



Associative synaptic plasticity creates dynamic persistent activity

Albert J. Wakhloo^{a,b,1} , David G. Clark^{a,c,d}, and L. F. Abbott^{a,c}

Edited by Moses V. Chao, New York University Langone Medical Center, New York, NY; received September 23, 2025; accepted August 24, 2026

In biological neural circuits, the dynamics of neurons and synapses are tightly coupled. We study the consequences of this coupling and show that it enables a form of working memory. In recurrent neural network models with ongoing Hebbian plasticity, we find that following oscillatory stimulation, neurons continue to oscillate long after the input is removed. This creates a dynamic form of memory that has no explicit storage or retrieval phases and that requires no prior knowledge of the input. We trace the mechanism of these “persistent oscillations” to an interaction between neurons and synapses that creates complex outlier eigenvalues of the connectivity matrix. This is shown both in simulation and analytically. We leverage this mechanistic understanding to generate persistent oscillations with prespecified dynamics, creating a dynamic analog of a classical Hopfield network. Our work demonstrates that coupling neuronal and synaptic dynamics enables novel forms of computation.

neural networks | synaptic plasticity | working memory | random matrix theory | nonlinear dynamics

Synaptic connections shape neuronal dynamics, and neuronal dynamics modify synapses through activity-dependent plasticity mechanisms. The dynamics of neurons and synapses are therefore tightly coupled. Nevertheless, classical neural network models treat these processes separately, with models of neuronal dynamics typically assuming fixed synapses (1–5) and models of learning and memory typically separating plasticity from the dynamics of the network (6–9). Most efforts to study the computational functions of joint neuronal-synaptic dynamics have limited their focus to purely presynaptic short-term plasticity (10–14). However, experimental work has described forms of bidirectional, associative plasticity over timescales not that different from those of purely neural systems (15–19). Thus, it is interesting to ask what unique network functions are enabled by ongoing associative plasticity (20–26).

In this paper, we study how recurrent neural networks with ongoing associative plasticity respond to external stimulation. In this model, introduced by Clark and Abbott (25), the synapses of a recurrent neural network evolve about random baseline strengths according to a Hebbian rule (Fig. 1A), and neurons and synapses fluctuate together on comparable timescales. Importantly, plasticity influences neuronal dynamics not by creating entirely new dynamical motifs that erase the prior weights, but by interacting with and sculpting a baseline connectivity. This requires the size of the synaptic fluctuations to be small—i.e., individual synapses remain close to their baseline strengths (Fig. 1B and C).

We study the response of such networks to a prototypical class of inputs, oscillations (27). Our main finding is that the interaction of recurrent neural dynamics with Hebbian synaptic dynamics creates oscillatory activity that outlasts the external stimulation for a time much longer than the intrinsic timescales of the system. This is unexpected, because Hebbian plasticity typically leads to fixed points, as in Hopfield models (6). We show that Hebbian associative plasticity rules can generate these more complex dynamical patterns as a result of the interaction between the static backbone connectivity in our model and ongoing plasticity.

1. Results

1.1. Recurrent-Network Model with Ongoing Hebbian Plasticity and External Drive. We study the recurrent-network model of Clark and Abbott (25) in which synapses fluctuate around fixed, random baselines. The dynamic variables comprise N neurons, with preactivations denoted by $x_i(t)$, together with N^2 plastic synaptic weights $A_{ij}(t)$, where i, j range from 1 to N . The neuronal variables evolve according to

$$\frac{dx_i(t)}{dt} = -x_i(t) + \sum_{j=1}^N [J_{ij} + A_{ij}(t)] \phi(x_j(t)) + b_i(t), \quad [1]$$

Significance

Neurons and synapses interact through activity-dependent plasticity mechanisms. Here, we study a recurrent neural network in which ongoing plastic weight updates make small adjustments to static background connections. We show that the interaction between plasticity and static connections creates a self-reinforcing form of memory such that transient stimulation generates long-lasting persistent activity. This persistent activity requires no prior learning of input statistics. Instead, ongoing synaptic plasticity anchors stimuli to preexisting motifs in the synaptic connections. Our work therefore describes a form of working memory in which synaptic plasticity interacts with, rather than overrides, preexisting structure.

Author affiliations: ^aCenter for Theoretical Neuroscience, Zuckerman Institute, Columbia University, New York, NY 10027; ^bCenter for Computational Neuroscience, Flatiron Institute, New York, NY 10010; ^cKavli Institute for Brain Science, Columbia University, New York, NY 10027; and ^dKempner Institute for the Study of Natural and Artificial Intelligence, Harvard University, Boston, MA 02134

Author contributions: A.J.W., D.G.C., and L.F.A. designed research; performed research; analyzed data; and wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Copyright © 2026 the Author(s). Published by PNAS. This article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](#).

¹To whom correspondence may be addressed. Email: ajw2232@cumc.columbia.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2526156123/-DCSupplemental>.

Published September 29, 2026.




Fig. 1. (A) Schematic of our modeling framework. (B) Neuronal dynamics of two example units, x_i and x_j , from a network of $N = 500$ units during the stimulation period ($t < 0$) and following input cessation ($t > 0$). Here and in the remainder of the paper, we measure time relative to the moment the input is turned off (gray dashed line) and in units of the neuronal time constant of the network. (C) Dynamics of the corresponding synaptic weights between these two units, $J_{ij} + A_{ij}(t)$ and $J_{ij} + A_{ij}(t)$. Fluctuations in the synaptic weights ($J_{ij} + A_{ij}(t)$) (Inset, gray line) are small relative to the baseline weights (J_{ij}) (Inset, dashed line).

where the variables J_{ij} are fixed random weights drawn from a Gaussian with mean zero and SD $g/\sqrt{N}$ (although our results are expected to apply for non-Gaussian weight distributions as well). These weights form the static backbone of the connectivity, while the dynamic variables $A_{ij}(t)$ are fluctuations due to ongoing plasticity. These are denoted in matrix form by $\mathbf{J}$ and $\mathbf{A}(t)$. The function $\phi(\cdot)$ is a saturating nonlinearity mapping preactivations to shifted and normalized firing rates, which we set to a hyperbolic tangent, $\phi(\cdot) = \tanh(\cdot)$. In addition to the recurrent input from other units, each unit receives an external input $h_i(t)$, defined below.

The time-varying component of the connectivity evolves according to a “pre-times-post” rate-based Hebbian rule,

$$p \frac{dA_{ij}(t)}{dt} = -A_{ij}(t) + \frac{k}{N} \phi(x_i(t)) \phi(x_j(t)). \quad [2]$$

Here, p sets the timescale of the synaptic dynamics relative to the neuronal dynamics, and k determines the strength and sign of the plasticity. We consider $k \geq 0$, corresponding to a Hebbian (rather than anti-Hebbian) rule, and $p > 1$, so that synapses evolve more slowly than neurons. Importantly, while the static background connections, J_{ij} , scale as $1/\sqrt{N}$, the plasticity-induced changes to these connections, $A_{ij}(t)$, scale as $1/N$. As we focus on the large- N regime, the plastic component is vanishingly small relative to the static component. Such minuscule changes in the synapses can, nevertheless, have dramatic effects on the network dynamics due to network activity aligning to low-rank structure present in $\mathbf{A}(t)$ (25).

