Training recurrent spiking neural networks in strong coupling regime

Christopher M. Kim and Carson C. Chow

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

Abstract

Recurrent neural networks can be trained to perform complex tasks. However, due to the large number of unconstrained synaptic connections, the recurrent connectivity that emerges from network training may not be biologically plausible and thus it is questionable as to whether they are applicable to actual neural processing underlying these tasks. To narrow this gap, we developed a training scheme that, in addition to achieving learning goals, respects the structural and dynamic properties of a standard cortical circuit model, i.e., strongly coupled excitatory-inhibitory spiking neural networks. By preserving the strong mean excitatory and inhibitory coupling of initial networks, we found that trained networks obeyed Dale's law without additional constraints, exhibited large trial-to-trial spiking variability, and operated in an inhibition-stabilized regime. We derived analytical estimates on how training and network parameters constrain the changes in mean synaptic strength during training. Our results demonstrate that training recurrent neural networks subject to strong coupling constraints can result in connectivity structure and dynamic regime relevant to cortical circuits.

Introduction

With recent advances in machine learning, recurrent neural networks are being applied to a wide range of problems in neuroscience as a tool for discovering dynamic mechanisms underlying biological neural circuits. The activity of network models trained either on tasks that animals perform or to directly generate neural activity have been shown to be consistent with neural recordings. Examples include context-dependent computation [1], neural responses in motor cortex [2], decision-making with robust transient activities [3], and flexible timing by temporal scaling [4]. However, due to the large number of parameters, the network connectivities are often not constrained to be biologically plausible, thus it remains unclear if and how biological neural circuits could implement the proposed dynamic mechanisms.

One approach for imposing biological constraints is to develop model architectures based on detailed connectivity structure of the neural system. The complete connectome is available for certain invertebrates such as Drosophila, and it is possible to build recurrent neural network models constrained by the detailed structure of the connectome [5–7]. In cortical regions of higher animals, however, obtaining a connectome and applying it for computational modeling purposes pose significant technical and conceptual challenges. Therefore, it is important to develop alternative methods for constraining the connectivity in a biologically plausible manner independently of the detailed connectome structure.

Recent studies have applied training methods to rate-based and spiking cortical circuits of excitatory and inhibitory neurons[8–11] but the focus was to preserve the excitatory-inhibitory structure (i.e. Dale's law) without constraining the synaptic strength to a strong coupling regime (but see [12, 13]) that exhibits chaotic dynamics [14], asynchronous spiking [15–17], spiking rate fluctuations [18], and inhibition-stabilization [19] as observed in neural recordings. Here, we show that by preserving the mean excitatory and inhibitory synaptic weights of a strongly coupled network, trained networks can exhibit large trial-to-trial variability in spiking activities, respect Dale's law without imposing strict constraints, and show a paradoxical phenomenon found in inhibition-stabilized networks. We provide analytical estimates on how the network and training parameters constrain the size of synaptic updates during network training. Our findings demonstrate that including synaptic constraints that maintain strong coupling

strength allows spiking networks to respect dynamic properties of standard cortical circuit models in addition to achieving the learning goal.

Results

1 Learning under excitatory and inhibitory synaptic constraints

The initial network consisted of excitatory and inhibitory theta neurons that had recurrent synaptic weights W_{ij} strong enough to generate spike trains with large trial-to-trial variability. Fano factor of spike counts in a short time window across trials starting with random initial conditions was close to 1 for each neuron indicating Poisson-like spiking activity (not shown). We trained the synaptic current $u_i(t)$ of each neuron $i = 1, ..., N$ to learn a sinusoidal wave $f_i(t)$ with random phase shift defined on time interval $t \in [0, T]$ and centered at the mean synaptic current the neuron received prior to training. Nonzero presynaptic connections $W_{ij}, j = 1, ..., N$ to each neuron i was trained so that $u_i(t)$ followed the target $f_i(t)$ in the trained time window $[0, T]$ when neurons were stimulated with a brief external input. We used recursive least squares [20] to train the recurrent connections with an objective function that minimized the error between the synaptic current and the target pattern. Importantly, in order to investigate the effects of constraining the excitatory and inhibitory synaptic weights, we considered two objective functions that regularized synaptic weights differently. The first objective function regularized the standard L_2-norm of the recurrent synaptic weights:

$$C_{FORCE} = \sum_{i=1}^{N} \Big[\int_0^T \big(f_i(t) - u_i(t)\big)^2 dt + \lambda \sum_j W_{ij}^2 \Big]. \tag{1}$$

The second objective function regularized the sums of excitatory and inhibitory synaptic weights, respectively, to each neuron i, in addition to the L_2-norm of the synaptic weights:

$$C_{ROWSUM} = \sum_i \Big[\int_0^T \big(f_i(t) - u_i(t)\big)^2 dt + \lambda \sum_j W_{ij}^2 + \mu\big(\sum_{j \in E} W_{ij}\big)^2 + \mu\big(\sum_{j \in I} W_{ij}\big)^2 \Big]. \tag{2}$$

The first scheme, which is equivalent to the algorithm of Sussillo and Abbott [21] and used previously in spiking networks [22], will be referred to as FORCE training. The second scheme is referred to as ROWSUM training since it regularizes the sum of excitatory and inhibitory weights, respectively, in each row of the connectivity matrix.

Figure 1 shows the network activity when recurrent connections were modified by two different training schemes. Figure 1a shows the excitatory (red), inhibitory (blue) and total (orange) synaptic currents of four example neurons after a successful FORCE training, and Figure 1b shows the synaptic currents after a successful ROWSUM training. Notably, if the excitatory and inhibitory connections were constrained during ROWSUM training to preserve the strong coupling, the excitatory and inhibitory currents exhibited large fluctuations and adjusted their amplitudes in order for the total current to track the target pattern (Fig. 1b). The large fluctuations of the excitatory and inhibitory synaptic currents were not observed in FORCE training. Instead, the total synaptic current closely followed the target patterns with little error (Fig. 1a).

