

Learning recurrent dynamics in spiking networks

Christopher M. Kim and Carson C. Chow

Laboratory of Biological Modeling, NIDDK/NIH, Bethesda, MD

Abstract

Spiking activity of neurons engaged in learning and performing a task show complex spatiotemporal dynamics. While the output of recurrent network models can learn to perform various tasks, the possible range of recurrent dynamics that emerge after learning remains unknown. Here we show that modifying the recurrent connectivity with a recursive least squares algorithm provides sufficient flexibility for synaptic and spiking rate dynamics of spiking networks to produce a wide range of spatiotemporal activity. We apply the training method to learn arbitrary firing patterns, stabilize irregular spiking activity of a balanced network, and reproduce the heterogeneous spiking rate patterns of cortical neurons engaged in motor planning and movement. We identify sufficient conditions for successful learning, characterize two types of learning errors, and assess the network capacity. Our findings show that synaptically-coupled recurrent spiking networks possess a vast computational capability that can support the diverse activity patterns in the brain.

Introduction

Neuronal populations exhibit diverse patterns of recurrent activity that can be highly irregular or well-structured when learning or performing a behavioral task¹⁻⁵. An open question is whether learning-induced synaptic rewiring is sufficient to give rise to the observed wide range of recurrent spiking dynamics that encodes and processes information.

It has been shown that a network of recurrently connected neuron models can be trained to perform complex motor and cognitive tasks. In this approach, synaptic connections to the outputs are rewired to generate a desired population-averaged signal, while the activity of individual neurons emerges in a self-organized way. The recurrent dynamics resulting from such learning schemes have harnessed chaotic temporally irregular activity of a network of rate-based neurons⁶ that is made repeatable either through direct feedback from the outputs^{7,8} or through training of the recurrent connections directly⁹. The resulting irregular yet stable dynamics provides a rich reservoir from which complex motor commands can be extracted by output neurons that sum over the network linearly^{10,11}.

Extending this idea to spiking networks poses a challenge because it is difficult to coordinate the spiking dynamics of many neurons, especially, if spike times are variable as in a balanced network. Some success has been achieved by training spiking networks directly with a feedback loop¹² or using a rate-based network as an intermediate step^{13,14}. A different top-down approach is to build networks that emit spikes optimally to correct the discrepancy between the actual and desired network outputs^{15,16}. This optimal coding strategy in a tightly balanced network can be learned with a local plasticity rule¹⁷ and is able to generate arbitrary network output^{18,19}. Although these studies demonstrate that network outputs can perform universal computations⁷, the possible repertoire of the recurrent activity has not been extensively explored. Various network models have been developed to explain certain important features of recurrent dynamics (e.g. fixed point attractor²⁰, synchrony²¹⁻²³, sequential activity^{3,4,24}, stable chaos²⁵, motor planning²⁶, and variable spiking²⁷⁻³⁰). However, we still lack an understanding of the scope of recurrent dynamics that spiking networks are capable of generating.

Here we show that a network of spiking neurons is capable of supporting arbitrarily complex coarse-grained recurrent dynamics provided the spatiotemporal patterns of the recurrent activity are diverse, the synaptic dynamics are fast, and the number of neurons in the network is large. We give a theoretical basis for how a network can learn and show various examples, which include stabilizing strong chaotic

rate fluctuations in balanced networks and constructing a recurrent network that reproduces the spiking rate patterns of a large number of cortical neurons involved in motor planning and movement. Our study suggests that individual neurons in a recurrent network have the capability to support near universal dynamics.

Results

1 Spiking networks can learn complex recurrent dynamics

We consider a network of N quadratic integrate-and-fire neurons that are recurrently connected with spike-activated synapses weighted by a connectivity matrix W . We show below that our results do not depend on the spiking mechanism. We focus on two measures of coarse-grained time-dependent neuron activity: 1) the synaptic drive $u_i(t)$ to neuron i which is given by the W -weighted sum of low-pass filtered incoming spike trains, and 2) the time-averaged spiking rate $R_i(t)$ of neuron i . The goal is to find a weight matrix W that can autonomously generate desired recurrent target dynamics when the network of spiking neurons connected by W is stimulated briefly with an external stimulus (**Fig. 1a**). The target dynamics are defined by a set of functions $f_1(t), f_2(t), \dots, f_N(t)$ on a time interval $[0, T]$. Learning of the recurrent connectivity W is considered successful if $u_i(t)$ or $R_i(t)$ evoked by the stimulus matches the target functions $f_i(t)$ over the time interval $[0, T]$ for all neurons $i = 1, 2, \dots, N$.

Previous studies have shown that recurrently connected rate units can learn chaotic trajectories that the initial network could already generate⁹, trajectories from a different network³¹, and sequential activity derived from imaging data³². Our study expands these results by showing that recurrent dynamics of spiking networks can be trained and the repertoire of recurrent dynamics that can be encoded is vast. To train the recurrent connectivity, we modified the Recursive Least Squares (RLS) algorithm developed in rate models, which minimizes a quadratic cost function between the activity measure and the target together with a quadratic regularization term^{8,9,33} (see Methods).

As a first example, we trained the network to produce synaptic drive patterns that matched a set of sine functions and the spiking rate to match the positive part of the same sine functions. The initial connectivity matrix has connection probability $p = 0.3$ and the coupling strength is drawn from a Normal distribution with mean 0 and standard deviation σ . Prior to training, the synaptic drive fluctuates irregularly, but as soon as the RLS algorithm is instantiated, the synaptic drives follow the target with small error; rapid changes in W quickly adjust the recurrent dynamics towards the target⁸ (**Fig. 1b, c**). As a result, the population spike trains exhibit reproducible patterns across training trials. A brief stimulus precedes each training session to reset the network to a specific state. If the training is successful, the trained response can be elicited whenever the same stimulus is applied regardless of the network state. We were able to train a network of rate-based neurons to learn arbitrarily complex target patterns using the same learning scheme (**Supplementary Fig. S1**).

Training the spiking rate is more challenging than training the synaptic drive because small changes in recurrent connectivity do not immediately affect the spiking activity if the effect is below the spike-threshold. Therefore, the spike trains may not follow the desired spiking rate pattern during the early stage of training, and the population spike trains no longer appear similar across training trials (**Fig. 1d**). This is also reflected in relatively small changes in recurrent connectivity and the substantially larger number of training runs required to produce desired spiking patterns (**Fig. 1e**). However, by only applying the training when the total input to a neuron is suprathreshold, the spiking rate can be trained to reproduce the target patterns. The correlation between the actual filtered spike trains and the target spiking rate increases gradually as the training progresses.

In contrast to training the network read-out, where it is proposed that the initial network needs to be at the “edge of chaos” to learn successfully^{8,12,14,34,35}, the recurrent connectivity can learn to produce the desired recurrent dynamics regardless of the initial network dynamics and connectivity. Even when the initial network has no synaptic connections, the brief stimulus preceding the training session is sufficient to build a fully functioning recurrent connectivity that captures the target dynamics. The RLS algorithm can grow new synapses or tune existing ones as long as some of the neurons become active after the initial stimulus (**Supplementary Fig. S2**).

Learning was not limited to one set of targets; the same network was able to learn multiple sets of targets. We trained the network to follow two independent sets of targets, where each target function was a sine function with random frequency. Every neuron in the network learned both activity patterns after




Figure 1: Synaptic drive and spiking rate of neurons in a recurrent network can learn complex patterns. **(a)** Schematic of network training. Blue square represents the external stimulus that elicits the desired response. Black curves represent target output for each neuron. Red arrows represent recurrent connectivity that is trained to produce desired target patterns. **(b)** Synaptic drive of 10 sample neurons before, during and after training. Pre-training is followed by multiple training trials. An external stimulus (blue) is applied prior to training for 100 ms. Synaptic drive (black) is trained to follow the target (red). If the training is successful, the same external stimulus can elicit the desired response. Bottom shows the spike rates of 100 neurons. **(c)** Top, The Pearson correlation between the actual synaptic drive and the target output during training trials. Bottom, The matrix (Frobenius) norm of changes in recurrent connectivity normalized to initial connectivity during training. **(d)** Filtered spike train of 10 neurons before, during and after training. As in (b), external stimulus (blue) is applied immediately before training trials. Filtered spike train (black) learns to follow the target spiking rate (red) with large errors during the early trials. Applying the stimulus to a successfully trained network elicits the desired spiking rate patterns in every neuron. **(e)** Top, Same as in (c) but measures the correlation between filtered spike trains and target outputs. Bottom, Same as in (c).

training, and, when stimulated with the appropriate cue, the network recapitulated the specified trained pattern of recurrent dynamics, regardless of initial activity. The synaptic drive and the spiking rate were both able to learn multiple target patterns (**Fig. 2**).

1.1 Learning arbitrary patterns of activity

To demonstrate that spiking networks can encode recurrent dynamics that have arbitrary spatiotemporal patterns, we considered targets generated from various families of functions; examples include complex periodic functions, chaotic trajectories, and Ornstein-Uhlenbeck (OU) noise. We randomly selected N different target patterns from one of the families to create a set of heterogeneous targets, and trained the




Figure 2: Learning multiple target patterns. **(a)** The synaptic drive of neurons learns two different target outputs. Blue stimulus evokes the first set of target outputs (red) and the green stimulus evokes the second set of target outputs (red). **(b)** The spiking rate of individual neurons learns two different target outputs.

synaptic drive of a network consisting of N neurons to learn the target dynamics.

As we will show more rigorously in Section 2, we identified two sufficient conditions on the dynamical state and spatiotemporal structure of target dynamics that ensure a wide repertoire of recurrent dynamics can be learned. The first is a “quasi-static” condition that stipulates that the dynamical time scale of target patterns must be slow enough compared to the synaptic time scale and average spiking rate. The second is a “heterogeneity” condition that requires the spatiotemporal structure of target patterns to be diverse enough. We will quantify these conditions below. The target patterns considered in **Figure 3** had slow temporal dynamics in comparison to the synaptic time constant ($\tau_s = 20$ ms) and the patterns were selected randomly to promote diverse structure. After training each neuron’s synaptic drive to produce the respective target pattern, the synaptic drive of every neuron in the network followed its target.

To verify the quasi-static condition, we compared the actual and a quasi-static approximation of the spiking rate and synaptic drive. The spiking rates of neurons were approximated using the current-to-rate transfer function with time-dependent synaptic input, and the synaptic drive was approximated by a weighted sum of the presynaptic neurons’ spiking rates. We elicited the trained patterns over multiple trials starting at random initial conditions to calculate the trial-averaged spiking rates. The quasi-static approximations of the synaptic drive and spiking rate closely matched the actual synaptic drive (**Fig. 3a**) and trial-averaged spiking rates (**Fig. 3b**).

