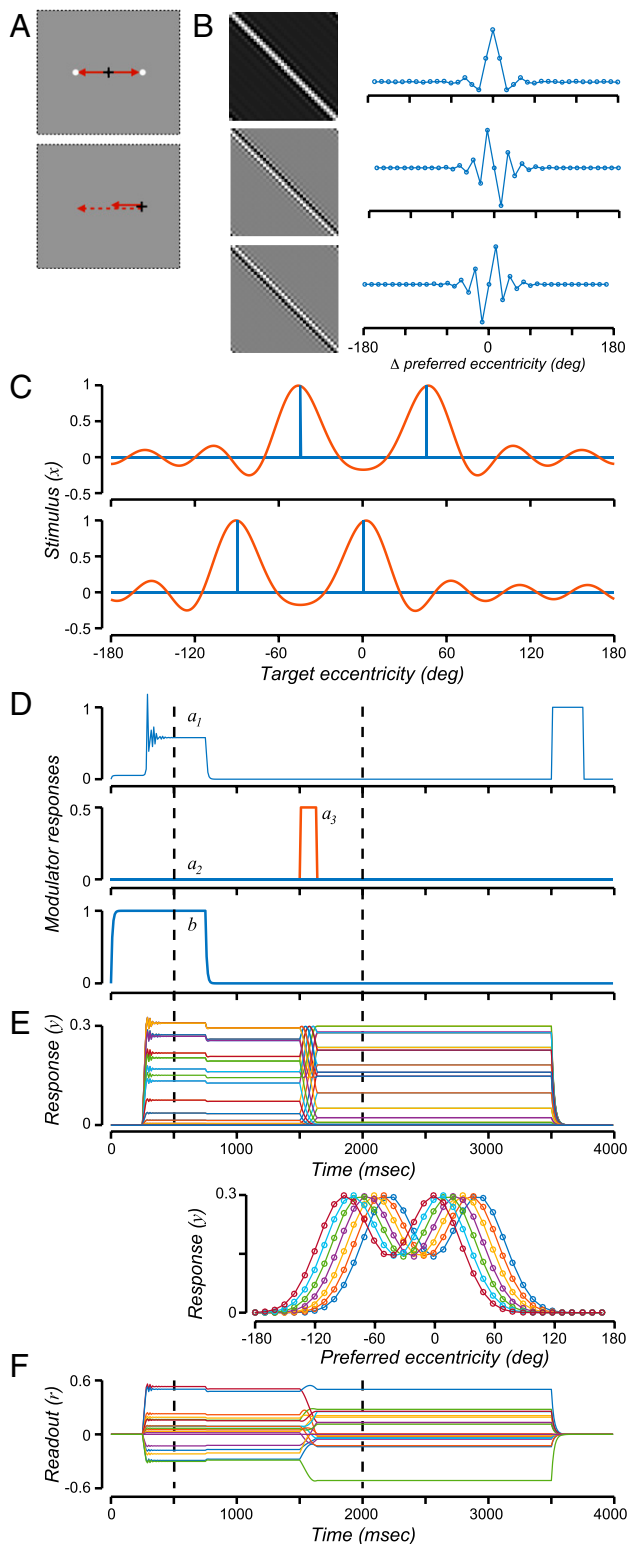


Fig. 6. Spatial updating. (A) Double-step saccade task. (A, Top) Targets presented. (A, Bottom) After eye movement to target 1. White dots, targets. Black cross-hairs, eye position. Solid red lines, planned eye movements without updating. Dashed red line, planned eye movement after updating. (B) Recurrent weight matrices. (Top) Recurrent weight matrix corresponding to modulator a_1 for maintaining a representation of the target locations. (Middle and Bottom) Recurrent weight matrices corresponding to modulators a_2 and a_3 for updating the representation with leftward and rightward eye movements, respectively. (C) Input stimulus and reconstructed stimulus. Blue, input stimulus (x) corresponding to the 2 target positions. Orange, reconstructed stimulus,

oscillate at correspondingly different frequencies (Fig. 5C). The frequencies of the various components corresponding to each of the eigenvalues, for this particular recurrent weight matrix, included a number of other frequencies in addition to the 1- and 8-Hz components shown in the figure. Motor control signals with any arbitrary phase, for each of the frequency components, can be generated by shifting the input drive (Fig. 5E). That way, all combinations of amplitudes, frequencies, and phases can be generated just by changing the spatial pattern of premotor activity, with a fixed, linear readout. This dovetails with experimental evidence demonstrating that the function of motor preparation is to set the initial conditions that generate the desired movement (61–63), and that complex movements are based on a library of motor primitives (64, 65).

The readout for open-loop control is, in general, a linear sum of the responses r^2 . The readout matrix for short-term memory, in the preceding sustained activity circuit (Fig. 2), comprised eigenvectors of the recurrent weight matrix to ensure that the input was recovered during the delay period. However, recovering the input is not the goal for open-loop control. Rather, a sum of the (co)sinusoidal basis functions was used to generate motor control signals for ballistic (open-loop) movements.

ORGaNICs may also generate more complicated control signals. The basis functions are damped oscillators when the modulators are greater than 0 but equal to one another ($a = b$) and constant over time, and when the input is constant over time. If the input is time-varying, then the responses depend on a linear combination of the inputs and the basis functions, and the responses may be used for closed-loop control. If the modulators are also time-varying, and different for each neuron, then the responses may exhibit a wide range of dynamics, with the capability (by analogy with LSTMs) of solving relatively sophisticated tasks (see Introduction for references).

Manipulation: Spatial Updating. A simulation of the double-step saccade task illustrates how ORGaNICs can both maintain and manipulate information over time (Fig. 6). In this task, 2 targets are shown while a subject is fixating the center of a screen (Fig. 6A, Upper). A pair of eye movements are then made in sequence to each of the 2 targets. Eye movements are represented in the brain using retinotopic, that is, eye-centered, coordinates (Fig. 6A, Upper, red lines). Consequently, after making the first eye movement, the plan for the second eye movement must be updated (Fig. 6A, Lower; the solid red line copied from the upper panel no longer points to the second target). This is done by combining a representation of the target location with a copy of the neural signals that control the eye muscles (i.e., corollary discharge) to update the planned eye movement (Fig. 6A, Lower, dashed red line).