Figures 1c and 1d show the spiking activity of the same neurons shown in Figures 1a and 1b, respectively, across multiple trials starting from random initial conditions. The synaptic current resulting from the FORCE training led to highly regular spikes with little variability across trials. In contrast, the synaptic current resulting from the ROWSUM training led to irregular spikes and large trial-to-trial variability. The trial-averaged spiking rates from two training schemes both matched the desired rates (Figs. 1e and 1f). However, the Fano factor that measures the cross-trial variability was significantly different (Figs. 1e and 1f); the mean Fano factor of all neurons from the FORCE and ROWSUM trainings were approximately 0.5 and 0.8, respectively.

It was shown previously that FORCE training can learn desired spiking rate patterns from data [22]. Here, we investigated if the ROWSUM training that constrained the excitatory and inhibitory synaptic weights can be applied in a similar manner to generate target rate patterns with large trial-to-trial variability in individual neurons' spiking activity. We used the f-I curve of the spiking neuron model in order

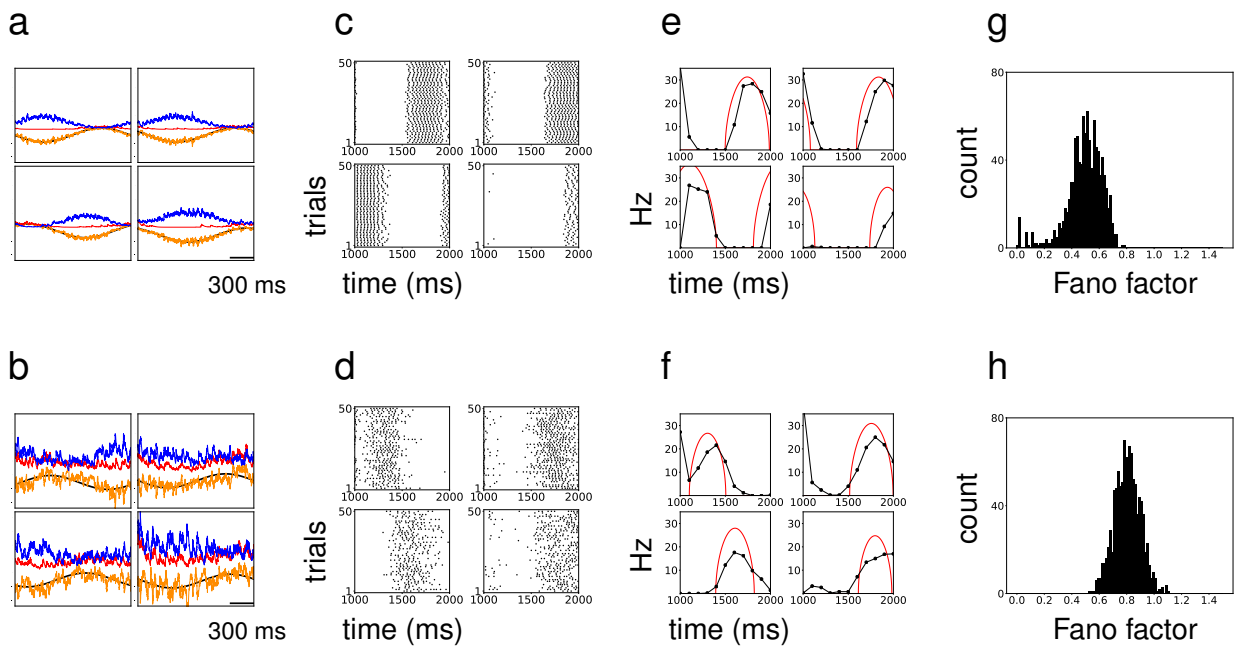

Figure 1: Comparison of FORCE- (top row) and ROWSUM-trained (bottom row) network activities. **(a)** Synaptic current of four example neurons from a FORCE-trained network; excitatory (red), inhibitory (blue) and total (orange) recurrent synaptic current; target pattern (black). **(b)** Synaptic current of example neurons from a ROWSUM-trained network. **(c-d)** Spike trains from the same neurons shown in (a) and (b) across trials starting at random initial conditions. **(e-f)** Trial-averaged spiking rate of neurons binned in 100 ms window (black) follows the target spiking rate (red). **(g-h)** Distribution of Fano factors of all neurons. Spike counts across trials in a 100 ms time window were used to compute Fano factor and averaged over 1 sec training window to obtain a time-averaged Fano factor for each neuron.

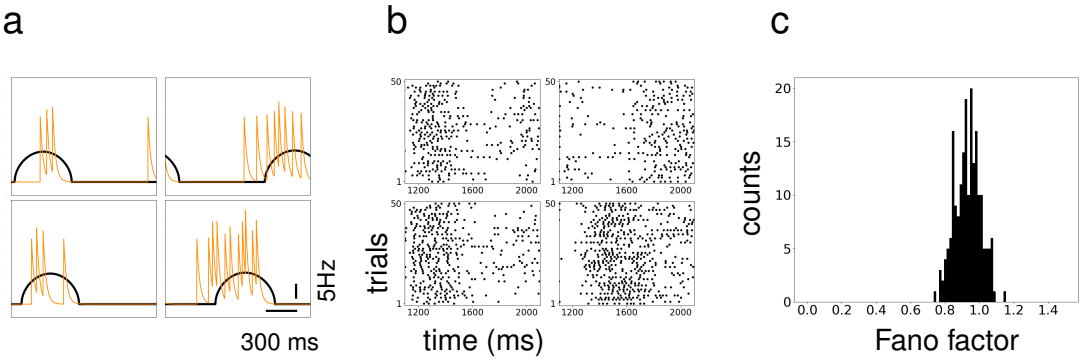

Figure 2: Generating target rate patterns under constrained excitatory and inhibitory synaptic connections. **(a)** Spiking response of four example neurons after ROWSUM training. **(b)** Large trial-to-trial variability is observed across trials in the spike trains of the same neurons shown in (a) when the target patterns are evoked starting from random initial conditions. **(c)** Distribution of the Fano factor of all neurons; spiking statistics is close to Poisson.

