



Characteristics of sequential activity in networks with temporally asymmetric Hebbian learning

Maxwell Gillett^{a,b}, Ulises Pereira^{c,1}, and Nicolas Brunel^{a,b,c,d,2}

^aDepartment of Neurobiology, The University of Chicago, Chicago, IL 60637; ^bDepartment of Neurobiology, Duke University, Durham, NC 27708; ^cDepartment of Statistics, The University of Chicago, Chicago, IL 60637; and ^dDepartment of Physics, Duke University, Durham, NC 27708

Edited by Terrence J. Sejnowski, Salk Institute for Biological Studies, La Jolla, CA, and approved October 16, 2020 (received for review October 24, 2019)

Sequential activity has been observed in multiple neuronal circuits across species, neural structures, and behaviors. It has been hypothesized that sequences could arise from learning processes. However, it is still unclear whether biologically plausible synaptic plasticity rules can organize neuronal activity to form sequences whose statistics match experimental observations. Here, we investigate temporally asymmetric Hebbian rules in sparsely connected recurrent rate networks and develop a theory of the transient sequential activity observed after learning. These rules transform a sequence of random input patterns into synaptic weight updates. After learning, recalled sequential activity is reflected in the transient correlation of network activity with each of the stored input patterns. Using mean-field theory, we derive a low-dimensional description of the network dynamics and compute the storage capacity of these networks. Multiple temporal characteristics of the recalled sequential activity are consistent with experimental observations. We find that the degree of sparseness of the recalled sequences can be controlled by nonlinearities in the learning rule. Furthermore, sequences maintain robust decoding, but display highly labile dynamics, when synaptic connectivity is continuously modified due to noise or storage of other patterns, similar to recent observations in hippocampus and parietal cortex. Finally, we demonstrate that our results also hold in recurrent networks of spiking neurons with separate excitatory and inhibitory populations.

sequences | recurrent networks | Hebbian plasticity

Sequential activity has been reported across a wide range of neural systems and behavioral contexts, where it plays a critical role in temporal information encoding. Sequences can encode choice-selective information (1), the timing of motor actions (2), planned or recalled trajectories through the environment (3), and elapsed time (4–6). This diversity in function is also reflected at the level of neuronal activity. Sequences occur at varying timescales, from those lasting tens of milliseconds during hippocampal sharp-wave ripples, to those spanning several seconds in the striatum (7, 8). Sequential activity also varies in temporal sparsity. In nucleus HVC of zebra finch, highly precise sequential activity is present during song production, where many neurons fire only a single short burst during a syllable (9). In primate motor cortex, single neurons are typically active throughout a whole reach movement, but with heterogeneous and rich dynamics (10).

Numerous models have explored how networks with specific synaptic connectivity can generate sequential activity (11–20). These models span a wide range of single neuron models, from binary to spiking, and a wide range of synaptic connectivities. A large class of models employs temporally asymmetric Hebbian (TAH) learning rules to generate a synaptic connectivity necessary for sequence retrieval. In these models, a sequence of random input patterns are presented to the network, and a Hebbian learning rule transforms the resulting patterns of activity into synaptic weight updates. In networks of binary neurons, TAH learning together with synchronous update dynamics lead to sequence storage and retrieval (21). With asynchronous

dynamics, additional mechanisms such as delays are needed to generate robust sequence retrieval (13).

Another longstanding and influential class of model has centered on the “synfire chain” architecture (22) in networks of spiking neurons, in which discrete pools of neurons are connected in a chain-like, feedforward fashion, and a barrage of synchronous activity in the first pool propagates down connected pools, separated in time by synaptic delays (22, 23). Theoretical studies have shown that synfire chains can self-organize from initially unstructured connectivity through an appropriate choice of plasticity rule and stimulus (16, 19). Other studies have shown that these chains produce sequential activity when embedded into recurrent networks if appropriate recurrent and feedback connections are introduced among pools (24–26).

Despite decades of research on this topic, the relationships between network parameters and experimentally observed features of sequential activity are still poorly understood. In particular, few models account for the temporal statistics of sequential activity during retrieval and across multiple recording sessions. In many sequences, a greater proportion of neurons encode for earlier times, and the tuning width of recruited neurons increases in time (5, 27, 28). Neural sequences are also not static, but change over the timescales of days, even when controlling for the same environment and task constraints (29, 30). TAH learning provides an appealing framework for how sequences of neuronal activity are embedded in the connectivity matrix. While TAH learning has been extensively explored in binary neuron

Significance

Sequential activity is a prominent feature of many neural systems, in multiple behavioral contexts. Here, we investigate how Hebbian rules lead to storage and recall of random sequences of inputs in both rate and spiking recurrent networks. In the case of the simplest (bilinear) rule, we characterize extensively the regions in parameter space that allow sequence retrieval and compute analytically the storage capacity of the network. We show that nonlinearities in the learning rule can lead to sparse sequences and find that sequences maintain robust decoding but display highly labile dynamics to continuous changes in the connectivity matrix, similar to recent observations in hippocampus and parietal cortex.

Author contributions: M.G., U.P., and N.B. designed research; M.G., U.P., and N.B. performed analytical calculations, M.G. performed numerical simulations of the model; M.G. analyzed the results; and M.G. and N.B. wrote the paper with assistance from U.P.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

This open access article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

¹Present address: Center for Neural Science, New York University, New York, NY 10003.

²To whom correspondence may be addressed. Email: nicolas.brunel@duke.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.1918674117/-DCSupplemental>.

First published November 11, 2020.

network models, it remains an open question as to whether these learning rules can produce sequential activity in more realistic network models and also reproduce the observed temporal statistics of this activity. It is also unclear whether a theory can be developed to understand the resulting activity, as has been done with networks of binary neurons. Rate networks provide a useful framework for investigating these questions, as they balance analytical tractability and biological realism.

We start by constructing a theory of transient activity that can be used to predict the sequence capacity of these networks and focus on two features of sequences that have been described in the experimental literature: temporal sparsity and the selectivity of single-unit activity. We find that several temporal characteristics of sequential activity are reproduced by the network model, including the temporal distribution of firing-rate peak widths and times. Introduction of a nonlinearity to the learning rule allows for the generation of temporally sparse sequences. We also find that storing multiple sequences yields selective single-unit activity resembling that found during two-alternative forced-choice tasks (1). Consistent with experiment, we find that the population activity of these sequences changes over time in the face of ongoing synaptic changes, while decoding performance is preserved (29, 30). Finally, we demonstrate that our findings also hold in networks of excitatory and inhibitory leaky-integrate and fire (LIF) neurons, where learning takes place only in excitatory-to-excitatory connections.

Results

We have studied the ability of large networks of sparsely connected rate units to learn sequential activity from random, unstructured input, using analytical and numerical methods. We focus initially on a network model composed of N neurons, each described by a firing rate r_i ($i = 1, \dots, N$) that evolves according to standard rate equations (31)

$$\tau \frac{dr_i}{dt} = -r_i + \phi(h_i), \quad h_i = \sum_{j \neq i} J_{ij} r_j, \quad [1]$$

where τ is the time constant of rate dynamics, h_i is the total synaptic input to neuron i , and J_{ij} is the connectivity matrix. ϕ is a sigmoidal neuronal transfer function,

$$\phi(h) = \frac{r_{\max}}{2} \left(1 + \operatorname{erf} \left(\frac{h - \theta}{\sqrt{2}\sigma} \right) \right), \quad [2]$$

which is described by three parameters: r_{\max} , the maximal firing rate (set to one in most of the following); θ , the input at which the firing rate is half its maximal value; and σ , a parameter which is inversely proportional to the slope (gain) of the transfer function. Note that when $\sigma \rightarrow 0$, the transfer function becomes a step function: $\phi(h) = 1$ for $h > \theta$, 0 otherwise. Note also that the specific shape of the transfer function (Eq. 2) was chosen for the sake of analytical tractability, but our results are robust to other choices of transfer function (see Fig. 7).

Learning Rule. We assume that the synaptic connectivity matrix is the result of a learning process through which a set of sequences of random inputs presented to the network has been stored using a TAH learning rule. Specifically, we consider two different types of input, corresponding to a continuous and discretized sequence. In the continuous-input scenario, S input sequences $\eta_i^s(t)$ of duration T are defined as realizations of independent Ornstein–Uhlenbeck (OU) processes, with zero mean, SD σ_{OU} , and correlation time constant τ_{OU} , where $i = 1, \dots, N$ is the neuronal index and where $s = 1, \dots, S$ is the sequence index. The strength J_{ij} of a synaptic connection from neuron j to neuron i is modified according to a TAH learning rule, that

associates presynaptic inputs at time t with postsynaptic inputs at time $t + \Delta_t$, $\Delta J_{ij} \propto f(\eta_i^s(t + \Delta_t))g(\eta_j^s(t))$, where the functions f and g describe how postsynaptic and presynaptic inputs affect synaptic strength, respectively, and Δ_t is a temporal offset. After presentation of all S sequences, the connectivity matrix is given by

$$J_{ij} = A \frac{c_{ij}}{K} \sum_{s=1}^S \int_{t=0}^{T-\Delta_t} f(\eta_i^s(t + \Delta_t))g(\eta_j^s(t))dt, \quad [3]$$

where c_{ij} is a matrix of independent and identically distributed (i.i.d.) Bernoulli random variables ($c_{ij} = 1, 0$ with probabilities $c, 1 - c$), describing the “structural” connectivity matrix (i.e., whether a connection from neuron j to neuron i exists or not), $K = Nc$ represents the average in-degree of a neuron, and A controls the strength of the recurrent connectivity.