To examine how the heterogeneity of target patterns may facilitate learning, we created sets of target patterns where the fraction of randomly generated targets was varied systematically. For non-random targets, we used the same target pattern repeatedly. Networks trained to learn target patterns with strong heterogeneity showed that a network is able to encode target patterns with high accuracy if there is a large fraction of random targets (**Fig. 3c**). Networks that are trained on too many repeated target patterns failed to learn. Beyond a certain fraction of random patterns, including additional patterns did not improve the performance, suggesting that the set of basis functions was over-complete. We probed the stability of over-complete networks under neuron loss by eliminating all the synaptic connections from a fraction of neurons. A network was first trained to learn target outputs where all the patterns were selected randomly (i.e. fraction of random targets equals 1) to ensure that the target patterns form a set of a redundant basis functions. Then, we elicited the trained patterns after removing a fraction of neurons from the network, which entails eliminating all the synaptic connections from the lost neurons. A trained network with 5% neuron loss was able to generate the trained patterns perfectly, 10% neuron loss resulted in a mild degradation of network response, and trained patterns completely disappeared after 40% neuron loss (**Fig. 3d**).

The target dynamics considered in **Figure 3** had population spiking rates of 9.1 Hz (periodic), 7.2 Hz (chaotic) and 12.1 Hz (OU) within the training window. To examine how population activity may influence learning, we trained networks to learn target patterns whose average amplitude was reduced gradually across target sets. The networks were able to learn when the population spiking rate of the target dynamics was as low as 1.5 Hz. However, the performance deteriorated as the population spiking rate decreased further (**Supplementary Fig. 2**). To demonstrate that learning does not depend on the spiking mechanism, we trained the synaptic drive of spiking networks using different neuron models. A network of leaky integrate-and-fire neurons, as well as a network of Izhikevich neurons whose neuron parameters were tuned to have five different firing patterns, successfully learned complex synaptic drive patterns (**Supplementary Fig. 3**).




Figure 3: Quasi-static and heterogeneous patterns can be learned. Example target patterns include complex periodic functions (product of sines with random frequencies), chaotic rate units (obtained from a random network of rate units), and OU noise (obtained by low pass filtering white noise with time constant 100 ms). **(a)** Target patterns (red) overlaid with actual synaptic drive (black) of a trained network. Quasi-static prediction (equation (1)) of synaptic drive (blue). **(b)** Spike trains of trained neurons elicited multiple trials, trial-averaged spiking rate calculated by the average number of spikes in 50 ms time bins (black), and predicted spiking rate (blue). **(c)** Performance of trained network as a function of the fraction of randomly selected targets. **(d)** Network response from a trained network after removing all the synaptic connections from 5%, 10% and 40% of randomly selected neurons in the network.

1.2 Stabilizing rate fluctuations in balanced networks

A random network with balanced excitation and inhibition is a canonical model for a cortical circuit that produces asynchronous single unit activity^{6,27–29,36}. The chaotic activity of balanced rate models⁶ has been harnessed to accomplish complex tasks by including a feedback loop⁸, stabilizing chaotic trajectories⁹ or introducing low-rank structure to the connectivity matrix³⁷. Balanced spiking networks have been shown to possess similar capabilities^{12–14,16,35}, but it is unknown if it is possible to make use of the heterogeneous fluctuations of the spiking rate in the strong coupling regime³⁶. Here, we extend the work of Laje and Buonomano⁹ to spiking networks and show that strongly fluctuating single neuron activities can be turned into dynamic attractors by adjusting the recurrent connectivity.

We take a network with randomly connected excitatory and inhibitory neurons that respects Dale's Law. Prior to training, the synaptic and spiking activity of individual neurons show large variations across trials because small discrepancies in the initial network state lead to rapid divergence of network dynamics. When simulated with two different initial conditions, the synaptic drive to neurons deviates

strongly from each other (**Fig. 4a**), and the spiking activity of single neurons is uncorrelated across trials and the trial-averaged spiking rate has little temporal structure (**Fig. 4b**). The balanced network exhibited sensitivity to small perturbation; the microstate of two identically prepared networks diverged rapidly if one spike is deleted from one of the networks (**Fig. 4c**). It has been previously questioned as to whether the chaotic nature of the balanced state could be utilized to perform reliable computations^{25,38}.

As in Laje and Buonomano⁹, we sought to tame the chaotic trajectories of single neuron activities when the coupling strength is strong enough to induce large and irregular spiking rate fluctuations in time and across neurons³⁶. We initiate the untrained network with random initial conditions to harvest innate synaptic activity, i.e. a set of synaptic trajectories that the network already knows how to generate. Then, the recurrent connectivity is trained so that the synaptic drive of every neuron in the network follows the innate pattern when stimulated with an external stimulus. To respect Dale's Law, the RLS learning rule is modified such that it does not update synaptic connections if there are changes in their signs.

After training, the synaptic drive to every neuron in the network is able to track the innate trajectories in response to the external stimulus within the trained window and diverge from the target pattern outside the trained window (**Fig. 4d**). When the trained network is stimulated to evoke the target patterns, the trial-averaged spiking rate develops a temporal structure that is not present in the untrained network (**Fig. 4e**). To verify the reliability of learned spiking patterns, we simulated the trained network twice with identical initial conditions but deleted one spike 200 ms after evoking the trained response from one of the simulations. Within the trained window, the relative deviation of the microstate is markedly small in comparison to the deviation observed in the untrained network. Outside the trained window, however, two networks diverge rapidly again, which demonstrates that training the recurrent connectivity creates an attracting flux tube around what used to be chaotic spike sequences²⁵ (**Fig. 4f, g**). Analyzing the eigenvalue spectrum of the recurrent connectivity reveals that the distribution of eigenvalues shifts towards zero and the spectral radius decreases as a result of training, which is consistent with the more stable network dynamics found in trained networks.

To demonstrate that learning the innate trajectories works well when a balanced network satisfies the quasi-static condition, we scanned the coupling strength J (see Figure Methods for definition of J) and synaptic time constant τ_s over a wide range and evaluated the accuracy of the quasi-static approximation in untrained networks. We find that increasing either J or τ_s promotes strong fluctuations in spiking rates^{36,39}, hence improving the quasi-static approximation (**Fig. 4i**). Learning performance was correlated with adherence to the quasi-static approximation, resulting in better performance for strong coupling and long synaptic time constants.

1.3 Generating an ensemble of *in-vivo* spiking patterns

We next investigated if the training method applies to actual spike recordings of a large number of neurons. In a previous study, a network of rate units was trained to match sequential activity imaged from posterior parietal cortex as a possible mechanism for short-term memory^{3,32}. Here, we aimed to construct recurrent spiking networks that capture heterogeneous spiking activity of cortical neurons involved in motor planning and movement^{1,2,5}.

The *in-vivo* spiking data was obtained from the publicly available data of Li et al.⁵, where they recorded the spike trains of a large number of neurons from the anterior lateral motor cortex of mice engaged in planning and executing directed licking over multiple trials. We compiled the trial-average spiking rate of $N_{\text{cor}} = 227$ cortical neurons from their data set⁴⁰, and trained a recurrent network model to reproduce the spiking rate patterns of all the N_{cor} neurons autonomously in response to a brief external stimulus. We only trained the recurrent connectivity and did not alter single neuron dynamics or external inputs.

First, we tested if a recurrent network of size N_{cor} is able to generate the spiking rate patterns of the same number of cortical neurons. This network model assumes that the spiking patterns of N_{cor} cortical neurons can be self-generated within a recurrent network. After training, the spiking rate of neuron models captured the overall trend of spiking rate, but not the rapid changes that may be pertinent to the short term memory and motor response (**Fig. 5b**). We hypothesized that the discrepancy may be attributed to other sources of input to the neurons not included in the model, such as recurrent input from other neurons in the local population or input from other areas of the brain, or the neuron dynamics that cannot be captured by our neuron model. We thus sought to improve the performance by adding N_{aux} auxiliary neurons to the recurrent network to mimic the spiking activity of unobserved neurons in the




Figure 4: Learning innate activity in balanced networks. **(a)** Synaptic drive of sample neurons starting at random initial conditions in response to external stimulus prior to training. **(b)** Spike raster of sample neurons evoked by the same stimulus over multiple trials with random initial conditions. **(c)** Single spike perturbation of an untrained network. **(d)-(f)** Synaptic drive, multi-trial spiking response and single spike perturbation in a trained network. **(g)** The average phase deviation of theta neurons due to single spike perturbation. **(h)** Left, distribution of eigenvalues of the recurrent connectivity before and after training as a function their absolute values. Right, Eigenvalue spectrum of the recurrent connectivity; gray circle has unit radius. **(i)** The accuracy of quasi-static approximation in untrained networks and the performance of trained networks as a function of coupling strength J and synaptic time constant τ_s . Color bar shows the Pearson correlation between predicted and actual synaptic drive in untrained networks (left) and innate and actual synaptic drive in trained networks (right).

local population, and trained the recurrent connectivity of a network of size $N_{\text{cor}} + N_{\text{aux}}$ (**Fig. 5a**). The auxiliary neurons were trained to follow spiking rate patterns obtained from an OU process and provided heterogeneity to the overall population activity patterns. When $N_{\text{aux}}/N_{\text{cor}} \geq 2$, the spiking patterns of neuron models accurately fit that of cortical neurons (**Fig. 5c**), and the population activity of all N_{cor} cortical neurons was well captured by the network model (**Fig. 5d**). The fit to cortical activity improved gradually as a function of the fraction of auxiliary neurons in the network due to increased heterogeneity in the target patterns (**Fig. 5e**)

To verify that the cortical neurons in the network model are not simply driven by the feedforward inputs from the auxiliary neurons, we randomly shuffled a fraction of recurrent connections between cortical neurons after a successful training. The fit to cortical data deteriorated as the fraction of shuffled synaptic connections between cortical neurons was increased, which confirmed that the recurrent connections between the cortical neurons played a role in generating the spiking patterns (**Fig. 5f**).




Figure 5: Generating *in-vivo* spiking activity in a subnetwork of a recurrent network. **(a)** Network schematic showing cortical (black) and auxiliary (white) neuron models trained to follow the spiking rate patterns of cortical neurons and target patterns derived from OU noise, respectively. Multi-trial spike sequences of sample cortical and auxiliary neurons in a successfully trained network. **(b)** Trial-averaged spiking rate of cortical neurons (red) and neuron models (black) when no auxiliary neurons are included. **(c)** Trial-averaged spiking rate of cortical and auxiliary neuron models when $N_{aux}/N_{cor} = 2$. **(d)** Spiking rate of all the cortical neurons from the data (left) and the recurrent network model (right) trained with $N_{aux}/N_{cor} = 2$. **(e)** The fit to cortical dynamics improves as the number of auxiliary neurons increases. **(f)** Random shuffling of synaptic connections between cortical neuron models degrades the fit to cortical data. Error bars show the standard deviation of results from 10 trials.

2 Sufficient conditions for learning

Here we quantify the sufficient conditions the target patterns need to satisfy in order to be successfully encoded in a network. The first condition is that the target patterns must be sufficiently slow compared to the dynamical time scale of both neurons and synapses, such that targets can be considered constant (quasi-static) on a short time interval. As we show below, this is not overly restrictive. In terms of network dynamics, the quasi-static condition implies that the synaptic and neuron dynamics operate as if in a stationary state even though the stationary values change as the network activity evolves in time. In this quasi-static state, we can use a mean field description of the spiking dynamics to derive a self-consistent equation that captures the time-dependent synaptic and spiking activity of neurons^{22,36,41} (see

Supplementary Notes). Under the quasi-static approximation, the synaptic drive satisfies

$$U_i(t) = \sum_{j=1}^N W_{ij} \phi(U_j(t) + I_j), \quad (1)$$

and the spiking rate $R_i = \phi(U_i + I_i)$ satisfies

$$R_i(t) = \phi\left(\sum_{j=1}^N W_{ij} R_j(t)\right), \quad (2)$$

where ϕ is the current-to-rate transfer (i.e. gain) function and I_i is a constant external input.