to approximate the neuron's firing rate, and included the regularization terms that constrained the mean excitatory and inhibitory synaptic weights of each neuron (see [22] for details).

Figure 2a shows that the spiking response of four example neurons from a single trial followed the target patterns after training under synaptic constraints. When the trained response was elicited over multiple trials starting with random initial conditions, the spike trains of individual neurons exhibited strong trial-to-trial variability (Fig. 2b) where the Fano factor measuring the cross-trial variability of neurons was close to 1 (Fig. 2c).

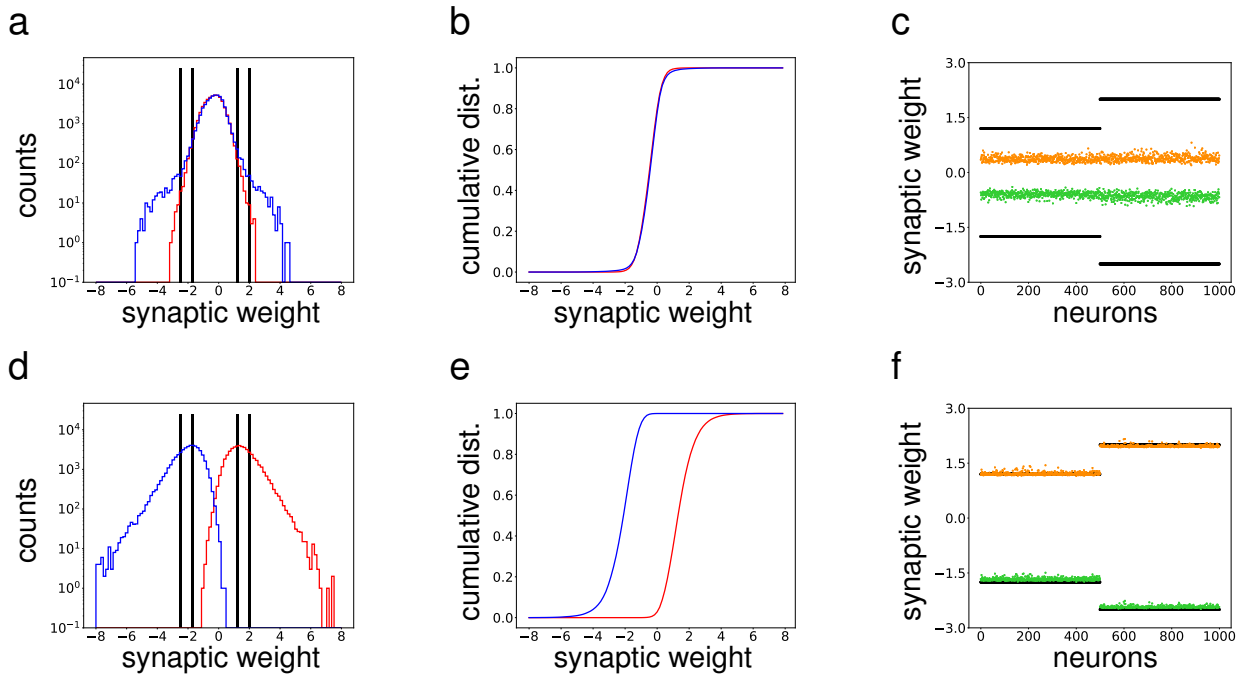

Figure 3: Connectivity of FORCE- and ROWSUM-trained networks. **(a)** Distribution of synaptic weights after FORCE training. Red (blue) shows the synapses that are excitatory (inibitory) connections in the initial network. Black shows the distribution of initial synaptic weights. **(b)** Cumulative distribution of FORCE-trained synapses. Same color code as in (a). **(c)** Mean excitatory (orange) and inhibitory (green) synaptic weights each neuron in the trained network receive is compared to those of the initial network. Black lines with positive (negative) values show the mean excitatory (inhibitory) synaptic weights to neurons in the initial network. **(d)-(f)** Same plots as in (a)-(c) for a ROWSUM-trained network.

2 Preserving the mean excitatory and inhibitory synaptic weights

We examined the trained synaptic weights to better understand the underlying network structure that led to different levels of variability in neuron activity in FORCE- and ROWSUM-trained networks. We found that the excitatory and inhibitory synaptic weights of the initial network were both redistributed around zero by the FORCE learning (Fig. 3a), and their cumulative distributions showed that approximately half of the synapses changed their signs after training, i.e., violated Dale's law (Fig. 3b). This shift in synaptic weight distributions towards zero can be explained by the L_2-regularization term in Eq. 1 that prefers small synaptic weights, and as a consequence, the mean excitatory and inhibitory synaptic weights each neuron received after FORCE training were significantly smaller than those of the initial network (Fig. 3c). These findings suggested that the network connectivity resulting from FORCE training was degenerate and highly tuned to generate the target dynamics with minimal error.

