



Multitasking via baseline control in recurrent neural networks

Shun Ogawa^{a,1}, Francesco Fumarola^{a,1}, and Luca Mazzucato^{b,c,2}

Edited by Terrence Sejnowski, Salk Institute for Biological Studies, La Jolla, CA; received March 29, 2023; accepted May 31, 2023

Changes in behavioral state, such as arousal and movements, strongly affect neural activity in sensory areas, and can be modeled as long-range projections regulating the mean and variance of baseline input currents. What are the computational benefits of these baseline modulations? We investigate this question within a brain-inspired framework for reservoir computing, where we vary the quenched baseline inputs to a recurrent neural network with random couplings. We found that baseline modulations control the dynamical phase of the reservoir network, unlocking a vast repertoire of network phases. We uncovered a number of bistable phases exhibiting the simultaneous coexistence of fixed points and chaos, of two fixed points, and of weak and strong chaos. We identified several phenomena, including noise-driven enhancement of chaos and ergodicity breaking; neural hysteresis, whereby transitions across a phase boundary retain the memory of the preceding phase. In each bistable phase, the reservoir performs a different binary decision-making task. Fast switching between different tasks can be controlled by adjusting the baseline input mean and variance. Moreover, we found that the reservoir network achieves optimal memory performance at any first-order phase boundary. In summary, baseline control enables multitasking without any optimization of the network couplings, opening directions for brain-inspired artificial intelligence and providing an interpretation for the ubiquitously observed behavioral modulations of cortical activity.

recurrent neural networks | mean field theory | reservoir computing | multitasking | decision-making

The activity of neurons across cortical areas is strongly modulated by changes in behavioral state such as arousal (1, 2), movements (3–6), and task-engagement (7). Intracellular recordings showed that these behavioral modulations are mediated by a change of baseline synaptic currents, likely originating from the thalamus and other subcortical areas (8, 9). Such baseline modulations exert strong effects on neural activity explaining up to 50% of its variance across cortical areas, a much larger effect compared to the task-related modulations (4–6). The functional role of these baseline modulations differs across experiments and areas, with arousal- or locomotion-induced improvement of visual (3, 10–13) and gustatory processing (2, 14), but degradation of auditory processing (15–17).

We aim to shed light on the potential role of baseline modulations on cortical activity within the framework of reservoir computing, a powerful tool based on recurrent neural networks (RNNs) with random couplings (Fig. 1). Random RNNs can recapitulate different dynamical phases observed in cortical circuits, such as silent or chaotic activity (18), fixed points (19), and the balanced regime (20) and provide a simple explanation for task selectivity features (21) and the heterogeneity of timescales (22) observed in cortical neurons. Random RNNs can achieve optimal performance in memory tasks when poised at a critical point either by fine-tuning their random couplings (23) or their noisy input (24).

Following recent theoretical (2, 11) and experimental studies (4, 25), we modeled the effect of changes in an animal's behavioral state as changes in the mean and across-neurons variance of the constant baseline input currents to an RNN (Fig. 1A). We found that baseline modulations steer the network activity to continuously interpolate between a large set of dynamical phases (Fig. 1B). Beyond known phases, such as fixed points and chaos, baseline modulations unlocked ergodicity-breaking phases, where the network activity can switch between weak and strong chaos, between a fixed point and chaos, or between two fixed points, depending on the initial conditions. All these different phases were continuously connected and achieved without any training or fine-tuning of synaptic couplings. We found an effect where an increase in quenched noise can induce chaos. When interpolating adiabatically between phases via baseline modulations, we found a manifestation of the phenomenon of neural hysteresis, whereby the network activity retains a memory of the path followed in phase space (Fig. 1C). We found that

Significance

Changes in an animal's behavioral state, such as arousal and movements, induce complex modulations of the baseline input currents to sensory areas, eliciting sensory modality-specific effects. A simple computational principle explaining the effects of baseline modulations on recurrent cortical circuits is lacking. We investigate the benefits of baseline modulations using a reservoir computing approach in recurrent neural networks with random couplings. Baseline modulations unlock a set of network phases and phenomena, including chaos enhancement, neural hysteresis, and ergodicity breaking. Strikingly, baseline modulations enable reservoir networks to perform multiple tasks, without any optimization of the network couplings. Baseline control of network dynamics opens directions for brain-inspired artificial intelligence and sheds light on behavioral modulations of cortical activity.

Author contributions: L.M. designed research; S.O., F.F., and L.M. performed research; S.O., F.F., and L.M. contributed new reagents/analytic tools; S.O., F.F., and L.M. analyzed data; and S.O., F.F., and L.M. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Copyright © 2023 the Author(s). Published by PNAS. This article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](#).

¹S.O. and F.F. contributed equally to this work. They have hence left RIKEN.

²To whom correspondence may be addressed. Email: lmazzuca@uoregon.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2304394120/-/DCSupplemental>.

Published August 7, 2023.




Fig. 1. Summary of main results. (A) Random neural network where the baseline input current b_i to the i -th neuron is drawn from a normal distribution with mean and variance μ and σ^2 . (B) Network phase diagram for varying μ and σ^2 shows four phases: fixed point (blue); chaos (cyan); bistable phase with coexistence of two fixed points (green); bistable phase with coexistence of fixed point and chaos (brown). (C) Neural hysteresis: Adiabatic changes in baseline variance $\sigma(t)^2$ lead to discontinuous transitions crossing over phase boundaries, retaining memory of the previous phase (blue: network simulations; red: exact DMFT calculation; y-axis: mean activity M). (D) Optimal memory capacity is achieved by varying σ across a phase boundary where the largest Lyapunov exponent (LLE) crosses zero. (E) Baseline control of multitasking. The two bistable network phases (chaos/fixed-point and double fixed-point phases: brown and green, respectively; see B) can be harnessed by a reservoir network to perform two different tasks: a delayed two-alternative forced-choice task (2AFC) in the double FP phase and a delayed go/no-go task (G/NG) in the chaos/FP phase. *Bottom*: Six trials, alternating 2AFC and G/NG blocks (green and red lines represent task rule onset), where in each block, stimuli from two classes are presented (blue/cyan and orange/brown color-shaded intervals represent the two classes for each task). After a delay, the decision outcome is read out (dot-dashed lines; pink lines: representative activity of four neurons).

baseline modulations can achieve optimal memory performance by poising the activity at any phase boundary where a Lyapunov exponent vanishes (Fig. 1D).

Crucially, our theory uncovered two computational principles in reservoir computing. First, the network can perform a different binary decision-making task in each of the bistable phases. Second, the network can achieve multitasking by simply varying the input baseline without any optimization of network weights (Fig. 1E). More generally, our theory shows that baseline modulations unlock a much richer dynamical phase portrait for RNNs than previously known. Baseline control represents a simple and efficient way for a reservoir network to flexibly toggle its dynamical regime to achieve flexible computations and multitasking. Our results thus suggest an important computational role for behavioral modulations of neural activity, whereby they might allow cortical circuits to flexibly adjust the cognitive task they perform to rapidly adapt to different contexts such as switching rapidly between multiple tasks.