The example circuit in Fig. 6 received 2 types of inputs: 1) the target locations at the beginning of the trial (Fig. 6C, Top, blue) and 2) a corollary discharge of the impending eye movement. The targets were assumed to be along the horizontal meridian of the visual field. There were again 36 neurons, but unlike the preceding examples each neuron responded selectively to a different eccentricity along the horizontal meridian of the visual field (i.e., degrees of visual angle away from fixation), not different polar angles around fixation at a fixed eccentricity. The encoding matrix W_{zx} was analogous to that in the preceding examples, but the neurons were selective for target eccentricity

computed as a weighted sum of the readout (F). (Top) Before eye movement to target 1. (Bottom) After eye movement to target 1. (D) Modulator responses. (Top) a_1 . (Middle) a_2 (blue) and a_3 (orange). (Bottom) b . (E) Output responses (y). (Top) Time course of activity, with different colors corresponding to different neurons. (Bottom) Responses for each of several time points (different colors correspond to different time points) while updating the neural representation of the target locations. (F) Readout (r). Dashed vertical lines in D–F correspond to the snapshots in A.

instead of polar angle. Readout and reconstruction were the same as for the sustained activity circuit (Fig. 2).

What distinguishes this example circuit from the preceding examples is that there were 3 recurrent weight matrices (Fig. 6B), the first for maintaining a representation of the target locations (Fig. 6B, Top), the second for changing the representation with leftward eye movements (Fig. 6B, Middle), and the third for changing the representation with rightward eye movements (Fig. 6B, Bottom). As in the preceding examples, the modulators were the same for each neuron in the circuit. Consequently, we can modify Eq. 1:

$$\tau_y \frac{dy}{dt} = -y + \left(\frac{b^+}{1+b^+} \right) z + \left(\frac{1}{1+a_1^+} \right) \hat{y}_1 + \left(\frac{a_2^+}{1+a_2^+} \right) \hat{y}_2 + \left(\frac{a_3^+}{1+a_3^+} \right) \hat{y}_3, \quad [6]$$

$$\hat{y}_k = \mathbf{W}_{y_k} y,$$

where the subscript k indexes over the 3 recurrent weight matrices. The first recurrent weight matrix was identical to that in the sustained activity circuit (Fig. 2B). The second recurrent weight matrix was a discrete approximation to the derivative of the responses (SI Appendix), and third was the negative derivative matrix (i.e., the second and third recurrent matrices differed from one another by a factor of -1). To accommodate 2 dimensions of eye movements, the input drive would depend on 2-dimensional response fields tiling the visual field, and the recurrent drive would depend on 5 recurrent weight matrices, one to maintain the current eye position, a pair for the horizontal component of movements, and another pair for the vertical component (or likewise a pair for the polar angle component of movements and another pair for the radial component).

The modulators were used to encode and update a representation of the target locations (Fig. 6D). As in the preceding examples, the responses followed the input drive at the beginning of the simulated trial because the input modulator was set to $b = 1$ (via \mathbf{W}_{bx} in Eq. 2) by the cue indicating the beginning of the trial. The value of b then switched to be small ($= 0$) before the targets were extinguished, so the output responses exhibited sustained activity that represented the original target locations (Fig. 6C, Top, orange). The modulator a_1 was responsible for recurrent normalization, as in the preceding example circuits. The modulator a_3 was nonzero for a period of time beginning just prior to the eye movement (Fig. 6D, Middle, orange). The amplitude of a_3 and duration of time during which it was nonzero determined the magnitude of updating, that is, corresponding to the amplitude of the impending saccade (for an eye movement in the opposite direction, the amplitude of a_2 , instead of a_3 , would have been nonzero). Finally, the value of the recurrent modulator was set to $a_1 \approx 1$ (via \mathbf{W}_{ax} in Eq. 2) by the cue indicating the end of the trial, causing the output responses to be extinguished.

The output responses exhibited a traveling wave of activity across the topographic map of target locations during the period of time when the neural representation of the targets was updated (Fig. 6E). The readout (Fig. 6F) encoded the 2 target locations, both before and after updating. The readout and decoding matrices were identical to those in the sustained activity circuit (Fig. 2). Preceding the eye movement, the original target locations were reconstructed from the readout (Fig. 6C, Top, orange curve). After the eye movement, the updated target locations were reconstructed (Fig. 6C, Bottom).

Manipulation: Time Warping and Time Reversal. A challenge for models of motor control is to generate movements at different speeds, for example playing a piece of piano music, generating speech (66), or generating birdsong (67) at different tempos. Likewise, a challenge for models of sensory processing is that perception must be tolerant with respect to compression or dilation of temporal signals, for example listening to fast vs. slow

speech (68). A possible mechanism for time warping is to scale the time constants of the neurons (69), all by the same factor, which scales the oscillation frequencies by the inverse of that scale factor (Eq. 5). A fixed value for the scale factor would handle linear time rescaling in which the entire input (and/or output) signal is compressed or dilated accordingly. A neural circuit might compute a time-varying value for the scale factor, based on the inputs and/or outputs, to handle time-varying time warping.

Here, we offer a different mechanism for time warping (also time reversal), making use of the modulators. An example open-loop motor control circuit was implemented that enabled time warping and time reversal (Fig. 7). The encoding matrix and the recurrent matrix were identical to those in the spatial updating example (Fig. 6). The a_1 and b modulators were also the same as in the spatial updating example, but the time courses of the other 2 modulators a_2 and a_3 were different (Fig. 7A). The readout was the same as that in the motor control circuit (Fig. 5), summing across the components r^{Σ} . The input was chosen to drive all of the eigenvectors with randomly chosen amplitudes and phases. Different values of the a_2 and a_3 modulators generated control signals that were time-warped and/or time-reversed. Increasing the modulator response from 1 to 5/3 caused the readout to increase in tempo by 25% (compare Fig. 7B and C); tempo was proportional to $a_2/(1+a_2)$. A time-varying modulator generated time-varying time warping. The circuit exhibited these phenomena because the responses exhibited oscillating traveling waves (Fig. 5B). The readout was a sum of these traveling waves, and the speed of the traveling waves was controlled by the modulators (SI Appendix). When a_3 (instead of a_2) was nonzero, the readout was time reversed (compare Fig. 7B and D) because the traveling waves of activity moved in the opposite direction.