In the case $\tau_{OU} \ll \Delta_t$, $\sigma_{OU} = 1$, $T/\Delta_t = P$, if we discretize the OU process with a timestep of Δ_t , then Eq. 3 becomes:

$$J_{ij} = A \frac{c_{ij}}{K} \sum_{s=1}^S \sum_{\mu=1}^{P-1} f(\xi_i^{s,\mu+1})g(\xi_j^{s,\mu}), \quad [4]$$

where memorized sequences are now composed of P random i.i.d. Gaussian input patterns $\xi_i^{s,\mu}$ ($i = 1, \dots, N$, $\mu = 1, \dots, P$). This procedure is equivalent to a continuous-time learning process where each input is active for a duration Δ_t , and in Eq. 3, corresponds to random uncorrelated inputs $\eta_i^s(t) = \xi_i^{s,\mu(t)}$, where $\mu(t) = \lfloor t/\Delta_t \rfloor$ (Fig. 1A). For simplicity, when investigating discretized input, we only explore the case in which the discretization timestep of the input matches the temporal offset of the learning rule, Δ_t .

Note that this learning rule makes a number of simplifications, for the sake of analytical tractability: 1) Plasticity is assumed to depend only on external inputs to the network, and recurrent synaptic inputs are neglected during the learning phase. 2) Inputs are assumed to be discrete in time, and the learning rule is assumed to associate a presynaptic input at a given time, with postsynaptic input at the next time step. In other words, the temporal structure in the input should be matched with the prepost delay maximizing synaptic potentiation. Continuous-time sequences in which inputs change on faster timescales than this delay can also be successfully stored and retrieved (Fig. 1, *Left*). We note that the temporal asymmetry is consistent with both spike-timing dependent plasticity (STDP), that operates on timescales of tens of milliseconds (32), and behavioral timescale synaptic plasticity that operates on timescales of seconds (33).

Retrieval of Stored Sequences. We first ask the question of whether such a network can recall the stored sequences and characterize the phenomenology of the retrieved sequences. We start by considering a network that stores a single sequence of length T (corresponding to continuous input), or P i.i.d. Gaussian patterns, $\xi_i^\mu \sim \mathcal{N}(0, 1)$ for all $i = 1, \dots, N$, $\mu = 1, \dots, P$ (corresponding to discretized input), using a bilinear learning rule ($f(x) = x$ and $g(x) = x$). Without loss of generality, we choose $A = 1$, as this parameter can be absorbed into the transfer function by defining new parameters $\theta^* = \frac{\theta}{A}$ and $\sigma^* = \frac{\sigma}{A}$. The Gaussian assumption can be justified for a network in which neurons receive a large number of weakly correlated inputs, since the sum of such inputs is expected to be close to Gaussian due to the central-limit theorem. The bilinearity assumption is made for the sake of analytical tractability and will be relaxed below.

In Fig. 1B, we show the retrieval of a single sequence of 16 input patterns from the perspective of a single unit. The gray line corresponds to the sequence of firing rates driven by input patterns for this particular neuron (i.e., $\phi(\eta_i(t))$, or $\phi(\xi_i^\mu)$, for $\mu = 1, \dots, 16$). The solid red line displays the unit’s activity




Fig. 1. Retrieval of a stored sequence in the continuous (*Left*) and discrete (*Right*) cases. (*A, Left*) Continuous input to a single neuron, drawn from an OU process with $\tau_{OU} = 4$ ms and $\sigma_{OU} = 1$. (*A, Right*) Sequence of 16 discrete input patterns ξ_i^t presented to a particular neuron i . A.u., arbitrary units. (*B*) Dark red line: dynamics of the firing rate of the neuron during recall of the sequence, following initialization to the first pattern, $r_i = \phi(\eta_i^1(0))$ (*Left*) or $r_i = \phi(\xi_i^1)$ (*Right*). Gray line: input patterns passed through the neuronal nonlinearity. Note that the dynamics during recall produces a smooth approximation of the sequence. Light red line: dynamics following initialization to a noisy version of the first pattern. (*C*) Dynamics of four additional representative neurons, showing the diversity of temporal profiles of activity. (*D*) Overlap of network activity with each stored pattern (solid lines). In the continuous case, pattern values were sampled every 10 ms. Dashed gray line: average squared rates, M . (*E*) Correlation of network activity with each stored pattern (or pattern sampled every 10 ms in the continuous case). Parameters: $N = 40,000$, $c = 0.005$, $\tau = 10$ ms, $\sigma = 0.1$, $\theta = 0.14$ (*Left*) or $\theta = 0.22$ (*Right*).

following initialization to the first pattern in that sequence, after learning has taken place. Note that for this learning rule, the network dynamics do not reproduce exactly the stored input sequence, but, rather, produce a dynamical trajectory that is correlated with that input sequence. In Fig. 1*B*, we have specifically selected a unit whose activity transits close to the stored patterns during recall. In Fig. 1*C*, the activity of several random units is displayed during sequence retrieval. As each unit experiences an uncorrelated random sequence of inputs during learning, activity is highly heterogeneous, and units differ in the degree to which they are correlated with input patterns (*SI Appendix*, Fig. S1). The distribution of firing rates during recall depends on the parameters of the neural transfer function, ϕ . If the threshold θ is large, then the distribution of rates is unimodal, with a peak around zero, and few neurons firing at high rates. If the threshold is zero, the distribution of firing rates is bimodal, with increasing density at the minimal and maximal rates as the inverse gain parameter σ is decreased (*SI Appendix*, Fig. S11). Retrieval is also highly robust to noise—a random perturbation of the initial conditions of magnitude 75% of the pattern itself leads to sequence retrieval, as shown in Fig. 1*B* (compare the two solid lines, corresponding to unperturbed and perturbed initial conditions, respectively—see also *SI Appendix* for more details about noise robustness).

A natural quantity to measure sequence retrieval at the population level is the overlap (dot product) of the instantaneous firing rates with a given stored pattern in the sequence. For discretized input, this quantity is given by $m_i(t) = \frac{1}{N} \sum_{i=1}^N r_i(t) \xi_i^t$. To measure overlaps in networks storing continuous input, we sampled pattern values every 10ms from the continuous input (Fig. 1*D* and *E, Left*; *SI Appendix*). Overlaps can be thought of as linear decoders of the activity of the network. Another quantity is the (Pearson) correlation between network activity and the stored pattern, which for the bilinear rule corresponds to the

overlap divided by the SD of the firing rates across the population. In the general case, we consider the correlation between the transformed presynaptic input $g(\xi^t)$ and network activity $r(t)$ (*SI Appendix*). Overlaps and correlations are shown in Fig. 1*D* and *E*, respectively, showing that all patterns in the sequence are successively retrieved, with an approximately constant peak overlap (~ 0.1) and correlation (~ 0.4). Note that the maximum achievable overlap and correlation for the bilinear rule are 0.388 and 0.825, respectively, following retrieval from initialization to the first pattern (*SI Appendix*). In the remainder of the paper, we investigate the dynamics of sequential activity arising from stored discretized input.

Temporal Characteristics of the Retrieved Sequence. We next consider the temporal characteristics of retrieved sequences. The speed of retrieval of a sequence is defined as the inverse of the interval between the peaks in the overlaps with successive patterns in the sequence. We find that the network retrieves the sequence at a speed of approximately one pattern retrieved per time constant τ . For instance, in Fig. 1, in which the time constant of the units is 10 ms, the whole sequence of 16 patterns is retrieved in approximately 150 ms (the first pattern was retrieved at time 0). Empirically, we find that the time it takes to retrieve the whole sequence is approximately invariant to the number of sequences stored and the parameters of the rate-transfer function, and is well predicted by $t_{\text{retrieval}} = \tau(P - 1)$ (*SI Appendix*, Fig. S2). When storing discretized input, the retrieval time (and connectivity in Eq. 4) does not depend on the temporal offset of the learning rule Δ_t , as the discretization timestep of the stored input is matched to this interval. For continuous sequences of duration T , the time to retrieve stored sequences is proportional to $\tau T / \Delta_t$ —i.e., the retrieval time is again proportional to the time constant of the rate dynamics τ , but is inversely proportional to the temporal offset of the learning rule Δ_t .