Results

We model our local cortical circuit as a recurrent neuronal network (RNN) of N neurons where the synaptic couplings are randomly drawn from a Gaussian distribution of mean J_0/N and variance $1/N$ (Fig. 2A). We choose a positive definite neuronal transfer function $\phi(x) = 1/(1 + \exp(2g(x - \theta))) = (1 + \tanh(g(x - \theta)))/2$ with threshold θ_0 . Every neuron in our model receives a constant external synaptic input b_i drawn from a Gaussian distribution with mean μ and variance σ^2 . This baseline represents the afferent projections to the local cortical circuit originating from other areas. Following experimental

(8, 9, 25) and theoretical studies (2, 11), we modeled behavioral modulations as a change in the baseline statistics (mean μ and variance σ^2) of synaptic inputs b_i to the local circuit, induced by long-range projections carrying information about the behavioral state of the animal, or other contextual modulations (2, 11). Because the characteristic timescale of behavioral modulations is typically much slower than a circuit's stimulus responses, we approximate the effects of such modulations as the quenched inputs b_i . Importantly, these baseline modulations are constant, time-independent offsets of the input current to each neuron and represent *quenched* input noise.

Baseline Control of the Network Dynamical Phases. We found that by varying the values of baseline mean and variance μ, σ^2 , one can access a large library of network phases (Fig. 2B and C). The first two phases are generalizations of the fixed point and the chaotic phase which were previously reported in ref. 26. Strikingly, we found a number of phases including “bistable” phases where the network activity can reach two different dynamical branches for the same values of recurrent couplings and baseline input, depending on the initial conditions. In the network of Fig. 2B, the bistable phases are of two different kinds, with coexistence of either a fixed point and chaos (brown) or two fixed points (green). Whereas in the monostable phase, the network Landau potential has one global minimum, in the bistable phases, it exhibits two local minima, each one defining the basin of attraction of the initial conditions leading to each of the two bistable branches. Depending on the network parameters (J_0, g), we found networks with up to five different phases, including a bistable phase featuring the coexistence of strong and weak chaos (SI Appendix). Each monostable phase




Fig. 2. Baseline control of the network dynamical phase. (A) Random neural network where the baseline input current b_i to the i -th neuron is drawn from a normal distribution $\mathcal{N}(\mu, \sigma^2)$ (red). (B) *Left*: Network phase diagram, obtained by varying the mean μ and variance σ^2 of the baseline input, shows four phases: fixed point (blue); chaos (cyan); bistable phase with coexistence of two fixed points (green); bistable phase with coexistence of fixed point and chaos (brown). *Top Right*: Multicritical point. *Bottom Right*: Schematic of the Landau potential along a phase space trajectory (black arrow in *Inset*) from a stable phase with a single fixed point (blue circle), to a bistable phase with coexistence of fixed point and chaos (blue and cyan circles), to a stable phase with chaos (cyan circle). (C) Positive (*Left*) and negative (*Right*) largest Lyapunov exponents (in the bistable phases both LLE coexist); (D) Order parameters in each phase: Mean network activity (*Top*); autocorrelation (*Middle*); local stability (*Bottom*). Representative network activity in the different phases. *Insets*: Order parameters (Autocovariance C_0 , C_∞ and mean activity M). Network parameters: $J_0 = 0.5$, $g = 5$, $\theta_0 = 1$.

and each branch of a bistable phase can be captured in terms of the network order parameters LLE, M , and C (Fig. 2C), representing, respectively, the largest Lyapunov exponent LLE and the mean M and variance C of the activity obtained from the self-consistent dynamic mean field equations (*Materials and Methods*).

The variance of the activity includes a contribution σ^2 from the quenched baseline input and a recurrent contribution. A useful characterization of the network dynamical phase is obtained when considering the population-averaged autocorrelation function $c(t)$ at lag t , in particular, its zero lag value $c(0) = C$, the network variance, and its asymptotic value for large lag $c(\infty)$. The network is at a fixed point if $c(t)$ does not depend on time (i.e., $c(\infty) = c(0) = C$), while it is in a chaotic phase if $c(0) > c(\infty)$, in which case the LLE is positive. Finally, a value of $c(\infty) > 0$ signals a nonzero mean activity driven by the quenched variance in the baseline input.

Noise-Induced Enhancement of Chaos. Exploring the features of baseline modulations revealed a surprising phenomenon, whereby increasing the variance of the quenched input can enhance chaos. This phenomenon can be understood from a mean field perspective by considering how the baseline and the recurrent synaptic inputs interact with the single cell transfer function to determine the operating point of the network dynamics (Fig. 3; see ref. 27 for additional details). To illustrate this phenomenon, we first revisit the known case of noise-driven suppression of chaos realized in a circuit with quenched inputs and a zero-centered transfer function (Fig. 3A), which can be obtained when the mean baseline is set equal to the threshold $\mu = \theta_0$ (see ref. 28 for a case where they both vanish). On

general grounds, one expects the network phase to be chaotic whenever a large fraction of the synaptic input distribution is concentrated in the high gain region of the transfer function, defined as the region where the gradient of the transfer function ϕ is large. In this region, $\phi'(x)^2$ is of order one, leading to a large LLE (see ref. 29 and *Materials and Methods*, Eq. 6). The distribution of synaptic inputs has mean M , which is centered at the threshold, and some nonzero variance C , obtained self-consistently from Eqs. 3 and 4. For zero baseline variance, the network exhibits chaotic activity (case 1), as a large fraction of the synaptic inputs have access to the high gain region of the transfer function. When turning on a quenched baseline variance σ^2 , the synaptic input increases its variance by a value proportional to σ^2 . For larger values of the baseline variance σ^2 , the fraction of synaptic inputs in the high gain region progressively shrinks and for large enough variance chaos is suppressed (case 2).

In the case where $\mu < \theta_0$, the transfer function is not zero-centered, and noise-driven enhancement of chaos can occur (Fig. 3B). For low baseline variance σ^2 , the network is in the fixed point regime as a small fraction of synaptic inputs has access to the high-gain region (case a). Increasing the baseline variance σ^2 leads to a transition into a chaotic phase, as a progressively larger fraction of synaptic inputs has access to the high gain region. At some large enough variance, though, the fraction of synaptic inputs in the high-gain region starts decreasing again, and eventually, this leads to a transition to the fixed point phase. This chaos enhancement can be achieved either by passing through an intermediate bistable phase (black arrow in Fig. 2B) or by inducing a direct transition from a fixed point to a chaotic phase at lower values of the mean baseline μ (direct transition from blue to cyan at $\mu \sim 0.5$, Fig. 3B). This chaos enhancement has a number of striking




Fig. 3. Noise-driven modulations of chaos. (A) When the transfer function input is zero-centered (i.e., the baseline mean μ equals the threshold θ_0). The chaotic phase 1) occurs when a large fraction of the synaptic input distribution (pink curve) lies within the high gain region (yellow-shaded area, where $\phi'(x)^2 \sim \mathcal{O}(1)$) of the transfer function (blue curve). Increasing the input quenched variance 2) reduces this fraction, suppressing chaos. (B) When the transfer function is nonzero-centered ($\mu < \theta_0$), for low (a) and high (c) quenched input variance network activity is at a fixed point because the high-gain region receives a small synaptic input fraction; this fraction is maximized at intermediate quenched input variance (b), enhancing chaos. Network parameters: $g = 5$, $\theta_0 = 1$, $J_0 = 0$. Panel A: $\phi(x) = \tanh(x - \theta_0)$, $\mu = 1$; cases (1, 2) $\sigma = (0.1, 0.6)$. Panel B: $\phi(x) = 1/(1 + \exp(2g(x - \theta))) = (1 + \tanh(g(x - \theta)))/2$, $\mu = 0.5$; case (a, b, c): $\sigma = (0.2, 0.4, 0.9)$.

consequences, such as baseline control of optimal performance and neural hysteresis, which we will examine in the next sections. While previous studies showed that an increase in the *temporal* noise (e.g., white noise inputs) always leads to suppression of chaos (24, 28–30), we found that quenched noise unlocks a much richer set of phenomena.