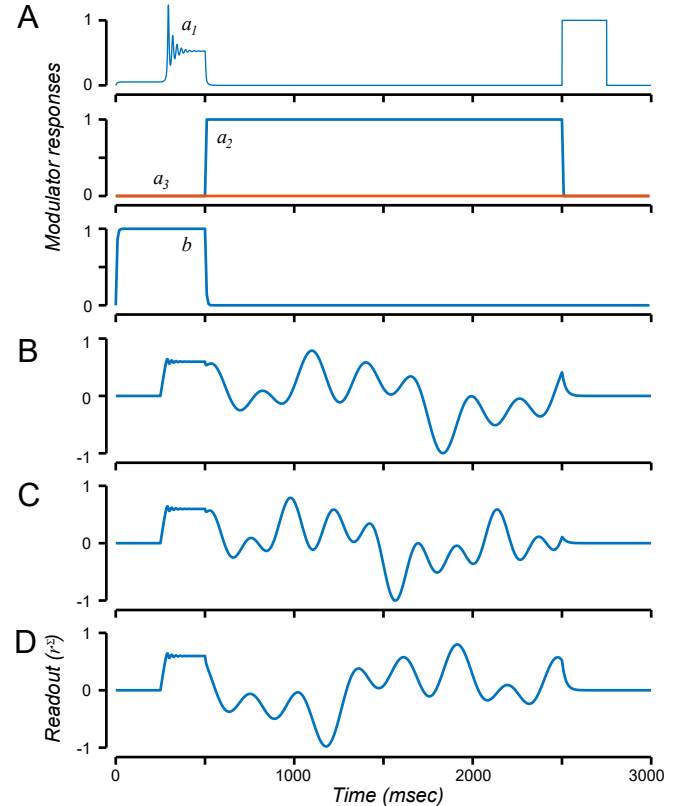


Fig. 7. Time warping and time reversal. (A) Modulator responses. (B) Readout for $a_2 = 1$ and $a_3 = 0$. (C) Time-warped readout for $a_2 = 5/3$ and $a_3 = 0$. (D) Time-reversed readout for $a_2 = 0$ and $a_3 = 1$.

Discussion

We developed a theoretical framework for neural dynamics called ORGaNICs and applied it to simulate key phenomena of working memory and motor control. We demonstrated the following results. 1) Working memory: ORGaNICs can simulate delay-period activity with complex dynamics, including sequential activity and traveling waves of activity, to maintain and manipulate information over time. Derivative-like recurrent connectivity, in particular, generated traveling waves of activity. We propose that these traveling waves play a role in circuit function to manipulate and update internal models. 2) Motor control: The exact same circuits (with the same synaptic weights) were used to generate signals with complex motor dynamics, by converting spatial patterns of premotor activity to temporal profiles of motor control activity. Different spatial patterns of premotor activity evoked different motor control dynamics. These circuits were controlled to manipulate (e.g., time-warp) the motor dynamics. 3) Normalization: Recurrent normalization, via the recurrent modulator, ensured stability over time and robustness with respect to perturbations of synaptic weights. 4) Mechanism: ORGaNICs can be implemented with a simplified biophysical (equivalent electrical circuit) model of pyramidal cells (see *SI Appendix* and ref. 46). There is considerable flexibility in the formulation of ORGaNICs, with different variants corresponding to different hypothesized neural circuits (*SI Appendix*). We demonstrated all of the above results with 2 circuits; the first circuit generated the simulation results in Figs. 2–5 and the second one generated Figs. 6 and 7, noting that the first circuit is equivalent to a special case of the second one.

Because they are generalizations of LSTMs, ORGaNICs can solve tasks that are much more sophisticated than the typical delayed-response tasks used in most cognitive psychology and neuroscience experiments. Indeed, although this is not an ML paper, we note that ORGaNICs may offer computational advantages compared to varieties of LSTMs that are commonly used in ML applications (see *SI Appendix* and ref. 46).

This theoretical framework, of course, includes components previously proposed in the computational/theoretical neuroscience literature, and the ML literature, that have achieved some of the same goals (70–87). However, with ORGaNICs we show that a single unified circuit architecture captures key neurophysiological phenomena associated with sensory, cognitive, and motor functions, each of which has been modeled separately in the previously published literature. Unlike linear recurrent neural networks, the modulators in ORGaNICs introduce nonlinearities (analogous to the gates in LSTMs) that can perform multiple functions including handling long-term dependencies and providing robustness via normalization (discussed below). Unlike most nonlinear recurrent neural nets, ORGaNICs are mathematically tractable, making it possible to derive concrete, quantitative predictions that can be fit to experimental measurements. The theory is tractable when the modulators are constant, that is, during each successive phase of a behavioral task. In addition, the responses of the normalization circuit follow the normalization equation (Eq. 3) exactly, so that this circuit makes predictions that are identical to those of the normalization model, thereby preserving all of the desirable features of that model, which has been fit to hundreds of experimental datasets. In classic work on neural fields (88–90), by contrast, diverse patterns of activity are accomplished by biasing a nonlinear network to different operating points, each having a different solution that can be approximated by local linearization. Here, we start with a linear dynamical system that is fully tractable, characterized by the eigenvalues and eigenvectors of the linear system, but also limited to only those patterns of activity that can be expressed as linear sums of the eigenvectors. We circumnavigate this limitation with the modulators that shape solutions to dynamically change the eigenstructure of the linear system; for each choice of values for the modulators, we have a different linear system. Unlike black-box ML approaches, ORGaNICs provide insight; for example, we understand exactly when and how it is possible to reconstruct an input by reading out the responses during the delay period of a working memory task and how to

generate motor control signals with complex dynamics (see *SI Appendix* for derivations). ML algorithms are particularly useful for computing solutions to optimization problems (e.g., model fitting via gradient descent), and we plan to use ML implementations of ORGaNICs to fit experimental data. ML approaches can also provide inspiration for neuroscience theories (and vice versa), like the links presented here between ORGaNICs and LSTMs. Left open in the current paper is how the weights in the various weight matrices emerge through development and/or learning. We engineered the weights to demonstrate the computational capabilities of this theoretical framework and to illustrate that the theory can reproduce neurobiological phenomena (although ORGaNICs are compatible with modified versions of ML algorithms; see *SI Appendix*). Some of the previously published literature (cited above) focuses on learning. However, having the right circuit architecture is a prerequisite for developing an accurate model of learning.

We propose that ORGaNICs can serve as a unifying theoretical framework for neural dynamics, a canonical computational motif based on recurrent amplification, gated integration, reset, and controlling the effective time constant. Rethinking cortical computation in these terms should have widespread implications, some of which are elucidated in the paragraphs that follow (see also *SI Appendix*).