While the sequence is retrieved at an approximately constant speed, several other features of retrieval depend strongly on the time within the sequence. In particular, the peak overlap magnitudes decay as a function of time within the sequence, but the width increases as $\tau\sqrt{l}$, where $l > 1$ denotes the index of the overlap (Fig. 2B and *SI Appendix*). Similarly, neurons at the beginning of a sequence are sharply tuned, and those toward the end have broader profiles (Fig. 2A). This is consistent with experimental findings in rat CA3 (5), medial prefrontal cortex (27) and monkey lateral prefrontal cortex (28). Another commonly reported experimental observation involves the temporal ordering of these peaks. Across these same areas and species, a larger number of neurons appear to encode for earlier parts of the sequence than for later segments, resulting in sequences that are not uniformly represented in time by neuronal activity (5, 27, 28). A similar pattern of overrepresentation is found in the model. The cumulative density of firing-rate peak times deviates from a uniform distribution, while the density of the overlap peak times are roughly uniform in time (Fig. 2C). This is due to a combination of two factors: Both initial overlaps and average activity are larger initially, leading to a larger probability that neurons with large values of ξ s for the initial patterns in the sequence will develop a peak.

A Low-Dimensional Description of Dynamics. To better understand the properties of these networks, we employed mean-field theory to analyze how the overlaps in a typical network realization evolve in time. To this end, we defined several “order parameters.” The first set of order parameters measures the typical overlap of network activity in time with the μ -th pattern: $q_\mu(t) = \mathbb{E}(\xi^\mu r(t))$, where the average is taken over the statistics of the patterns. Note that these are distinct from the previously defined overlaps m , which are empirically measured values for a particular realization of sequential activity patterns. However, for the large, sparse networks considered here, these quantities describe any typical realization of random input patterns and connectivity (*SI Appendix*). We also define the average squared rate $M(t) = \mathbb{E}(r^2(t))$, the two time autocorrelation function $C(u, t) = \mathbb{E}(r(u)r(t))$, and the network memory load as the number of patterns stored per average synaptic in-degree—i.e., $\alpha = \frac{S(P-1)}{K}$.

Framing the network in these terms yields a low-dimensional description in which the time-dependent evolution of the overlaps can be written as an effective delay line system (*SI Appendix*). In this formulation, the activation level of each overlap is driven by a gain-modulated version of the previous one in

the sequence, where the gain depends on the sequential load, rate variability, and norm of the overlaps

$$\tau \frac{dq_\mu}{dt} = -q_\mu + q_{\mu-1} G \left(\sum_{\nu=1}^{P-1} q_\nu^2(t) + \alpha M \right) \quad \mu = 2, \dots, P, \quad [5]$$

and the gain function G is given by

$$G(x) \equiv \frac{1}{\sqrt{2\pi(\sigma^2 + x)}} \exp\left(-\frac{\theta^2}{2(\sigma^2 + x)}\right). \quad [6]$$

The dynamics of M and C are given by coupled integro-differential equations (*SI Appendix*). With a constant gain G below one, recall of a sequence would decay to zero, as each overlap would become increasingly less effective at driving the next one in the sequence. Conversely, a constant gain above one would eventually result in runaway growth (*SI Appendix*). In Eq. 5, the gain changes in time due to its dependence on the norm of \bar{q} and M . During successful retrieval, the gain transiently rises to a value that is slightly larger than one, allowing for overlap peaks to maintain an approximately constant value (*SI Appendix*, Fig. S3B). We find that there are two different regimes for successful retrieval that depend on the shape of G (and, therefore, the transfer-function parameters): One in which sequences of arbitrarily small initial overlap are retrieved and another in which a finite initial overlap is required (*SI Appendix*, Fig. S3A and C).

The predictions from mean-field theory agree closely with the full network simulations. Fig. 3A shows the dynamics of a network in which two sequences of identical length are stored. At time 0, we initialize to the first pattern of the first sequence (corresponding overlaps shown in red). At 250 ms, we present an input lasting 10 ms corresponding to the first pattern of the second (blue) sequence. The solid lines show the overlap of the full network activity with each of the stored input patterns. The predicted overlaps from simulating mean-field equations are shown in dashed lines. The average squared rate is also predicted well by the theory, as is the two time autocorrelation function (*SI Appendix*, Fig. S4).

Storage Capacity of the Network. We next asked how the properties of retrieved sequences depend on the memory load α and what is the maximal storage capacity of the network, defined as the largest value of α for which sequences can be retrieved




Fig. 2. Temporal characteristics of a retrieved sequence. (A) Distribution of single-neuron peak widths, defined as continuous firing intervals occurring one SD above the single-neuron time-averaged firing rate. Black error bars denote mean and SD of widths within each 10-ms interval. Dashed red curve is best-fit trend using scaled square root function (*SI Appendix*). (B) Green dots correspond to observed overlap widths, defined as the weighted sample SD of $m_i(t)$ (*SI Appendix*). Red dots display the analytically predicted overlap width of $\tau\sqrt{l}$, where $\tau = 10$ ms and $l \geq 1$ (*SI Appendix*). (C) Cumulative percentage of peak times for single neurons (blue) and overlaps (gray). The dashed black line represents a uniform distribution. All parameters are as in Fig. 1.

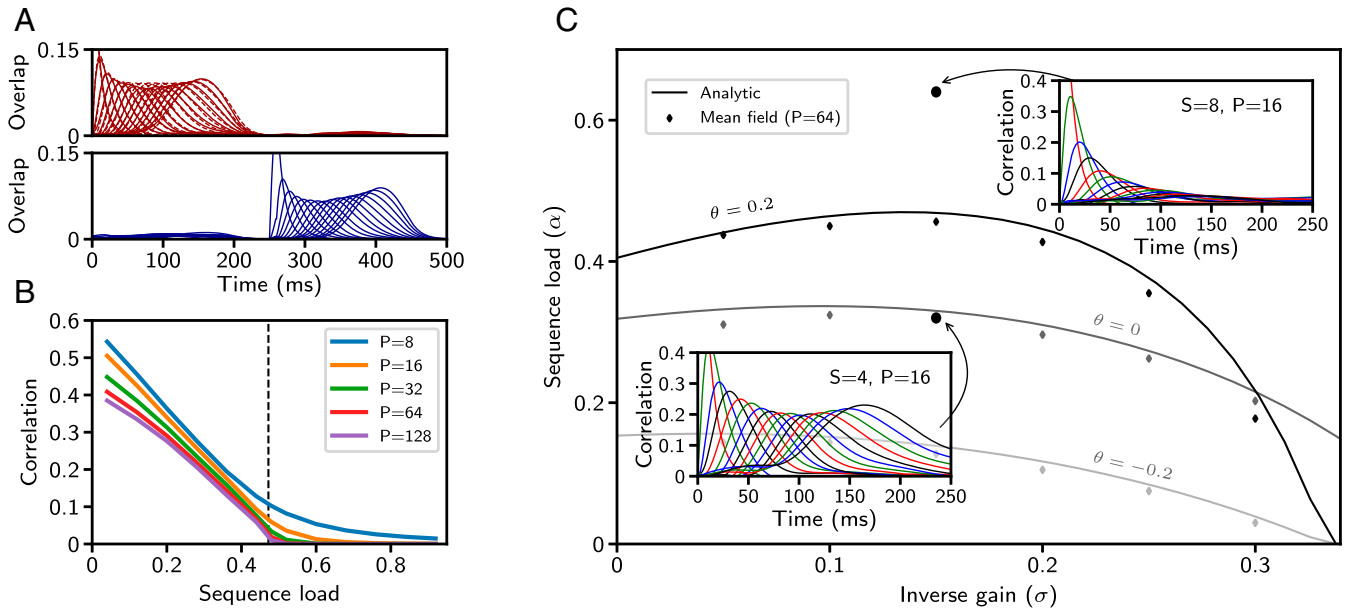


Fig. 3. Sequence capacity for Gaussian patterns. (A) Overlaps of two discrete sequences. Solid lines are full network simulations; dashed lines are simulations of the mean-field equations. (B) The maximal correlation with the final pattern in the stored sequence obtained from simulating mean-field equations, as a function of sequence load, for parameters corresponding to Fig. 1 ($\theta = 0.22$ and $\sigma = 0.1$). The vertical dashed line marks the predicted capacity for $\theta = 0.22$ and $\sigma = 0.1$. (C) Storage capacity as a function of the gain of the neural transfer function for three values of θ (all other parameters fixed). Solid curves: storage capacity computed from *SI Appendix*, Eqs. 39 and 40. Symbols: storage capacity, computed from simulations of mean-field equations (*SI Appendix*). C, Insets display representative overlap dynamics for network parameters above and below the capacity curve corresponding to $\theta = 0.2$ (solid lines within Insets are full network simulations). Parameters for A are as in Fig. 1, except $S = 2$.

successfully. Fig. 3B shows the peak correlation value that the network attained with the final pattern in a sequence as a function of the sequence load, obtained by using the mean-field equations. It shows that the correlation with the last pattern decreases with α until it reaches, for long enough sequences, a value close to zero at $\alpha \approx 0.47$. We calculated analytically this sequence capacity α_c and found that it is implicitly determined by solving the relation $G(\alpha_c M_c) = 1$ for α_c , where M_c is the average squared firing rate at capacity (*SI Appendix*). At capacity, for a fixed number of incoming connections K , the network can store one long patterned sequence of length $K\alpha_c$, or any number of sequences that collectively sum to this length (*SI Appendix*, Fig. S6).