The advantage of operating in a quasi-static state is that both measures of network activity become conducive to learning new patterns. First, equation (1) is closed in terms of U , which implies that training the synaptic drive is equivalent to training a rate-based network. Second, the RLS algorithm can efficiently optimize the recurrent connectivity W , thanks to the linearity of equation (1) in W , while the synaptic drive closely follows the target patterns as shown in **Fig. 1b**. The spiking rate also provides a closed description of the network activity, as described in equation (2). However, due to nonlinearity in W , it learns only when the total input to a neuron is supra-threshold, i.e. the gradient of ϕ must be positive. For this reason, the learning error cannot be controlled as tightly as the synaptic drive and requires additional trials for successful learning as shown in **Fig. 1d**.

The second condition requires the target patterns to be sufficiently heterogeneous in time and across neurons. Such complexity allows the ensemble of spiking activity to have a rich spatiotemporal structure to generate the desired activity patterns of every neuron within the network. In the perspective of “reservoir computing”^{8,11,42}, every neuron in a recurrent network is considered to be a read-out, and, at the same time, part of the reservoir that is collectively used to produce desired patterns in single neurons. The heterogeneity condition is equivalent to having a set of complete (or over-complete) basis functions, i.e. $\phi(U_j + I_j)$, $j = 1, \dots, N$ in equation (1) and R_j , $j = 1, \dots, N$ in equation (2), to generate the target patterns, i.e. the left hand side of equations (1) and (2). The two conditions are not necessarily independent. Heterogeneous targets also foster asynchronous spiking activity that support quasi-static dynamics.

We note that, although equations (1) and (2) describe the dynamical state in which the learning works well, merely finding W that satisfies one of the equations does not guarantee that a spiking network with recurrent connectivity W will produce the target dynamics. The recurrent connectivity W needs to be trained iteratively as the network dynamics unfold in time to ensure that the target dynamics can be generated in a stable manner⁸. Moreover, the equations explain how the target dynamics can be learned, but does not address how recurrent spiking networks are able to encode such dynamic patterns.

2.1 Characterizing learning error

Learning errors can be classified into two categories. There are *tracking* errors, which arise because the target is not a solution of the true spiking network dynamics and *sampling* errors, which arise from encoding a continuous function with a finite number of spikes. We quantified these learning errors as a function of the network and target time scales. The intrinsic time scale of spiking network dynamics is the synaptic decay constant τ_s , and the time scale of target dynamics is the decay constant τ_c of OU noise. We used target patterns generated from OU noise since the trajectories have a predetermined time scale and their spatio-temporal patterns are sufficiently heterogeneous.

We systematically varied τ_s and τ_c from fast AMPA-like (~ 1 ms) to slow NMDA-like synaptic transmission (~ 100 ms) and trained the synaptic drive of networks with synaptic time scale τ_s to learn a set of OU trajectories with time scale τ_c . The parameter scan reveals a learning regime, where the networks successfully encode the target patterns, and two error-dominant regimes. The tracking error is prevalent when synapses are slow in comparison to target patterns, and the sampling error dominates when the synapse is fast (**Fig. 6a**).

A network with a synaptic decay time $\tau_s = 200$ ms fails to follow rapid changes in the target patterns, but still captures the overall shape, when the target patterns have a faster time scale $\tau_c = 100$ ms (**Fig. 6b, Tracking error**). This prototypical example shows that the synaptic dynamics are not fast enough to encode the target dynamics in the tracking error regime. With a faster synapse $\tau_s = 30$ ms, the synaptic




Figure 6: Sampling and tracking errors. Synaptic drive was trained to learn 1 s long trajectories generated from OU noise with decay time τ_c . **(a)** Performance of networks of size $N = 500$ as a function of synaptic decay time τ_s and target decay time τ_c . **(b)** Examples of trained networks whose responses show sampling error, tracking error, and successful learning. The target trajectories are identical and $\tau_c = 100$ ms. **(c)** Inverted “U”-shaped curve as a function of synaptic decay time. Error bars show the s.d. of five trained networks of size $N = 500$. **(d)** Inverted “U”-shaped curve for networks of sizes $N = 500$ and 1000 for $\tau_c = 100$ ms. **(e)** Network performance shown as a function of τ_s/τ_c where the range of τ_s is from 30 ms to 500 ms and the range of τ_c is from 1 ms to 500 ms and $N = 1000$. **(f)** Network performance shown as a function of $1/\sqrt{N\tau_s}$ where the range of τ_s is from 1 ms to 30 ms, the range of N is from 500 to 1000 and $\tau_c = 100$ ms.

drive is able to learn the identical target trajectories with high accuracy (**Fig. 6b, Learning**). However, when the synapse is too fast $\tau_s = 5$ ms, the synaptic drive fluctuates around the target trajectories with high frequency (**Fig. 6b, Sampling error**). This is a typical network response in the sampling error regime since discrete spikes with narrow width and large amplitude are summed to “sample” the target synaptic activity.

To better understand how network parameters determine the learning errors, we mathematically analyzed the errors assuming that (1) target dynamics can be encoded if the quasi-static condition holds, and (2) the mean field description of the target dynamics is accurate (see **Supplementary Notes**). The learning errors were characterized as a deviation of these assumptions from the actual spiking network dynamics. We found that the tracking errors ϵ_{track} are substantial if the quasi-static condition is not valid, i.e. synapses are not fast enough for spiking networks to encode targets, and the sampling errors ϵ_{sample} occur if the mean field description becomes inaccurate, i.e. discrete representation of targets in terms of spikes deviates from their continuous representation in terms of spiking rates. The errors are estimated to scale with

$$\epsilon_{\text{track}} \sim \tau_s/\tau_c, \quad \epsilon_{\text{sample}} \sim 1/\sqrt{\tau_s N}, \quad (3)$$

which imply that tracking error can be controlled as long as synapses are relatively faster than target




Figure 7: Capacity as a function of network size. **(a)** Performance of trained networks as a function of target length T for networks of size $N = 500$ and 1000 . Target patterns were generated from OU noise with decay time $\tau_c = 100$ ms. **(b)** Networks of fixed sizes trained on a range of target length and correlations. Color bar shows the Pearson correlation between target and actual synaptic drive. The black lines show the function $T_{\max} = \tilde{T}_{\max} \tau_c$ where \tilde{T}_{\max} was fitted to minimize the least square error between the linear function and maximal target length T_{\max} that can be successfully learned at each τ_c . **(c)** Learning capacity \tilde{T}_{\max} shown as a function of network size.

patterns, and the sampling error can be controlled by either increasing τ_s to stretch the width of individual spikes or increasing N to encode the targets with more input spikes. The error estimates reveal the versatility of recurrent spiking networks to encode arbitrary patterns since ϵ_{track} can be reduced by tuning τ_s to be small enough and ϵ_{sample} can be reduced by increasing N to be large enough.

We examined the performance of trained networks to verify if the theoretical results can explain the learning errors. The learning curve, as a function of τ_s , has an inverted U-shape when both types of errors are present (**Fig. 6c, d**). Successful learning occurs in an optimal range of τ_s , and, consistent with the error analysis, the performance decreases monotonically with τ_s on the right branch due to the tracking error while increases monotonically with τ_s on the left branch due to the sampling error. The tracking error is reduced if target patterns are slowed down from $\tau_c = 50$ ms to $\tau_c = 200$ ms, hence decrease the ratio τ_s/τ_c . Then, the learning curve becomes sigmoidal, and the performance remains high even when τ_s is in the slow NMDA regime (**Fig. 6c**). On the other hand, the sampling error is reduced if the network size is increased from $N = 500$ to 1500 , which lifts the left branch of the learning curve (**Fig. 6d**). Note that when two error regimes are well separated, changes in target time scale τ_c does not affect ϵ_{sample} , and changes in network size N does not affect ϵ_{sample} , as predicted.

Finally, we condensed the training results over a wide range of target time scales in the tracking error regime (**Fig. 6e**), and similarly condensed the training results over different network sizes in the sampling error regime (**Fig. 6f**) to demonstrate that τ_s/τ_c and $N\tau_s$ explain the overall performance in the tracking and sampling error regimes, respectively.

2.2 Learning capacity increases with network size

It has been shown that a recurrent rate network’s capability to encode target patterns deteriorates as a function of the length of time⁹, but increase in network size can enhance its storage capacity^{32,43,44}. Consistent with these results, we find that the performance of recurrent spiking networks to learn complex trajectories decreases with target length and improves with network size (**Fig. 7a**).

To assess the storage capacity of spiking networks, we evaluated the maximal target length that can be encoded in a network as a function of network size. It was necessary to define the target length in terms of its “effective length” to account for the fact that target patterns with the same length may have different effective length due to their temporal structures; for instance, OU noise with short temporal correlation times has more structure to be learned than a constant function. For target trajectories generated from an OU process with decay time τ_c , we rescaled the target length T with respect to τ_c and defined the effective length $\tilde{T} = T/\tau_c$. The capacity of a network is the maximal \tilde{T} that can be successfully encoded in a network.

To estimate the maximal \tilde{T} , we trained networks of fixed size to learn OU trajectories while varying T and τ_c (each panel in **Fig. 7b**). Then, for each τ_c , we find the maximal target length T_{\max} that can be learned successfully, and estimate the maximal \tilde{T} by finding a constant \tilde{T}_{\max} that best fits the line $T_{\max} = \tilde{T}_{\max} \tau_c$ to training results (black lines in **Fig. 7b**). **Figure 7c** shows that the learning capacity

\tilde{T}_{\max} increases monotonically with the network size.

We can illustrate why this may be the case with a simple albeit non-rigorous counting argument. Successful learning is achieved for the synaptic drive when equation (1) is satisfied. If we artificially discretize time into P “quasi-static” bins then we can consider $U_i(t)$ as a $N \times P$ matrix that satisfies a matrix equation $U = WV$, where we treat $V \equiv \phi(U + I)$ as an independent matrix. If we operate on the matrix equation with V^T we find that a solution to this artificial system is possible if the matrix VV^T is invertible, which is possible if the rank of VV^T is full. This argument then shows how increasing pattern heterogeneity, which makes the rows of V less correlated, is beneficial for learning, and that learning capability will decline as P increases, with a steep decline for $P > N$. If we ascribe quasi-static bin to some fraction of the pattern correlation time then P will scale with the length of the time interval. In this way, we can intuitively visualize the temporal storage capacity.

Discussion

We have shown that the synaptic drive and spiking rate of a network of spiking neurons can be trained to follow any arbitrary spatiotemporal patterns. The necessary ingredients for learning are that the spike train inputs to a neuron are weakly correlated (i.e. heterogeneous target patterns), the synapses are fast enough (i.e. small tracking error), and the network is large enough (i.e. small sampling error and large capacity). We demonstrated that (1) a balanced network consisting of excitatory and inhibitory neurons can learn to track its strongly fluctuating innate synaptic trajectories, and (2) a recurrent spiking network can learn to reproduce the spiking rate patterns of an ensemble of cortical neurons involved in motor planning and movement.