In contrast to Clark and Abbott (25), we introduce a nonzero external input $h_i(t)$ given by

$$h_i(t) = \begin{cases} I \cos(2\pi f t + \theta_i) & t < 0, \\ 0 & t \geq 0, \end{cases} \quad [3]$$

with amplitude I , frequency f , and neuron-specific phase θ_i . We choose the phases randomly for each unit following

Rajan et al. (27), though our results are qualitatively unchanged when all phases are equal ($\theta_i = \theta$).

1.2. Behavior of the Externally Driven, Plastic Network. We first study the behavior of this model when the input is active, $t < 0$. This extends the phenomenology studied by Rajan et al. (27), which considered static synapses, to the model of Clark and Abbott (25), which includes plastic synapses. Neuronal traces are visualized in Fig. 2 A, *i–iv* ($t < 0$). When both the input amplitude, I , and plasticity strength, k , are small, neurons show chaotic fluctuations with superimposed oscillations (Fig. 2 A, *i*). As the plasticity strength k increases, the chaotic activity progressively slows (Fig. 2 A, *ii*). This occurs because plastic synapses create a positive feedback loop that slows the neuronal dynamics and drags the network toward fixed points (25). In this state, external input produces small oscillations atop the nearly static activity. When the input amplitude I increases sufficiently, the network transitions from this “small-oscillation” regime to a state that contains both a stimulus-driven component and, due to the recurrent connectivity, an intrinsically generated chaotic component (Fig. 2 A, *iii*). As I continues to increase, the neurons eventually become perfectly entrained to the input, eliminating all traces of chaos (Fig. 2 A, *iv* and Materials and Methods).

1.3. Persistent Oscillations Following Input Cessation. The four regimes shown in (Fig. 2 A) exhibit distinct behaviors following input cessation at $t = 0$ (Fig. 2 A, *i–iv*, $t \geq 0$). Most notably, the network can continue to oscillate after the input is removed (Fig. 2 A, *iii*). The lifetime of these oscillations can exceed the synaptic timescale, p , by an order of magnitude (Figs. 2 B and 3 A and Materials and Methods). The fact that these oscillations persist for much longer than the intrinsic timescales of the system indicates that they reflect a self-renewing process. We demonstrate this in SI Appendix, Fig. S1 by setting $k = 0$ at $t = 0$ so that synapses continue to influence neurons, and continue to decay, but neurons no longer influence synapses. In this case, persistent oscillations are curtailed. These results show that persistent oscillations involve oscillatory neural activity “feeding back into itself” through ongoing synaptic weight changes.

We now describe the dependence of persistent oscillations on system parameters. Here, and for the rest of the paper, we focus on the regime in which synapses are slower than neurons ($p > 1$) and where, without plasticity, the activity would be chaotic ($g > 1$) (1, 25). As shown in Fig. 3 B, the lifetime of persistent oscillations grows supralinearly with p . For these oscillations to occur, the input amplitude, I , must be large enough to avoid the small-oscillation regime, but small enough to avoid eliminating the influence of the recurrent connections on the dynamics (Figs. 2 A, *ii* and 3 A). That is, the activity must express a mixture of input- and recurrent-driven features for persistent oscillations to occur.

The fact that persistent oscillations are eliminated for large I shows that this phenomenon is not the result of a simple delayed feedback loop between neurons and synapses. Indeed, one would expect the opposite trend if this were the mechanism driving persistent oscillations. We also note that the parameter regimes in which persistent oscillations do or do not occur are distinct from those defining the phases of chaos or oscillatory entrainment described in ref. 27 (Fig. 2 B).*

*That persistent oscillations can occur even when the activity is fully oscillatory is likely because full entrainment does not imply that the recurrence plays no role. In particular, near this phase transition, the recurrence still plays a large role in determining the amplitude and phases of the units, so the neural activity expresses features driven by both the inputs and the recurrence, leading to persistent oscillations.




Fig. 2. Overview of model behavior. (A) Example traces at four different points of the parameter space. In each subplot, we stimulate a network for 200 time units, after which point all inputs are halted. The duration of the stimulation is not essential, provided it is sufficiently long (SI Appendix Fig. S2). Each line denotes a single unit from the same network. (B) Lifetime of persistent oscillations as a function of k and l for networks with $p = 25$, $g = 1.3$, and $f = 0.1$. The black line marks the transition from chaos to perfectly oscillatory activity during stimulation (Materials and Methods). Each point in the (k, l) plane corresponds to an average over 50 networks of size $N = 500$.

We next consider the relationship between the frequency of the input and the frequency of persistent oscillations (Materials and Methods). As shown in Fig. 3C, there is a preferred frequency band within which oscillations persist. Inside this band, the frequency of the oscillations grows with the frequency of the input. As the input frequency approaches the upper edge of the band, the likelihood of generating persistent oscillations rapidly falls. These findings imply that when oscillations persist, they reflect the structure of the previously present inputs, with faster inputs leading to faster persistent oscillations.

We expect these trends in the lifetime and frequency of persistent oscillations to hold in the $N \rightarrow \infty$ limit. While it is possible to derive a dynamical mean field theory that describes this limit exactly (25), we found that the resulting equations were prohibitively difficult to solve numerically. Instead, we show in Fig. 3 that all trends discussed above are identical in networks with $N = 1,000$ and $N = 5,000$.

1.4. Complex Outlier Eigenvalues Underlie Persistent Oscillations. How does the connectivity matrix generate oscillations in the absence of inputs? We begin by considering the eigenvalues of the matrix $\mathbf{J} + \mathbf{A}(t)$ at the time the input is turned off, $t = 0$. We expect these eigenvalues to be predictive of the nonlinear dynamics of the system, with large, complex eigenvalues leading to oscillations.

This is indeed the case. In the parameter regime where oscillations persist, the connectivity matrix develops a pair of complex-conjugate outlier eigenvalues (Fig. 4A, i, outliers). The distribution of the remaining eigenvalues is a uniform disk in the complex plane of radius g , which is the eigenvalue distribution of an i.i.d. random matrix—i.e., the eigenvalue distribution for $\mathbf{J}$ alone, without plastic modifications (Fig. 4A and B, i, disk). In the small-oscillation regime in which oscillations do not persist, the connectivity develops a single large real outlier eigenvalue, reflecting the neurons being continuously dragged toward fixed points by the synapses (Fig. 4B, i).

Of course, the connectivity is actually dynamic, and so too are the outlier eigenvalues. The dynamics of the outlier eigenvalues reflect the behavior of the network after the input is turned off. As oscillations persist (Fig. 4A, ii), the imaginary component of the eigenvalues gradually falls to zero while the real part grows

more positive (Fig. 4A, iii). During this time, the oscillations progressively slow, eventually stopping altogether around the time the outlier becomes real. In SI Appendix, Fig. S5, we show that the instantaneous network frequency closely tracks the imaginary component of the eigenvalue outlier through the entire period following input cessation, providing further evidence that outliers drive persistent oscillations.