In Fig. 3C, the analytically computed capacity curve is shown as a function of the inverse gain parameter of the rate-transfer function σ and for three values of θ (solid lines). Below a specified capacity curve, a sequence can be retrieved, and above it decays to zero. A small inverse gain corresponds to a transfer function with a very steep slope, whereas a large inverse gain produces a more shallow slope. The parameter θ determines the threshold of input required to drive a unit to half its maximum firing rate r_{\max} , where a larger θ corresponds to a higher threshold. The storage capacity obtained from simulating the mean-field equations (symbols) agrees well with the analytical result (solid lines).

For Gaussian patterns, we find that capacity is maximal for transfer functions with a positive threshold $\theta_c^+ = 1/\sqrt{2\pi e} \approx 0.24$ (*SI Appendix*, Figs. S3A and 5). The capacity decreases smoothly for $\theta < \theta_c^+$, but drops abruptly to zero above θ_c^+ . This abrupt drop is due to the behavior of the gain function G as parameters σ and θ are varied (*SI Appendix*, Fig. S3C). For $\theta < \theta_c^+$ (*SI Appendix*, Fig. S3A, region F), G has a region of stability (range of values of x such that $G(x) \geq 1$) that is bounded away from zero. For $\alpha < \alpha_c$, and for appropriate initial conditions, the gain eventually stabilizes around one: Either an initially large gain will decrease or an initially small gain will grow. As θ

crosses θ_c^+ , however, this region of stability disappears, resulting in zero capacity $\alpha_c = 0$. Capacity decreases continuously in other parameter directions. In particular, as θ approaches its lower bound θ_c^- , and as σ approaches its upper bound σ_c , the capacity continuously decreases to zero.

Nonlinear Learning Rules Produce Sparse Sequences. We have explored so far the storage of Gaussian sequences and found that the stored sequences can be robustly retrieved. Activity at the population level is heterogeneous in time (Fig. 1B), with many units maintaining high levels of activity throughout recall. However, many neural systems exhibit more temporally sparse sequential activity (9). While the average activity level of the network can be controlled to some degree by the neuronal threshold θ , retrieved sequences are nonsparse, even for the largest values of θ , for which the capacity is nonzero. Altering the statistics of the stored patterns will, in turn, change the statistics of neuronal activity during retrieval, which could lead to more sparse activity.

We therefore revisited Eq. 4 and explored nonlinear presynaptic and postsynaptic transformations ($g(x)$ and $f(x)$, respectively). We focus on a simplified functional form, similar to one shown to fit learning rules inferred from in vivo cortical data (34) and shown to maximize storage capacity in networks storing fixed-point attractors (35). For each synaptic terminal (pre and post), we apply a transformation that binarizes the activity patterns ξ into high and low values according to a threshold (Fig. 4A):

$$f(x) = \begin{cases} q_f & \text{if } x > x_f \\ q_f - 1 & \text{if } x \leq x_f, \end{cases} \quad [7]$$

$$g(x) = \begin{cases} q_g & \text{if } x > x_g \\ q_g - 1 & \text{if } x \leq x_g. \end{cases} \quad [8]$$

In the following, we choose $x_f = x_g$ and set q_g such that $\int Dz g(z) = 0$. This fixes the mean connection strength to zero,

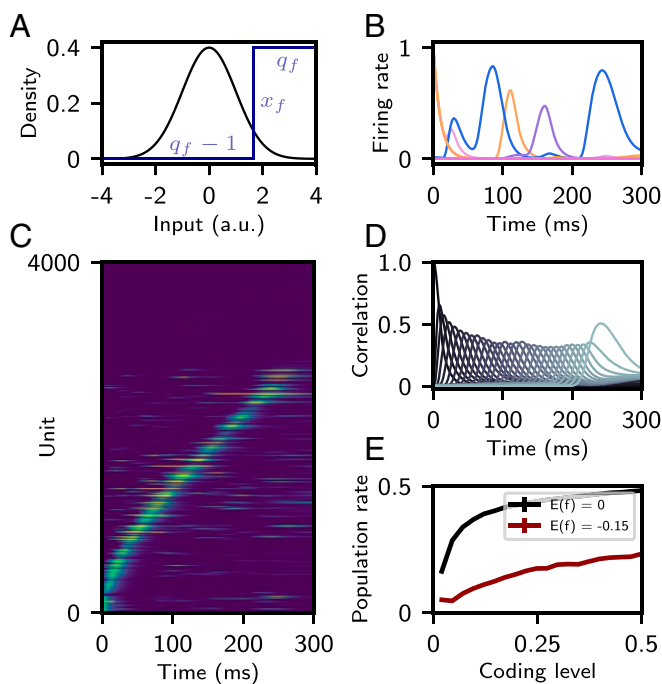


Fig. 4. Sparse sequences with a nonlinear learning rule. (A) Probability density of Gaussian input pattern ξ (black). Step function f binarizing the input patterns prior to storage in the connectivity matrix (blue) is shown. (B) Firing rate of several representative units as a function of time. (C) Firing rates of 4,000 neurons (out of 40,000) as a function of time, with “silent” neurons shown on top and active neurons on the bottom, sorted by time of peak firing rate. (D) Correlation of network activity with each stored pattern. (E) Average population rate as a function of “coding level” (probability that input pattern is above x_f). The average of $f(x)$ is maintained by varying q_f with x_f . In red, the average of $f(x)$ is constrained to -0.15 , the value in A–D. In black, the average of $f(x)$ is fixed to zero. All other parameters are as in A–D. For A–D, parameters of the learning rule were $x_f = 1.645$, $x_g = 1.645$, $q_f = 0.8$, and $q_g = 0.95$. $S = 1$, $P = 30$.

thereby preventing the mean total synaptic input from growing with the number of patterns stored, as was also the case with the bilinear learning rule. The parameter q_f is chosen such that $\int Dz f(z) \leq 0$ (34, 35).

High thresholds (x_f , x_g) lead to a low “coding level” (the probability that an input pattern is above x_f) of the binarized representation of the patterns stored in the connectivity matrix. We find that storage of such patterns leads to sparse sequential activity at a population level (Fig. 4 B and C) and strong transient correlations with these patterns during retrieval (Fig. 4D). Quantitative analysis of this activity reveals that both single-neuron peak and overlap widths are largely constant in time, unlike retrieval following a bilinear learning rule. Single-neuron peak times are still biased toward the start of the sequence, but to a lesser degree than with the bilinear rule (SI Appendix, Fig. S7). As a result of this sparse activity, overlap values are smaller than those resulting from the bilinear rule, but correlations are of similar magnitude, due to smaller pattern and firing rate variance.

This sparse activity is reflected in the average population firing rate, which is $0.04 \cdot r_{\max}$ in Fig. 4 B–D, compared to Fig. 1, where it was $0.11 \cdot r_{\max}$. Compared to networks storing Gaussian patterned sequences, this is a consequence of fewer neurons firing, and not due to a uniform decrease of single-unit activity. As the coding level decreases, the average population firing rate decreases (Fig. 4E). The magnitude of this decrease is dependent on the average of $f(x)$ (holding all other parameters equal) and

becomes smaller as this average becomes more negative. Given a low coding level, we expect many neurons to be silent during recall. As can be seen in Fig. 4C, for this choice of parameters, roughly 25% of neurons display no activity. The sequence capacity of these networks is larger than those resulting from the bilinear rule and increases as the coding level decreases (SI Appendix, Fig. S9).

Diverse Selectivity Properties Emerge from Learning Random Sequences. Experimentally observed sequential activity often displays some degree of diversity in selectivity properties. For instance, in posterior parietal cortex (PPC), mice show choice-selective sequential activity during two-alternative force-choice T-maze tasks. In the task, mice are briefly cued at the start of a trial before running down a virtual track that is identical across trial contexts. At the end of the track, they must turn either left or right according to the remembered cue to receive reward (1). Within a single recording session and across trials, many neurons display a preference for firing at a specific location during a single-choice context and are silent or weakly active otherwise. Other neurons fire at the same interval during both choice contexts, at different intervals, or are not modulated by the task (1). Similar activity has been described in rat CA1 hippocampal neurons during the short delay period preceding a different two-alternative figure-eight maze (4). It is an open question as to whether these different types of selectivity are a signature of task-specific mechanisms or simply the expected by-product of storing uncorrelated random inputs. To investigate this question, we stored two random uncorrelated sequences, each corresponding to the “left” and “right” target decisions in the T-maze task. The patterns in these sequences were transformed according to the nonlinear learning rule described in the previous section. We find that the same qualitative heterogeneity of response type exists in our network model. In Fig. 5, subpopulations of identified right-preferring, left-preferring, and nonspecific units (i.e., units that fire at similar times in both sequences) are plotted during a single recall trial of each sequence. Note that while the model reproduces qualitatively the diversity of selectivity properties seen in the data, quantitatively, the fraction of nonspecific neurons found in the model ($\sim 1\%$) is much smaller than that found in experiment ($\sim 10\%$ of putative imaged cells in ref. 1). This discrepancy might be due to different levels of sparsity or to correlations in the two sequences stored in mouse PPC.