Ergodicity Breaking in the Bistable Phases. The network activity in a bistable phase exhibits dynamical breaking of ergodicity. To illustrate this effect, we consider a network with fixed baseline mean μ at different values of σ (Fig. 4). At intermediate values of σ the network is in the bistable phase featuring a coexistence of a fixed point attractor and chaos, while at low and high values, the network is in the monostable fixed point phase and the chaotic phase, respectively. In the bistable phase, the network dynamics converge to either a fixed point attractor or to a chaotic attractor, depending on the initial conditions (Fig. 4A). These two branches are characterized by a negative (fixed point)

or a positive (chaos) LLE, respectively, and by branch-specific values for the network order parameters (C, M , Fig. 4B). We quantified ergodicity breaking in terms of the average distance $\langle d(T) \rangle$ between temporal trajectories (starting from different initial conditions, or between different replicas) over an epoch T (Fig. 4C). Monostable phases (fixed point or chaos) are ergodic and $\langle d(T) \rangle$ converges to C_∞ at large $T \rightarrow \infty$, since the network activity eventually explores all possible configurations (in the chaotic phase, the decay is typically slower than in a phase with a single attractor). The network breaks ergodicity when $\langle d(T) \rangle$ does not decay to C_∞ but rather it monotonically increases to reach nonzero late time values larger than C_∞ . In this case, depending on the initial conditions, there are two basins at a finite distance from each other. We found that the network is nonergodic in all the bistable phases, although each one of these phases retains specific values of the order parameters.

The library of bistable phases induced by changes in the baseline statistics includes all the phases in Fig. 4 and, remarkably, a previously unobserved phase exhibiting the coexistence of two chaotic phases Fig. 4D. This double chaos phase features a weak chaotic branch with small positive LLE and slow dynamics, and a strong chaotic branch with large positive LLE and fast dynamics. We found that this double chaos phase occurs for large g and it exhibits important computational properties that we investigate below.

Neural Hysteresis Retains Memory of Network Phase Trajectories. What are the effects of adiabatic changes in baseline statistics on the network dynamics? We sought to elucidate the effects of slow baseline changes, by driving the network with time-varying values of $\sigma(t)$ for fixed μ , describing a closed loop (Fig. 4E and F). We found that the network order parameters C, M changed discontinuously across phase boundaries, signaling a phase transition. When the baseline trajectory crosses the phase boundary from a stable phase (with a single LLE) to a bistable phase (with two branches, each characterized by its own LLE), the network activity in the bistable phase lies on either of the two branches, characterized by two separate basins of attractions (Fig. 4). The rules governing which of the two branches will be reached are determined by a hysteresis effect. We found that the network activity in the bistable phase retained a memory of the dynamical branch that it occupied before crossing the phase boundary. In the particular example of Fig. 4F, when crossing the boundary from the monostable fixed point to the bistable phase, the activity will persist on the fixed point branch of the bistable phase, whose negative LLE is continuously connected with the fixed point phase. For larger values of $\sigma(t)$, the network will eventually enter the monostable chaotic phase, where the LLE discontinuously jumps to a very large value. Vice versa, when inverting the time-varying trajectory in phase space by slowly decreasing the $\sigma(t)$ from the monostable chaotic phase into the bistable phase, the network will persist on the chaotic branch of the latter, whose positive LLE is continuously connected to the monostable chaotic phase. Eventually, for lower $\sigma(t)$, the network falls back into the fixed point phase where the LLE discontinuously jumps from large positive to negative values. Thus, when crossing phase boundaries adiabatically, the network will choose the branch of the bistable phase whose LLE is continuously connected to the previous phase.

Neural hysteresis occurs not just in the fixed point/chaos bistable phase, but also in the double fixed point and double chaos bistable phases. When crossing boundaries between two adjacent bistable phases, more complex hysteresis profiles can




Fig. 4. Ergodicity breaking in the bistable phase. (A) *Top*: Representative trials from initial conditions leading to the fixed point (*Right*) or chaotic attractor (*Left*) within the bistable chaos/fixed point phase. *Bottom*: 10 representative trajectories of network activity in the same bistable phase starting from different initial conditions (five leading to the chaotic attractor, black circles; five leading to the fixed point, black crosses; only three initial conditions per phase are shown; dashed circles represent the positions of the chaotic and fixed point attractors, respectively). The first two principal components of the set of all trajectories (PCs) are shown. The activity in both examples is captured by the mean and variance as shown in Fig. 2D). (B) For increasing values of σ , a cross-over from a monostable fixed point phase (*Left*) to a bistable phase fixed point/chaos (*Middle*) to a monostable chaotic phase (*Right*) is revealed by the order parameters (LLE: largest Lyapunov exponent; M: mean activity; C: mean autocorrelation; 1S: 1-replica stability). In the bistable phase, the fixed point and chaotic branches exhibit different order parameters. (C) Average distance between replica trajectories $\langle d \rangle$ reveals ergodicity breaking: in the monostable fixed point (blue) and chaotic (cyan) phases $\langle d \rangle$ asymptotes to C_∞ , but in the bistable phase (brown) it asymptotes to a value larger than C_∞ , representing the average distance between the basins of attraction of the two branches. (D) Example of a cross-over from a monostable fixed point phase (blue), to a bistable phase fixed point/chaos (brown) to a bistable weak/strong chaos phase (red), to a monostable chaotic phase (cyan), as revealed by the order parameters (same as panel B). Neural hysteresis. (E) Slow changes in baseline variance $\sigma(t)$ lead to discontinuous transitions in the network order parameters M, C (left: temporal profile of M, C, σ). (F) Crossing-over phase boundaries by a time-varying $\sigma(t)$ retains memory of the previous phase (blue: network simulations; red: exact DMFT calculation). Network parameters: panel A: $J_0 = 0.5, \theta_0 = 1, \mu = 0.54, g = 5, \sigma = 0.1$; panel B: $J_0 = 0.5, g = 6, \theta_0 = 1, \mu = 0.5$; panel C: same as Fig. (2B) and $\mu = 0.5, 0.6, 0.7$; panel D: $J_0 = 0.5, \theta = 1, \mu = 0.5, g = 18$; panels E and F: $J_0 = 0.5, g = 12, \theta_0 = 1, \mu = 0.5, \sigma(t) = \sigma_0 + \sigma_1 \sin(\pi t/T)$ with $T = 2048, \sigma_0 = 0.17, \sigma_1 = 0.025$.

occur. For example, when crossing into the double chaos phase (with fast/slow chaotic branches, Fig. 4D), from the fixed point branch of the fixed point/chaotic phase, the network dynamics will lie on the slow chaotic branch, whose positive but small LLE is continuously connected to the fixed point branch of the previous bistable phase. However, when crossing into the double chaos phase from the chaotic branch of the fixed point/chaotic bistable phase, the network dynamics will persist on the fast chaotic branch, whose large positive LLE is continuously connected to the chaotic branch of the fixed point/chaotic bistable phase. We then examined the relevance of neural hysteresis for controlling the network performance in a memory task.