The locations of the outlier eigenvalues largely determine the lifetime and frequency of the persistent oscillations. To show this, we sweep over a large number of model parameters and plot the lifetime and frequency of the oscillations as a function of the real and imaginary parts of the outlier at the time the input is turned off (Fig. 4C and D). In general, oscillations have a longer lifetime when the outlier has a moderately large real part and have both a higher frequency and a longer lifetime when the outlier has a large imaginary part. That is, the imaginary component affects both the lifetime and the frequency of the oscillations, while the real part controls only the lifetime.

We now use this mechanism to generate persistent oscillations that never decay. To do this, we freeze all synaptic dynamics at $t = 0$, halting both Hebbian updates and synaptic decay (effectively sending $p \rightarrow \infty$ as in ref. 25). This fixes the synaptic weights and eigenvalue spectrum at their poststimulation values. With plasticity frozen, oscillations induced by the outlier eigenvalues persist indefinitely (Fig. 4E).

The emergence of complex-conjugate outlier eigenvalues is surprising given that the Hebbian rule produces symmetric updates to the weights ($\mathbf{A}(t) = \mathbf{A}^T(t)$; Eq. 2), and symmetric matrices have purely real eigenvalues. Moreover, for any low-rank, symmetric matrix $\mathbf{H} = \mathbf{H}^T$ that does not depend on $\mathbf{J}$, the combined matrix $\mathbf{J} + \mathbf{H}$ can possess only real outlier eigenvalues in the large- N limit (SI Appendix, section 2.1; note that, in our model, the plastic matrix $\mathbf{A}(t)$ is approximately low-rank (25) due to p being finite) (28). These outliers, if they exist, are simply the (real) eigenvalues of $\mathbf{H}$ (SI Appendix, section 2.1).

Thus, for complex outliers to emerge, the plastic matrix cannot be uncorrelated with the static background matrix $\mathbf{J}$. Because $\mathbf{A}(t)$ is shaped by activity-dependent plasticity, complex outliers and persistent oscillations can only occur when the activity expresses features driven by both the external inputs and the recurrent background connectivity, $\mathbf{J}$. We visualize this in Fig. 4A, iv,




Fig. 3. Scaling of lifetime and frequency with model parameters. (A and B) Lifetime of persistent oscillations as a function of input amplitude I and synaptic time constant p for networks with $N = 1,000$ (dashed) and $N = 5,000$ (solid). Each point represents an average over 400 simulations for the small networks and over 30 for the large networks. (C) Frequency of the initial phase of the persistent oscillation as a function of the frequency of the inputs. We define this “ringing frequency” by the first peak in the power spectrum of the neuronal dynamics (*Materials and Methods*). When no peak is detected, we set the frequency to 0. (D–F) In each subplot, we plot the activity of four example units from three different networks with increasing values of I and with all other parameters held fixed (first 100 time units of stimulation not shown). When I is too low (panel D; $I = 0.15$), the network is in the small oscillation regime. When it is moderately large (E; $I = 0.75$), the network expresses a mixture of recurrent- and input-driven features, and oscillations persist for a long time. Finally, as I continues to increase (F; $I = 1.2$), the activity becomes increasingly dominated by the inputs, and the lifetime of the oscillations gradually decreases.

which demonstrates that in the persistent oscillation regime, $A(t)$ is aligned with the subspaces of eigenvectors of J corresponding to complex eigenvalues (*Materials and Methods*).

In the following section, we demonstrate this idea by first showing that a matrix built from eigenvectors of J is sufficient for producing complex outliers and thus oscillations, and then designing inputs that generate such a matrix via Hebbian plasticity during the stimulation period.

1.5. Generating Targeted Persistent Oscillations. We now address how the interaction between the static backbone, ongoing plasticity, and oscillatory inputs can generate complex outlier eigenvalues in the connectivity matrix. We first define a “target” matrix $\hat{A}$ in terms of eigenvectors of J that, when added to J , generates complex outlier eigenvalues and thus oscillations. We then find a set of inputs that allow Hebbian plasticity to generate such a matrix during the stimulation period, that is, stimulation leading to $A(0) \approx \hat{A}$. This approach is similar to that of Hopfield networks in which targeted inputs cause a network to express a

pattern of activity embedded in the synapses. Here, however, the memory-related activity is dynamic, rather than static, and memory storage is part of the dynamics of the model.

The target matrix is rank-two, real, and symmetric, given by

$$\hat{A} = 2\alpha \operatorname{Re}\{\psi\psi^\dagger\}, \quad [4]$$

where ψ is a complex eigenvector of J , $\dagger$ denotes the complex conjugate transpose, and α is a real scalar. We take ψ to be a right eigenvector of J , that is, $J\psi = \eta\psi$, with eigenvalue η . In *SI Appendix, section 2.2*, we show that for any choice of J (i.e., not necessarily random), the combined matrix $J + \hat{A}$ develops an eigenvalue pair located at

$$\lambda = \operatorname{Re}(\eta) \pm i\sqrt{\operatorname{Im}(\eta)^2 - \alpha^2|\psi^\top\psi|^2} + \alpha. \quad [5]$$

In the large- N limit with J i.i.d. Gaussian, $|\psi^\top\psi|^2 \ll 1$, so that the above equation simplifies to

$$\lambda = \eta + \alpha, \quad [6]$$




Fig. 4. Eigenvalue outliers drive persistent oscillations. (A and B) Behavior of outlier eigenvalues when the network is in the persistent oscillation regime (A), vs. when it is in the small oscillation/slow dynamics regime (B). (i) Eigenvalues of $J + A(t)$ at $t = 0$, when the input is halted. The bulk of the eigenvalue distribution fills a disk of radius g (gray line), which is what one would expect from a random matrix with variance g^2/N , while the outliers are produced by the plastic matrix. (ii) Dynamics of 5 example units, $x_i(t)$, from the corresponding network (first 100 time units not shown). (iii) Dynamic trajectory of the real and imaginary parts of the largest outlier eigenvalue of the network. (iv) Overlaps of the primary eigenvectors of the plastic $A(t)$ matrix with each eigenvector of J at the time the input is halted. Here we color each eigenvalue of J by the magnitude of the corresponding eigenvector's alignment with $A(0)$ (Materials and Methods). In the persistent oscillation regime, the plastic weights align with eigenvectors of J with complex eigenvalues. We return to this point in Section 1.6. (C and D) The lifetime and frequency of persistent oscillations as a function of the eigenvalue outlier, with the expected disk-shaped eigenvalue bulk shaded in gray. We simulated a set of networks over a range of parameters and plotted the resulting eigenvalue outliers, coloring by either lifetime (C) or persistent oscillation frequency (D) and binning the final values (Materials and Methods). (E) Freezable persistent oscillations. Here, we halt the ongoing plasticity at $t = 0$, creating persistent oscillations that never decay. This is a dynamic variant of freezable chaos—see ref. 25.

together with its conjugate. In other words, the matrix $\hat{A}$ translates the eigenvalue η along the real axis; this translation is rightward for $\alpha > 0$.