Sustained delay-period activity and sequential activity are opposite sides of the same coin. ORGaNICs, a straightforward extension of leaky neural integrators and neural oscillators, provide a unified theoretical framework for sustained activity (Fig. 2), oscillatory activity (*SI Appendix*, Fig. S4), and sequential activity (Fig. 4), just by changing the recurrent weight matrix. Indeed, ORGaNICs can switch between these different behaviors. The spatial updating circuit, for example, exhibits sustained activity during the delay periods and sequential activity coincident with the eye movement (Fig. 6). The modulators a_2 and a_3 do the job of toggling between sustained and sequential. We assert that complicated dynamics is the norm, to support manipulation as well as maintenance (e.g., Fig. 6).

ORGaNICs can be used to generate motor control signals, with the very same circuits used to model working memory, just by changing the readout. The circuits convert spatial patterns of input (premotor) activity to temporal profiles of output (motor control) activity. Different spatial patterns of premotor activity evoke motor control outputs with different temporal response dynamics (e.g., as in Figs. 5 and 7), and the modulators provide a means for manipulating (time warping and time reversal) the dynamics (Fig. 7).

ORGaNICs are applicable also to models of sensory integration (e.g., integrating corollary discharge in Fig. 6) and sensory processing (e.g., with normalization as in Fig. 3). ORGaNICs may be stacked in layers such that the inputs to one ORGaNIC are the outputs from one or more other ORGaNICs. Particular stacked architectures encompass convolutional neural nets (i.e., deep nets) as a special case: specifically, when the encoding/embedding weight matrices are convolutional and when the modulators are large ($a_j = b_j \gg 0$) such that the output responses from each layer are dominated by the input drive to that layer. Consequently, working memory, motor control, sensory processing (including prediction over time; see *SI Appendix* and ref. 46), and possibly other cognitive functions (in addition to working memory, such as cognitive control, for example controlling attention) may all share a common canonical computational foundation.

Derivative-like recurrent connectivity (55) simulates sequential activity and traveling waves of activity (Figs. 4–7), and we propose that these traveling waves play a particular role in circuit function. Weight matrices with derivative-like weights are a mainstay of feed-forward models of sensory processing (91, 92), but the contribution of derivative-like weights in recurrent connectivity has been underappreciated. Traveling waves are ubiquitous in cortical activity, but their functional role has remained a mystery (93). We used recurrent weight matrices based on derivatives (i.e., the difference in activity between nearby neurons) to evoke traveling waves of activity that functioned to support manipulation. The traveling

waves served to transform spatial patterns of premotor activity to temporal patterns of motor control activity (Figs. 5 and 7) or to update internal models (working memory representations) whether or not there was an overt movement (Fig. 6).

Why do some neural circuits exhibit sustained activity while others exhibit sequential activity, and what are the relative advantages or disadvantages of each? Sustained activity circuits are useful for short-term memory (i.e., maintenance), but not for other cognitive functions that require manipulation and control. For sustained-activity circuits, a simple linear readout of the responses can be used to reconstruct the input drive (and to approximately reconstruct the input stimulus), at any point in time during a delay period (Fig. 2). In addition, sustained-activity circuits are likely to be more robust than sequential-activity circuits, because all of the components share the same dynamics. Sequential-activity circuits, on the other hand, offer much more flexibility. The same circuit, with the same fixed recurrent weight matrix and the same fixed encoding matrix, can support multiple different functions just by changing the readout. For example, the sequential-activity circuit (Fig. 4) and the motor-control circuit (Fig. 5) were identical except for the readout. For the sequential-activity circuit (Fig. 4), a (nonlinear) modulus readout generated an output that was constant over time (i.e., to support maintenance). For the motor-control circuit (Fig. 5), a linear readout was used to generate control signals as sums of (co)sinusoidal basis functions with various different frequencies and phases. Likewise, the spatial-updating circuit (Fig. 6) and the time-warping/time-reversal circuit (Fig. 7) were identical. This circuit can be used to perform working memory (maintenance and manipulation), and the same circuit (without changing the encoding or recurrent weights) can be used to execute movements with complex dynamics. One way to implement this, for example, would be to have 2 different brain areas with stereotypical intrinsic circuitry (i.e., identical recurrent weights) that support 2 different functions with different readouts. Indeed, there is experimental evidence that different brain areas support different functions with similar circuits, for example parietal areas underlying working memory maintenance and PFC areas underlying motor planning (94). Alternatively, the output from a single circuit could innervate 2 different brain areas, one of which performs the first readout and the other of which performs the second readout, or a single brain area might switch between 2 different readouts (e.g., using a gating mechanism analogous to the modulators in ORGaNICs), corresponding to different behavioral states, without changing the intrinsic connectivity within the circuit. This makes biological sense. Rather than having to change everything (the encoding matrix, the recurrent matrix, the modulators, and the readout), you need only change one thing (the readout matrix) to enable a wide variety of functions. This is not possible with recurrent weight matrices that exhibit sustained activity, simply because there is only a single mode of dynamics (constant over time).

The modulators perform multiple functions and can be implemented with a variety of circuit, cellular, and synaptic mechanisms. The time-varying values of the modulators determine the state of the circuit by controlling the recurrent gain and effective time constant of each neuron in the circuit. The multiple functions of the modulators include normalization (Fig. 3), maintenance (Figs. 2–7), controlling pattern generators (Figs. 5 and 7), gated integration/updates (Fig. 6), time warping and time reversal (Fig. 7), reset (Figs. 2–7), controlling the effective time constant (SI Appendix, Fig. S1), controlling the relative contributions of bottom-up versus top-down connections (95), representing and weighting the reliability of sensory evidence (likelihood) and internal model (prior, expectation) for inference, prediction over time, and multisensory integration (95). ORGaNICs may have multiple recurrent weight matrices, each multiplied by different recurrent modulators, to perform combinations of these functions (Eq. 6 and Figs. 6 and 7). Some of the modulator functions need to be fast and selective (e.g., normalization), likely implemented in local circuits. A variety of mechanisms have been hypothesized for adjusting the gain of local circuits (96–98). Some modulator

functions might depend on thalamocortical loops (20, 99–101). Other modulator functions are relatively nonselective and evolve relatively slowly over time and may be implemented with neuromodulators (102–105).