Baseline Control of Multitasking. In any of the bistable phases, our reservoir network can perform binary decision-making tasks by equipping it with a linear readout (Fig. 5A). The two possible outcomes of the binary decision are represented by the two branches of a bistable phase and the linear readout is proportional to the mean activity (leveraging the fact that different branches of a bistable phase have different mean activity M), reporting the outcome of the binary decision in each trial. In each bistable phase, stimuli are drawn from two classes, associated with the two choices available to the reservoir (Fig. 5A), and are presented for a

short interval, nudging the network activity toward either branch of a bistable phase via the neural hysteresis mechanism explained in Fig. 4 E and F. For example, in the bistable chaos/fixed-point phase (brown region in Fig. 5A), one class of stimuli transiently nudges the network activity toward the chaotic phase (cyan region in Fig. 5A), such that after stimulus offset, the network settles into the chaotic branch of the chaos/fixed-point phase. The second class of stimuli transiently nudges the network activity toward the single fixed-point phase (blue region in Fig. 5A), such that after stimulus offset the network settles into the fixed point branch of the chaos/fixed-point phase. In the representative simulated session in Fig. 5B, the network is performing two trials of the Go/No-Go (G/NG) task and reports the correct choice in response to either stimuli after a delay period. In this neuroscience-inspired task, the network is interpreted as a model of the motor cortex, and the chaotic and the fixed-point branches are interpreted, respectively, as the animal performing a movement (Go) in response to one class of stimuli (e.g., a monkey releasing a bar), and withdrawing that movement (No-Go) in response to the other class of stimuli (31).

The next step is to model task-switching by leveraging the repertoire of multiple bistable phases (Fig. 5C). By changing the baseline input mean and variance, we can interpolate between




Fig. 5. Baseline control of multitasking. (A) The two bistable network phases (chaos/fixed-point and double fixed-point phases: brown and green, respectively; same as Fig. 2B) can be harnessed by a reservoir network to perform two different tasks: a delayed two-alternative forced-choice task (2AFC) in the double FP phase and a delayed go/no-go task (G/NG) in the chaos/FP phase. (B) *Top*: Experimental design for two representative trials of the G/NG task: Task rules are implemented by sustained values of task-specific baseline μ, σ . Stimuli are represented by transient changes in baseline mean μ during a short sample epoch (100 ms). Following a delay epoch (200 ms), the network decision outcome is extracted via a linear readout z (the z-scored mean activity). *Bottom*: Neural mechanism of decision-making along the hysteresis loop (circles and letters mark time points in the two representative trials at the top, projected onto the plane with LLEs as functions of the momentary input baseline). (C) Representative session with eight trials, alternating 2AFC and G/NG blocks (green and red lines represent task rule onset). In each block, stimuli from two classes are presented (blue/cyan and orange/brown color-shaded intervals represent the two classes for each task). After a delay, the decision outcome is read out (dot-dashed lines). *Top*: Representative activity of four neurons. *Bottom*: Network readout reports the stimulus class in either task from network activity in 2AFC and G/NG tasks: Positive or negative readout values represent s_R, s_L or s_G, s_{NG} stimulus classes, respectively. An additional linear readout reports task rule from network activity (positive and negative values for G/NG and 2AFC tasks, respectively; linear discriminant between task rules). (D) Readouts during the task. (E) Baseline values during the task (see panel A for comparison). (F) Network activity mean and variance (see Fig. 2D for comparison). Network parameters as in Fig. 2.

different bistable phases and therefore obtain a reservoir network that performs multiple tasks. We illustrate this ability by showing how our reservoir network can quickly switch between the delayed two-alternative forced choice task (2AFC) and the delayed G/NG, two classic paradigms commonly used in systems neuroscience (31, 32). Each task rule is represented by a sustained value of the baseline (μ, σ), which may change before trial onset, signaling a change in task starting in the upcoming trial. In the double fixed-point phase, the network performs a delayed 2AFC task, whereby each stimulus class is associated with one of the two fixed points attractors. In this neuroscience-inspired task, the network is interpreted as a model of the premotor cortex, and the two fixed points represent attractors which hold in working memory during the delay period the two choices available to the animal (e.g., licking the left or right water spout) in response to the two classes of stimuli (32).

In a representative session featuring task switching every two trials, the network decision-making performance was perfect (all 8 stimuli were correctly classified in Fig. 5C; in a longer session with 100 trials, 50 per task, yielded perfect performance in both tasks, respectively). The network time-varying baseline and order parameters reveal that the neural hysteresis mechanism underlies the binary decision-making tasks in each bistable phase (Fig. 5D–F). The network multitasking repertoire may vary depending on the set of bistable phases available for given values of the random coupling variance g , including the double chaos bistable phase in Fig. 4D. A striking feature of our framework is that the reservoir is performing the task without any weight optimization, contrary to the typical multitasking scenarios where RNNs are trained to perform multiple tasks via a costly weight optimization via gradient descent (33).

Baseline Control of Optimal Memory Capacity. A classic result in the theory of random neural networks is that, by fine-tuning the recurrent couplings at the “edge of chaos,” one can achieve optimal performance in a memory task, where the network activity maintains for a very long time a memory of stimuli presented sequentially (23). This was achieved by fine-tuning the network recurrent couplings to values close to the transition between fixed point and chaos, which is a metabolically costly and slow procedure typically requiring synaptic plasticity. Is it possible to achieve optimal memory capacity without changing

the recurrent couplings? We found that baseline control can achieve optimal memory capacity by simply adjusting the mean and variance of the baseline input distribution, without requiring any change in the recurrent couplings (Fig. 6).

We first derived an analytical formula for the memory capacity in the vicinity of a second-order phase transition boundary

$$\mathcal{M} \sim \frac{1}{1 - \langle \phi'^\alpha \phi'^\beta \rangle}, \quad [1]$$

where α, β are replica indices. Optimal memory capacity is achieved close to a phase boundary, and its features are qualitatively different depending on whether the phases separated by the boundary are monostable or bistable. At a boundary between two monostable phases, where the activity transitions between a fixed point and chaotic phase, optimal memory capacity is achieved at the edge of chaos. For fixed values of the recurrent couplings (Fig. 6A), one can easily achieve optimal memory capacity by adiabatically changing either the mean or the variance of the baseline. This external modulation thus sets the network at the edge of chaos, in the region where memory capacity is maximized, via baseline control, without any change in the recurrent couplings. Around a phase boundary involving a bistable phase, the optimal performance region can be reached by making use of the neural hysteresis phenomenon. We illustrate this intriguing scenario in the case of the transition from a bistable fixed point/chaos branch to a bistable double chaos branch (Fig. 6B). Optimal performance is achieved only on the branch of the bistable phase transition which undergoes a second-order phase transition (i.e., the branch whose LLE crosses zero). In this specific case, then, we can reach optimal performance on the lower branch of the LLE curve, describing the transition between the weak chaotic branch of the double chaos phase to the fixed point branch of the fixed point/chaos phase. Because of the neural hysteresis, achieving the optimal performance region requires first initializing the network on the lower LLE branch (on either side of the transition) and then adiabatically controlling the baseline to reach the desired point. The phase boundaries where only first-order phase transitions occur (i.e., no branch exhibits an LLE that crosses zero) do not lead to optimal memory capacity. For example, in Fig. 4B, neither the upper nor lower branch of the transition between a monostable fixed point phase to a bistable fixed point/chaos phase lead to large memory capacity, since no




Fig. 6. Baseline control of optimal memory capacity. (A) Two representative trajectories in baseline (μ, σ) space (Left: green and orange lines) allow to reach a phase transition where the LLE crosses zero (Top) and memory capacity is optimized (Bottom). (B) In a transition between bistable phases, memory capacity is optimized by a baseline trajectory whose branch exhibits an LLE that crosses zero at the phase boundary (orange curve); the branch with positive LLE (blue curve) does not maximize memory capacity. Network parameters: Panel A, same as Fig. 2A; panel B, same as Fig. 4D.