Furthermore, we provide a stimulation protocol to create a plastic matrix similar to $\hat{A}$ above. If the neuronal activations were perfectly sinusoidal with $\phi(x(t)) = \text{Re}\{\psi e^{i2\pi ft}\}$ and $p \gg 1/f$, Eq. 2 shows that the plastic matrix $A(t)$ would converge to the $\hat{A}$ of Eq. 4 with $\alpha = k/4N$. Of course, the actual neuronal variables are nonlinear, and it is in general not possible to force them to be exact sinusoids. Nevertheless, we can stimulate the network with such perfectly sinusoidal inputs given by

$$b(t) = I\sqrt{N} \text{Re}\{\psi e^{i2\pi ft}\}. \quad [7]$$

The factor of $\sqrt{N}$ ensures that these inputs are of order one for the conventional normalization $\sum_i |\psi_i|^2 = 1$. These sinusoidal

inputs have phases that are aligned with the chosen eigenvectors of the static connectivity matrix, J .

As shown in Fig. 5A, this eigenvector-aligned input selectively amplifies specific complex-conjugate eigenvalues, translating them rightward along the real axis and creating a complex outlier pair that drives persistent oscillations. Fig. 5B and C illustrate the relationship between the lifetime and frequency of the persistent oscillations and the target eigenvalue η . The lifetime is controlled by both the real and imaginary parts of η , while the frequency is controlled by its imaginary part alone. Using this eigenvalue-targeting method, we can generate persistent oscillatory activity with controlled lifetime and frequency characteristics.

The relationship between system parameters and persistent oscillations is different in this eigenvector-aligned case than in the random-phase case explored previously. Most notably, with eigenvector-aligned inputs, we can use arbitrarily large input amplitudes and still elicit persistent oscillations. This key




Fig. 5. Generating structured persistent oscillations with targeted inputs. (A) Trajectory of a targeted eigenvalue. Here, we target an eigenvalue located at $\lambda = 1 + \frac{4}{5}i$ and track its trajectory during the stimulation period (blue) and the stimulus off period (red). As in previous figures, t denotes the time from the moment the stimulation is halted. (B and C) Lifetime and ringing frequency as a function of the targeted eigenvalue. The real part affects the lifetime, while the imaginary part affects both the lifetime and the frequency of the evoked oscillation.

difference arises because the stimulus is now explicitly correlated with the random recurrent connectivity. Consequently, even when network activity is input-dominated, the plastic matrix $A(t)$ still develops correlations with the random backbone, J .

We can show analytically that, with eigenvector-aligned inputs and taking p , I , and $1/f$ large with $p \gg 1/f$, the resulting eigenvalue outliers following a sufficiently long stimulation period are approximated by

$$\lambda \approx \eta + \frac{4k}{\pi^2}, \quad [8]$$

and its conjugate $\bar{\lambda}$ (SI Appendix, section 3.1), if $|\lambda| > g$ (otherwise no outliers exist). The coefficient $4/\pi^2$ results from the effects of the nonlinearity $\phi(\cdot)$. This expression indicates that the eigenvalue outlier—and thus the lifetime and frequency of the targeted oscillation—depends only on the plasticity strength k and the targeted eigenvalue η . Eigenvector-aligned inputs therefore enable persistent oscillations with arbitrarily large input amplitudes I , while also making these oscillations insensitive to the exact input frequency f . Instead, the oscillations depend on the targeted connectivity mode encoded in the pattern of input weights across neurons.

These results demonstrate that targeted inputs decouple stimulus-evoked dynamics from the persistent dynamics that follow. What matters is the “spatial” pattern of the inputs across neurons (i.e., their alignment with a particular eigenvector), not their temporal structure. We show an extreme example of this in SI Appendix, Fig. S4, where we evoke persistent oscillations with a single biphasic pulse rather than sustained oscillatory drive. This behavior differs from the random-phase case, in which input amplitude and frequency affect the persistent oscillations that follow (Fig. 3).

In the next section, we analyze the random-phase case and show that it too involves eigenvector alignment, unifying the two

cases. We note, however, that eigenvector alignment may not be the only route to complex outliers; other low-rank perturbations of J could produce them as well, a possibility we leave to future work.

1.6. Analytical Solution for Random Phases: Outlier Location and Eigenvector Alignment. How do complex outlier eigenvalues emerge in response to random-phase oscillatory input (Fig. 4), and is this mechanism related to the eigenvector alignment for targeted inputs described above? To answer these questions, we obtain an analytical approximation for the plastic matrix $A(0)$ and its resulting outlier eigenvalues as a function of the model parameters. This approximation reveals that the random-phase case also involves alignment between $A(0)$ and the eigenvectors of J , unifying the two cases.

To motivate our analysis, note that the key features of the dynamics during stimulation are that $x_i(t)$ is 1) driven by the input $h_i(t)$, and 2) also influenced by J (Fig. 2). We reasoned that these features could be captured by a linearized model in which we make the approximation $\phi(x_i(t)) \approx \gamma x_i(t)$, with a gain parameter γ that is set heuristically (SI Appendix, section 3.2).

As described in SI Appendix, section 3.2, this approach yields the following approximation for the dynamics of the activations:

$$\phi(x(t)) \approx \sqrt{N} I \text{Re}\{\Psi e^{i2\pi f t}\}, \quad [9]$$

$$\text{where } \Psi = N^{-1/2}(\nu I_N - J)^{-1} e^{i\theta}. \quad [10]$$

Here, $e^{i\theta}$ is a vector with elements $e^{i\theta_i}$ with θ_i being the phase of the input to neuron i , I_N is the identity matrix, and ν is a complex scalar that depends on system parameters (see Materials and Methods for the full expression).

Expanding the above equation in powers of J , we find that the parameter ν measures the degree to which the neuronal dynamics are influenced by the recurrent connectivity. For large values of $|\nu|$, the variables $\phi(x_i(t))$ are dominated by the inputs, while for $|\nu|$ close to g —i.e., for ν close to an eigenvalue of J —the recurrence plays a large role in shaping the dynamics.

Under this approximation, the plastic matrix converges to a static matrix during the stimulation period. Following stimulation for a time much longer than the synaptic timescale p , and given synapses slower than the inputs $p \gg 1/f$, the plastic matrix is

$$A(0) \approx 2\alpha \text{Re}\{\Psi\Psi^\dagger\}, \quad [11]$$

where $\alpha = kI^2/4$. Note the similarity of this expression and Eq. 4, the difference being that, in this case, Ψ is not an exact eigenvector of J . Nevertheless, this expression incorporates information about both the recurrent connectivity, J , and the input structure both through Ψ and because the value of ν depends on the input’s amplitude and frequency (Materials and Methods).

We calculate the outlier eigenvalues of the matrix $J + F(J)$ for a class of $F(J)$ that includes the approximate $A(0)$ given by Eq. 11. Using a diagrammatic method, we show that if $|\nu| > g$, then the outlier eigenvalues are (SI Appendix, section 2.3)

$$\lambda = \frac{g^2 \nu}{|\nu|^2} + \frac{\alpha}{|\nu|^2 - g^2}, \quad [12]$$

and its complex conjugate, unless $|\lambda| \leq g$, in which case no outliers exist. Note that while the steps leading to the formula for $A(0)$ in Eq. 11 are approximate, given such a matrix, the above equation for the outlier eigenvalue pair is exact in the large- N

limit. Eq. 12 is similar to Eq. 6 with v playing the role of η ; both have an α -dependent term contributing to the real part of the eigenvalue.