Recurrent normalization, as implemented with ORGaNICs (Fig. 3), is consonant with the idea that normalization operates via recurrent amplification, that is, that weak inputs are strongly amplified but that strong inputs are only weakly amplified. Several hypotheses for the recurrent circuits underlying normalization have been proposed (50, 51, 96, 106–108), but most of them are inconsistent with experimental observations suggesting that normalization is implemented via recurrent amplification (109–114). ORGaNICs offer a family of dynamical systems models of normalization, each of which comprises coupled neural integrators to implement normalization via recurrent amplification (SI Appendix). When the input drive is constant over time, the circuit achieves an asymptotic stable state in which the output responses follow the normalization equation exactly (Eq. 3).

There is a critical need for developing behavioral tasks that animal models are capable of learning, and that involve both maintaining and manipulating information over time. ORGaNICs (and LSTMs) manage long-term dependencies between sensory inputs at different times, using a combination of gated integration and reset. Typical delayed-response tasks like the memory-guided saccade task are appropriate for studying what psychologists call “short-term memory,” but they are weak probes for studying working memory (115–118), because those tasks do not involve manipulation of information over time. Behavioral tasks that are popular in studies of decision making involve integration of noisy sensory information (30, 32) or integration of probabilistic cues (119). Variants of these tasks (31, 34) might be used to test the gated integration and reset functionality of ORGaNICs. The antisaccade task (120–123) and the double-step saccade task (124–126) might also be used, with delay periods, to test the theory and to characterize how cortical circuits manage long-term dependencies.

Finally, the theory motivates a variety of experiments, some examples of which are as follows. First, the theory predicts that the modulators change the effective time constant and recurrent gain of a PFC neuron. Experimental evidence suggests that the modulatory responses are computed in the thalamus (2, 20, 99). Consequently, manipulating the responses of these thalamic neurons (e.g., via optogenetics) should have a particular impact on both the time constant and recurrent gain of cortical neurons. Second, the specific biophysical implementation (SI Appendix, Fig. S6) predicts that the soma and basal dendrites share input drive, but with opposite sign. This would, of course, have to be implemented with inhibitory interneurons. Third, the theory predicts that that neural activity underlying motor control and working memory is normalized. Normalization might be measured in motor cortex by comparing activity when making each of 2 simple movements vs. the combination of those movements simultaneously, or by comparing activity in one subpopulation of neurons with and without optogenetic stimulation of a separate subpopulation of neurons. Normalization might be measured in working memory circuits by comparing activity when maintaining one item versus multiple items during a delay period (127–129). Fourth, following previous research (130), a model based on ORGaNICs may be fit to behavioral and neurophysiological measurements of working memory. Trial-to-trial variability of behavioral performance during a working memory task has been shown to be linked with trial-to-trial variability in delay-period activity. These data might be fit by adding noise to the responses and/or synaptic weights, leading to drift in activity ratios during a delay period. Fifth, as noted above, variants of sensory integration tasks might be used to test the gated integration and reset functionality of ORGaNICs, and variants of the antisaccade and double-step saccade tasks might also be used, with delay periods, to characterize how cortical circuits manage long-term dependencies.

ACKNOWLEDGMENTS. We thank Mike Landy, Eero Simoncelli, Charlie Burlingham, Gyuri Buzsaki, Mike Long, Kenway Louie, Roozbeh Kiani, E. J. Chichilnisky, and Lindsey Glickfeld for comments and discussion.