LLE on either branch of the intermediate bistable phase crosses zero. Nevertheless, it is always possible to reach a different second-order phase boundary from any point in (μ, σ) space by following an appropriate adiabatic trajectory in the baseline, where optimal memory capacity can be achieved (Fig. 6A). Therefore, one can achieve baseline control of optimal performance via neural hysteresis.

Discussion

We presented a brain-inspired framework for reservoir computing where we controlled the dynamical phase of a recurrent neural network by modulating the mean and quenched variance of its baseline inputs. Baseline modulations revealed a host of phenomena. First, we found that they can set the operating point of the network activity by controlling whether synaptic inputs overlap with the high gain region of the transfer function. A manifestation of this effect is a noise-induced enhancement of chaos. Second, baseline modulations unlocked access to a large repertoire of network phases. On top of the known fixed point and chaotic ones, we uncovered three bistable phases, where the network activity breaks ergodicity and exhibits the simultaneous coexistence of a fixed point and chaos, of two different fixed points, and weak and strong chaos. By driving the network with adiabatic changes in the baseline statistics one can toggle between the different phases, charting a trajectory in phase space. These trajectories exhibited a manifestation of the phenomenon of neural hysteresis, whereby adiabatic transitions across a phase boundary retain the memory of the adiabatic trajectory. Moreover, we showed that baseline control can achieve optimal performance in a memory task at a second-order phase boundary without any fine-tuning of the network recurrent couplings. In the bistable phases, we showed that the reservoir can perform different decision-making tasks, leveraging neural hysteresis and ergodicity breaking. Strikingly, we found that by simply varying the network baseline the reservoir can perform multiple tasks without any weight optimization. Our work provides a conceptual framework to achieve flexible performance and multitasking via the simple neural mechanism of baseline control, paving the way for an approach to reservoir computing.

Noise-Induced Enhancement of Chaos. Previous theoretical work found a noise-induced suppression of chaos in random neural networks driven by time-varying inputs both in discrete time (29) and continuous time (22, 24, 28, 30, 34). In previous cases, featuring a mean synaptic input centered in the middle of the high-gain region of the transfer function, suppression of chaos occurs because an increase in the variance drives the network away from the chaotic regime. In contrast, we found that, when the baseline statistics sets the mean synaptic input away from the center of the high gain region, one can induce a transition from fixed point to chaos at intermediate values of the variance (Fig. 3). Larger values of the variance eventually suppress chaos, such that a nonmonotonic dependence of the Lyapunov exponent on the baseline variance or mean can be realized. This is an example of noise-induced chaos in recurrent neural networks with additive interactions, although a similar phenomenon was recently found in networks with gated recurrent units (35) (for the logistic map see ref. 36). We believe that noise-induced modulation of chaos in discrete time networks is similar for both quenched and dynamical noise (24) since the LLE and the edge of chaos are the same for both cases. We speculate that introducing a leak term and generalizing our results to a continuous time system may induce a dynamical suppression of

chaos on general grounds, based on the memory effect. Another interesting direction is to drive the network with dynamical noise at different values of the baseline input and investigate its effect on the different monostable and bistable phases we uncovered via baseline modulation.

Optimal Sequential Memory. Previous studies showed that optimal performance in random networks can be achieved by either tuning the recurrent couplings at the edge of chaos (23) or by driving the network with noisy input tuned to a particular amplitude (24). Both those methods requires simple tuning of two hyperparameters [mean and variance of the random couplings (23) or noise (24) distribution], as in our model. It would be interesting to compare these alternative methods, test whether any of them is realized in cortical circuits and develop optimization algorithms to learn their parameters.

Comparison with Other Multitasking Frameworks. Humans learn to perform new cognitive tasks by directly following instructions, without any training at all (37). On the other hand, brain-inspired RNNs can be trained to perform multiple tasks by optimizing their recurrent weights via gradient descent (33, 38). This optimization procedure is costly, scaling as the square of the network size, and typically requires thousands or millions of training epochs to achieve good task performance; moreover, their maintenance is biologically implausible, as it requires a mechanism to fine-tune the value of the recurrent weights. Recent work showed that RNNs trained to perform a library of tasks via gradient descent can then quickly learn a new task by reutilizing learned computational motifs, such as learned fixed points or line attractors (38, 39). Here, we took a different approach to multitasking by interpreting the reservoir's own dynamical phases as a library of 'innate' computational motifs. Each of the multiple bistable phases already present with random recurrent couplings was shown to implement a different binary choice, relying on the combination of their ergodicity breaking and neural hysteresis property. Task rules were implemented as values of the baseline input mean and variance (Fig. 5). Unlike previous studies, our approach does not require any training of recurrent weights, thus avoiding the issues listed above. A limitation of our approach is that only a small number of bistable phases are available, and therefore, the expressivity of the reservoir is not large as the one achieved by trained RNNs (33). It is tantalizing to speculate that by combining our reservoir approach with some limited weight optimization one could learn a larger variety of computational motifs and lead to a more biologically plausible theory of multitasking RNNs.

Information Processing Capabilities and Bistability. Bistable phases with coexistence of fixed points and chaos were previously reported in recurrent networks with random couplings (40) and with gated recurrent units (35). We generalized this to a set of bistable phases featuring the coexistence of two fixed points and, remarkably, two chaotic attractors with slow and fast chaos, respectively. This is a report of a doubly chaotic phase in recurrent neural networks. Are there any information processing benefits of the double chaos phase? Neural activity unfolding within the weakly chaotic branch of this bistable phase has large sequential memory capacity, as the Fisher information diverges at the edge of chaos. On the other hand, the strongly chaotic branch erases memory fast. In this doubly chaotic phase, the network's information processing ability can be changed drastically by switching between the two branches, for example,

via an external pulse. It would be tantalizing to explore the computational capabilities of these bistable phases unlocked by baseline modulation. Here, we only considered homogeneous inputs where the baseline statistics is the same for all network neurons. Although, one may consider a more general setup with heterogeneous inputs, where different neural populations receive baseline modulations with different statistics. The simplest such possibility would be the ability to perform different tasks by gating in and out specific subpopulations, driving them with negative input. This is a promising direction for multitasking, and we leave it for future work.

Evidence for Baseline Modulations in Brain Circuits. In biologically plausible models of cortical circuits based on spiking networks, it was previously shown that increasing the baseline quenched variance leads to improved performance. This mechanism was shown to explain the improvement of sensory processing observed in the visual cortex during locomotion (11) and in the gustatory cortex with general expectation (2). In these studies, the effect of locomotion or expectation was modeled as a change in the constant baseline input to each neuron realizing an increase in the input quenched variance. This model was consistent with the physiological observation of the heterogeneous neuronal responses to changes in behavioral state, comprising a mix of enhanced and suppressed firing rate responses [during locomotion (3, 11, 25), movements (4–6), or expectation (14, 41)]. Intracellular recordings showed that these modulations are mediated by a change of baseline synaptic currents, likely originating from subcortical areas (8, 9). Because the effects of these changes in behavioral state on neural activity unfolded over a slower timescale (a few seconds) compared to the typical information processing speed in neural circuits (subsecond), we modeled them as constant baseline changes, captured by changes in the mean and variance of the distribution of input currents. Our results provide an interpretation of these phenomena, leading to the hypothesis that they could enable cortical circuits to adapt their operating regimes to changing demands.