In the *SI Appendix*, section 3.2, we derive an additional condition describing the transition between the small oscillation and persistent oscillation regimes (*Materials and Methods*, Eq. 19). This condition implies $|v| > g$ whenever oscillations persist, justifying our application of the above formula. We confirmed that Eq. 12 approximates the outliers observed in the simulations of Section 1.4 (*SI Appendix*, Fig. S8A) and that the transition boundary approximates what we observed in Fig. 2 (*SI Appendix*, Fig. S9B). We emphasize that although this approximation captures the qualitative features of this model, it does not yield quantitatively exact results, particularly near the boundary between the small and persistent oscillation regimes (Fig. 2).

We now trace the phenomenology of persistent oscillations for random-phase input back to our formula for the eigenvalue outliers. First, we note that, when $I \gg g$, k —i.e., when the inputs dominate the recurrence—we find that $|\lambda| < g$ (*Materials and Methods*), so that no outliers exist. Conversely, when I is excessively small, we find that our calculation correctly predicts that the network will be in the small oscillation regime (Eq. 19). This accounts for the observed amplitude dependence of persistent oscillations: the amplitude must be weak enough to allow for dynamics to reflect contributions of the recurrence, but strong enough to avoid the small oscillations regime (Figs. 2 A, ii and 3).

Turning to the input frequency, we find that outlier eigenvalues only form within a preferred input frequency band. Within this band, the imaginary component of this eigenvalue outlier generally increases with the input frequency, which explains why faster inputs generally lead to faster persistent oscillations (Fig. 3 C and *SI Appendix*, Fig. S9A). These results emphasize that in this setting, the network learns to generate persistent oscillations that reflect the structure of the inputs, in contrast to the targeted inputs of the previous section.

Although the vector Ψ is not an eigenvector of J for $|v| > g$, it overlaps with a small number of J eigenvectors (*SI Appendix*, section 3.3). To show this, we calculate the normalized overlap:

$$Q = \left\langle \frac{|\Psi^\dagger J \Psi|}{\|\Psi\| \|J \Psi\|} \right\rangle_{J, \theta}, \quad [13]$$

which is 1 when Ψ is an eigenvector of J and is 0 for a random Ψ which is uncorrelated with J . We show that $Q \rightarrow 1^-$ in the limit $|v| \rightarrow g^+$, so that the vector Ψ becomes an eigenvector of J (*SI Appendix*, section 3.3). For intermediate values of Q , Ψ is partially aligned with an eigenvector of J ; such intermediate values of Q are found in the parameter regime in which oscillations persist (*SI Appendix*, Fig. S9C). Numerically, we validate this by measuring the alignment between the leading eigenvectors of A and each eigenvector of J (*Materials and Methods*). As shown in Fig. 4 A–B, iv, we find that this projection localizes around a few complex eigenvalues in the persistent oscillation regime. Thus, the plastic matrix and neuronal activity aligns itself with eigenvectors of the static background, J , in the persistent oscillation regime.

This is similar in spirit to what we found for targeted inputs with the key difference that the eigenvalue-eigenvector pair to which the network aligns itself is determined by the frequency and amplitude of the input. In a sense, the stimulus causes the network to express a pattern of activity, already embedded in the background connectivity, possessing dynamical properties that

match those of the input. That is, the network anchors the input to preexisting modes of the connectivity.

2. Discussion

We have shown that an ongoing Hebbian plasticity mechanism can learn to generate persistent oscillatory patterns of activity in response to both random and structured input. Classical models of working memory invoke Hebbian plasticity to explain persistent *static* activity that bridges temporal delays in working memory tasks (11–13, 29–31). In these models, it is assumed that neurons with similar tuning profiles have excitatory recurrent connections between them, and this is justified by an appeal to Hebbian plasticity. Subsequent stimulus presentation leads to reverberating excitation between similarly tuned neurons, stabilizing the evoked activity patterns (30, 32). Our work differs from these models in two important ways. First, the patterns of persistent activity in our model can be learned “on the fly,” in response to arbitrary input, rather than being burned into a prespecified connectivity scheme. Second, our model produces oscillatory rather than static persistent activity patterns. Many experimental studies report complex dynamical activity patterns following stimulus cessation during working memory tasks that do not align neatly with the picture of sustained firing (33–37). Our results suggest that such complex persistent patterns could be generated by Hebbian plasticity processes that underlie canonical working memory models.

The vast majority of theoretical models of learning involve strong plasticity introducing entirely new dynamical patterns into a neural network (6–9). In our model, plasticity affects ongoing neuronal dynamics by manipulating a dynamical reservoir provided by the static backbone connectivity. Previous work suggests that this could be the operating regime of many biological networks. For example, studies using brain–computer interfaces have found that neural circuits in the motor cortex can only learn patterns of activity that lie within a constrained dynamical repertoire (38, 39), and studies in the hippocampus and sensory cortex have found that correlations in the activity patterns between neurons remain remarkably stable, even after periods of learning or across different task conditions (40–44). These results suggest that plasticity and learning occur by adjusting preexisting connectivity structures, rather than creating entirely new dynamical motifs from scratch. Our results emphasize that even microscopic changes in synaptic weights can have dramatic effects on neuronal dynamics and that interactions between preexisting structure and ongoing plasticity can lead to novel forms of neural computation.

In our model, plasticity exerts a significant impact on network dynamics even though the size of the weight updates is small. Indeed, the elements of the plastic matrix A_{ij} scale as $1/N$, while the individual baseline connections J_{ij} scale as $1/\sqrt{N}$. These small plastic updates can still influence neuronal dynamics when activity aligns with the low-rank structure of the plastic matrix (25). For example, low-rank matrices with elements scaling as $1/N$ can produce order-one eigenvalue outliers when added to J (*SI Appendix*, section 2). In general, across different learning rules, there should be a relationship between the rank and the scale of plasticity-induced connectivity, an interesting topic for future study.

How might the persistent oscillations we describe relate to studies of oscillatory brain activity? Previous studies have reported persistent oscillations in neural activity following exposure to a periodic stimulus (45–49), and other studies have found that

single neurons can phase lock to oscillations in the local field potential during working memory tasks (50–52). These findings have sometimes been interpreted as evidence that the circuits being studied have an intrinsic tendency toward oscillation (45, 48). Our results suggest that the situation may be more complex. Our networks have no a-priori oscillatory structure embedded in them, and in the random-phase case, no relationship to the stimulus. Yet, ongoing plasticity leads to persistent oscillations within preferred input frequency bands (Fig. 3C). A clear signature of such plasticity-induced oscillations is a decaying frequency over time (SI Appendix, Fig. S5). This pattern cannot be explained by simple underdamped oscillator models. Observing such a slowing down in biological neural networks would suggest that plasticity induces persistent oscillations.