Our scheme works because the network quickly enters a quasi-static state where the instantaneous firing rate of a neuron is a fixed function of the inputs. Learning fails if the synaptic time scale is slow compared to the time scale of the target, in which case the quasi-static condition is violated and the tracking error becomes large. There is a trade-off between tracking error and sampling noise; fast synapse can decrease the tracking error, but it also increases the sampling noise. Increasing the network size can decrease sampling noise without affecting the tracking error. Therefore, by adjusting the synaptic time scale and network size, it becomes possible to learn arbitrarily complex recurrent dynamics.

An important structural property of our network model is that the synaptic inputs are summed linearly, which allows the synaptic activity to be trained using a recursive form of linear regression⁸. Linear summation of synaptic inputs is a standard assumption for many spiking network models^{21,22,27–29} and there is physiological evidence that linear summation is prevalent^{45,46}. Training the spiking rate, on the other hand, does not take full advantage of the linear synapse due to the nonlinear current-to-transfer function. The network is capable of following a wide repertoire of patterns because even though the network dynamics are highly nonlinear, the system effectively reduces to a linear system for learning. Moreover, learning capacity can be estimated using a simple solvability condition for a linear system. However, nonlinear dendritic processing has been widely observed^{47,48} and may have computational consequences^{14,49,50}. It requires further investigation to find out whether a recurrent network with nonlinear synapses can be trained to learn arbitrary recurrent dynamics.

We note that our training scheme does not teach precise spike timings; it either trains the spiking rate directly or trains the synaptic drive which in turn trains the spiking rate. The spike times can be divergent from instance to instance, hence our learning scheme supports rate coding as opposed to spike coding; we demonstrated that the same learning scheme can be used to train the recurrent dynamics of rate-based networks. We find that spike trains that have temporally irregular structure across neurons enhance the rate coding scheme by providing sufficient computational complexity to encode the target dynamics. All neurons in the network can be trained to follow the same synaptic drive patterns as long as there is sufficient heterogeneity, e.g. noisy external input, that decorrelates the spike trains of the neurons (**Supplementary Fig. S5**).

Although our results confirm that recurrent spiking networks have the capability to generate a wide range of repertoire of recurrent dynamics, it is unlikely that a biological network is using this particular learning scheme. The learning rule derived from recursive least squares algorithm is very effective but is nonlocal in time, i.e. it uses the activity of all presynaptic neurons within the train time window to update synaptic weights. Moreover, each neuron in the network is assigned with a target signal and the synaptic connections are updated at a fast time scale as the error function is computed in a supervised

manner. It would be of interest to find out whether more biologically plausible learning schemes, such as reward-based learning⁵¹⁻⁵³, can lead to similar performance.

Our study provided conditions under which a spiking network can learn a wide range of target dynamics. Previous studies investigated the mathematical relationship between the patterns of stored fixed points and the recurrent connectivity in simple network models^{54,55}. It remains an open question to classify possible target patterns that can be encoded in a spiking network and identify the class of connectivity matrices that generates the target patterns in a dynamically stable manner.

Methods

1 Network of spiking neurons

We considered a network of N randomly and sparsely connected quadratic integrate-and-fire neurons given by

$$\tau \dot{v}_i = I_i(t) + u_i(t) + v_i^2 \quad (4)$$

where v_i is a dimensionless variable representing membrane potential, $I_i(t)$ is an applied input, $u_i(t)$ is the total synaptic drive the neuron receives from other neurons in the recurrent network, and $\tau = 10$ ms is a neuron time constant. The threshold to spiking is zero input. For negative total input, the neuron is at rest and for positive input, v_i will go to infinity or “blow up” in finite time from any initial condition. The neuron is considered to spike at $v_i = \infty$ whereupon it is reset to $-\infty$ ^{56,57}.

To simulate the dynamics of quadratic integrate-and-fire neurons, we used its phase representation, i.e. theta neuron model, that can be derived by a simple change of variables, $v_i = \tan(\theta_i/2)$; its dynamics are governed by

$$\tau \dot{\theta}_i = 1 - \cos \theta_i + (I_i(t) + u_i(t))(1 + \cos \theta_i), \quad (5)$$

where a spike is emitted when $\theta(t) = \pi$. The synaptic drive to a neuron obeys

$$\tau_s \dot{u}_i(t) = -u_i(t) + \sum_{j=1}^N W_{ij} s_j(t), \quad (6)$$

where $s_j(t) = \sum_{t_j^k < t} \delta(t - t_j^k)$ is the spike train neuron j generates up to time t , and τ_s is a synaptic time constant.

The recurrent connectivity W_{ij} describes the synaptic coupling from neuron j to neuron i . It can be any real matrix but in many of the simulations we use a random matrix with connection probability p , and the coupling strength of non-zero elements is modeled differently for different figures.

2 Training recurrent dynamics

To train the synaptic and spiking rate dynamics of individual neurons, it is more convenient to divide the synaptic drive equation (6) into two parts; one that isolates the spike train of single neuron and computes its synaptic filtering

$$\tau_s \dot{r}_i(t) = -r_i(t) + s_i(t), \quad (7)$$

and the other that combines all the presynaptic neurons’ spiking activity and computes the synaptic drive

$$u_i(t) = \sum_{j=1}^N W_{ij} r_j(t). \quad (8)$$

The synaptic drive u_i and the filtered spike train r_i are two measures of spiking activity that have been trained in this study. Note that equations (7) and (8) generate synaptic dynamics that are equivalent to equation (6).

Training procedure. We select N target trajectories $f_1(t), \dots, f_N(t)$ of length T ms for a recurrent network consisting of N neurons. We train either the synaptic drive or spiking rate of individual neuron

i to follow the target $f_i(t)$ over time interval $[0, T]$ for all $i = 1, \dots, N$. External stimulus I_i with amplitude sampled uniformly from $[-1, 1]$ is applied to neuron i for all $i = 1, 2, \dots, N$ for 100 ms immediately preceding the training to situate the network at a specific state. During training, the recurrent connectivity W is updated every Δt ms using a learning rule described below in order to steer the network dynamics towards the target dynamics. The training is repeated multiple times until changes in the recurrent connectivity stabilize.

2.1 Training synaptic drive

We modify the recurrent learning rule that was developed to train the recurrent connectivity of a network of rate units using Recursive Least Squares (RLS) algorithm^{8,9}. Our learning rule extends the RLS algorithm to training the recurrent connectivity of spiking networks.

When learning the synaptic drive patterns, the objective is to find recurrent connectivity W that minimizes the cost function

$$C[W] = \int_0^T \frac{1}{2} \|\mathbf{f}(t) - \mathbf{u}(t)\|_{L_2}^2 dt + \frac{\lambda}{2} \|W\|_{L_2}^2, \quad (9)$$

which measures the mean-square error between the targets and the synaptic drive over the time interval $[0, T]$ plus a quadratic regularization term. To derive the learning rule, we use equation (8) to express \mathbf{u} as a function of W , view the synaptic connections W_{i1}, \dots, W_{iN} to neuron i to be the read-out weights that determine the synaptic drive u_i , and apply the learning rule to the row vectors of W . To keep the recurrent connectivity sparse, learning occurs only on synaptic connections that are non-zero prior to training (see **Supplementary Notes** for details).

Let $\mathbf{w}_i(t)$ be the reduced row vector of $W(t)$ consisting of elements that have non-zero connections to neuron i prior to training. Similarly, let $\mathbf{r}_i(t)$ be a (column) vector of filtered spikes of presynaptic neurons that have non-zero connections to neuron i . The synaptic update to neuron i is

$$\mathbf{w}_i(t)^T = \mathbf{w}_i(t - \Delta t)^T + e_i(t)P(t)\mathbf{r}_i(t), \quad (10)$$

where the error term is

$$e_i(t) = f_i(t) - \mathbf{w}_i(t - \Delta t)\mathbf{r}_i(t) \quad (11)$$

and the inverse of the correlation matrix of filtered spike trains is

$$P(t) = P(t - \Delta t) - \frac{P(t - \Delta t)\mathbf{r}_i(t)\mathbf{r}_i(t)^T P(t - \Delta t)}{1 + \mathbf{r}_i(t)^T P(t - \Delta t)\mathbf{r}_i(t)}, \quad P(0) = \lambda^{-1}I. \quad (12)$$

Finally, $W(t)$ is obtained by concatenating the row vectors $\mathbf{w}_i(t)$, $i = 1, \dots, N$.

2.2 Training spiking rate

To train the spiking rate of neurons, we approximate the spike train $s_i(t)$ of neuron i with its spiking rate $\phi(u_i(t) + I_i)$ where ϕ is the current-to-rate transfer function of theta neuron model. For constant input,

$$\phi_1(x) = \pi^{-1} \sqrt{[x]_+} \quad \text{where} \quad [x]_+ = \max(x, 0), \quad (13)$$

and for noisy input

$$\phi_2(x) = \frac{1}{\pi} \sqrt{c \log(1 + e^{x/c})}. \quad (14)$$

Since ϕ_2 is a good approximation of ϕ_1 and has a smooth transition around $x = 0$, we used $\phi \equiv \phi_2$ with $c = 0.1$ ⁵⁸. The objective is to find recurrent connectivity W that minimizes the cost function

$$C[W] = \int_0^T \frac{1}{2} \|\mathbf{f}(t) - \phi(W\mathbf{r}(t) + \mathbf{I})\|_{L_2}^2 dt + \frac{\lambda}{2} \|W\|_{L_2}^2. \quad (15)$$

If we define \mathbf{w}_i and \mathbf{r}_i as before, we can derive the following synaptic update to neuron i

$$\mathbf{w}_i^T(t) = \mathbf{w}_i^T(t - \Delta t) + e_i(t)P(t)\tilde{\mathbf{r}}_i(t), \quad (16)$$

where the error term is

$$e_i(t) = f_i(t) - \phi(\mathbf{w}_i(t - \Delta t)\mathbf{r}_i(t) + I_i) \quad (17)$$

and

$$P(t) = P(t - \Delta t) - \frac{P(t - \Delta t)\tilde{\mathbf{r}}_i(t)\tilde{\mathbf{r}}_i(t)^T P(t - \Delta t)}{1 + \tilde{\mathbf{r}}_i(t)^T P(t - \Delta t)\tilde{\mathbf{r}}_i(t)}, \quad P^0 = \lambda^{-1}I. \quad (18)$$

(see **Supplementary Notes** for details). Note that the nonlinear effects of the transfer function is included in

$$\tilde{\mathbf{r}}_i(t) = \phi'(u_i(t) + I_i)\mathbf{r}_i(t), \quad (19)$$

which scales the spiking activity of neuron i by its gain function ϕ' .

As before, $W(t)$ is obtained by concatenating the row vectors $\mathbf{w}_i(t)$, $i = 1, \dots, N$.

3 Simulation parameters

Computer code. Example code is available at <http://github.com/chrisrkkim/SpikeLearning>

Figure 1. A network of $N = 200$ neurons was connected randomly with probability $p = 0.3$ and the coupling strength was drawn from a Normal distribution with mean 0 and standard deviation σ/\sqrt{Np} with $\sigma = 4$. In addition, the average of all non-zero synaptic connections to a neuron was subtracted from the connections to the neuron such that the summed coupling strength was precisely zero. Networks with balanced excitatory and inhibitory connections produced highly fluctuating synaptic and spiking activity in all neurons. The synaptic decay time was $\tau_s = 20$ ms.