Baseline Modulations and Gain Modulation. The effect of the baseline modulations on network dynamics highlighted in this study can be understood in terms of changes in the network effective transfer function $\Phi_{\text{eff}}(x) = \int Dz \phi(\sqrt{C}z + \mu + x)$, where Dz is a standard Gaussian measure, and C is the self-consistent variance of the activity, giving the self-consistent equation for the mean rate $M = \Phi_{\text{eff}}(M)$ (*Materials and Methods*). Baseline modulations lead to changes in the slope of the effective transfer function, a relationship previously derived in spiking networks (2). This is consistent with experimental observations that changes in behavioral states are mediated by gain modulation, as observed at the level of single cells (1) as well as populations (11). Alternative mechanisms for gain modulation include changes in the background synaptic currents controlling the single-cell conductances (42), which are not captured by our rate-based model.

Ergodicity Breaking. We found ergodicity breaking in network dynamics occurring in a series of bistable phases, which include phases with two fixed points, with a fixed point and chaos, and with weak/strong realizations of chaos. Ergodicity breaking was recently reported independently in a dynamically balanced neural network of inhibitory units in ref. 43. The origin of the ergodicity breaking in these two models is different. While in our case it

is driven by heterogeneity, or disorder, in the input baseline, in ref. 43, it is caused by an overrepresentation of symmetric connections, leading to non-Gaussian inputs for each neuron as a consequence. Moreover, while we relied on DMFT to prove the existence of bistability, Berlemont and Mongillo (43) applied the cavity method to reveal a large number of metastable states.

Neural Hysteresis. A prediction of our model is that baseline modulations may induce neural hysteresis when crossing a bistable phase boundary. Hysteresis is a universal phenomenon observed in many domains of physics. Hysteresis in neural networks was first observed in the presence of recurrent inhibition (44, 45) and later confirmed in visual areas in vitro (46). In the Wilson–Cowan model (47), hysteresis was observed in the transitions between fixed points. In our case, hysteresis occurs in the transition between different network phases including chaotic and fixed-point regimes. Our results suggest a potential way to examine the existence of hysteresis in brain circuits, within the assumption that increasing baseline variance represents increasing values of a continuous behavioral modulation such as arousal [e.g., measured by pupil size (48)]. A potential signature of hysteresis could be detected whether the autocorrelation time of neural activity at a specific arousal level exhibited a strong dependence on whether arousal levels decreased from very high levels or increased from very low levels. We leave this interesting direction for future work.

Materials and Methods

Random Neural Network Model. Our discrete time neural network model with top-down control, illustrated in Fig. 2, is governed by the dynamical equation

$$x_{i,t+1} = \sum_{j=1}^N J_{ij} \phi(x_{j,t}) + b_i + \eta_t. \quad [2]$$

Here, b_i is quenched Gaussian noise with mean μ , and variance σ^2 , η_t is a possible time-dependent external stimulus (relevant for the sequential memory task below). The synaptic couplings J_{ij} are drawn from a normal distribution with mean J_0/N , and variance is g^2/N ; the scaling $1/N$ guarantees the existence of the large N limit. We will assume $\mu > 0$ in accordance with the fact that long-range projections are typically mediated by pyramidal cells. The activation function $\phi(x) = \frac{1}{2} [\tanh(x - \theta_0) + 1]$ is positive definite and biologically plausible as it incorporates both a soft rectification and thresholding. Indeed, the activation function ϕ satisfies $\phi(x) \approx 0$ when $x \ll \theta_0$ and $\phi(x) \approx 1$ when $x \gg \theta_0$.

For this model, the measure of the path integral is

$$\mathcal{D}x = \prod_{i=1}^N \mathcal{D}x_i, \quad \mathcal{D}x_i = \sum_{t \in \mathbb{Z}} dx_{i,t}.$$

We apply dynamical mean field theory (DMFT) as described in ref. 27. The aim of DMFT is to obtain the single body density functional $P_1(x)$ or equivalently its moment generating functional, averaged over the randomness of the synaptic connections and the external noise in the infinite population limit $N \rightarrow \infty$. That is,

$$P_1(x_{1,t}) \equiv \int \langle P_N(x) \rangle_{\zeta, J} \prod_{i=2}^N \mathcal{D}x_i,$$

or its characteristic function,

$$Z_1(l_{1,t}) = \int e^{i \sum_t l_{1,t} x_{1,t}} P_1(x_{1,t}) \mathcal{D}x_1,$$

where $P_N(x) \mathcal{D}x$ is the N -body density functional given by

$$P_N(x) = \prod_{i=1}^N \prod_t \delta(x_{i,t+1} - l_{i,t} - \eta_t - b_i),$$

where $l_{i,t} = \sum_{j=1}^N J_{ij} \phi(x_{j,t})$. Using the expression of the Dirac δ function as $\delta(x) = (2\pi)^{-1} \int e^{i\tilde{x}x} d\tilde{x}$, and the saddle point method (18, 24, 27), we derive the single-body density function, whose detail is shown in *SI Appendix*.

Order Parameters. The order parameters of the model are the population mean and variance at equilibrium of the single neuron activity $\langle x_{i,t} \rangle$. A rigorous derivation of self-consistent equations for these two quantities requires dynamical mean field theory (*SI Appendix*); a heuristic argument for them can be sketched as follows. Averaging Eq. 2 in the absence of external input yields

$$\langle x_{i,t+1} \rangle = \sum_{j=1}^N \langle J_{ij} \phi(x_{j,t}) \rangle.$$

Neglecting correlation between the random variables J_{ij} and $x_{j,t}$ on the right-hand side, and using the statistical invariance under permutation of neuron labels to drop cell indices, we obtain $\langle x_{t+1} \rangle = J \langle \phi(x_t) \rangle$. Focusing now on the stationary regime, where the distribution of x_{t+1} and x_t are identical, and assuming them to be Gaussian with mean M and variance C , leads to

$$M = J \int \frac{dx}{\sqrt{2\pi}} e^{-x^2/2} \phi(\sqrt{C}x + M). \quad [3]$$

Taking the second moment of Eq. 2, without neglecting the variance of the quenched disorder, term, and deploying once again the same assumptions yields

$$C = \sigma^2 + g^2 \int \frac{dx}{\sqrt{2\pi}} e^{-x^2/2} \phi(\sqrt{C}x + M)^2. \quad [4]$$

Stability of the systems with single or double replicas is checked by computing the linear response or by checking that the hessian matrix is positive definite (18, 49). In *SI Appendix*, the dynamical mean-field theory approach is rigorously developed to derive two dynamical equations for the mean-field momenta. The stationary limit of those equations is found to correspond to Eqs. 3 and 4, thus confirming the heuristic result.

Largest Lyapunov Exponent. The Lyapunov exponent of a dynamical system is a quantity that characterizes the rate of separation of infinitesimally close trajectories. Quantitatively, two trajectories in phase space with an initial separation vector diverge (provided that the divergence can be treated within the linearized approximation) at an exponential rate given, and the Lyapunov exponent governs this exponential growth. The LLE for a discrete-time dynamical system is defined as

$$\lambda_{\max} = \lim_{\tau \rightarrow \infty} \lim_{\|x_t^1 - x_t^2\| \rightarrow 0} \frac{1}{2\tau} \ln \frac{\langle |x_{t+\tau}^1 - x_{t+\tau}^2|^2 \rangle}{\langle |x_t^1 - x_t^2|^2 \rangle}, \quad [5]$$

which indicates how the two orbits, or replicas, get to be far from each other. In the N body picture, when $N \rightarrow \infty$, we find (50) (see *SI Appendix* for a derivation):

$$\lambda_{\text{LLE}} = \frac{1}{2} \ln \langle \phi'(x)^2 \rangle = \frac{1}{2} \ln \int \phi'(\sqrt{C}x + M)^2 dx. \quad [6]$$

Here, C and M are the self-consistent solutions to the dynamical mean-field equation [3] and [4]. In the monostable phases, a single LLE exists since a single solution to these equations can be found. In the bistable phases, two different solutions for C and M exist, depending on the initial conditions for the mean-field equations, corresponding to the two basins of attraction of the two branches. The two solutions in turn yield two different LLE via Eq. 6.