1. C. F. Jacobsen, Functions of frontal association area in primates. *Arch. Neurol. Psychiatry* **33**, 558–569 (1935).
2. J. M. Fuster, G. E. Alexander, Neuron activity related to short-term memory. *Science* **173**, 652–654 (1971).
3. J. M. Fuster, Unit activity in prefrontal cortex during delayed-response performance: Neuronal correlates of transient memory. *J. Neurophysiol.* **36**, 61–78 (1973).
4. C. Constantinidis *et al.*, Persistent spiking activity underlies working memory. *J. Neurosci.* **38**, 7020–7028 (2018).
5. J. W. Gnadt, R. A. Andersen, Memory related motor planning activity in posterior parietal cortex of macaque. *Exp. Brain Res.* **70**, 216–220 (1988).
6. S. Funahashi, C. J. Bruce, P. S. Goldman-Rakic, Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J. Neurophysiol.* **61**, 331–349 (1989).
7. P. S. Goldman-Rakic, Cellular basis of working memory. *Neuron* **14**, 477–485 (1995).
8. C. Constantinidis, G. V. Williams, P. S. Goldman-Rakic, A role for inhibition in shaping the temporal flow of information in prefrontal cortex. *Nat. Neurosci.* **5**, 175–180 (2002).
9. D. Schluppeck, C. E. Curtis, P. W. Glimcher, D. J. Heeger, Sustained activity in topographic areas of human posterior parietal cortex during memory-guided saccades. *J. Neurosci.* **26**, 5098–5108 (2006).
10. E. K. Miller, C. A. Erickson, R. Desimone, Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *J. Neurosci.* **16**, 5154–5167 (1996).
11. R. Romo, C. D. Brody, A. Hernández, L. Lemus, Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature* **399**, 470–473 (1999).
12. C. R. Hussar, T. Pasternak, Memory-guided sensory comparisons in the prefrontal cortex: Contribution of putative pyramidal cells and interneurons. *J. Neurosci.* **32**, 2747–2761 (2012).
13. M. J. Goard, G. N. Pho, J. Woodson, M. Sur, Distinct roles of visual, parietal, and frontal motor cortices in memory-guided sensorimotor decisions. *eLife* **5**, e13764 (2016).
14. X. J. Wang, Synaptic basis of cortical persistent activity: The importance of NMDA receptors to working memory. *J. Neurosci.* **19**, 9587–9603 (1999).
15. A. Compte, N. Brunel, P. S. Goldman-Rakic, X. J. Wang, Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cereb. Cortex* **10**, 910–923 (2000).
16. E. H. Baeg *et al.*, Dynamics of population code for working memory in the prefrontal cortex. *Neuron* **40**, 177–188 (2003).
17. S. Fujisawa, A. Amarasingham, M. T. Harrison, G. Buzsáki, Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nat. Neurosci.* **11**, 823–833 (2008).
18. E. Pastalkova, V. Itskov, A. Amarasingham, G. Buzsáki, Internally generated cell assembly sequences in the rat hippocampus. *Science* **321**, 1322–1327 (2008).
19. C. D. Harvey, P. Coen, D. W. Tank, Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* **484**, 62–68 (2012).
20. L. I. Schmitt *et al.*, Thalamic amplification of cortical connectivity sustains attentional control. *Nature* **545**, 219–223 (2017).
21. M. Lundqvist, P. Herman, E. K. Miller, Working memory: Delay activity, yes! Persistent activity? Maybe not. *J. Neurosci.* **38**, 7013–7019 (2018).
22. C. D. Brody, A. Hernández, A. Zainos, R. Romo, Timing and neural encoding of somatosensory parametric working memory in macaque prefrontal cortex. *Cereb. Cortex* **13**, 1196–1207 (2003).
23. C. K. Machens, R. Romo, C. D. Brody, Flexible control of mutual inhibition: A neural model of two-interval discrimination. *Science* **307**, 1121–1124 (2005).
24. M. Shafi *et al.*, Variability in neuronal activity in primate cortex during working memory tasks. *Neuroscience* **146**, 1082–1108 (2007).
25. D. A. Markowitz, C. E. Curtis, B. Pesaran, Multiple component networks support working memory in prefrontal cortex. *Proc. Natl. Acad. Sci. U.S.A.* **112**, 11084–11089 (2015).
26. D. Kobak *et al.*, Demixed principal component analysis of neural population data. *eLife* **5**, e10989 (2016).
27. J. D. Murray *et al.*, Stable population coding for working memory coexists with heterogeneous neural dynamics in prefrontal cortex. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 394–399 (2017).
28. B. Pesaran, J. S. Pezaris, M. Sahani, P. P. Mitra, R. A. Andersen, Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nat. Neurosci.* **5**, 805–811 (2002).
29. M. Lundqvist *et al.*, Gamma and beta bursts underlie working memory. *Neuron* **90**, 152–164 (2016).
30. M. N. Shadlen, W. T. Newsome, Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J. Neurophysiol.* **86**, 1916–1936 (2001).
31. J. I. Gold, M. N. Shadlen, The influence of behavioral context on the representation of a perceptual decision in developing oculomotor commands. *J. Neurosci.* **23**, 632–651 (2003).
32. B. W. Brunton, M. M. Botvinick, C. D. Brody, Rats and humans can optimally accumulate evidence for decision-making. *Science* **340**, 95–98 (2013).
33. R. Kiani, A. K. Churchland, M. N. Shadlen, Integration of direction cues is invariant to the temporal gap between them. *J. Neurosci.* **33**, 16483–16489 (2013).
34. B. A. Purcell, R. Kiani, Hierarchical decision processes that operate over distinct timescales underlie choice and changes in strategy. *Proc. Natl. Acad. Sci. U.S.A.* **113**, E4531–E4540 (2016).
35. E. Aksay, G. Gamkrelidze, H. S. Seung, R. Baker, D. W. Tank, In vivo intracellular recording and perturbation of persistent activity in a neural integrator. *Nat. Neurosci.* **4**, 184–193 (2001).