In the spirit of classical models of associative memory, we used targeted inputs to generate structured persistent oscillations (6). More precisely, we aligned the phases of the oscillating inputs to a specific eigenvector of the background connectivity matrix, $\mathbf{J}$. This process is similar to classical Hopfield models of associative memory in which a set of random patterns are embedded into a connectivity matrix and then evoked using pattern-aligned stimuli. In this way, the static backbone $\mathbf{J}$ encodes a “dynamical repertoire” of patterns that can be evoked through stimulation. We present evidence in SI Appendix, Fig. S6 that this form of memory is extensive in the number of units.

The key difference between our model and Hopfield networks is that the synapses of our model are not held fixed. It is this feature that allows us to generate oscillatory, rather than static, memory patterns. Moreover, unlike Hebbian models for sequence storage (8, 9), our networks use ongoing plasticity to generate memory-related activity patterns, rather than discrete learning epochs. That is, we do not require learning episodes in which a pattern is written into the network prior to stimulation. Ongoing plasticity allows the network to generate persistent oscillations in response to inputs that the network has never seen before by anchoring the inputs to preexisting modes of the static connectivity (Section 1.6).

The Hebbian mechanism for persistent oscillations we studied may provide various functional advantages. Previous experimental work suggests that neuromodulators and other factors may gate or modify the strength of ongoing plasticity processes (53). A previous study (25) examined static memory in the same network model we analyzed here. In a particular parameter regime in which the full network dynamics are chaotic, if synapses are frozen by halting plasticity, the network remains at a fixed point very close to the state it was in at the time of the freezing. In other words, any static state of the network can be held or remembered indefinitely by halting synaptic dynamics. Here we extend this form of memory across a richer dynamic range. Driving a similar network evokes a range of oscillations that persist indefinitely in the absence of input if synaptic plasticity is halted, freezing the synapses (Fig. 4E). This is essentially a form of input-driven reservoir computing in which a particular mode of the rich dynamics inherent in random network connectivity can be extracted by an input and then remembered indefinitely through the gating of plasticity. Such a computation could be useful for the maintenance of a dynamically extended memory or the generation of a rhythmic motor command. As another example of modulating plasticity, in SI Appendix, Fig. S4, we altered k to induce persistent oscillations in networks that otherwise lie far outside the parameter regime in which this would otherwise occur.

Drawing from previous models of working memory (12), it would be interesting to design classes of $\mathbf{J}$ matrices from which

ongoing plasticity can pull out more complex dynamical motifs through targeted stimulation. It may be that by choosing $\mathbf{J}$ appropriately, rather than randomly, one can design a network in which plasticity “pulls out” the appropriate patterns when cued with suitable transient stimulation. We leave this possibility to future work.

Overall, our work demonstrates that understanding memory-related neural activity may require modeling synaptic and neuronal dynamics together. Separating synaptic from neuronal dynamics, while convenient, can obscure basic functions of a circuit.

3. Materials and Methods

3.1. Simulation of Plastic Networks. We simulated the full neuronal-synaptic dynamics in Eqs. 1 and 2 using a fifth order Runge-Kutta (RK) solver with an adaptive step size. Specifically, we used the GPU-accelerated implementation of Tsitouras’ 5/4 method in Diffax (54). We set the absolute and relative tolerance of the solver to 10^{-6} . In SI Appendix, Fig. S3, we show that over a small subset of the parameters examined in the above figures, the lifetime and frequencies one obtains are left unchanged after increasing this precision by two orders of magnitude.

Unless stated otherwise, networks were run for 700 time units with the input lasting for the first 200 time units and with the parameters $p = 25$, $g = 1.3$, $f = 0.1$. In all analyses considered in the main text, the inputs h_i were present from the start point of the simulation. As shown in SI Appendix, Fig. S2, this was not an essential choice. So long as an initial strong stimulation is given to the network to break it out of the small oscillation regime (Fig. 2A, ii), the exact choice of initialization appeared to have little to no effect. Moreover, we show in SI Appendix, Fig. S2 the lifetime of the persistent oscillations is unaffected by the amount of time the network is stimulated provided that the stimulation time is sufficiently long, a factor which scales with p .

3.2. Chaos to Entrainment Transition. To calculate the entrainment to chaos transition line depicted in Fig. 2B, we used a bisection approach. For each value of k , we calculated the critical I above which more than half of a batch of 9 networks were fully entrained. We simulated networks of size $N = 500$ for 5,000 time units, with hyperparameters $g = 1.3$, $f = 0.1$, and $p = 25$. Networks were considered fully entrained when the average autocovariance function of the preactivations, $\Delta(t - t') = \langle x(t)x(t') \rangle$, had a sinusoidal shape with no signs of decay between peaks. Numerically, we required that each peak of the autocovariance was within a distance of $\epsilon = 10^{-3}$ from all others. After obtaining critical I values for each k , we smoothed the resulting transition line using a cubic spline.

3.3. Lifetime and Frequency of Persistent Oscillations. We used the power spectrum of the network activity to measure the lifetime of persistent oscillations. Specifically, we calculated the average power spectrum of the preactivations over windows of width of 40 time units. A network was defined as oscillating at a certain time if the power spectrum of the preactivations within that window had a significant peak, defined as a point in the power spectrum that was 10^{-1} greater than adjacent points. We defined the lifetime of persistent oscillation as the time it took for the power spectrum to stop having any significant peaks. We show an example of this in SI Appendix, Fig. S3.

We used a similar approach to measure the frequency of the persistent oscillations. Specifically, we calculated the power spectrum of the preactivations over a slightly smaller window of size 37.5 time units. The frequency of the ringing was then defined as the slowest frequency that had a significant peak in the power spectrum, where a significant peak was defined as above. If no such peak was detected, we simply set the ringing frequency to 0 in Figs. 3–5.

3.4. Lifetime and Frequency as a Function of Outliers. We obtained the results shown in Fig. 4C and D by considering a small three-dimensional grid of I , k , and f values. For each combination of parameter values, we simulated a batch of 50 networks, each of size $N = 500$, and calculated the eigenvalues

of $\mathbf{J} + \mathbf{A}$ at the time the input was halted, as well as the lifetime and frequency of the resulting persistent oscillation. To obtain the results in Fig. 4, we first formed a two-dimensional grid with the minimal and maximal x and y set equal to the minimum and maximum of the real and imaginary parts of the observed outlier eigenvalues, respectively. We then used Nadaraya–Watson smoothing with a Gaussian kernel to interpolate the lifetime and frequency of the persistent oscillations as a function of the outlier eigenvalue. Finally, because the outlier eigenvalue distribution was highly nonconvex, we masked out points on the grid which were further than a distance of 0.1 from any observed eigenvalue. The results shown in Fig. 4 were generated by binning the resulting expected lifetime and frequencies.