Changes in Synaptic Connectivity Preserve Collective Sequence Retrieval while Profoundly Changing Single-Neuron Dynamics. While sequential activity in PPC is stereotyped across multiple trials in a single recording session, it changes significantly across multiday recordings (30). In a single recording session, many neurons display a strong peak of activity at one point in time during the task. Across consecutive days of recordings, however, a substantial fraction of these peaks are either gained or lost. Critically, information about trial type is not lost over multiple recording sessions, as reflected by above chance decoding performance of trial type (left vs. right) in subsets of cells (30). Similarly, place cells in CA1 of hippocampus form stable sequences within a single recording session, but exhibit large changes in single neuron place fields across days, even in the same environment (29). As in PPC, a total remapping is not observed, and information about the position of the animal is preserved at the population level across time (29).

We explored the possibility that these changes in single neuron selectivity might be due to ongoing changes in synaptic connectivity, consistent with spine dynamics observed in cortex and hippocampus (36, 37). To probe whether dynamic reorganization of sequential activity is consistent with such ongoing synaptic dynamics, we introduced random weight perturbations to the synaptic connectivity. We simulated the network over multiple




Fig. 5. Diverse selectivity properties emerge after learning two random input sequences. Single-trial raster plots display units that are active only during the right trial (*Left*), left trial (*Center*), and at similar times during both trials (*Right*). Activity is sorted by the time of maximal activity, and this sorting is fixed across rows. Numbers indicate amount of neurons passing selection criteria (total neuron count is 40,000). $S = 2$, and all other parameters are as in Fig. 4.

days, where the weights on day n are defined as a sum of a fixed, sequential component (J_{ij}^{seq}) as in Eq. 4, and a random dynamic component (J_{ij}^{dyn}):

$$J_{ij}(n) = J_{ij}^{\text{dyn}}(n) + J_{ij}^{\text{seq}}, \quad [9]$$

$$J_{ij}^{\text{seq}} = A \frac{C_{ij}}{K} \sum_{s=1}^S \sum_{\mu=1}^{P-1} f(\xi_i^{s,\mu+1}) g(\xi_j^{s,\mu}), \quad [10]$$

$$J_{ij}^{\text{dyn}}(1) = \sigma_z c_{ij} z_{ij}^1, \quad [11]$$

$$J_{ij}^{\text{dyn}}(n) = \lambda J_{ij}^{\text{dyn}}(n-1) + \sqrt{1-\lambda^2} \sigma_z z_{ij}^n, \quad n > 1, \quad [12]$$

where λ controls the decay of random perturbations, σ_z controls their amplitude, and $z_{ij}^n \sim \mathcal{N}(0,1)$ are i.i.d. random Gaussian variables of zero mean and unit variance.

In Fig. 6, we examine the reorganization of activity during a 30-d period, for parameters such that the activity strongly decorrelates over the time course of a few days. In Fig. 6A, sorted sequential activity is plotted for retrieval on day 1, followed by that during the 30-d mark. We find that sequential activity shows progressive reorganization, with the sequence on day 30 being only weakly, but significantly correlated (with a Pearson correlation coefficient of ~ 0.4) with the sequence on day 1. Critically, the underlying stored sequence can be read out equally well on both days, as shown by the pattern-correlation values plotted below. In Fig. 6B, representative single-unit activity is shown on the first, second, and final simulated day for four neurons, showing the emergence and disappearance of localized activity peaks in multiple neurons. To measure how quickly sequential activity decorrelates across time, for each day we plotted the average correlation of single-unit activity profiles with those on day 1 or 30 (Fig. 6C). For the parameters chosen here, sequences decor-

relate in a few days, such that by day 15, the correlation has reached a stable baseline. This qualitatively matches the ensemble activity correlation times reported in mouse CA1 (38). The dependence of the timescale and degree of decorrelation on λ and σ_z is shown in *SI Appendix*, Fig. S8.

Sequence Retrieval in Excitatory-Inhibitory Spiking Networks. For simplicity, we had previously chosen to neglect two key features of neuronal circuits: the separation of excitation and inhibition, and action-potential generation. To investigate whether our results are valid in more realistic networks, we developed a procedure to map the dynamics of a single-population rate network onto a network of current-based LIF neurons with separate populations of excitatory and inhibitory units. To implement sequence learning in two population networks, we made the following assumptions: 1) Learning takes place only between excitatory-to-excitatory recurrent connections; 2) to implement the sign constraint in those connections, we apply to the Hebbian connectivity matrix (Eq. 4) a nonlinear synaptic transfer function that imposes a nonnegativity constraint—specifically, we use a simple rectification (*SI Appendix*); and 3) inhibition is faster than excitation. Under these assumptions, we reformulate the two-population network into an equivalent single-population network with two sets of connections: one representing the excitatory-to-excitatory inputs, and the other representing the inhibitory feedback. We compute the effective inhibitory drive that balances the average positive excitatory recurrent drive arising from rectification, and, in turn, excitatory-to-inhibitory and inhibitory-to-excitatory connection strengths in the original two-population model that achieve this balance.

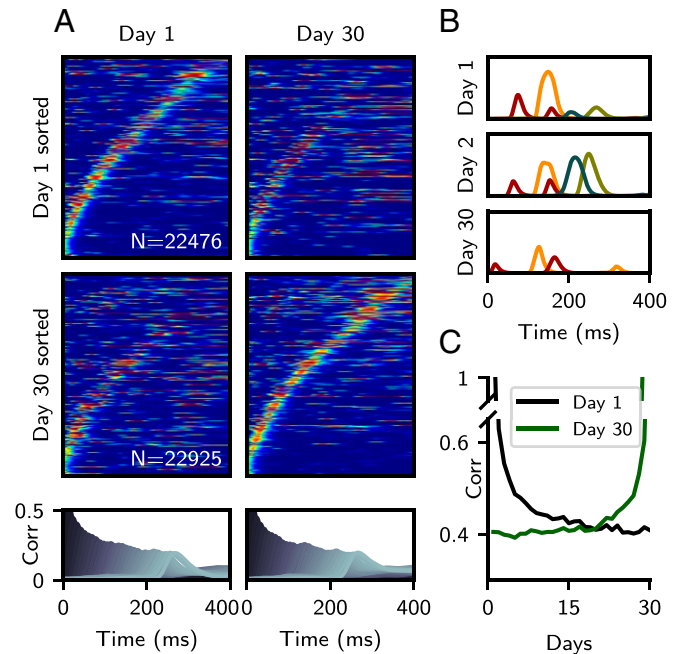


Fig. 6. Changes in synaptic connectivity preserve retrieval while changing single-neuron dynamics. (A) Sorted raster plots at the start and end of a 30-d simulation. (A, Top) Sorted activity on day 1. (A, Middle) Sorted activity on day 30. (A, Bottom) Pattern correlations computed on day 1 and 30. (B) Representative single-unit activity profiles of four neurons across days 1, 2, and 30. Color corresponds to unit identity. (C) Average correlation of activity profiles between day n , and either day 1 (black) or 30 (green). $\lambda = 0.85$, $\sigma_z = 0.03$, $S = 1$, and all other parameters as in Fig. 4. Note that the SD of the “noise” term in the connectivity matrix σ_z is larger than the SD of the fixed sequential component $\sigma_{\text{seq}} = 0.022$. The colormap scale for days 1 and 30 in A are the same as in Fig. 5.

Fig. 7 shows the results of applying this procedure in two steps (from a one-population to two-population rate model, and then from a two-population rate model to a two-population LIF network—see *SI Appendix* for details). In Fig. 7*A*, sequential activity is shown from a single-population rate network of size $N = 20,000$ with a rectified linear rate-transfer function. In Fig. 7*B*, we transform this population into excitatory units by rectifying their synaptic weights and add an additional population of $N = 5,000$ inhibitory units, with random sparse connections to and from the excitatory units. Finally, in Fig. 7*C*, we construct a current-based spiking network using parameters derived from LIF transfer functions. These LIF transfer functions were fit to a bounded region of the rectified linear transfer functions in Fig. 7*B* (*SI Appendix*). Connectivity matrices are identical for Fig. 7*B* and *C* up to a constant scaling factor. Raster plots of excitatory neurons reveal that activity is sparse, with individual neurons firing strongly for short periods of time. Inhibitory neurons fire at rates that are much less temporally modulated than those of excitatory neurons and inherit their fluctuations in activity from random connections from the excitatory neurons. Retrieval in the spiking network is robust to perturbations of the inhibitory feedback (*SI Appendix*, Fig. S12). Increasing by 5% or decreasing by 20% the strength of excitatory-to-inhibitory connections (J_{IE}) results in diminished or elevated global firing activity, respectively, but the sequence is retrieved with the same fidelity, as measured by pattern correlations.