For the ROWSUM training, we used the same network and training parameters as the FORCE training but included the additional penalty that constrained the sums of excitatory and inhibitory synaptic weights, respectively, to each neuron. We found that the excitatory and inhibitory weights of the trained synapses remained around their initial values (Fig. 3d), and their cumulative distributions showed that most synapses did not change their sign even though no additional constraints were imposed to enforce Dale's law (Fig. 3e). The mean excitatory and inhibitory synaptic weights each neuron received after training deviated little from the initial mean synaptic weights (Fig. 2f) suggesting that the ROWSUM-regularization on the excitatory and inhibitory synaptic weights in Eq. 2 allowed the network to preserve their mean weights during training.

These findings from the trained networks demonstrated that constraining the sums of excitatory and inhibitory synaptic weights, respectively, allowed the individual synapses to learn target patterns while keeping the mean synaptic strength close to the initial values. Since the excitatory and inhibitory synaptic weights of the initial network were strong enough to generate variable spiking activity, the trained network was able to inherit the dynamic property of the initial network.

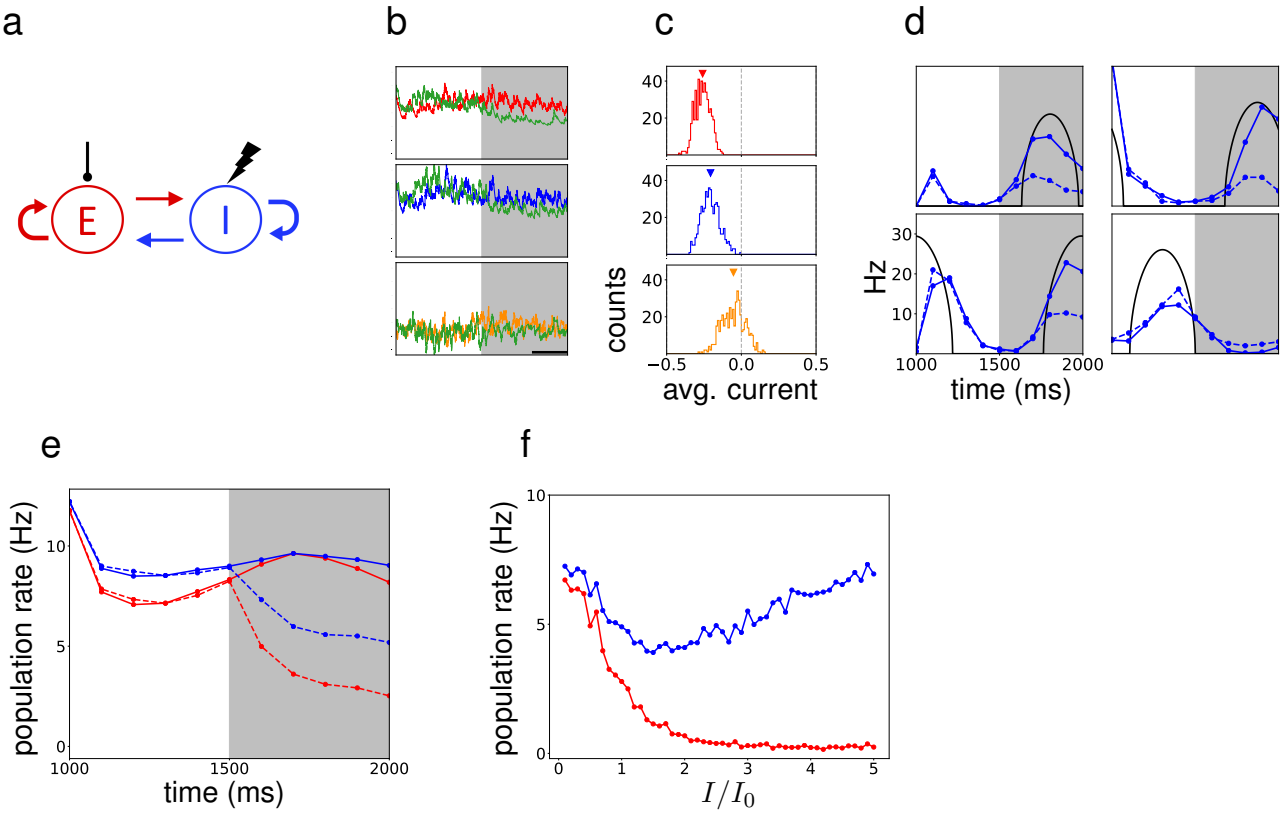

Figure 4: ROWSUM-trained networks exhibit dynamic features of inhibition-stabilization. Shaded region shows the time period during which inhibitory neurons receive external stimulus. **(a)** Schematic of network diagram where inhibitory neurons are excited by additional external stimulus. **(b)** Synaptic current of an inhibitory neuron; excitatory (red), inhibitory (blue) and total (orange) in an unstimulated network is compared to those (green) of a stimulated network. **(c)** Histogram of differences in time-averaged synaptic current to individual neurons before and after injecting additional stimulus to inhibitory neurons. Negative value implies reduction in average current after the additional stimulation. Color code same as in (b). Triangles show the mean of each distribution. **(d)** Trial-average firing rate of inhibitory neurons. Solid and dotted lines show the firing rate of unstimulated and stimulated neurons, respectively. **(e)** Firing rate of excitatory (red) and inhibitory (blue) populations in unstimulated (solid) and stimulated (dotted) networks. **(f)** Firing rate of excitatory (red) and inhibitory (blue) populations as a function of the stimulus intensity applied to inhibitory neurons. I_0 is the default external stimulus all neurons in the network receive.