The target functions for the synaptic drive (**Fig. 1b**) were sine waves $f(t) = A \sin(2\pi(t - T_0)/T_1)$ where the amplitude A , initial phase T_0 , and period T_1 were sampled uniformly from $[0.5, 1.5]$, $[0, 1000$ ms] and $[300$ ms, 1000 ms], respectively. We generated N distinct target functions of length $T = 1000$ ms. The target functions for the spiking rate (**Fig. 1d**) were $\pi^{-1}\sqrt{[f(t)]_+}$ where $f(t)$ were the same synaptic drive patterns that have been generated.

Immediately before each training loop, every neuron was stimulated for 50 ms with constant external stimulus that had random amplitude sampled from $[-1, 1]$. The same external stimulus was used across training loops. The recurrent connectivity was updated every $\Delta t = 2$ ms during training using the learning rule derived from RLS algorithm and the learning rate was $\lambda = 1$. After training, the network was stimulated with the external stimulus to evoke the trained patterns. The performance was measured by calculating the average Pearson correlation between target functions and the evoked network response.

Figure 2. The initial network and target functions were generated as in **Figure 1** using the same parameters, but now the target functions consisted of two sets of N sine waves. To learn two sets of target patterns, the training loops alternated between two patterns, and immediately before each training loop, every neuron was stimulated for 50 ms with constant external stimuli that had random amplitudes, using a different stimulus for each pattern. Each target pattern was trained for 100 loops (i.e. total 200 training loops), synaptic update was every $\Delta t = 2$ ms, and the learning rate was $\lambda = 10$. To evoke one of the target patterns after training, the network was stimulated with the external stimulus that was used to train that target pattern.

Figure 3. The network consisted of $N = 500$ neurons. The initial connectivity was sparsely connected with connection probability $p = 0.3$ and coupling strength was sampled from a Normal distribution with mean 0 and standard deviation σ/\sqrt{Np} with $\sigma = 1$. The synaptic decay time was $\tau_s = 20$ ms.

We considered three families of target functions with length $T = 1000$ ms. The complex periodic functions were defined as a product of two sine waves $f(t) = A \sin(2\pi(t - T_0)/T_1) \sin(2\pi(t - T_0)/T_2)$ where A , T_0 , T_1 and T_2 were sampled randomly from intervals $[0.5, 1.5]$, $[0, 1000$ ms], $[500$ ms, 1000 ms], and $[100$ ms, 500 ms], respectively. The chaotic rate activity was generated from a network of N randomly connected rate units, $\tau\dot{x}_i = -x_i + \sum_{j=1}^N M_{ij}h(x_j)$ where $\tau = 40$ ms, $h(x) = \pi^{-1}\sqrt{[x]_+}$ and

M_{ij} is non-zero with probability $p = 0.3$ and is drawn from Gaussian distribution with mean zero and standard deviation g/\sqrt{Np} with $g = 5$. The Ornstein-Uhlenbeck process was obtained by simulating, $\tau_c \dot{x} = -x + s\xi(t)$, N times with random initial conditions and different realizations of the white noise $\xi(t)$ satisfying $\langle \xi \rangle = 0$ and $\langle \xi(t)\xi(t') \rangle = \delta(t - t')$. The decay time constant was $\tau_c = 200$ ms, and the amplitude of target function was determined by $s = 0.3$.

The recurrent connectivity was updated every $\Delta t = 2$ ms during training, the learning rate was $\lambda = 1$, and the training loop was repeated 30 times.

Figure 4. A balanced network had two populations where the excitatory population consisted of $(1 - f)N$ neurons and the inhibitory population consisted of fN neurons with ratio $f = 0.2$ and network size $N = 1000$. Each neuron received $p(1 - f)N$ excitatory connections with strength J and pfN inhibitory connections with strength $-gJ$ from randomly selected excitatory and inhibitory neurons. The connection probability was set to $p = 0.1$ to have sparse connectivity. The relative strength of inhibition to excitation g was set to 5 so that the network was inhibition dominant²². In **Figure 4a-h**, the initial coupling strength $J = 6$ and synaptic decay time $\tau_s = 60$ ms were adjusted to be large enough, so that the synaptic drive and spiking rate of individual neurons fluctuated strongly and slowly prior to training.

After running the initial network that started at random initial conditions for 3 seconds, we recorded the synaptic drive of all neurons for 2 seconds to harvest target trajectories that are innate to the balanced network. Then, the synaptic drive was trained to learn the innate trajectories, where synaptic update occurred every 10 ms, learning rate was $\lambda = 10$ and training loop was repeated 40 times. To respect Dale's Law while training the network, we did not modify the synaptic connections if the synaptic update reversed the sign of original connections, either from excitatory to inhibitory or from inhibitory to excitatory. Moreover, the synaptic connections that attempted to change their signs were excluded in subsequent trainings. In **Figure 4h**, the initial and trained connectivity matrices were normalized by a factor $\sqrt{[(1 - f)J^2 + f(gJ)^2](1 - p)}$ so that the spectral radius of the initial connectivity matrix is approximately 1, then we plotted the eigenvalue spectrum of the normalized matrices.

In **Figure 4i**, the coupling strength J was scanned from 1 to 6 in increments of 0.25, and the synaptic decay time τ_s was scanned from 5 ms to 100 ms in increments of 5 ms. To measure the accuracy of quasi-static approximation in untrained networks, we simulated the network dynamics for each pair of J and τ_s , then calculated the average Pearson correlation between the predicted synaptic drive (equation (1)) and the actual synaptic drive. To measure the performance of trained networks, we repeated the training 10 times using different initial network configurations and innate trajectories, and calculated the Pearson correlation between the innate trajectories and the evoked network response for all 10 trainings. The heat map shows the best performance out of 10 trainings for each pair, J and τ_s .

Figure 5. The initial connectivity was sparsely connected with connection probability $p = 0.3$ and the coupling strength was sampled from a Normal distribution with mean 0 and standard deviation σ/\sqrt{Np} with $\sigma = 1$. The synaptic decay time was $\tau_s = 50$ ms. There were in total N neurons in the network model, of which N_{cor} neurons, called cortical neurons, were trained to learn the spiking rate patterns of cortical neurons, and N_{aux} neurons, called auxiliary neurons, were trained to learn trajectories generated from OU process.

We used the trial-averaged spiking rates of neurons recorded in the anterior lateral motor cortex of mice engaged in motor planning and movement that lasted 4600 ms⁵. The data was available from the website CRCNS.ORG⁴⁰. We selected $N_{\text{cor}} = 227$ neurons from the data set, whose average spiking rate during the behavioral task was greater than 5 Hz. Each cortical neuron in the network model was trained to learn the spiking rate pattern of one of the real cortical neurons.

To generate target rate functions for the auxiliary neurons, we simulated an OU process, $\tau_c \dot{x}(t) = -x(t) + s\xi(t)$, with $\tau_c = 800$ ms and $s = 0.1$, then converted into spiking rate $\phi([x(t)]_+)$ and low-pass filtered with decay time τ_s to make it smooth. Each auxiliary neuron was trained on 4600 ms-long target rate function that was generated with a random initial condition.

Figures 6, 7. Networks consisting of $N = 500$ neurons with no initial connections and synaptic decay time τ_s were trained to learn OU process with decay time τ_c and length T . In **Figure 6**, target length was fixed to $T = 1000$ ms while the time constants τ_s and τ_c were varied systematically from 10^0 ms to $5 \cdot 10^2$ ms in log-scale. The trainings were repeated 5 times for each pair of τ_s and τ_c to find the average performance. In **Figure 7**, the synaptic decay time was fixed to $\tau_s = 20$ ms and T was scanned from

250 ms to 5000 ms in increments of 250 ms, τ_c was scanned from 25 ms to 500 ms in increments of 25 ms, and N was scanned from 500 to 1000 in increments of 50.

To ensure that the network connectivity after training is sparse, synaptic learning occurred only on connections that were randomly selected with probability $p = 0.3$ prior to training. Recurrent connectivity was updated every $\Delta t = 2$ ms during training, learning rate was $\lambda = 1$, and training loop was repeated 30 times. The average Pearson correlation between the target functions and the evoked synaptic activity was calculated to measure the network performance after training.

Acknowledgments

This research was supported [in part] by the Intramural Research Program of the NIH, The National Institute of Diabetes and Digestive and Kidney Diseases (NIDDK).

References

- [1] Mark M Churchland and Krishna V Shenoy. Temporal complexity and heterogeneity of single-neuron activity in premotor and motor cortex. *Journal of neurophysiology*, 97(6):4235–4257, 2007.
- [2] Mark M Churchland, John P Cunningham, Matthew T Kaufman, Justin D Foster, Paul Nuyujukian, Stephen I Ryu, and Krishna V Shenoy. Neural population dynamics during reaching. *Nature*, 487(7405):51–56, 2012.
- [3] Christopher D Harvey, Philip Coen, and David W Tank. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature*, 484(7392):62, 2012.
- [4] Eva Pastalkova, Vladimir Itskov, Asohan Amarasingham, and György Buzsáki. Internally generated cell assembly sequences in the rat hippocampus. *Science*, 321(5894):1322–1327, 2008.
- [5] Nuo Li, Tsai-Wen Chen, Zengcai V Guo, Charles R Gerfen, and Karel Svoboda. A motor cortex circuit for motor planning and movement. *Nature*, 519(7541):51–56, 2015.
- [6] Sompolinsky, Crisanti, and Sommers. Chaos in random neural networks. *Physical review letters*, 61(3):259–262, jul 1988. ISSN 1079-7114. doi: 10.1103/PhysRevLett.61.259.
- [7] Wolfgang Maass, Prashant Joshi, and Eduardo D Sontag. Computational aspects of feedback in neural circuits. *PLoS computational biology*, 3(1):e165, 2007.
- [8] David Sussillo and L. F. Abbott. Generating Coherent Patterns of Activity from Chaotic Neural Networks. *Neuron*, 63(4):544–557, 2009. ISSN 08966273. doi: 10.1016/j.neuron.2009.07.018.
- [9] Rodrigo Laje and Dean V Buonomano. Robust timing and motor patterns by taming chaos in recurrent neural networks. *Nature Neuroscience*, 16(7):925–933, may 2013. ISSN 1097-6256. doi: 10.1038/nn.3405.
- [10] Dean V. Buonomano and Wolfgang Maass. State-dependent computations: spatiotemporal processing in cortical networks. *Nature Reviews Neuroscience*, 10(2):113–125, feb 2009. ISSN 1471-003X. doi: 10.1038/nrn2558.
- [11] Herbert Jaeger and Harald Haas. Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication. *science*, 304(5667):78–80, 2004.
- [12] Wilten Nicola and Claudia Clopath. Supervised learning in spiking neural networks with force training. *Nature communications*, 8(1):2208, 2017.
- [13] Brian DePasquale, Mark M. Churchland, and L. F. Abbott. Using Firing-Rate Dynamics to Train Recurrent Networks of Spiking Model Neurons. pages 1–17, 2016.
- [14] Dominik Thalmeier, Marvin Uhlmann, Hilbert J Kappen, and Raoul-Martin Memmesheimer. Learning universal computations with spikes. *PLoS Comput Biol*, 12(6):e1004895, 2016.