Discussion

We have investigated a simple TAH learning rule in rate and spiking networks models and developed a theory of the transient activity observed in these networks. Using this theory, we have computed analytically the storage capacity for sequences in the case of a bilinear learning rule and derived parameters of the neuronal transfer function (threshold and gain) that maximize capacity.

We have found that a variety of temporal features of population activity in the network model match experimental observations, including those made from single trials of activity and across multiday recording sessions. A bilinear learning rule produces sequences that qualitatively match recorded distributions of single-neuron preferred times and tuning widths in cortical

and hippocampal areas. Tuning widths increase over the course of retrieval, and activity peaks are more concentrated at earlier segments of retrieval (5, 27, 28). With a nonlinear learning rule, sequential activity is temporally sparse and more uniform in time, consistent with findings in nucleus HVC of zebra finch (39, 40). To mimic the effects of synaptic turnover and ongoing learning across multiday recording sessions, we continually perturbed synaptic connectivity across repeated trials of activity. We show that single-unit activity can dramatically reorganize over this time period while maintaining a stable readout of sequential activity at the population level, consistent with recent experimental findings (29, 30). Finally, we have developed a procedure to map the sequential activity in simple rate networks to excitatory–inhibitory spiking networks.

Comparison with Previous Models. A large body of work has explored the generation and development of sequential activity in network models (11–19, 41). The storage and retrieval of binary activity patterns using TAH learning rules has been studied in networks of both binary and rate units. In networks of binary neurons, with asymmetric connectivity and parallel updating, the activity transitions instantaneously from one pattern to the next at each time step. With asynchronous updating, synaptic currents induced by the asymmetric component in the connectivity need to be endowed with delays in order to stabilize pattern transitions (13, 42). In contrast to previously studied rate and binary models, the present rate network neither transitions immediately between patterns nor dwells in pure pattern states for fixed periods of time. Instead, network activity during retrieval smoothly evolves in time in such a way as to become transiently correlated successively with each of the patterns in the sequence.

Previous studies of network of spiking neurons have heavily focused on the reproduction of temporally sparse synfire chain activity (15, 16, 44). However, a growing number of observations indicate that, outside specialized neural systems such as nucleus HVC, sequences are rarely consistent with such a simplified description. Contrary to previous work, our approach seeks to provide a unified framework to account for the diversity of observed dynamics. Recent modeling efforts to reproduce quantitatively experimentally observed dynamics have used supervised learning algorithms, in which initially randomly

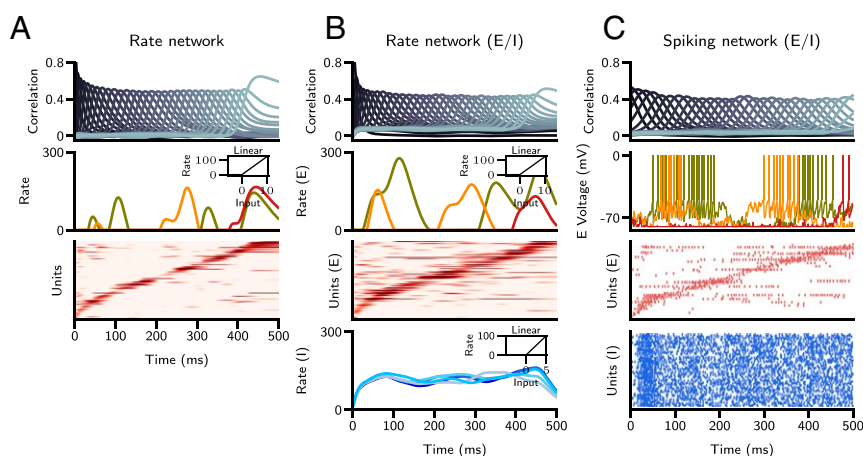


Fig. 7. Sequence retrieval in excitatory–inhibitory (E-I) networks. (A) Single-population rate network with threshold-linear transfer function. (A, Top) Correlation of network activity with stored patterns. (A, Middle) Firing rates of three representative units, where color corresponds to unit identity. A, Middle, Inset displays transfer function. (A, Bottom) Raster plot of units, sorted by time of peak firing rate. (B) Two-population rate network; panels are as in A. The lowermost graph in B shows the firing rate of representative inhibitory units. (C) Two-population LIF spiking network. (C, Top) As in A and B. (C, Middle) Voltage traces of three representative excitatory units. (C, Bottom) Raster plots of excitatory (sorted) and inhibitory units. Note that the same realizations of random sequences are stored in the three networks, and parameters of the E-I networks are computed so as to match the characteristics of sequences in the single-population rate network. In A and B, every 100th neuron is plotted in the raster plots for clarity, and silent neurons ($\sim 25\%$ of the population) are excluded. In C, every 200th neuron is plotted in the excitatory raster, and every 50th neuron in the inhibitory raster.

connected networks are trained to reproduce exactly activity patterns observed in experiment (18, 44). However, these learning rules are nonlocal, and there are currently no known biophysical mechanisms that would allow them to be implemented in brain networks. Here, we have shown that a simpler and more biophysically realistic learning rule is sufficient to reproduce many experimentally observed features of sequential activity.

Storage Capacity. The analytical methods used in this paper are generalizations of mean-field methods used extensively both in networks of binary neurons (48–51) and in networks of rate units (35, 52). Using these methods, we were able to compute analytically the storage capacity of sequences of independent Gaussian-distributed patterns of activity in rate models using a TAH learning rule. This allowed us to characterize how the storage capacity depends on the parameters of the transfer function and to obtain the threshold and gain of the transfer function that optimizes this storage capacity. We found, as in fixed-point attractor networks (35, 48, 53), and networks of binary neurons storing sequences (49, 50), that capacity scales linearly with the mean number of connections per neuron. The storage capacity of our model for the bilinear rule is of the same order of magnitude as storage capacities previously computed in networks of binary neurons, both in sparsely connected (49) and fully connected networks (50). While our results are derived in the large N limit, we have shown that it predicts well the measured capacity for finite-size networks. We also computed numerically the capacity for networks with nonlinear learning rules and found that the capacity increases when sequences become sparser, similar to what is found in networks storing fixed-point attractor states (35, 54, 55).

Storage and Retrieval. In the present work, we have used a simplified version of a TAH rule in which only nearest-neighbor pairs in the sequence of uncorrelated patterns modify the weights. This learning rule assumes that network activity is “clamped” during learning by external inputs, and that recurrent inputs do not interfere with such learning. The temporally discrete input considered here can be seen as an approximation of a continuous input sequence (Fig. 14). Alternatively, it can describe situations in which the inputs are discrete in time and/or learning in the presence of a strong global oscillation, where the time difference between consecutive input patterns would correspond to the period of the oscillation. We found for the bilinear rule that, following learning, the speed of retrieval (i.e., the inverse time between recall of consecutive patterns) matches the time constant of the rate dynamics. Moreover, we showed empirically that this speed does not depend on the number of stored patterns, nor on the particulars of the neural transfer function. This suggests that if patterns are presented during learning at this same timescale (i.e., $\Delta t = \tau$), one that is consistent with STDP, then they can be retrieved at the same speed as they were presented, without any additional mechanisms (56). If patterns are presented on slower timescales, then retrieval of input patterns could proceed on a faster speed than in which they were originally presented. Such a mechanism could underlie the phenomenon of hippocampal replay during sharp-wave ripple events, in which previously experienced sequences of place cell-activity can reactivate on significantly temporally compressed timescales (57). Whether additional cell-intrinsic or network mechanisms allow for dynamic control of retrieval speed remains an open question.

Retrieval in the present network model bears similarities to that of the functionally feedforward linear-rate networks analyzed by Goldman (58), in which recurrent connectivity was constructed by using an orthogonal random matrix to rotate a feedforward chain of unity synaptic weights. Retrieval was ini-

tiated by briefly stimulating an input corresponding to the first orthogonal pattern in the chain. The retrieval dynamics in those networks correspond exactly to those of the overlap dynamics for the bilinear rule in the presently studied network if a constant gain of $G = 1$ is assumed (*SI Appendix*). The key difference in our model is that the gain G is dynamic (Eq. 6), as it depends on network activity and overlaps with the stored patterns (Eq. 5). When retrieval is successful, network dynamics ensure that it maintains a value slightly above one throughout retrieval of the full sequence (see *SI Appendix* for detailed discussion), without any need for fine tuning.

We also found that retrieval is highly robust to initial Gaussian perturbations. Contrary to expectation, storage of additional nonretrieved patterns/sequences confers significantly increased robustness against a perturbation of the same magnitude (*SI Appendix*, Fig. S10). This can be understood by examining the mean-field equations and noticing that a larger α provides a greater range of convergence in initial overlap values to a gain that reaches an equilibrium value around one (see *SI Appendix* for detailed discussion).