3 Synaptically constrained networks operate in inhibition-stabilized regime

Next we investigated if other aspects of the initial network dynamics were inherited by the trained network when the mean excitatory and inhibitory synaptic weights were preserved. The initial network was set up to be an inhibition-stabilized network (ISN), where strong recurrent excitatory connections were stabilized by inhibitory feedback without which the network would exhibit run-away excitation [19, 23]. Since the ROWSUM training preserved the mean synaptic strength of the initial network, we investigated if, in addition to learning target activity patterns, the trained network preserved unique features of initial network's population dynamics characterized by the strong excitation and feedback inhibition.

An ISN exhibits a paradoxical phenomenon where an external stimulus that excites inhibitory neurons decreases their firing rates. Although the stimulus can transiently increase the firing rates of inhibitory neurons, this activity inhibits excitatory neurons that in turn reduce recurrent excitatory input to inhibitory neurons. The reduction in recurrent excitation is larger than the external stimulus resulting in a net decrease in excitatory input to inhibitory neurons. Such large reduction in recurrent excitation occurs because strongly connected excitatory neurons respond to changes in input with large gain [19, 23].

To test if an ROWSUM-trained network exhibited the dynamic features of ISN, we first evoked the learned activity patterns in the trained network, and then injected external stimulus that excited all

inhibitory neurons (Fig. 4a). The external stimulus was injected during the second half of the target patterns (shown as shaded region in Fig. 4) in order to compare its effect with the unstimulated network. We found that the recurrent excitatory input to inhibitory neurons was reduced when the inhibitory neurons were stimulated (Fig. 4b, top panel) indicating reduced excitatory population activity. Similarly, the recurrent inhibitory input to inhibitory neurons was also reduced (Fig. 4b, middle panel) indicating reduced inhibitory population activity. Although the reduced excitation and inhibition had mixed effects on the total synaptic current to inhibitory neurons (i.e. approximately 25% of inhibitory neurons ended up receiving additional excitatory input due to the stimulation; Fig. 4b, bottom panel), the net effect averaged over all inhibitory neurons was inhibition (Fig. 4c).

The effects of exciting inhibitory neurons were observed more clearly in trial-averaged firing rates in individual neurons and population rates. Without additional stimulus to inhibitory neurons, the trial-averaged firing rates followed the target rate patterns as trained (Fig. 4c). However, when the external stimulus was injected to excite the inhibitory neurons, their firing rates decreased compared to those of unstimulated neurons and a similar effect was found in both the excitatory and inhibitory population firing rates (Fig. 4d).

We found that this paradoxical phenomenon occurred only within a limited range of stimulus intensity, which is consistent with theoretical analysis of ISNs. As the external stimulus to inhibitory neurons increased, the excitatory population firing rate decreased gradually, and once the excitatory population activity was fully suppressed the firing rate of inhibitory neurons started increasing with the external stimulus (Fig. 4e). In this strong stimulus regime, the network was effectively reduced to a single inhibitory population, therefore the paradoxical phenomenon was no longer observed.

4 Analytical upper bound on synaptic weight updates

We further inquired how the network and training parameters in the ROWSUM-training constrained the size of synaptic weight updates. In particular, we sought to derive an analytical upper bound on the changes in the sums of excitatory and inhibitory synaptic weights, respectively. For the analysis, we considered a vector of synaptic connections $\mathbf{w}$ to an arbitrary neuron i in the network whose weights were modified by the RLS algorithm under the ROWSUM-regularization. Elements of $\mathbf{w}$ are denoted by $w_j = W_{ij}$ for $j = 1, ..., K$ where K is the total number of synapses modified under ROWSUM-regularization. Then, the sums of excitatory and inhibitory synaptic weights are defined, respectively, by

$$z_\alpha = \mathbf{w} \cdot \mathbf{1}_\alpha \quad \text{for} \quad \alpha = E, I. \tag{3}$$

Here, $\mathbf{1}_E$ denotes a vector $(1, ..., 1, 0, ..., 0)$ where the number of 1's is equal to the number of excitatory synapses K_E modified under ROWSUM-regularization and the rest of $K_I = K - K_E$ elements is 0. $\mathbf{1}_I = (0, ..., 0, 1, ..., 1)$ is defined similarly for the inhibitory synapses modified under the ROWSUM-regularization.

From the RLS algorithm (Eq. 17 in Methods), the changes in the sums of excitatory and inhibitory synaptic weights at each synaptic update obey

$$\delta z_\alpha^n = z_\alpha^n - z_\alpha^{n-1} = e^n \boldsymbol{\sigma}_\alpha^n \cdot \mathbf{r}^n \quad \text{for} \quad \alpha = E, I \tag{4}$$

where

$$\boldsymbol{\sigma}_\alpha^n = P^n \mathbf{1}_\alpha \tag{5}$$

and

$$P^n = \left[\sum_{t=1}^{n} \mathbf{r}^t (\mathbf{r}^t)' + \lambda I + \mu \mathbf{1}_E \mathbf{1}_E' + \mu \mathbf{1}_I \mathbf{1}_I' \right]^{-1}. \tag{6}$$

Since $\boldsymbol{\sigma}_\alpha^n$ was the only term in the synaptic update rule (Eq. 4) that depended on training and network parameters, it constrains the upper bound on synaptic weight changes. In simulations, the elements of $\boldsymbol{\sigma}_\alpha^n$ fluctuated in the early phase of training but stabilized and decreased monotonically as the training iteration n increased (Fig. 5a). Based on this observation, we derived an approximation of $\boldsymbol{\sigma}_\alpha^n$ for large n and obtained an estimate on its upper bound.