Distance between Replicas. Let us define the mean activity in the replica α (corresponding to some initial conditions $x_i^\alpha(0)$) as

$$\bar{x}_i^\alpha(T) = \frac{1}{T} \int_0^T x_i^\alpha(t) dt.$$

We then define the distance between replicas as (43)

$$d_{\alpha\beta}^2(T) = \frac{1}{N} \sum_{i=1}^N [\bar{x}_i^\alpha(T) - \bar{x}_i^\beta(T)]^2,$$

and its average $\langle d \rangle = \frac{1}{n^2} \sum_{a,b=1}^n d_{ab}(T)$, as used in the visualization of Fig. 4C.

Multitasking Readouts. Network readouts z in each task were chosen as z-scored mean network. The task readout was chosen as a projection on the linear discriminant direction maximizing separability of the two tasks from the activity immediately preceding stimulus presentation (min-maxed as well).

Memory Capacity. Following refs. 51 and 52, we define the memory capacity of a dynamical system for an observer in possession of an unbiased estimator for the mean, who can therefore remove the mean values from all the time series he records. Moreover, we would like the resulting memory capacity to be zero when the linear readout is dominated by a constant baseline value because nothing can be learned from a readout independent of the input. Adopting therefore the mean-removed formula, we find for the memory capacity $\mathcal{M}$ in the neighborhood of the second-order phase transition boundary

$$\mathcal{M} \sim \frac{1}{1 - \langle \phi'^\alpha \phi'^\beta \rangle}. \quad [7]$$

To derive this formula, we proceed along the same lines as in Ref. 24, considering the input signal u_t as $u_t = \frac{1}{N} \sum_i \xi_{i,t}$ and trying to reconstruct the input $u(t_0)$ with the sparse linear readout $\sum_{j=1}^K w_j x_{j,t}$ with $O(K) < O(\sqrt{N})$. The memory curve C_τ and capacity C_M are given respectively by the determinant coefficient which measures how well the readout neurons reconstruct the past input $u(t - \tau)$ correctly, and their sum (52)

$$C_\tau = \frac{\sum_{i,j=1}^K \text{Cov}_t(u_t, x_{i,t+\tau}) \text{Cov}_t(x_{i,t}, x_{j,t})^{-1} \text{Cov}_t(u_t, x_{j,t+\tau})}{\text{Var}_t(u_t)}, \quad [8]$$

$$C_M = \sum_\tau C_\tau,$$

where

$$\text{Cov}_t(u_t, v_{t+\tau}) = \lim_{T \rightarrow \infty} \left[\frac{1}{T} \sum_{t=1}^T u_t v_{t+\tau} - \left(\frac{1}{T} \sum_{t=1}^T u_t \right) \left(\frac{1}{T} \sum_{s=1}^T v_{s+\tau} \right) \right],$$

and $\text{Var}_t(u_t)$ is computed in the same manner. The readout is sparse, so that the covariance $\text{Cov}_t(x_i(t), x_j(t))$ becomes diagonal in the infinite population limit $N \rightarrow \infty$ (23). Moreover, we deal with the steady state so that this term is constant with respect to time. The detail of the derivation is exhibited in Supplement.

Data, Materials, and Software Availability. Jupyter notebooks reproducing the main figures can be found at <https://github.com/mazzulab/multitasking> (53). All study data are included in the article and/or *SI Appendix*.

ACKNOWLEDGMENTS. We would like to thank Enrico Rinaldi for advice on the numerics and Taro Toyozumi and Łukasz Kuśmierczak for discussions. S.O. and F.F. were partially supported by RIKEN Center for Brain Science. L.M. was supported by National Institute of Neurological Disorders and Stroke grant R01-NS118461 and by National Institute on Drug Abuse grant R01-DA055439 (CRCNS). An open license has been selected upon submission.

Author affiliations: ^aLaboratory for Neural Computation and Adaptation, RIKEN Center for Brain Science, Wako, Saitama 351-0198, Japan; ^bDepartment of Biology, Institute of Neuroscience, University of Oregon, Eugene, OR 97403; and ^cDepartment of Mathematics, Institute of Neuroscience, University of Oregon, Eugene, OR 97403