36. R. H. Hahnloser, A. A. Kozhevnikov, M. S. Fee, An ultra-sparse code underlies the generation of neural sequences in a songbird. *Nature* **419**, 65–70 (2002).
37. A. A. Kozhevnikov, M. S. Fee, Singing-related activity of identified HVC neurons in the zebra finch. *J. Neurophysiol.* **97**, 4271–4283 (2007).
38. D. Z. Jin, N. Fujii, A. M. Graybiel, Neural representation of time in cortico-basal ganglia circuits. *Proc. Natl. Acad. Sci. U.S.A.* **106**, 19156–19161 (2009).
39. M. A. Long, D. Z. Jin, M. S. Fee, Support for a synaptic chain model of neuronal sequence generation. *Nature* **468**, 394–399 (2010).
40. S. Hochreiter, J. Schmidhuber, Long short-term memory. *Neural Comput.* **9**, 1735–1780 (1997).
41. A. Graves, Generating sequences with recurrent neural networks. arXiv:1308.0850 (4 August 2013).
42. A. Graves, A.-r. Mohamed, G. Hinton, “Speech recognition with deep recurrent neural networks” in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing (IEEE, 2013)*, pp. 6645–6649.
43. K. Cho *et al.*, Learning phrase representations using RNN encoder-decoder for statistical machine translation. arXiv:1406.1078 (3 June 2014).
44. A. Graves, G. Wayne, I. Danihelka, Neural Turing machines. arXiv:1410.5401 (20 October 2014).
45. I. Sutskever, O. Vinyals, Q. V. Le, “Sequence to sequence learning with neural networks” in *Advances in Neural Information Processing Systems (MIT Press, 2014)*, vol. 2, pp. 3104–3112.
46. D. J. Heeger, W. E. Mackey, ORGaNICS: A theory of working memory in brains and machines. arXiv:1803.06288 (16 March 2018).
47. D. J. Heeger, W. E. Mackey, ORGaNICS: A canonical neural circuit computation. bioRxiv:10.1101/506337 (26 December 2018).
48. D. J. Heeger, Supplemental Material for “Oscillatory recurrent gated neural integrator circuits (ORGaNICS), a unifying theoretical framework for neural dynamics.” NYU Faculty Digital Archive. <https://archive.nyu.edu/handle/2451/60439>. Deposited 10 October 2019.
49. D. J. Heeger, Half-squaring in responses of cat striate cells. *Vis. Neurosci.* **9**, 427–443 (1992).
50. D. J. Heeger, Normalization of cell responses in cat striate cortex. *Vis. Neurosci.* **9**, 181–197 (1992).
51. M. Carandini, D. J. Heeger, Normalization as a canonical neural computation. *Nat. Rev. Neurosci.* **13**, 51–62 (2011).
52. H. S. Seung, How the brain keeps the eyes still. *Proc. Natl. Acad. Sci. U.S.A.* **93**, 13339–13344 (1996).
53. A. A. Koulakov, S. Raghavachari, A. Kepecs, J. E. Lisman, Model for a robust neural integrator. *Nat. Neurosci.* **5**, 775–782 (2002).
54. C. D. Brody, R. Romo, A. Kepecs, Basic mechanisms for graded persistent activity: Discrete attractors, continuous attractors, and dynamic representations. *Curr. Opin. Neurobiol.* **13**, 204–211 (2003).
55. K. Zhang, Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: A theory. *J. Neurosci.* **16**, 2112–2126 (1996).
56. M. Abeles, *Corticonics: Neural Circuits of the Cerebral Cortex* (Cambridge University Press, 1991).
57. M. Abeles, H. Bergman, E. Margalit, E. Vaadia, Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *J. Neurophysiol.* **70**, 1629–1638 (1993).
58. E. Bienenstock, A model of neocortex. *Network* **6**, 179–224 (1995).
59. M. Herrmann, J. Hertz, A. Prügel-Bennett, Analysis of synfire chains. *Network* **6**, 403–414 (1995).
60. E. M. Izhikevich, Polychronization: Computation with spikes. *Neural Comput.* **18**, 245–282 (2006).
61. M. M. Churchland, J. P. Cunningham, M. T. Kaufman, S. I. Ryu, K. V. Shenoy, Cortical preparatory activity: Representation of movement or first cog in a dynamical machine? *Neuron* **68**, 387–400 (2010).
62. K. V. Shenoy, M. Sahani, M. M. Churchland, Cortical control of arm movements: A dynamical systems perspective. *Annu. Rev. Neurosci.* **36**, 337–359 (2013).
63. A. A. Russo *et al.*, Motor cortex embeds muscle-like commands in an untangled population response. *Neuron* **97**, 953–966.e8 (2018).
64. K. A. Thoroughman, R. Shadmehr, Learning of action through adaptive combination of motor primitives. *Nature* **407**, 742–747 (2000).
65. S. F. Giszter, Motor primitives—New data and future questions. *Curr. Opin. Neurobiol.* **33**, 156–165 (2015).
66. M. A. Long *et al.*, Functional segregation of cortical regions underlying speech timing and articulation. *Neuron* **89**, 1187–1193 (2016).
67. M. A. Long, M. S. Fee, Using temperature to analyse temporal dynamics in the songbird motor pathway. *Nature* **456**, 189–194 (2008).
68. Y. Lerner, C. J. Honey, M. Katkov, U. Hasson, Temporal scaling of neural responses to compressed and dilated natural speech. *J. Neurophysiol.* **111**, 2433–2444 (2014).
69. R. Güttig, H. Sompolinsky, Time-warp-invariant neuronal processing. *PLoS Biol.* **7**, e1000141 (2009).
70. S. Grossberg, Nonlinear neural networks: Principles, mechanisms, and architectures. *Neural Netw.* **1**, 17–61 (1988).
71. B. A. Olshausen, C. H. Anderson, D. C. Van Essen, A neurobiological model of visual attention and invariant pattern recognition based on dynamic routing of information. *J. Neurosci.* **13**, 4700–4719 (1993).
72. R. C. O'Reilly, M. J. Frank, Making working memory work: A computational model of learning in the prefrontal cortex and basal ganglia. *Neural Comput.* **18**, 283–328 (2006).
73. G. Mongillo, O. Barak, M. Tsodyks, Synaptic theory of working memory. *Science* **319**, 1543–1546 (2008).
74. M. S. Goldman, Memory without feedback in a neural network. *Neuron* **61**, 621–634 (2009).