- [15] Martin Boerlin, Christian K Machens, and Sophie Denève. Predictive coding of dynamical variables in balanced spiking networks. *PLoS computational biology*, 9(11):e1003258, 2013.
- [16] Sophie Denève and Christian K Machens. Efficient codes and balanced networks. *Nature neuroscience*, 19(3):375, 2016.
- [17] Wieland Brendel, Ralph Bourdoukan, Pietro Vertechi, Christian K Machens, and Sophie Denève. Learning to represent signals spike by spike. *arXiv preprint arXiv:1703.03777*, 2017.
- [18] Ralph Bourdoukan and Sophie Deneve. Enforcing balance allows local supervised learning in spiking recurrent networks. In *Advances in Neural Information Processing Systems*, pages 982–990, 2015.
- [19] Sophie Denève, Alireza Alemi, and Ralph Bourdoukan. The brain as an efficient and robust adaptive learner. *Neuron*, 94(5):969–977, 2017.
- [20] John J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8):2554–2558, 1982.
- [21] Xiao-Jing Wang and György Buzsáki. Gamma oscillation by synaptic inhibition in a hippocampal interneuronal network model. *Journal of neuroscience*, 16(20):6402–6413, 1996.
- [22] Nicolas Brunel. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of computational neuroscience*, 8(3):183–208, 2000.
- [23] Nancy Kopell and Bard Ermentrout. Chemical and electrical synapses perform complementary roles in the synchronization of interneuronal networks. *Proceedings of the national academy of sciences of the United States of America*, 101(43):15482–15487, 2004.
- [24] R Ben-Yishai, R Lev Bar-Or, and H Sompolinsky. Theory of orientation tuning in visual cortex. *Proceedings of the National Academy of Sciences*, 92(9):3844–3848, 1995.
- [25] Michael Monteforte and Fred Wolf. Dynamic flux tubes form reservoirs of stability in neuronal circuits. *Physical Review X*, 2(4):041007, 2012.
- [26] David Sussillo, Mark M Churchland, Matthew T Kaufman, and Krishna V Shenoy. A neural network that finds a naturalistic solution for the production of muscle activity. *Nature neuroscience*, 18(7):1025–1033, 2015.
- [27] Carl van Vreeswijk and Haim Sompolinsky. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, 274(5293):1724, 1996.
- [28] Alfonso Renart, Jaime De La Rocha, Peter Bartho, Liad Hollender, Néstor Parga, Alex Reyes, and Kenneth D Harris. The asynchronous state in cortical circuits. *science*, 327(5965):587–590, 2010.
- [29] Robert Rosenbaum, Matthew A Smith, Adam Kohn, Jonathan E Rubin, and Brent Doiron. The spatial structure of correlated neuronal variability. *Nature neuroscience*, 20(1):107, 2017.
- [30] Carsen Stringer, Marius Pachitariu, Nicholas A Steinmetz, Michael Okun, Peter Bartho, Kenneth D Harris, Maneesh Sahani, and Nicholas A Lesica. Inhibitory control of correlated intrinsic variability in cortical networks. *Elife*, 5:e19695, 2016.
- [31] Brian DePasquale, Christopher J Cueva, Kanaka Rajan, LF Abbott, et al. full-force: A target-based method for training recurrent networks. *PloS one*, 13(2):e0191527, 2018.
- [32] Kanaka Rajan, Christopher D. Harvey, and David W. Tank. Recurrent Network Models of Sequence Generation and Memory. *Neuron*, 90(1):128–142, 2016. doi: 10.1016/J.NEURON.2016.02.009.
- [33] Simon Haykin. *Adaptive Filter Theory (3rd Ed.)*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1996. ISBN 0-13-322760-X.

- [34] Nils Bertschinger and Thomas Natschläger. Real-time computation at the edge of chaos in recurrent neural networks. *Neural computation*, 16(7):1413–1436, 2004.
- [35] L F Abbott, Brian DePasquale, and Raoul-martin Memmesheimer. Building functional networks of spiking model neurons. *Nature Neuroscience*, 19(3):350–355, feb 2016. ISSN 1097-6256. doi: 10.1038/nn.4241.
- [36] Srdjan Ostojic. Two types of asynchronous activity in networks of excitatory and inhibitory spiking neurons. *Nature neuroscience*, 17(4):594–600, 2014.
- [37] Francesca Mastrogiuseppe and Srdjan Ostojic. Linking connectivity, dynamics and computations in recurrent neural networks. *arXiv preprint arXiv:1711.09672*, 2017.
- [38] Michael London, Arnd Roth, Lisa Beeren, Michael Häusser, and Peter E Latham. Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex. *Nature*, 466(7302): 123–127, 2010.
- [39] Omri Harish and David Hansel. Asynchronous Rate Chaos in Spiking Neuronal Circuits. *PLoS computational biology*, 11(7):e1004266, 2015. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1004266.
- [40] Nuo Li, Tsai-Wen Chen, Zengcai V. Guo, Charles R. Gerfen, and Karel Svoboda. Extracellular recordings from anterior lateral motor cortex (alm) neurons of adult mice performing a tactile decision behavior. *crcns.org*. 2014.
- [41] Michael A Buice and Carson C Chow. Dynamic finite size effects in spiking neural networks. *PLoS computational biology*, 9(1):e1002872, 2013.
- [42] Wolfgang Maass, Thomas Natschläger, and Henry Markram. Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural computation*, 14(11):2531–2560, 2002.
- [43] Herbert Jaeger. *Short term memory in echo state networks*, volume 5. GMD-German National Research Institute for Computer Science, 2001.
- [44] Olivia L White, Daniel D Lee, and Haim Sompolinsky. Short-term memory in orthogonal neural networks. *Physical review letters*, 92(14):148102, 2004.
- [45] Sydney Cash and Rafael Yuste. Input summation by cultured pyramidal neurons is linear and position-independent. *Journal of Neuroscience*, 18(1):10–15, 1998.
- [46] Sydney Cash and Rafael Yuste. Linear summation of excitatory inputs by ca1 pyramidal neurons. *Neuron*, 22(2):383–394, 1999.
- [47] Sonia Gasparini and Jeffrey C Magee. State-dependent dendritic computation in hippocampal ca1 pyramidal neurons. *Journal of Neuroscience*, 26(7):2088–2100, 2006.
- [48] Thomas Nevian, Matthew E Larkum, Alon Polsky, and Jackie Schiller. Properties of basal dendrites of layer 5 pyramidal neurons: a direct patch-clamp recording study. *Nature neuroscience*, 10(2): 206, 2007.
- [49] Raoul-Martin Memmesheimer. Quantitative prediction of intermittent high-frequency oscillations in neural networks with supralinear dendritic interactions. *Proceedings of the National Academy of Sciences*, 107(24):11092–11097, 2010.
- [50] Raoul-Martin Memmesheimer and Marc Timme. Non-additive coupling enables propagation of synchronous spiking activity in purely random networks. *PLoS computational biology*, 8(4): e1002384, 2012.
- [51] Ila R Fiete and H Sebastian Seung. Gradient learning in spiking neural networks by dynamic perturbation of conductances. *Physical review letters*, 97(4):048104, 2006.

- [52] Gregor M Hoerzer, Robert Legenstein, and Wolfgang Maass. Emergence of complex computational structures from chaotic neural networks through reward-modulated hebbian learning. *Cerebral cortex*, 24(3):677–690, 2012.
- [53] Thomas Miconi. Biologically plausible learning in recurrent neural networks reproduces neural dynamics observed during cognitive tasks. *Elife*, 6, 2017.
- [54] Carina Curto, Anda Degeratu, and Vladimir Itskov. Encoding binary neural codes in networks of threshold-linear neurons. *Neural computation*, 25(11):2858–2903, 2013.
- [55] Nicolas Brunel. Is cortical connectivity optimized for storing information. *Nature neuroscience*, 19(5):749–755, 2016.
- [56] Bard Ermentrout. Type i membranes, phase resetting curves, and synchrony. *Neural computation*, 8(5):979–1001, 1996.
- [57] Peter E Latham, BJ Richmond, PG Nelson, and S Nirenberg. Intrinsic dynamics in neuronal networks. i. theory. *Journal of neurophysiology*, 83(2):808–827, 2000.
- [58] Nicolas Brunel and Peter E Latham. Firing rate of the noisy quadratic integrate-and-fire neuron. *Neural Computation*, 15(10):2281–2306, 2003.
- [59] Eugene M Izhikevich. Simple model of spiking neurons. *IEEE Transactions on neural networks*, 14(6):1569–1572, 2003.

Supplementary notes

1 Training recurrent dynamics

Here, we derive the synaptic update rules for the synaptic drive and spiking rate trainings, (10) and (16). We use RLS algorithm³³ to learn target functions $f_i(t)$, $i = 1, 2, \dots, N$ defined on a time interval $[0, T]$, and the synaptic update occurs at evenly spaced time points, $0 = t^0 \leq t^1 \leq \dots \leq t^K = T$.

In the following derivation, super-script k on a variable X_i^k implies that X is evaluated at t^k , and the sub-script i implies that X pertains to neuron i .

1.1 Training synaptic drive

The cost function measures the discrepancy between the target functions $f_i(t)$ and the synaptic drive $u_i(t)$ for all $i = 1, \dots, N$ at discrete time points t_0, \dots, t_K ,

$$C[W] = \frac{1}{2} \sum_{k=0}^K \|\mathbf{f}^k - \mathbf{w}^k\|_{L_2}^2 + \frac{\lambda}{2} \|W\|_{L_2}^2. \quad (20)$$

The Recursive Least Squares (RLS) algorithm solves the problem iteratively by finding a solution W^n to (20) at t^n and updating the solution at next time step t^{n+1} . We do not directly find the entire matrix W^n , but find each row of W^n , i.e. synaptic connections to each neuron i that minimize the discrepancy between u_i and f_i , then simply combine them to obtain W^n .