1. M. J. McGinley *et al.*, Waking state: Rapid variations modulate neural and behavioral responses. *Neuron* **87**, 1143–1161 (2015).
2. L. Mazzucato, G. La Camera, A. Fontanini, Expectation-induced modulation of metastable activity underlies faster coding of sensory stimuli. *Nat. Neurosci.* **22**, 787–796 (2019).
3. C. M. Niell, M. P. Stryker, Modulation of visual responses by behavioral state in mouse visual cortex. *Neuron* **65**, 472–479 (2010).
4. C. Stringer *et al.*, Spontaneous behaviors drive multidimensional, brain-wide population activity. *bioRxiv* [Preprint] (2018). <https://doi.org/10.1101/306019> (Accessed 4 January 2023).
5. S. Musall, M. T. Kaufman, A. L. Juvinet, S. Gluf, A. K. Churchland, Single-trial neural dynamics are dominated by richly varied movements. *bioRxiv* [Preprint] (2019). <https://doi.org/10.1101/308288> (Accessed 4 January 2023).
6. D. B. Salkoff, E. Zagha, E. McCarthy, D. A. McCormick, Movement and performance explain widespread cortical activity in a visual detection task. *Cerebral Cortex* **30**, 421–437 (2020).
7. A. Fontanini, D. B. Katz, Behavioral states, network states, and sensory response variability. *J. Neurophysiol.* **100**, 1160–8 (2008).
8. P. O. Polack, J. Friedman, P. Golshani, Cellular mechanisms of brain state-dependent gain modulation in visual cortex. *Nat. Neurosci.* **16**, 1331 (2013).
9. D. B. Nestvogel, D. A. McCormick, Visual thalamocortical mechanisms of waking state-dependent activity and alpha oscillations. *Neuron* **110**, 120–138 (2022).
10. M. C. Dardlat, M. P. Stryker, Locomotion enhances neural encoding of visual stimuli in mouse v1. *J. Neurosci.* **37**, 3764–3775 (2017).
11. D. Wyrick, L. Mazzucato, State-dependent regulation of cortical processing speed via gain modulation. *J. Neurosci.* **41**, 3988–4005 (2021).
12. P. A. Lin, S. K. Asinof, N. J. Edwards, J. S. Isaacson, Arousal regulates frequency tuning in primary auditory cortex. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 25304–25310 (2019).
13. C. R. Heller, Z. P. Schwartz, D. Saderi, S. V. David, Selective effects of arousal on population coding of natural sounds in auditory cortex. *bioRxiv* [Preprint] (2020). <https://doi.org/10.1101/2020.08.31.276584> (Accessed 4 January 2023).
14. C. L. Samuelsen, M. P. Gardner, A. Fontanini, Effects of cue-triggered expectation on cortical processing of taste. *Neuron* **74**, 410–422 (2012).
15. D. M. Schneider, A. Nelson, R. Mooney, A synaptic and circuit basis for corollary discharge in the auditory cortex. *Nature* **513**, 189–194 (2014).
16. J. Bigelow, R. J. Morrill, J. Dekloe, A. R. Hasenstaub, Movement and VIP interneuron activation differentially modulate encoding in mouse auditory cortex. *eNeuro* **6**, ENEURO.0164-19.2019 (2019).
17. I. Yavorska, M. Wehr, Effects of locomotion in auditory cortex are not mediated by the VIP network. *Front. Neural Circuits* **15**, 24 (2021).
18. A. Crisanti, H. Sompolinsky, Path integral approach to random neural networks. *Phys. Rev. E* **98**, 062120 (2018).
19. M. Stern, H. Sompolinsky, L. Abbott, Dynamics of random neural networks with bistable units. *Phys. Rev. E* **90**, 062710 (2014).
20. C. Van Vreeswijk, H. Sompolinsky, Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* **274**, 1724–1726 (1996).
21. A. Sederberg, I. Nemenman, Randomly connected networks generate emergent selectivity and predict decoding properties of large populations of neurons. *PLoS Comput. Biol.* **16**, e1007875 (2020).
22. M. Stern, N. Istrate, L. Mazzucato, A reservoir of timescales in random neural networks. *arXiv* [Preprint] (2021). <http://arxiv.org/abs/2110.09165> (Accessed 4 January 2023).
23. T. Toyozumi, L. Abbott, Beyond the edge of chaos: Amplification and temporal integration by recurrent networks in the chaotic regime. *Phys. Rev. E* **84**, 051908 (2011).
24. J. Schuecker, S. Goedeke, M. Helias, Optimal sequence memory in driven random networks. *Phys. Rev. X* **8**, 041029 (2018).
25. M. Dipoppa *et al.*, Vision and locomotion shape the interactions between neuron types in mouse visual cortex. *Neuron* **98**, 602–615 (2018).
26. H. Sompolinsky, A. Crisanti, H. J. Sommers, Chaos in random neural networks. *Phys. Rev. Lett.* **61**, 259 (1988).
27. M. Helias, D. Dahmen, *Statistical Field Theory for Neural Networks* (Springer, 2020).
28. K. Rajan, L. Abbott, H. Sompolinsky, Stimulus-dependent suppression of chaos in recurrent neural networks. *Phys. Rev. E* **82**, 011903 (2010).
29. L. Molgedey, J. Schuchhardt, H. G. Schuster, Suppressing chaos in neural networks by noise. *Phys. Rev. Lett.* **69**, 3717 (1992).
30. L. F. Abbott, K. Rajan, H. Sompolinsky, Interactions between Intrinsic and Stimulus-Evoked Activity in Recurrent Neural Networks in M. Ding, D. Ginzburg, Eds. (Oxford University Press, 2011).
31. J. R. Hollerman, L. Tremblay, W. Schultz, Influence of reward expectation on behavior-related neuronal activity in primate striatum. *J. Neurophysiol.* **80**, 947–963 (1998).
32. H. K. Inagaki, L. Fontolan, S. Romani, K. Svoboda, Discrete attractor dynamics underlies persistent activity in the frontal cortex. *Nature* **566**, 212 (2019).
33. G. R. Yang, M. R. Joglekar, H. F. Song, W. T. Newsome, X. J. Wang, Task representations in neural networks trained to perform many cognitive tasks. *Nat. Neurosci.* **22**, 297–306 (2019).
34. R. Engelken, A. Ingrosso, R. Khajeh, S. Goedeke, L. Abbott, Input correlations impede suppression of chaos and learning in balanced rate networks. *arXiv* [Preprint] (2022). <http://arxiv.org/abs/2201.09916> (Accessed 4 January 2023).
35. K. Krishnamurthy, T. Can, D. J. Schwab, Theory of gating in recurrent neural networks. *Phys. Rev. X* **12**, 011011 (2022).
36. G. Mayer-Kress, H. Haken, The influence of noise on the logistic model. *J. Stat. Phys.* **26**, 149 (1981).
37. M. W. Cole, J. A. Etzel, J. M. Zacks, W. Schneider, T. S. Braver, Rapid transfer of abstract rules to novel contexts in human lateral prefrontal cortex. *Front. Hum. Neurosci.* **5**, 142 (2011).
38. L. Driscoll, K. Shenoy, D. Sussillo, Flexible multitask computation in recurrent networks utilizes shared dynamical motifs. *bioRxiv* [Preprint] (2022). <https://doi.org/10.1101/2022.08.15.503870> (Accessed 4 January 2023).
39. C. D. Márton, L. Gagnon, G. Lajoie, K. Rajan, Efficient and robust multi-task learning in the brain with modular latent primitives. *arXiv* [Preprint] (2021). <http://arxiv.org/abs/2105.14108> (Accessed 4 January 2023).
40. Ł. Kuśmierczak, S. Ogawa, T. Toyozumi, Edge of chaos and avalanches in neural networks with heavy-tailed synaptic weight distribution. *Phys. Rev. Lett.* **125**, 028101 (2020).
41. R. Vincis, A. Fontanini, Associative learning changes cross-modal representations in the gustatory cortex. *eLife* **5**, e16420 (2016).
42. F. S. Chance, L. F. Abbott, A. D. Reyes, Gain modulation from background synaptic input. *Neuron* **35**, 773–782 (2002).
43. K. Berlemont, G. Mongillo, Glassy phase in dynamically-balanced neural networks. *bioRxiv* [Preprint] (2022). <https://doi.org/10.1101/2022.03.14.484348> (Accessed 4 January 2023).
44. B. G. Cragg, H. N. V. Temperley, Memory: The analogy with ferromagnetic hysteresis. *Brain* **78**, 304–316 (1955).
45. D. R. Smith, C. H. Davidson, Maintained activity in neural nets. *J. ACM (JACM)* **9**, 268–279 (1962).
46. D. Fender, B. Julesz, Extension of panum's fusional area in binocularly stabilized vision. *JOSA* **57**, 819–830 (1967).
47. H. R. Wilson, J. D. Cowan, Excitatory and inhibitory interactions in localized populations of model neurons. *Biophys. J.* **12**, 1–24 (1972).
48. M. J. McGinley, S. V. David, D. A. McCormick, Cortical membrane potential signature of optimal states for sensory signal detection. *Neuron* **87**, 179–192 (2015).
49. J. R. de Almeida, D. J. Thouless, Stability of the Sherrington–Kirkpatrick solution of a spin glass model. *J. Phys. A: Math. General* **11**, 983 (1978).
50. B. Cessac, Increase in complexity in random neural networks. *J. Phys. I*, 409–432 (1995).
51. J. Dambre, D. Verstraeten, B. Schrauwen, S. Massar, Information processing capacity of dynamical systems. *Sci. Rep.* **2**, 1–7 (2012).
52. H. Jaeger, "Short term memory in echo state networks" in *GMD-Report 152, GMD-German National Research Institute for Computer Science* (Citeseer, 2002). <http://www.faculty.jacobs-university.de/hjaeger/pubs/STMEchoStatesTechRep.pdf> Accessed 4 January 2023.
53. S. Ogawa, F. Fumarola, L. Mazzucato, Multitasking RNN. Github. <https://github.com/mazzucato/multitasking>. Deposited 17 July 2023.