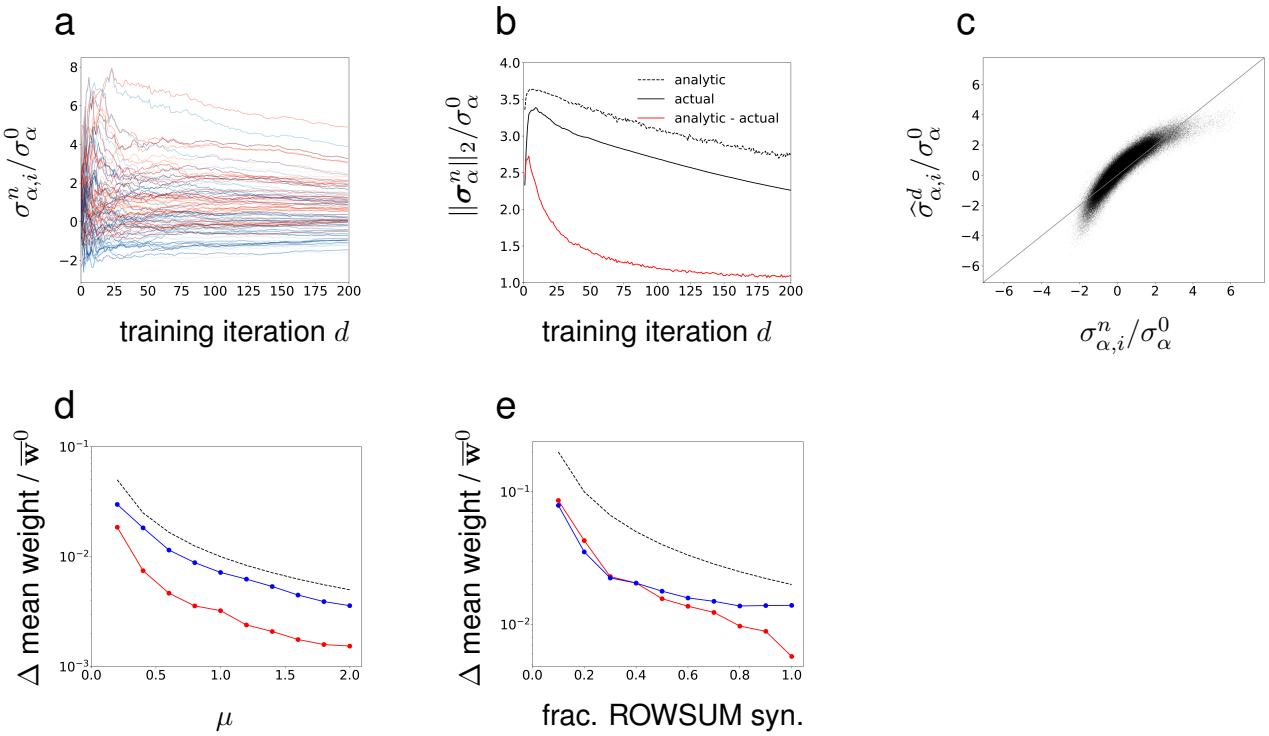

Figure 5: Analytical upper bound on synaptic updates. **(a)** Elements in $\boldsymbol{\sigma}^n_\alpha$ decrease monotonically if training iteration d becomes large. $\boldsymbol{\sigma}^n_\alpha$ of sample neurons are shown (red, excitatory; blue, inhibitory). $\sigma^0_\alpha = (\lambda + \mu K_\alpha)^{-1}$ denotes the initial value. **(b)** Comparison of the actual $\boldsymbol{\sigma}^n_\alpha$ and the analytical estimate $\widehat{\boldsymbol{\sigma}}^d_\alpha$ as a function of training iterations d. **(c)** Comparison of the actual values $\boldsymbol{\sigma}^n_\alpha$ and the analytical estimates $\widehat{\boldsymbol{\sigma}}^d_\alpha$ at training iteration $d = 200$. Scatter plot shows vector elements for all neurons. Gray shows the identity line. **(d)** Total changes in the mean synaptic weights $\sum_{t=1}^n \overline{\delta z^t_\alpha}/\overline{\mathbf{w}}^0$ after training as a function of μ. $\overline{\mathbf{w}}^0$ is the average synaptic strength of the initial network. Black, analytical estimate $O(\mu^{-1})$; Red, average excitatory weights; Blue, average inhibitory weights. **(e)** Total changes in the mean synaptic weights $\sum_{t=1}^n \overline{\delta z^t_\alpha}/\overline{\mathbf{w}}^0$ after training as a function of the fraction of ROWSUM-regularized synapses. Color code is same as (d); Black, analytical estimate $O(K_\alpha^{-1})$. In (d) and (e), networks were trained for $d = 100$ iterations.

The analysis of $\boldsymbol{\sigma}^n_\alpha$ presented below revealed that (1) the initial vector $\boldsymbol{\sigma}^0_\alpha = (\lambda + \mu K_\alpha)^{-1}\mathbf{1}_\alpha$ for $\alpha = E, I$ can be made arbitrarily small if the product μK_α of ROWSUM-penalty μ and the number of modified synapses K_α is large (see Methods Section 3.1), (2) $\boldsymbol{\sigma}^n_\alpha$ is bounded by $\boldsymbol{\sigma}^0_\alpha$ uniformly in $n \geq 1$ up to a uniform constant (see Methods Section 3.2), and (3) the update size of mean synaptic weights $\overline{\delta z^n_\alpha} = K_\alpha^{-1}\delta z^n_\alpha$ has an upper bound with a scaling factor $O(1/(\mu K_\alpha))$ thanks to (1) and (2) (see Methods Section 3.3). This analysis result showed that the product of μ and K_α tightly constrained the mean synaptic updates because $\boldsymbol{\sigma}^0_\alpha$ effectively served as a uniform upper bound on $\boldsymbol{\sigma}^n_\alpha$ that persisted throughout the training iterations n.