To find the i^{th} row of W^n , we denote it by \mathbf{w}_i^n and rewrite the cost function for neuron i that evaluates the discrepancy between $f_i(t)$ and $u_i(t)$ on a time interval $[0, t^n]$,

$$C[\mathbf{w}_i^n] = \frac{1}{2} \sum_{k=0}^n (f_i^k - \mathbf{w}_i^n \cdot \mathbf{r}^k)^2 + \frac{\lambda}{2} \|\mathbf{w}_i^n\|_{L_2}^2. \quad (21)$$

Calculating the gradient and setting it to 0, we obtain

$$0 = \nabla_{\mathbf{w}_i^n} C = - \sum_{k=1}^n (\hat{u}_i^k - \mathbf{w}_i^n \cdot \mathbf{r}^k) \mathbf{r}^k + \lambda \mathbf{w}_i^n$$

We express the equation concisely as follows.

$$\begin{aligned} [R^n + \lambda I] \mathbf{w}_i^n &= \mathbf{q}^n \\ R^n &= \sum_{k=1}^n \mathbf{r}^k (\mathbf{r}^k)^T, \quad \mathbf{q}^n = \sum_{k=1}^n \hat{u}_i^k \mathbf{r}^k. \end{aligned} \quad (22)$$

To find \mathbf{w}_i^n iteratively, we rewrite equation (22) up to t^{n-1} ,

$$[R^{n-1} + \lambda I] \mathbf{w}_i^{n-1} = \mathbf{q}^{n-1}, \quad (23)$$

and subtract equations (22) and (23) to obtain

$$[R^n + \lambda I] [\mathbf{w}_i^n - \mathbf{w}_i^{n-1}] + \mathbf{r}^n (\mathbf{r}^n)^T \mathbf{w}_i^{n-1} = \hat{u}_i^n \mathbf{r}^n. \quad (24)$$

The update rule for \mathbf{w}_i^n is then given by

$$\mathbf{w}_i^n = \mathbf{w}_i^{n-1} + e_i^n [R^n + \lambda I]^{-1} \mathbf{r}^n, \quad (25)$$

where the error term is

$$e_i^n = f_i^n - \mathbf{r}^n \cdot \mathbf{w}_i^{n-1}. \quad (26)$$

The matrix inverse $P^n = [R^n + \lambda I]^{-1}$ can be computed iteratively

$$P^n = P^{n-1} - \frac{P^{n-1} \mathbf{r}^n (\mathbf{r}^n)^T P^{n-1}}{1 + (\mathbf{r}^n)^T P^{n-1} \mathbf{r}^n}, \quad P^0 = \lambda^{-1} I,$$

using the matrix identity

$$(A + \mathbf{r} \mathbf{r}^T)^{-1} = A^{-1} - \frac{A^{-1} \mathbf{r} \mathbf{r}^T A^{-1}}{1 + \mathbf{r}^T A^{-1} \mathbf{r}}.$$

1.2 Training spiking rate

To train the spiking rate of neurons, we approximate the spike train $s_i(t)$ of neuron i with its spiking rate $\phi(u_i(t) + I_i)$ where ϕ is the current-to-rate transfer function of theta neuron model. For constant input,

$$\phi_1(x) = \pi^{-1} \sqrt{[x]_+} \quad \text{where} \quad [x]_+ = \max(x, 0), \quad (27)$$

and for noisy input

$$\phi_2(x) = \frac{1}{\pi} \sqrt{c \log(1 + e^{x/c})}. \quad (28)$$

Since ϕ_2 is a good approximation of ϕ_1 and has a smooth transition around $x = 0$, we used $\phi \equiv \phi_2$ with $c = 0.1^{58}$.

If the synaptic update occurs at discrete time points, t^0, \dots, t^K , the objective is to find recurrent connectivity W that minimizes the cost function

$$C[W] = \frac{1}{2} \sum_{k=0}^K \|\mathbf{f}^k(t) - \phi(W \mathbf{r}^k(t) + \mathbf{I})\|_{L_2}^2 + \frac{\lambda}{2} \|W\|_{L_2}^2. \quad (29)$$

As in training the synaptic drive, we optimize the following cost function to train each row of W^n that evaluates the discrepancy between the spiking rate of neuron i and the target spiking rate f_i over a time interval $[0, t_n]$,

$$C[\mathbf{w}_i^n] = \frac{1}{2} \sum_{k=1}^n (f_i^k - \phi(\mathbf{w}_i^n \cdot \mathbf{r}^k + I_i^k))^2 + \frac{\lambda}{2} |\mathbf{w}_i^n|^2. \quad (30)$$

Calculating the gradient and setting it to zero, we obtain

$$0 = \nabla_{\mathbf{w}_i^n} C = - \sum_{k=1}^n [f_i^k - \phi(\mathbf{w}_i^n \cdot \mathbf{r}^k + I_i^k)] \tilde{\mathbf{r}}_i^k + \lambda \mathbf{w}_i^n. \quad (31)$$

where

$$\tilde{\mathbf{r}}_i^k = \phi'(u_i^k + I_i^k) \mathbf{r}^k \quad (32)$$

is the vector of filtered spike trains scaled by the gain of neuron i . Note that when evaluating ϕ' in equation (32), we use the approximation $u_i^k \approx \mathbf{w}_i^n \cdot \mathbf{r}^k$ to avoid introducing nonlinear functions of \mathbf{w}_i^n .

To find an update rule for \mathbf{w}_i^n , we rewrite equation (31) up to t_{n-1} ,

$$0 = - \sum_{k=1}^{n-1} [f_i^k - \phi(\mathbf{w}_i^{n-1} \cdot \mathbf{r}^k + I_i^k)] \tilde{\mathbf{r}}_i^k + \lambda \mathbf{w}_i^{n-1}, \quad (33)$$

and subtract equations (31) and (33) and obtain

$$0 = \sum_{k=1}^n [\phi(\mathbf{w}_i^n \cdot \mathbf{r}^k + I_i^k) - \phi(\mathbf{w}_i^{n-1} \cdot \mathbf{r}^k + I_i^k)] \tilde{\mathbf{r}}_i^k - [f_i^n - \phi(\mathbf{w}_i^{n-1} \cdot \mathbf{r}^n + I_i^n)] \tilde{\mathbf{r}}_i^k + \lambda [\mathbf{w}_i^n - \mathbf{w}_i^{n-1}]. \quad (34)$$

Since \mathbf{w}_i^{n-1} is updated by small increment, we can approximate the first line in equation (34),

$$\phi(\mathbf{w}_i^n \cdot \mathbf{r}^k + I_i^k) - \phi(\mathbf{w}_i^{n-1} \cdot \mathbf{r}^k + I_i^k) \approx [\mathbf{w}_i^n - \mathbf{w}_i^{n-1}] \cdot \tilde{\mathbf{r}}_i^k \quad (35)$$

where we use the approximation $u_i^k \approx \mathbf{w}_i^n \cdot \mathbf{r}^k$ as before to evaluate the derivative ϕ' . Substituting equation (35) to equation (34), we obtain the update rule

$$\mathbf{w}_i^n = \mathbf{w}_i^{n-1} + e_i^n [R^n + \lambda I]^{-1} \tilde{\mathbf{r}}_i^n, \quad (36)$$

where the error is

$$e_i^n = f_i^n - \phi(\mathbf{w}_i^{n-1} \cdot \mathbf{r}^n + I_i^n), \quad (37)$$

and the correlation matrix of the normalized spiking activity is

$$R^n = \sum_{k=1}^n \tilde{\mathbf{r}}_i^k (\tilde{\mathbf{r}}_i^k)^T. \quad (38)$$

As shown above, the matrix inverse $P^n = [R^n + \lambda I]^{-1}$ can be computed iteratively,

$$P^n = P^{n-1} - \frac{P^{n-1} \tilde{\mathbf{r}}_i^n (\tilde{\mathbf{r}}_i^n)^T P^{n-1}}{1 + (\tilde{\mathbf{r}}_i^n)^T P^{n-1} \tilde{\mathbf{r}}_i^n}, \quad P^0 = \lambda^{-1} I.$$

2 Mean field description of the quasi-static dynamics

We say that a network is in a quasi-static state if the synaptic drive to a neuron changes sufficiently slower than the dynamical time scale of neurons and synapses. Here, we use a formalism developed by Buice and Chow⁴¹ and derive equations (1) and (2), which provide a mean field description of the synaptic and spiking rate dynamics of neurons in the quasi-static state.

First, we recast single neuron dynamic equation (5) in terms of the empirical distribution of neuron's phase $\eta_i(\theta, t) = \delta(\theta_i(t) - \theta)$. Since the number of neurons in the network is conserved, we can write the Klimontovich equation for the phase distribution

$$\partial_t \eta_i(\theta, t) + \partial_\theta [F(\theta, u_i + I_i) \eta_i(\theta, t)] = 0 \quad (39)$$

where $F(\theta, I) = 1 - \cos \theta + I(1 + \cos \theta)$. The synaptic drive equation (6) can be written in the form

$$\tau_s \dot{u}_i(t) = -u_i(t) + 2 \sum_{j=1}^N W_{ij} \eta_j(\pi, t) \quad (40)$$

since $s_j(t) = \eta_j(\pi, t) \dot{\theta}|_{\theta_j=\pi}$ and $\dot{\theta}_j|_{\theta_j=\pi} = 2$ for a theta neuron model. Equation (39), together with (40), fully describes the network dynamics.

Next, to obtain a mean field description of the spiking dynamics, we take the ensemble average prepared with different initial conditions and ignore the contribution of higher order moments resulting from nonlinear terms $\langle u_i \eta_i \rangle$. Then we obtain the mean field equation

$$\partial_t \rho(\theta, t) + \partial_\theta [F(\theta, U_i + I_i) \rho_i(\theta, t)] = 0 \quad (41)$$

$$\tau_s \dot{U}_i = -U_i + 2 \sum_{j=1}^N W_{ij} \rho_j(\pi, t). \quad (42)$$

where $\langle u_i \rangle = U_i$ and $\langle \eta_i \rangle = \rho_i$. We note that the mean field equations (41) and (42) provide a good description of the trained network dynamics because W learns over repeatedly training trials and starting at random initial conditions, to minimize the error between target trajectories and actual neuron activity.

Now, we assume that the temporal dynamics of synaptic drive and neuron phase can be suppressed in the quasi-static state,

$$\tau_s \dot{U}_i \approx 0, \quad \partial_t \rho \approx 0. \quad (43)$$

Substituting (43) to equation (41), but allowing $U(t)$ to be time-dependent, we obtain the quasi-static solution of phase density

$$\rho_i(\theta, t) = \frac{\sqrt{[U_i(t) + I_i]_+}}{\pi [1 - \cos \theta + (U_i(t) + I_i)(1 + \cos \theta)]}, \quad (44)$$

which has been normalized such that $\int_{-\pi}^{\pi} \rho_i(\theta) d\theta = 1$, and the spiking rate of a neuron is given by

$$\phi(U_i(t) + I_i) = 2\rho_i(\pi, t) = \sqrt{[U_i(t) + I_i]_+} / \pi, \quad (45)$$

the current-to-rate transfer function of a theta neuron model. Substituting (43) and (45) to equation (42), we obtain a quasi-static solution of the synaptic drive

$$U_i(t) = \sum_{j=1}^N W_{ij} \phi(U_j(t) + I_j). \quad (46)$$

If we define the spiking rate of a neuron as $R_i(t) = \phi(U_i + I_i)$, we immediately obtain

$$R_i(t) = \phi\left(\sum_{j=1}^N W_{ij} R_j + I_i\right). \quad (47)$$

3 Analysis of learning error

In this section, we identify and analyze two types of learning errors, assuming that for sufficiently heterogeneous targets, (1) the learning rule finds a recurrent connectivity W that can generate target patterns if the quasi-static condition holds, and (2) the mean field description of the spiking network dynamics is accurate due to the error function and repeated training trials. These assumptions imply that equations (46) and (47) hold for the target patterns $U_i(t)$ and the trained W . We show that learning errors arise when our assumptions become inaccurate, hence the network dynamics described by equations (46) and (47) deviate from the actual spiking network dynamics. As we will see, tracking error is prevalent if the target is not an exact solution of the mean field dynamics (i.e. quasi-static approximation fails), and the sampling error dominates if the discrete spikes do not accurately represent continuous targets (i.e. mean field approximation fails).