3.5. Eigenvalue Outlier and Transition Formulae. In this section, we present the full formulae for the analytical results described in Section 1.6. See *SI Appendix* for the full derivations. Recall that when an outlier eigenvalue pair forms in response to random-phase inputs, its value is given by

$$\lambda = \frac{g^2 \nu}{|\nu|^2} + \frac{\alpha}{|\nu|^2 - g^2}, \quad [14]$$

together with its conjugate. In this expression, the values of α and ν are given by

$$\alpha = \frac{kl^2}{4}, \quad \nu = \frac{1 + i\omega}{\gamma} - \frac{k\phi(l)}{4}. \quad [15]$$

The value of γ is

$$\gamma = \frac{\beta}{\xi} \left(\sqrt{1 + \frac{(1 + \omega^2)\xi}{\beta^2}} - 1 \right), \quad [16]$$

$$\xi = g^2 + \frac{1}{\phi(l)} l^2 - \beta^2, \quad [17]$$

$$\beta = \frac{k\phi(l)}{4}, \quad [18]$$

and $\omega = 2\pi f$. We plot the behavior of ν as a function of the input parameters, f and l in *SI Appendix, Fig. S9A*. Here, we make two observations based on the form of these equations. First, it is clear that the imaginary part of the outlier is determined by the term $g^2 \nu / |\nu|^2$. From the above expression, we can see that ν grows with l . This implies that the imaginary component of the eigenvalue will decay as inputs grow, explaining why we do not find persistent oscillations in response to strong inputs. Second, we note that the imaginary component of ν is given by $i2\pi f / \gamma$, and in many parameter regimes, $\text{Im}(\nu)$ grows with f (*SI Appendix, Fig. S9A*). Thus, faster inputs generally lead to outliers with larger imaginary part. This explains why the frequency of persistent oscillations increases with the frequency of the inputs, within some admissible band (Fig. 3C).

In addition to this formula for the eigenvalue outlier λ , we approximate the transition from the small oscillation regime to the persistent oscillation regime. As described in more detail in *SI Appendix*, the idea behind this calculation is to note that the constant component of the activity is typically large in the

small oscillation regime and 0 in the persistent oscillation and input dominated regimes (Fig. 2A). We therefore identify the small oscillation regime with the parameter configurations that admit a nonzero time-independent component. Under our approximations and for $\mathbf{J}$ i.i.d. random with mean zero and variance g^2/N , we find that this is equivalent to:

$$\frac{1}{\gamma} \leq g + \beta. \quad [19]$$

As shown in *SI Appendix, Fig. S9B*, this approximates the transition between these two regimes well. Moreover, since $|\nu| > \frac{1}{\gamma} - \beta$, this implies that in the persistent oscillation regime, $|\nu| > g$, which allows us to apply the results of *SI Appendix, section 2.3* to estimate the eigenvalue outlier.

3.6. Eigenvector Overlaps. In Fig. 4, *iv* we calculated the alignment between the plastic matrix, $\mathbf{A}(t)$, and the eigenvectors of the random backbone, $\mathbf{J}$, at the time the input is turned off, $t = 0$. To do this, we first performed a spectral decomposition of $\mathbf{A}(0)$, and we recorded the two eigenvectors with the largest eigenvalues, $\mathbf{u}, \mathbf{v} \in \mathbb{R}^N$. This was motivated by our analytical approximation in which $\mathbf{A}(0)$ is exactly rank two. Next, we calculated the principal angle between the subspace spanned by $\mathbf{u}$ and $\mathbf{v}$, compared with the subspace spanned by the real and imaginary parts of each eigenvector of $\mathbf{J}$. That is, given an eigenvector $\mathbf{e} = \mathbf{u}' + i\mathbf{v}'$ of $\mathbf{J}$, we calculated:

$$\text{Tr}(\mathbf{M}^T \mathbf{M}), \quad \text{where} \quad \mathbf{M} = \begin{pmatrix} \mathbf{u} \cdot \mathbf{u}' & \mathbf{u} \cdot \mathbf{v}' \\ \mathbf{v} \cdot \mathbf{u}' & \mathbf{v} \cdot \mathbf{v}' \end{pmatrix}. \quad [20]$$

In Fig. 4, *iv*, we color each eigenvalue of $\mathbf{J}$ by the corresponding overlap value. This measure captures the alignment between the leading eigenvectors of $\mathbf{A}(0)$ and each eigenvector of $\mathbf{J}$.

3.7. Targeted Oscillations. For the targeting results we describe in Section 1.5, we drew 5,000 different random backbone matrices $\mathbf{J}$ with $N = 500$. Then, for each random $\mathbf{J}$ matrix, we chose a random eigenvalue, formed oscillatory inputs following Eq. 7, and stimulated the networks for 200 time units, followed by a period without stimulation lasting 500 time units. For these analyses, we used $l = 6, k = 3, p = 25, f = 0.1$, and $g = 1.3$. After simulating the networks, we calculated the lifetime and frequency of the resulting persistent oscillation.

Data, Materials, and Software Availability. Code and simulation data to reproduce all results are available at <https://github.com/awakhloo/dynamic-memories> and <https://doi.org/10.5281/zenodo.21499172> (55).

ACKNOWLEDGMENTS. We thank members of the Aronov lab for feedback. This work was supported by the Gatsby Foundation (A.J.W., D.G.C., and L.F.A.), The Kavli Foundation (D.G.C. and L.F.A.), and NIH Grants T32NS064929 (D.G.C.) and T32MH126036 (A.J.W.).

- H. Sompolinsky, A. Crisanti, H. J. Sommers, Chaos in random neural networks. *Phys. Rev. Lett.* **61**, 259 (1988).
- J. Kadmon, H. Sompolinsky, Transition to chaos in random neuronal networks. *Phys. Rev. X* **5**, 041030 (2015).
- J. Aljadeff, M. Stern, T. Sharpee, Transition to chaos in random networks with cell-type-specific connectivity. *Phys. Rev. Lett.* **114**, 088101 (2015).
- F. Mastrogiuseppe, S. Ostojic, Linking connectivity, dynamics, and computations in low-rank recurrent neural networks. *Neuron* **99**, 609–623 (2018).
- D. G. Clark, L. Abbott, A. Litwin-Kumar, Dimension of activity in random neural networks. *Phys. Rev. Lett.* **131**, 118401 (2023).
- J. J. Hopfield, Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. U.S.A.* **79**, 2554–2558 (1982).
- D. Krotov, J. Hopfield, Large associative memory problem in neurobiology and machine learning. *arXiv [Preprint]* (2020). <https://arxiv.org/abs/2008.06996> (Accessed 23 September 2025).
- H. Sompolinsky, I. Kanter, Temporal association in asymmetric neural networks. *Phys. Rev. Lett.* **57**, 2861 (1986).
- M. Farrell, C. Pehlevan, Recall tempo of Hebbian sequences depends on the interplay of Hebbian kernel with tutor signal timing. *Proc. Natl. Acad. Sci. U.S.A.* **121**, e2309876121 (2024).
- L. F. Abbott, J. Varela, K. Sen, S. Nelson, Synaptic depression and cortical gain control. *Science* **275**, 221–224 (1997).
- O. Barak, M. Tsodyks, Persistent activity in neural networks with dynamic synapses. *PLoS Comput. Biol.* **3**, e35 (2007).
- G. Mongillo, O. Barak, M. Tsodyks, Synaptic theory of working memory. *Science* **319**, 1543–1546 (2008).
- Y. Mi, M. Katkov, M. Tsodyks, Synaptic correlates of working memory capacity. *Neuron* **93**, 323–330 (2017).
- D. V. Buonomano, W. Maass, State-dependent computations: Spatiotemporal processing in cortical networks. *Nat. Rev. Neurosci.* **10**, 113–125 (2009).
- A. Lansner, F. Fiebig, P. Herman, Fast Hebbian plasticity and working memory. *Curr. Opin. Neurobiol.* **83**, 102809 (2023).
- R. C. Malenka, R. A. Nicoll, NMDA-receptor-dependent synaptic plasticity: Multiple forms and mechanisms. *Trends Neurosci.* **16**, 521–527 (1993).
- R. C. Malenka, Postsynaptic factors control the duration of synaptic enhancement in area CA1 of the hippocampus. *Neuron* **6**, 53–60 (1991).
- B. Gustafsson, F. Asztely, E. Hanse, H. Wigström, Onset characteristics of long-term potentiation in the guinea-pig hippocampal CA1 region in vitro. *Eur. J. Neurosci.* **1**, 382–394 (1989).