An approximation of $\boldsymbol{\sigma}^n_\alpha$ was derived in a few steps. First, the spiking activity $\mathbf{r}^t$ during training was approximated by the firing rates $\bar{\mathbf{r}}^t$ induced by the target patterns and a noise term $\bar{\boldsymbol{\varepsilon}}^t$ that modeled the spiking fluctuations around the firing rates (see Methods Section 3.2 for details). Such an approximation was based on the observation that spiking activity stayed close to the target patterns during training due to the fast weight updates (e.g., every 10 ms) by the RLS algorithm.

Next, since the target patterns were learned repeatedly, firing rates $\bar{\mathbf{r}}^t$ were assumed to have a period equal to the target length (m-steps) and repeated d-times to cover the entire training duration, i.e., $n = d \cdot m$. Assuming such firing rate patterns, the correlation matrix was approximated by $\sum_{t=1}^n \mathbf{r}^t(\mathbf{r}^t)' \approx d\sum_{t=1}^m(\bar{\mathbf{r}}^t + \sqrt{d}\,\bar{\boldsymbol{\varepsilon}}^t)(\bar{\mathbf{r}}^t + \sqrt{d}\,\bar{\boldsymbol{\varepsilon}}^t)'$ and substituted into the definition of P^n (Eq. 6) to obtain

$$\widehat{\boldsymbol{\sigma}}^d_\alpha = [I_K - P^0\bar{R}(d^{-1}I_m + \bar{R}'P^0\bar{R} + d\,\mathcal{E}'P^0\mathcal{E})^{-1}\bar{R}']\boldsymbol{\sigma}^0_\alpha \tag{7}$$

where $\bar{R} = [\bar{\mathbf{r}}^1, ..., \bar{\mathbf{r}}^m]$ and $\bar{\mathcal{E}} = [\bar{\boldsymbol{\varepsilon}}^1, ..., \bar{\boldsymbol{\varepsilon}}^m]$. We verified numerically that $\boldsymbol{\sigma}^n_\alpha$ and $\widehat{\boldsymbol{\sigma}}^d_\alpha$ decrease mono-

tonically for large n or d, and $\boldsymbol{\sigma}_\alpha^n$ can be estimated accurately by $\widehat{\boldsymbol{\sigma}}_\alpha^d$ (Figs. 5b, c). Then, by monotonicity, there exists $\widehat{\boldsymbol{\sigma}}_\alpha^{d_0}$ for some constant d_0 that serves as a uniform upper bound on $\boldsymbol{\sigma}_\alpha^n$:

$$\|\boldsymbol{\sigma}_\alpha^n\|_2 \approx \|\widehat{\boldsymbol{\sigma}}_\alpha^d\|_2 \leq \|\widehat{\boldsymbol{\sigma}}_\alpha^{d_0}\|_2 \leq C_\alpha \sigma_\alpha^0 \quad \text{for} \quad d \gg d_0 \gg 1 \quad \text{and} \quad \alpha = E, I. \tag{8}$$

Moreover, the constant C_α is independent of not only n but also μ and K_α because P^0 converges elementwise when μ or K_α becomes large (see Methods Section 3.3 for details).

Finally, we applied the upper bound on $\boldsymbol{\sigma}_\alpha^n$ (Eq. 8) to the synaptic learning rule (Eq. 4) to estimate the size of mean synaptic weight updates:

$$|\overline{\delta z_\alpha^n}| = K_\alpha^{-1} |\delta z_\alpha^n| \leq |e^n| B_\alpha (\mu K_\alpha)^{-1} \quad \text{for} \quad n \gg 1 \quad \text{and} \quad \alpha = E, I, \tag{9}$$

where $B_\alpha = C_\alpha r_{max} q_\alpha^{-1}$, $r_{\max}$ is the maximal firing rate induced by the target patterns, $q_\alpha = K_\alpha / K$ is the proportion of $\alpha = E, I$ synapses relative to the total number of synapses K modified under ROWSUM-regularization, and e^n is the learning error.

The analytical bound obtained in Eq. 9 suggested two factors that determined the size of synaptic updates. First, the ROWSUM-penalty μ constrained the upper bound of $\overline{\delta z_\alpha^n}$ with a scaling factor μ^{-1}. To verify this anlaytical estimate, we set up an initial network and trained all the synaptic weights with different μ-values across training sessions. We evaluated the mean excitatory and inhibitory synaptic weights each neuron received, and then averaged over all neurons in the network. We found that the total changes in the mean excitatory and inhibitory weights after training $\sum_{t=1}^n \overline{\delta z_\alpha^t}$ decreased with μ consistently to the analytical estimate $O(\mu^{-1})$ (Fig. 5d).

Second, the upper bound on the size of mean synaptic weight updates $|\overline{\delta z_\alpha^n}|$ was inversely proportional to K_α if it was sufficiently large and $q_\alpha = K_\alpha / K$ remained constant. In other words, increasing the number of $\alpha = E, I$ synapses modified under the ROWSUM-regularization effectively strengthened the ROWSUM-penalty. To verify this analytical result, we set up an initial network consisting of $N = 1000$ neurons and connection probability $p = 0.1$. Among Np synapses each neuron received, a fraction $f = K/(Np)$ of synapses was trained under the ROWSUM-regularization while the rest $1 - f$ of the synapses were trained under the standard L_2-regularization. The fraction f of ROWSUM-regularized synapses was increased systematically, while the proportion $q_\alpha = K_\alpha / K$ of ROWSUM-modified $\alpha = E, I$ synapses were fixed to 0.5 for both $\alpha = E, I$. After training, we evaluated the average strength of the excitatory and inhibitory synapses trained under ROWSUM-regularization and confirmed that the total changes in average synaptic strength decreased consistently with the analytical estimate $O(K_\alpha^{-1})$ (Fig. 5e).