While we have made several simplifying assumptions in the interest of analytical tractability, many of these assumptions do not hold in real neural networks. Biological neurons receive spatially and temporally correlated input. Plasticity rules also operate over a range of temporal offsets, and not at a single discrete time interval, as was assumed here (59). The learning rule we considered here operates on external inputs and does not account for the recurrent inputs arising from pattern presentations during learning. Feedback from these inputs could dramatically alter firing rates during learning and require additional mechanisms to stabilize learning. However, we note that the learning rule described here could potentially be a good approximation of a firing-rate-based rule, in a scenario in which learning is gated by neuromodulators. In such a scenario, plasticity would be gated by neuromodulators that would also act to weaken recurrent inputs, ensuring that external inputs are the primary drivers of neuronal activity. During the retrieval phase, recurrent inputs would be strong enough to sustain retrieval without ongoing additional external inputs (see, e.g., ref. 60). Understanding how each of these properties impact the findings presented will be the subject of future work.

Temporal Characteristics of Activity. We have found two emergent features of retrieval with a bilinear learning rule that are consistent with experiment: a broadening of single-unit activity profiles with time and an overrepresentation of peaks at earlier times of the sequence. These features have been reported in hippocampal area CA3 (5), medial prefrontal cortex (27) and lateral prefrontal cortex (28). Our theory provides a simple explanation for both of these features. Mean-field theory analysis shows that the temporal profile of the overlaps control, to a large extent, the profiles of single units. Specifically, the mean input current of a single unit is a linear combination of overlaps, with coefficients equal to the pattern values shown during learning. For Gaussian patterns, the overlaps rise transiently through an effective delay-line system, where the activation of each overlap is fed as input to drive the next. As the overlaps decay on the same timescale that they rise, the effective decay time grows for inputs later in the sequence, resulting in more broadly tuned profiles. This is also reflected in the longer decay time in the two time autocorrelation function at later times (*SI Appendix*, Fig. S4B). The overrepresentation of peaks at earlier times can be accounted for by a decrease in overlap magnitude throughout recall, as fluctuations in activity earlier in recall are more likely to rise above threshold and establish a peak.

We have also investigated the effects of storing sparse patterns using a nonlinear learning rule. The form of this rule bears

similarities to that inferred from electrophysiological data in inferior temporal cortex (34) and has been shown to be optimal for storing fixed-point attractors in the space of separable, sigmoidal learning rules (35). Sparse patterns lead to sequence statistics that are consistent, instead, with song-specific sequential activity in nucleus HVC of zebra finch (39, 40).

Effects of Random Synaptic Fluctuations on Sequence Retrieval.

We showed that simulating the effects of random fluctuations of synaptic weights by adding random weight perturbations can account for sequences that are unstable at the single-unit level, but which still maintain information as a population, consistent with experiment (29, 30, 38). We assumed that synaptic connectivity has both a stable component (storing the memory of sequences) and an unstable component that decays over timescales of days. The fluctuations of synaptic connectivity described by Eqs. 9–12 are roughly consistent with observations of spine dynamics over long timescales (e.g., ref. 61). When both components are of similar magnitude, the unstable component interferes with recall, producing noise that causes the input current to a single neuron to fluctuate strongly from day to day. While this noise is of sufficient strength to produce temporally sparse sequential activity with activity peaks that depends on its realization, it is not strong enough to disrupt recall driven by the stable sequential weight component.

If synaptic weights had no stable component and performed an unbiased random walk starting from their initial values, then decoded sequential activity representing task variables would unsurprisingly be expected to dissipate after some period of elapsed time. Experimentally, sequential activity after several weeks still exhibits a significant correlation with the initial pattern of activity. It is still an open question whether this correlation decays to zero on longer timescales. Here, we assumed that there exists a stable component to the synaptic connectivity (perhaps due to repeated exposure, and consequent relearning, of the animal to the same sensory environment), while fluctuations, described by a simple OU process, either represent purely stochastic synaptic fluctuations or are due to learning of other information unrelated to the specific task under which sequential activity is measured.

From Rate to Spiking Network Model. While our initial results were derived in rate models that do not respect Dale's law, we have shown that it is possible to build a network of spiking excitatory and inhibitory neurons that stores and retrieves sequences in a qualitatively similar way as in the rate model. In the spiking network, the Hebbian rule operates in excitatory-to-excitatory connections only, while all connections involving inhibitory neurons are fixed. With this architecture, most excitatory neurons show strong temporal modulations in their firing rate, while inhibitory neurons show much smaller temporal modulations around an average firing rate. These firing patterns are consistent with observations in multiple neural systems of smaller temporal modulations in the firing of inhibitory neurons, compared to excitatory cells (9, 62, 63).

Overall, our paper demonstrates that a simple Hebbian learning rule leads to storage and retrieval of sequences in large networks, with a phenomenology that mimics multiple experimental observations in multiple neural systems. The next step will be to investigate how the connectivity matrix used in this study can be obtained through biophysically realistic online learning dynamics.

Materials and Methods

Details of the numerical and analytical methods can be found in *SI Appendix*, including the mean-field theory for the bilinear rule, construction of the excitatory–inhibitory rate network, and details of the LIF spiking network. Simulation and analysis details for all figures are documented in *SI Appendix, Supplemental Procedures*. Parameters for all simulations can be found in *SI Appendix, Parameter Tables*.

Software. Rate and spiking network simulations were performed by using custom C++ and Python routines with the aid of NumPy (64), SciPy (65), Numba (66), Cython (67), Ray (68), and Jupyter (69). Graphics were generated by using Matplotlib (70).

Data Availability. Code to generate the figures is available at GitHub, <https://www.github.com/maxjillett/hebbian.sequence.learning>.

ACKNOWLEDGMENTS. This work was supported by NIH Grant R01 EB022891 and Office of Naval Research Grant N00014-16-1-2327. We thank Henry Greenside and Harel Shouval for their comments on a previous version of the manuscript, and all members of the N.B. group for discussions.

- C. D. Harvey, P. Coen, D. W. Tank, Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* **484**, 62–68 (2012).
- J. Wang, D. Narain, E. A. Hosseini, M. Jazayeri, Flexible timing by temporal scaling of cortical responses. *Nat. Neurosci.* **21**, 102–110 (2018).
- A. K. Lee, M. A. Wilson, Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron* **36**, 1183–1194 (2002).
- E. Pastalkova, V. Itskov, A. Amarasingham, G. Buzsaki, Internally generated cell assembly sequences in the rat hippocampus. *Science* **321**, 1322–1327 (2008).
- D. M. Salz et al., Time cells in hippocampal area CA3. *J. Neurosci.* **36**, 7476–7484 (2016).
- J. G. Heys, D. A. Dombeck, Evidence for a subcircuit in medial entorhinal cortex representing elapsed time during immobility. *Nat. Neurosci.* **21**, 1574–1582 (2018).
- G. Buzsaki, Hippocampal sharp waves: Their origin and significance. *Brain Res.* **398**, 242–252 (1986).
- H. Akhlaghpour et al., Dissociated sequential activity and stimulus encoding in the dorsomedial striatum during spatial working memory. *eLife* **5**, e19507 (2016).
- R. H. Hahnloser, A. A. Kozhevnikov, M. S. Fee, An ultra-sparse code underlies the generation of neural sequences in a songbird. *Nature* **419**, 65–70 (2002).
- M. M. Churchland, K. V. Shenoy, Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. *J. Neurophysiol.* **97**, 4235–4257 (2007).
- S.-I. Amari, Learning patterns and pattern sequences by self-organizing nets of threshold elements. *IEEE Trans. Comput. C-21*, 1197–1206 (1972).
- A. Herz, B. Sulzer, R. Kühn, J. L. van Hemmen, The Hebb rule: Storing static and dynamic objects in an associative neural network. *Europhys. Lett.* **7**, 663–669 (1988).
- R. Kühn, J. L. Van Hemmen, "Temporal association" in *Models of Neural Networks I*, E. Domany, J. L. Van Hemmen, K. Schulten, Eds. (Springer, Berlin, 1995), pp. 221–288.
- M. Rabinovich et al., Dynamical encoding by networks of competing neuron groups: Winnerless competition. *Phys. Rev. Lett.* **87**, 068102 (2001).
- M. A. Long, D. Z. Jin, M. S. Fee, Support for a synaptic chain model of neuronal sequence generation. *Nature* **468**, 394–399 (2010).
- I. R. Fiete, W. Senn, C. Z. Wang, R. H. Hahnloser, Spike-time-dependent plasticity and heterosynaptic competition organize networks to produce long scale-free sequences of neural activity. *Neuron* **65**, 563–576 (2010).
- J. Cannon, N. Kopell, T. Gardner, J. Markowitz, Neural sequence generation using spatiotemporal patterns of inhibition. *PLoS Comput. Biol.* **11**, e1004581 (2015).
- K. Rajan, C. D. Harvey, D. W. Tank, Recurrent network models of sequence generation and memory. *Neuron* **90**, 128–142 (2016).
- N. R. Tannenbaum, Y. Burak, Shaping neural circuits by high order synaptic interactions. *PLoS Comput. Biol.* **12**, e1005056 (2016).
- J. M. Murray, G. S. Escola, Learning multiple variable-speed sequences in striatum via cortical tutoring. *eLife* **6**, e26084 (2017).
- M. Herrmann, J. A. Hertz, A. Prugel-Bennet, Analysis of synfire chains. *Netw. Comput. Neural Syst.* **6**, 403–414 (1995).
- M. Abeles, *Corticonics* (Cambridge University Press, New York, NY, 1991).
- M. Diesmann, M. O. Gewaltig, A. Aertsen, Stable propagation of synchronous spiking in cortical neural networks. *Nature* **402**, 529–533 (1999).
- Y. Aviel, E. Pavlov, M. Abeles, D. Horn, Synfire chain in a balanced network. *Neurocomputing* **44**, 285–292 (2002).
- A. Kumar, S. Rotter, A. Aertsen, Conditions for propagating synchronous spiking and asynchronous firing rates in a cortical network model. *J. Neurosci.* **28**, 5268–5280 (2008).
- N. Chenkov, H. Sprekeler, R. Kempter, Memory replay in balanced recurrent networks. *PLoS Comput. Biol.* **13**, e1005359 (2017).
- Z. Tiganj, M. W. Jung, J. Kim, M. W. Howard, Sequential firing codes for time in rodent medial prefrontal cortex. *Cerebr. Cortex* **27**, 5663–5671 (2017).
- Z. Tiganj, J. A. Cromer, J. E. Roy, E. K. Miller, M. W. Howard, Compressed timeline of recent experience in monkey lateral prefrontal cortex. *J. Cognit. Neurosci.* **30**, 935–950 (2018).
- Y. Ziv et al., Long-term dynamics of CA1 hippocampal place codes. *Nat. Neurosci.* **16**, 264–266 (2013).