19. M. A. Erickson, L. A. Maramba, J. Lisman, A single brief burst induces GluR1-dependent associative short-term potentiation: A potential mechanism for short-term memory. *J. Cogn. Neurosci.* **22**, 2530–2540 (2010).
20. F. Fiebig, A. Lansner, A spiking working memory model based on Hebbian short-term potentiation. *J. Neurosci.* **37**, 83–96 (2017).
21. S. G. Manohar, N. Zokaei, S. J. Fallon, T. P. Vogels, M. Husain, Neural mechanisms of attending to items in working memory. *Neurosci. Biobehav. Rev.* **101**, 1–12 (2019).
22. T. Micini, K. Stanley, J. Clune, "Differentiable plasticity: Training plastic neural networks with backpropagation" in *International Conference on Machine Learning* (PMLR, 2018), pp. 3559–3568.
23. A. Sandberg, J. Tegnér, A. Lansner, A working memory model based on fast Hebbian learning. *Network: Comput. Neural Syst.* **14**, 789 (2003).
24. J. Ba, G. E. Hinton, V. Mnih, J. Z. Leibo, C. Ionescu, Using fast weights to attend to the recent past. *Adv. Neural Inf. Process. Syst.* **29**, 4331–4339 (2016).
25. D. G. Clark, L. F. Abbott, Theory of coupled neuronal-synaptic dynamics. *Phys. Rev. X* **14**, 021001 (2024).
26. F. Schönsberg, D. Giana, Y. Chopra, M. Diamond, S. Goldt, Diverse perceptual biases emerge from Hebbian plasticity in a recurrent neural network model. *Neuron* **113**, 3673–3684.e6 (2025).
27. K. Rajan, L. Abbott, H. Sompolinsky, Stimulus-dependent suppression of chaos in recurrent neural networks. *Phys. Rev. E-Stat., Nonlinear, Soft Matter Phys.* **82**, 011903 (2010).
28. T. Tao, Outliers in the spectrum of IID matrices with bounded rank perturbations. *Probab. Theory Relat. Fields* **155**, 231–263 (2013).
29. D. J. Amit, N. Brunel, Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb. Cortex (New York, NY: 1991)* **7**, 237–252 (1997).
30. X. J. Wang, Synaptic reverberation underlying mnemonic persistent activity. *Trends Neurosci.* **24**, 455–463 (2001).
31. X. J. Wang, 50 years of mnemonic persistent activity: Quo vadis? *Trends Neurosci.* **44**, 888–902 (2021).
32. 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).
33. D. Mendoza-Halliday, S. Torres, J. C. Martinez-Trujillo, Sharp emergence of feature-selective sustained activity along the dorsal visual pathway. *Nat. Neurosci.* **17**, 1255–1262 (2014).
34. D. Zaksas, T. Pasternak, Directional signals in the prefrontal cortex and in area MT during a working memory for visual motion task. *J. Neurosci.* **26**, 11726–11742 (2006).
35. 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).
36. E. Baeg *et al.*, Dynamics of population code for working memory in the prefrontal cortex. *Neuron* **40**, 177–188 (2003).
37. J. Zhu *et al.*, Transient delay-period activity of agranular insular cortex controls working memory maintenance in learning novel tasks. *Neuron* **105**, 934–946 (2020).
38. P. T. Sadtler *et al.*, Neural constraints on learning. *Nature* **512**, 423–426 (2014).
39. M. D. Golub *et al.*, Learning by neural reassociation. *Nat. Neurosci.* **21**, 607–616 (2018).
40. G. Dragoi, S. Tonegawa, Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature* **469**, 397–401 (2011).
41. S. McKenzie *et al.*, Preexisting hippocampal network dynamics constrain optogenetically induced place fields. *Neuron* **109**, 1040–1054 (2021).
42. A. Luczak, P. Barthó, K. D. Harris, Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron* **62**, 413–425 (2009).
43. E. H. Park, S. Keeley, C. Savin, J. B. Ranck, A. A. Fenton, How the internally organized direction sense is used to navigate. *Neuron* **101**, 285–293 (2019).
44. A. A. Fenton, Remapping revisited: How the hippocampus represents different spaces. *Nat. Rev. Neurosci.* **25**, 428–448 (2024).
45. K. Duecker *et al.*, Challenges and approaches in the study of neural entrainment. *J. Neurosci.* **44**, e1234242024 (2024).
46. G. Sumbre, A. Muto, H. Baier, M.-m. Poo, Entrained rhythmic activities of neuronal ensembles as perceptual memory of time interval. *Nature* **456**, 102–106 (2008).
47. J. P. Lerousseau, A. Trébuchon, B. Morillon, D. Schön, Frequency selectivity of persistent cortical oscillatory responses to auditory rhythmic stimulation. *J. Neurosci.* **41**, 7991–8006 (2021).
48. S. Van Bree, E. Sohoglu, M. H. Davis, B. Zoefel, Sustained neural rhythms reveal endogenous oscillations supporting speech perception. *PLoS Biol.* **19**, e3001142 (2021).
49. L. T. Blanpain *et al.*, Multisensory flicker modulates widespread brain networks and reduces interictal epileptiform discharges. *Nat. Commun.* **15**, 3156 (2024).
50. H. Lee, G. V. Simpson, N. K. Logothetis, G. Rainer, Phase locking of single neuron activity to theta oscillations during working memory in monkey extrastriate visual cortex. *Neuron* **45**, 147–156 (2005).
51. J. Kamiński, A. Brzezicka, A. N. Mamelak, U. Rutishauser, Combined phase-rate coding by persistently active neurons as a mechanism for maintaining multiple items in working memory in humans. *Neuron* **106**, 256–264 (2020).
52. M. Siegel, M. R. Warden, E. K. Miller, Phase-dependent neuronal coding of objects in short-term memory. *Proc. Natl. Acad. Sci. U.S.A.* **106**, 21341–21346 (2009).
53. J. C. Magee, C. Grienberger, Synaptic plasticity forms and functions. *Annu. Rev. Neurosci.* **43**, 95–117 (2020).
54. P. Kidger, "On neural differential equations," PhD thesis, University of Oxford, Oxford (2021).
55. A. J. Wakhloo, *awakhloo/dynamic-memories: Zen* (2026). <https://doi.org/10.5281/zenodo.21499172>. Accessed 23 September 2025.