75. M. Lundqvist, A. Compte, A. Lansner, Bistable, irregular firing and population oscillations in a modular attractor memory network. *PLoS Comput. Biol.* **6**, e1000803 (2010).
76. M. Lundqvist, P. Herman, A. Lansner, Theta and gamma power increases and alpha/beta power decreases with memory load in an attractor network model. *J. Cogn. Neurosci.* **23**, 3008–3020 (2011).
77. S. Druckmann, D. B. Chklovskii, Neuronal circuits underlying persistent representations despite time varying activity. *Curr. Biol.* **22**, 2095–2103 (2012).
78. G. Hennequin, T. P. Vogels, W. Gerstner, Optimal control of transient dynamics in balanced networks supports generation of complex movements. *Neuron* **82**, 1394–1406 (2014).
79. D. Sussillo, M. M. Churchland, M. T. Kaufman, K. V. Shenoy, A neural network that finds a naturalistic solution for the production of muscle activity. *Nat. Neurosci.* **18**, 1025–1033 (2015).
80. K. Rajan, C. D. Harvey, D. W. Tank, Recurrent network models of sequence generation and memory. *Neuron* **90**, 128–142 (2016).
81. R. Costa, I. A. Assael, B. Shillingford, N. de Freitas, T. Vogels, "Cortical microcircuits as gated-recurrent neural networks" in *31st Annual Conference on Neural Information Processing Systems*, U. von Luxburg et al., Eds. (NIPS, La Jolla, CA, 2017), vol. 1, pp. 272–283.
82. J. M. Murray, G. S. Escola, Learning multiple variable-speed sequences in striatum via cortical tutoring. *eLife* **6**, e26084 (2017).
83. V. Goudar, D. V. Buonomano, Encoding sensory and motor patterns as time-invariant trajectories in recurrent neural networks. *eLife* **7**, e31134 (2018).
84. N. Kraynyukova, T. Tchumatchenko, Stabilized supralinear network can give rise to bistable, oscillatory, and persistent activity. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 3464–3469 (2018).
85. E. Orhan, W. J. Ma, A diverse range of factors affect the nature of neural representations underlying short-term memory. bioRxiv:10.1101/244707 (13 October 2018).
86. J. P. Stroud, M. A. Porter, G. Hennequin, T. P. Vogels, Motor primitives in space and time via targeted gain modulation in cortical networks. *Nat. Neurosci.* **21**, 1774–1783 (2018).
87. C. Tallec, Y. Ollivier, Can recurrent neural networks warp time? arXiv:1804.11188 (23 March 2018).
88. S. Amari, Dynamics of pattern formation in lateral-inhibition type neural fields. *Biol. Cybern.* **27**, 77–87 (1977).
89. B. Ermentrout, Neural networks as spatio-temporal pattern-forming systems. *Rep. Prog. Phys.* **61**, 353–430 (1998).
90. P. C. Bressloff, Spatiotemporal dynamics of continuum neural fields. *J. Phys. A Math. Theor.* **45**, 033001 (2011).
91. E. H. Adelson, J. R. Bergen, "The Plenoptic function and the elements of early vision" in *Computational Models of Visual Processing*, M. S. Landy, J. A. Movshon, Eds. (MIT Press, 1991), pp. 3–20.
92. E. P. Simoncelli, D. J. Heeger, A model of neuronal responses in visual area MT. *Vision Res.* **38**, 743–761 (1998).
93. L. Muller, F. Chavane, J. Reynolds, T. J. Sejnowski, Cortical travelling waves: Mechanisms and computational principles. *Nat. Rev. Neurosci.* **19**, 255–268 (2018).
94. W. E. Mackey, C. E. Curtis, Distinct contributions by frontal and parietal cortices support working memory. *Sci. Rep.* **7**, 6188 (2017).
95. D. J. Heeger, Theory of cortical function. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 1773–1782 (2017).
96. M. Carandini, D. J. Heeger, Summation and division by neurons in primate visual cortex. *Science* **264**, 1333–1336 (1994).
97. M. Carandini, D. J. Heeger, W. Senn, A synaptic explanation of suppression in visual cortex. *J. Neurosci.* **22**, 10053–10065 (2002).
98. F. S. Chance, L. F. Abbott, A. D. Reyes, Gain modulation from background synaptic input. *Neuron* **35**, 773–782 (2002).
99. Z. V. Guo et al., Maintenance of persistent activity in a frontal thalamocortical loop. *Nature* **545**, 181–186 (2017).
100. R. V. Rikhye, A. Gilra, M. M. Halassa, Thalamic regulation of switching between cortical representations enables cognitive flexibility. *Nat. Neurosci.* **21**, 1753–1763 (2018).
101. R. V. Rikhye, R. D. Wimmer, M. M. Halassa, Toward an integrative theory of thalamic function. *Annu. Rev. Neurosci.* **41**, 163–183 (2018).
102. K. Thurley, W. Senn, H. R. Lüscher, Dopamine increases the gain of the input-output response of rat prefrontal pyramidal neurons. *J. Neurophysiol.* **99**, 2985–2997 (2008).
103. V. Varga et al., Fast synaptic subcortical control of hippocampal circuits. *Science* **326**, 449–453 (2009).
104. E. Marder, Neuromodulation of neuronal circuits: Back to the future. *Neuron* **76**, 1–11 (2012).
105. K. Wei et al., Serotonin affects movement gain control in the spinal cord. *J. Neurosci.* **34**, 12690–12700 (2014).
106. H. Ozeki, I. M. Finn, E. S. Schaffer, K. D. Miller, D. Ferster, Inhibitory stabilization of the cortical network underlies visual surround suppression. *Neuron* **62**, 578–592 (2009).
107. K. Louie, T. LoFaro, R. Webb, P. W. Glimcher, Dynamic divisive normalization predicts time-varying value coding in decision-related circuits. *J. Neurosci.* **34**, 16046–16057 (2014).
108. D. B. Rubin, S. D. Van Hooser, K. D. Miller, The stabilized supralinear network: A unifying circuit motif underlying multi-input integration in sensory cortex. *Neuron* **85**, 402–417 (2015).
109. H. Adesnik, M. Scanziani, Lateral competition for cortical space by layer-specific horizontal circuits. *Nature* **464**, 1155–1160 (2010).
110. X. Huang, Y. M. Elyada, W. H. Bosking, T. Walker, D. Fitzpatrick, Optogenetic assessment of horizontal interactions in primary visual cortex. *J. Neurosci.* **34**, 4976–4990 (2014).
111. J. J. Nassi, M. C. Avery, A. H. Cetin, A. W. Roe, J. H. Reynolds, Optogenetic activation of normalization in alert macaque visual cortex. *Neuron* **86**, 1504–1517 (2015).
112. T. K. Sato, B. Haider, M. Häusser, M. Carandini, An excitatory basis for divisive normalization in visual cortex. *Nat. Neurosci.* **19**, 568–570 (2016).
113. H. Adesnik, Synaptic mechanisms of feature coding in the visual cortex of awake mice. *Neuron* **95**, 1147–1159.e4 (2017).
114. K. A. Bolding, K. M. Franks, Recurrent cortical circuits implement concentration-invariant odor coding. *Science* **361**, eaat6904 (2018).
115. R. C. Atkinson, R. M. Shiffrin, "Human memory: A proposed system and its control processes" in *Psychology of Learning and Motivation*, K.W. Spence, J. T. Spence, Eds. (Elsevier, 1968), vol. 2, pp. 89–195.
116. N. Cowan, *Attention and Memory: An Integrated Framework* (Oxford University Press, 1998).
117. N. Cowan, What are the differences between long-term, short-term, and working memory? *Prog. Brain Res.* **169**, 323–338 (2008).
118. B. R. Postle, The cognitive neuroscience of visual short-term memory. *Curr. Opin. Behav. Sci.* **1**, 40–46 (2015).
119. T. Yang, M. N. Shadlen, Probabilistic reasoning by neurons. *Nature* **447**, 1075–1080 (2007).
120. P. E. Hallett, Primary and secondary saccades to goals defined by instructions. *Vision Res.* **18**, 1279–1296 (1978).
121. S. Funahashi, M. V. Chafee, P. S. Goldman-Rakic, Prefrontal neuronal activity in rhesus monkeys performing a delayed anti-saccade task. *Nature* **365**, 753–756 (1993).
122. D. P. Munoz, S. Everling, Look away: The anti-saccade task and the voluntary control of eye movement. *Nat. Rev. Neurosci.* **5**, 218–228 (2004).
123. K. Johnson, S. Everling, Neural activity in monkey prefrontal cortex is modulated by task context and behavioral instruction during delayed-match-to-sample and conditional prosaccade-antisaccade tasks. *J. Cogn. Neurosci.* **18**, 749–765 (2006).
124. G. Westheimer, Eye movement responses to a horizontally moving visual stimulus. *AMA Arch. Ophthalmol.* **52**, 932–941 (1954).
125. W. Becker, R. Jürgens, An analysis of the saccadic system by means of double step stimuli. *Vision Res.* **19**, 967–983 (1979).
126. M. E. Goldberg, C. J. Bruce, Primate frontal eye fields. III. Maintenance of a spatially accurate saccade signal. *J. Neurophysiol.* **64**, 489–508 (1990).
127. W. J. Ma, W. Huang, No capacity limit in attentional tracking: Evidence for probabilistic inference under a resource constraint. *J. Vis.* **9**, 3.1–30 (2009).
128. Z. Wei, X. J. Wang, D. H. Wang, From distributed resources to limited slots in multiple-item working memory: A spiking network model with normalization. *J. Neurosci.* **32**, 11228–11240 (2012).
129. S. Keshvari, R. van den Berg, W. J. Ma, No evidence for an item limit in change detection. *PLoS Comput. Biol.* **9**, e1002927 (2013).
130. K. Wimmer, D. Q. Nykamp, C. Constantinidis, A. Compte, Bump attractor dynamics in prefrontal cortex explains behavioral precision in spatial working memory. *Nat. Neurosci.* **17**, 431–439 (2014).