30. L. N. Driscoll, N. L. Pettit, M. Minderer, S. N. Chettih, C. D. Harvey, Dynamic reorganization of neuronal activity patterns in parietal cortex. *Cell* **170**, 986–999 (2017).
31. J. J. Hopfield, Neurons with graded response have collective computational properties like those of two-state neurons. *Proc. Natl. Acad. Sci. U.S.A.* **81**, 3088–3092 (1984).
32. H. Markram, J. Lubke, M. Frotscher, B. Sakmann, Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* **275**, 213–215 (1997).
33. K. C. Bittner, A. D. Milstein, C. Grienberger, S. Romani, J. C. Magee, Behavioral time scale synaptic plasticity underlies CA1 place fields. *Science* **357**, 1033–1036 (2017).
34. S. Lim *et al.*, Inferring learning rules from distributions of firing rates in cortical neurons. *Nat. Neurosci.* **18**, 1804–1810 (2015).
35. U. Pereira, N. Brunel, Attractor dynamics in networks with learning rules inferred from in vivo data. *Neuron* **99**, 227–238 (2018).
36. A. Holtmaat, K. Svoboda, Experience-dependent structural synaptic plasticity in the mammalian brain. *Nat. Rev. Neurosci.* **10**, 647–658 (2009).
37. T. Pfeiffer *et al.*, Chronic 2P-STED imaging reveals high turnover of dendritic spines in the hippocampus in vivo. *eLife* **7**, e34700 (2018).
38. A. Rubin, N. Geva, L. Sheintuch, Y. Ziv, Hippocampal ensemble dynamics timestamp events in long-term memory. *eLife* **4**, e12247 (2015).
39. G. F. Lynch, T. S. Okubo, A. Hanuschkin, R. H. Hahnloser, M. S. Fee, Rhythmic continuous-time coding in the songbird analog of vocal motor cortex. *Neuron* **90**, 877–892 (2016).
40. M. A. Picardo *et al.*, Population-level representation of a temporal sequence underlying song production in the zebra finch. *Neuron* **90**, 866–876 (2016).
41. D. Kleinfeld, Sequential state generation by model neural networks. *Proc. Natl. Acad. Sci. U.S.A.* **83**, 9469–9473 (1986).
42. H. Sompolinsky, I. Kanter, Temporal association in asymmetric neural networks. *Phys. Rev. Lett.* **57**, 2861–2864 (1986).
43. D. Kleinfeld, H. Sompolinsky, Associative neural network model for the generation of temporal patterns. Theory and application to central pattern generators. *Biophys. J.* **54**, 1039–1051 (1988).
44. R. M. Memmesheimer, R. Rubin, B. P. Olveczky, H. Sompolinsky, Learning precisely timed spikes. *Neuron* **82**, 925–938 (2014).
45. A. Veliz-Cuba, H. Z. Shouval, K. Josić, Z. P. Kilpatrick, Networks that learn the precise timing of event sequences. *J. Comput. Neurosci.* **39**, 235–254 (2015).
46. P. J. Tully, H. Linden, M. H. Hennig, A. Lansner, Spike-based Bayesian-Hebbian learning of temporal sequences. *PLoS Comput. Biol.* **12**, e1004954 (2016).
47. J. D. Murray *et al.*, Stable population coding for working memory coexists with heterogeneous neural dynamics in prefrontal cortex. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 394–399 (2017).
48. B. Derrida, E. Gardner, A. Zippelius, An exactly solvable asymmetric neural network model. *Europhys. Lett.* **4**, 167–173 (1987).
49. H. Gutfreund, M. Mezard, Processing of temporal sequences in neural networks. *Phys. Rev. Lett.* **61**, 235–238 (1988).
50. A. During, A. C. C. Coolen, D. Sherrington, Phase diagram and storage capacity of sequence processing neural networks. *J. Phys. Math. Gen.* **31**, 8607–8621 (1998).
51. M. Kawamura, M. Okada, Transient dynamics for sequence processing neural networks. *J. Phys. Math. Gen.* **35**, 253–266 (2002).
52. B. Tirozzi, M. Tsodyks, Chaos in highly diluted neural networks. *Europhys. Lett.* **14**, 727–732 (1991).
53. D. J. Amit, H. Gutfreund, H. Sompolinsky, Storing infinite numbers of patterns in a spin-glass model of neural networks. *Phys. Rev. Lett.* **55**, 1530–1531 (1985).
54. M. Tsodyks, M. V. Feigel'man, The enhanced storage capacity in neural networks with low activity level. *Europhys. Lett.* **6**, 101–105 (1988).
55. M. V. Tsodyks, Associative memory in asymmetric diluted network with low level of activity. *Europhys. Lett.* **7**, 203–208 (1988).
56. W. Gerstner, R. Ritz, J. L. van Hemmen, Why spikes? Hebbian learning and retrieval of time-resolved excitation patterns. *Biol. Cybern.* **69**, 503–515 (1993).
57. Z. Nadasdy, H. Hirase, A. Czurko, J. Csicsvari, G. Buzsaki, Replay and time compression of recurring spike sequences in the hippocampus. *J. Neurosci.* **19**, 9497–9507 (1999).
58. M. S. Goldman, Memory without feedback in a neural network. *Neuron* **61**, 621–634 (2009).
59. L. F. Abbott, S. B. Nelson, Synaptic plasticity: Taming the beast. *Nat. Neurosci.* **3**, 1178–1183 (2000).
60. M. E. Hasselmo, The role of acetylcholine in learning and memory. *Curr. Opin. Neurobiol.* **16**, 710–715 (2006).
61. J. T. Trachtenberg *et al.*, Long-term in vivo imaging of experience-dependent synaptic plasticity in adult cortex. *Nature* **420**, 788–794 (2002).
62. G. Kosche, D. Vallentin, M. A. Long, Interplay of inhibition and excitation shapes a premotor neural sequence. *J. Neurosci.* **35**, 1217–1227 (2015).
63. C. Grienberger, A. D. Milstein, K. C. Bittner, S. Romani, J. C. Magee, Inhibitory suppression of heterogeneously tuned excitation enhances spatial coding in CA1 place cells. *Nat. Neurosci.* **20**, 417–426 (2017).
64. S. van der Walt, S. C. Colbert, G. Varoquaux, The NumPy array: A structure for efficient numerical computation. *Comput. Sci. Eng.* **13**, 22–30 (2011).
65. P. Virtanen *et al.*, SciPy 1.0: Fundamental algorithms for scientific computing in python. *Nat. Methods* **17**, 261–272 (2020).
66. S. K. Lam, A. Pitrou, S. Seibert, “Numba: A LLVM-based python JIT compiler” in *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC, LLVM'15* (Association for Computing Machinery, New York, NY, 2015), pp. 1–6.
67. S. Behnel *et al.*, Cython: The best of both worlds. *Comput. Sci. Eng.* **13**, 31–39 (2011).
68. P. Moritz *et al.*, “A distributed framework for emerging AI applications” in *Proceedings of the 13th USENIX Conference on Operating Systems Design and Implementation* (USENIX Association, 2018), pp. 561–577.
69. K. Thomas *et al.*, “Jupyter Notebooks – A publishing format for reproducible computational workflows” in *Positioning and Power in Academic Publishing: Players, Agents, and Agendas*, F. Loizides, B. Schmidt, eds. (IOS Press, Amsterdam, Netherlands, 2016), pp. 87–90.
70. J. D. Hunter, Matplotlib: A 2D graphics environment. *Comput. Sci. Eng.* **9**, 90–95 (2007).