Suppose we are trying to learn a target \hat{u}_i which obeys an Ornstein-Uhlenbeck process

$$\left(\tau_c \frac{d}{dt} + 1\right) \hat{u}_i = \xi_i(t) \quad (48)$$

on a time interval $0 < t < T$ where $\xi_i(t)$ are independent white noise with zero mean and variance σ^2 . The time constant τ_c determines the temporal correlation of a target trajectory. In order for perfect training, the target dynamics (48) needs to be compatible with the network dynamics (6); in other words, there must exist a recurrent connectivity W such that the following equation

$$\left(\tau_s \frac{d}{dt} + 1\right) \hat{u}_i(t) = \sum_{j=1}^N W_{ij} s[\hat{u}_j(t)] \quad (49)$$

obtained by substituting the solution of (48) into (6) must hold for $0 < t < T$. Here, $s[\hat{u}_j(t)]$ maps the synaptic drive $\hat{u}_j(t)$ to the entire spike train $s_j(t)$.

It is very difficult to find W that may solve equation (49) exactly since it requires fully understanding the solution space of a high dimensional system of nonlinear ordinary differential equations. Instead, we assume that the target patterns are quasi-static and the learning rule finds a recurrent connectivity W that satisfies

$$\hat{u}_i(t) = \sum_{j=1}^N W_{ij} \phi(\hat{u}_j(t)). \quad (50)$$

We then substitute equation (50) to equation (49) to estimate how the quasi-static mean field dynamics deviate from the actual spiking network dynamics. A straightforward calculation shows that

$$\hat{u}_i(t) - \sum_{j=1}^N W_{ij} \phi(\hat{u}_j(t)) + \epsilon_{\text{track}} + \epsilon_{\text{sample}} = 0 \quad (51)$$

where we define the tracking and sampling errors as

$$\epsilon_{\text{track}} = \tau_s \frac{d\hat{u}_i}{dt} \quad (52)$$

and

$$\epsilon_{\text{sample}} = \sum_{j=1}^N W_{ij} (\phi(\hat{u}_j(t)) - s[\hat{u}_j(t)]) \quad (53)$$

on the time interval $0 < t < T$.

Tracking error. From its definition, ϵ_{track} captures the deviation of the quasi-static solution (50) from the exact solution of the mean field description obtained when $\epsilon_{\text{sample}} = 0$. ϵ_{track} becomes large if the quasi-static condition (43) fails and, in such network state, the synaptic dynamic is not able to “track” the target patterns, thus learning is obstructed. In the following, we estimate ϵ_{track} in terms of two time scales τ_s and τ_c .

First, we take the Fourier transform of equation (52) and obtain

$$\mathcal{F}[\epsilon_{\text{track}}](\omega) = i\tau_s\omega \cdot \mathcal{F}[\hat{u}](\omega). \quad (54)$$

Next, normalize $\mathcal{F}[\epsilon_{\text{track}}]$ with respect to $\mathcal{F}[\hat{u}]$ to estimate the tracking error for target patterns with different amplitudes, then compute the power of normalized tracking error.

$$\frac{1}{\Omega} \int_0^{\Omega} \left\| \frac{\mathcal{F}[\epsilon_{\text{track}}]}{\mathcal{F}[\hat{u}]} \right\|^2 d\omega = \frac{1}{2} \tau_s \Omega \Big|_{\Omega=\Omega_c} = \frac{1}{4\pi} \frac{\tau_s}{\tau_c} \quad (55)$$

where $\Omega_c = 1/(2\pi\tau_c)$ is the cut-off frequency of the power spectrum of a Gaussian process, $S_{GP}(\omega) = \sigma^2\tau_c^2/(1 + 4\pi^2\tau_c^2\omega^2)$. Thus, the tracking error scales with τ_s/τ_c .

Sampling error. ϵ_{sample} captures how the actual representation of target patterns in terms of spikes deviates from their continuous representation in terms of rate functions. In the following, we estimate ϵ_{sample} in terms of τ_s and N under the assumption that the continuous representation provides an accurate description of the target patterns.

We low-pass filtered ϵ_{sample} to estimate the sampling error since the synaptic drive (i.e. the target variable in this estimate) is a W weighted sum of filtered spikes with width that scales with τ_s . If the spike trains of neurons are uncorrelated (i.e. cross product terms are negligible),

$$\text{Var}[\epsilon_{\text{sample}}^{\text{filtered}}] = \sum_{j=1}^N W_{ij}^2 \langle (\bar{r}_j - r_j(t))^2 \rangle \quad (56)$$

where $r_j(t)$ is the filtered spike train and $\bar{r}_j = \langle r_j(t) \rangle = \frac{1}{\Delta t} \int_{t_k}^{t_k+\Delta t} r_j(s) ds$ is the empirical estimate of mean spiking rate on a short time interval.

First, we calculate the fluctuation of filtered spike trains under the assumption that a neuron generates spikes sparsely, hence the filtered spikes are non-overlapping. Let $s_j(t) = \sum_k \delta(t - t_j^k)$ be a spike train of neuron j and the filtered spike train $r_j(t) = \frac{1}{\tau_s} \sum_k \exp(-(t - t_j^k)/\tau_s) H(t - t_j^k)$. Then, the rate fluctuation of neuron j is

$$\langle (r_j(t) - \bar{r}_j)^2 \rangle = \langle r_j^2(t) \rangle - \bar{r}_j^2 \quad (57)$$

$$= \frac{1}{\tau_s^2} \sum_k \langle \exp(-2(t - t_j^k)/\tau_s) H(t - t_j^k) \rangle - \bar{r}_j^2 \quad (58)$$

$$= \bar{r}_j \left(\frac{1}{2\tau_s} - \bar{r}_j \right) \quad (59)$$

where k is summed over the average number of spikes, $\bar{r}_j \Delta t$, generated in the time interval of length Δt .

Next, to estimate the effect of network size on the sampling error, we examined equation (50) and observed that $O(W) \sim 1/N$. This follows from that, for pre-determined target patterns, $O(U), O(\phi(U)) \sim 1$ regardless of the network size, hence $O(W)$ must scale with $1/N$ in order for both sides of the equation to be compatible. If the network is dense, i.e. the number of synaptic connections to a neuron is pN on average, then the sampling error scales as follows.

$$O\left(\text{Var}[\epsilon_{\text{sample}}^{\text{filtered}}]\right) \sim \sum_{j=1}^N O(W_{ij}^2) O(\langle (\bar{r}_j - r_j(t))^2 \rangle) \sim \frac{1}{\tau_s N} \quad (60)$$

Supplementary figures




Figure S1: Learning arbitrarily complex target patterns in a network of rate-based neurons. The network dynamics obey $\tau \dot{x}_i = -x_i + \sum_{j=1}^N W_{ij} r_j + I_i$ where $r_j = \tanh(x_j)$. The synaptic current x_i to every neuron in the network was trained to follow complex periodic functions $f(t) = A \sin(2\pi(t - T_0)/T_1) \sin(2\pi(t - T_0)/T_2)$ where the initial phase T_0 and frequencies T_1, T_2 were selected randomly. The elements of initial connectivity matrix W_{ij} were drawn from a Gaussian distribution with mean zero and standard deviation σ/\sqrt{Np} where $\sigma = 2$ was strong enough to induce chaotic dynamics; Network size $N = 500$, connection probability between neurons $p = 0.3$, and time constant $\tau = 10$ ms. External input I_i with constant random amplitude was applied to each neuron for 50 ms (blue) and was set to zero elsewhere. **(a)** Before training, the network is in chaotic regime and the synaptic current (black) of individual neurons fluctuates irregularly. **(b)** After learning to follow the target trajectories (red), the synaptic current tracks the target pattern closely in response to the external stimulus.




Figure S2: Training a network that has no initial connections. The coupling strength of the initial recurrent connectivity is zero, and, prior to training, no synaptic or spiking activity appears beyond the first few hundred milliseconds. **(a)** Training synaptic drive patterns using the RLS algorithm. Black curves show the actual synaptic drive of 10 neurons and red curves show the target outputs. Blue shows the 100 ms external stimulus. **(b)** Correlation between synaptic drive and target function (top) and the Frobenius norm of changes in recurrent connectivity normalized to initial connectivity during training (bottom). **(c)-(d)** Same as in (a) and (b), but spiking rate patterns are trained.




Figure S3: Learning target patterns with low population spiking rate. The synaptic drive of networks consisting of 500 neurons were trained to learn complex periodic functions $f(t) = A \sin(2\pi(t - T_0)/T_1) \sin(2\pi(t - T_0)/T_2)$ where the initial phase T_0 and frequencies T_1, T_2 were selected randomly from [500 ms, 1000 ms]. **(a)** The amplitude $A = 0.1$, resulting in population spiking rate 2.8 Hz in trained window. **(b)** The amplitude $A = 0.05$, resulting in population spiking rate 1.5 Hz in trained window. **(c)** The amplitude $A = 0.01$, resulting in population spiking rate 0.01 Hz in trained window and learning fails.




Figure S4: Learning recurrent dynamics with leaky integrate-and-fire and Izhikevich neuron models. Synaptic drive of a network of spiking neurons were trained to follow 1000 ms long targets $f(t) = A \sin(2\pi(t - T_0)/T_1) \sin(2\pi(t - T_0)/T_2)$ where T_0, T_1 and T_2 were selected uniformly from the interval [500 ms, 1000 ms]. **(a)** Network consisted of $N = 200$ leaky integrate-and-fire neuron models, whose membrane potential obeys $\dot{v}_i = -(v_i - I_i)/\tau + u_i$ with a time constant $\tau = 10$ ms; the neuron spikes when v_i exceeds spike threshold $v_{thr} = -50$ mV then v_i is reset to $v_{res} = -65$ mV. Red curves show the target pattern and black curves show the voltage trace and synaptic drive of a trained network. **(b)** Spike rastergram of a trained leaky integrate-and-fire neuron network generating the synaptic drive patterns. **(c)** Network consisted of $N = 200$ Izhikevich neurons, whose dynamics are described by two equations $\dot{v}_i = 0.04v_i^2 + 5v_i + 140 - w_i + I_i + u_i$ and $\dot{w}_i = a(bv_i - w_i)$; the neuron spikes when v_i exceeds 30 mV, then v_i is reset to c and w_i is reset to $w_i + d$. Neuron parameters a, b, c and d were selected as in the original study⁵⁹ so that there were equal numbers of regular spiking, intrinsic bursting, chattering, fast spiking and low threshold spiking neurons. Synaptic current u_i is modeled as in equations (6) for all neuron models with synaptic decay time $\tau_s = 30$ ms. Red curves show the target patterns and black curves show the voltage trace and synaptic drive of a trained network. **(d)** Spike rastergram of a trained Izhikevich neuron network showing the trained response of different cell types.




Figure S5: Synaptic drive of a network of neurons is trained to learn an identical sine wave while external noise generated independently from OU process is injected to individual neurons. The same external noise (gray curves) is applied repeatedly during and after training. **(a)-(b)** The amplitude of external noise is varied from (a) low, (b) medium to (c) high. The target sine wave is shown in red and the synaptic drive of neurons are shown in black. The raster plot in (c) shows the ensemble of spike trains of a successfully trained network with strong external noise.