5 Variance of synaptic weight distribution

So far, we demonstrated that by including the additional μ-penalty, trained networks preserved the mean excitatory and inhibitory synaptic weights of the initial network. To investigate the role of L_2-regularization in controlling the synaptic weight distribution, we fixed the μ-penalty and systematically varied λ that penalized the L_2-regularizer. Since large λ constrains every synapse to its initial value while decreasing λ eases this constraint, we expected that the variance of synaptic distribution will increase if λ becomes small and μ-penalty is large enough to keep the mean synaptic weights fixed.

As we varied λ over a wide range, the variance of synaptic weights indeed decreased gradually with λ (Fig. 6a). The distributions of excitatory and inhibitory synaptic weights in trained networks that used different λ values, $0.005, 1.0$ and 10.0, are shown in Figure 6b. When λ became large, the strong penalty constrained the synaptic weight distribution to narrow range and did not allow the network to explore the solution space that can generate the target activity patterns resulting in low performance (Fig. 6c).

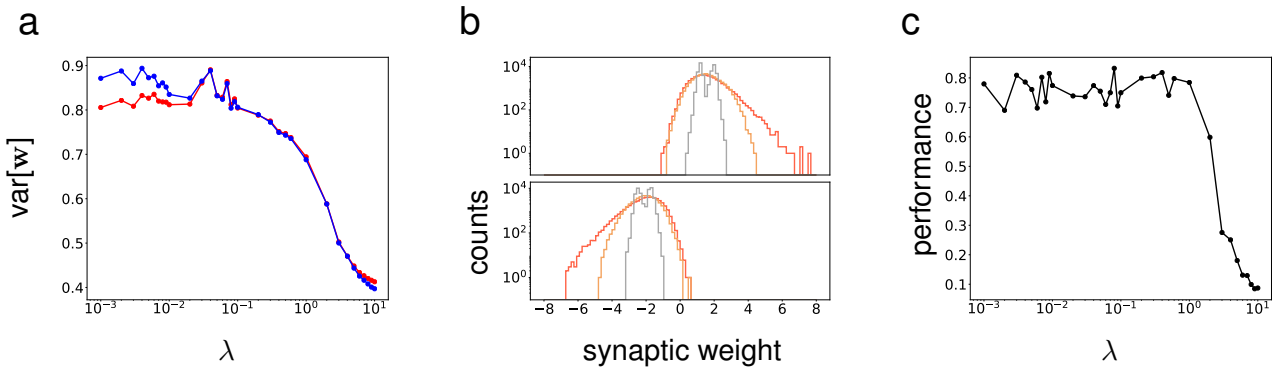

Figure 6: Variance of synaptic weight distribution in ROWSUM-trained networks. **(a)** The variance of synaptic weight distribution in networks trained with different λ-values. **(b)** Distribution of excitatory (top) and inhibitory (bottom) synaptic weights in networks trained with $\lambda = 0.005$ (red), 1.0 (orange) and 10 (gray). **(c)** Performance of ROWSUM-trained networks in generating target activity patterns. The correlation between trial-averaged-firing rate and target rate of each neuron were averaged over all neurons to measure performance.

Discussion

Recurrent neural networks can learn to perform cognitive tasks [1–4] and generate complex activity patterns [22, 24, 25]. However, due to the large number of connections, the connectivity structure that achieves the learning goal is nonunique [2, 8]. Therefore, it is important to further constrain the recurrent connectivity structure in order to identify solutions relevant to biological neural circuits. In this study, we showed that trained networks can maintain strong excitatory and inhibitory synaptic strength and, consequently, exhibit dynamic features of strongly coupled cortical circuit models, such as stochastic spiking and inhibition-stabilization, if the mean synaptic strength of the strongly coupled initial network was preserved throughout training.

In a related work, Ingrosso and Abbott [12] developed a training scheme that maintained strong synaptic strength in trained networks by setting up a stroungly coupled excitatory-inhibitory network and penalizing the deviation of synaptic weights from their initial values. Specifically, they considered a regularizer $R_i^{\text{syn}} = \sum_j (W_{ij} - W_{ij}^0)^2$ for each neuron i that constrained the L_2-distance of individual synaptic weights between the initial and trained networks. The main difference to our work is that the regularizer $R_i^{\text{sum}} = (\sum_{j\in E} W_{ij})^2 + (\sum_{j\in I} W_{ij})^2$ considered in our study constrained the sums of excitatory and inhibitory weights, respectively, but not the individual synaptic weights. In fact, including R_i^{syn} to the cost function is equivalent to adding the standard L_2-regularizer $R_i^{\text{stan}} = \sum_j W_{ij}^2$ if the synaptic update rule is derived through the RLS algorithm. This is because the effects of subtracting constants W_{ij}^0 from W_{ij} are nullified when the gradients at two time points are compared to derive a synaptic update rule (see Methods Section 2). Bounded coordinate descent algorithm implemented by Ingrosso and Abbott did not suffer from this issue, and they showed that regularizing the weights with R_i^{syn}, but not R_i^{stan}, allowed them to preserve the strong synaptic strength of initial networks. For the RLS algorithm, strengthing the L_2-regularizer R_i^{stan} will keep the weights close to their initial values; however, as demonstrated in Fig. 6, strong L_2-penalty narrows the solution space the network can explore, thus degrades the performance level. Our study demonstrated that constraining the mean synaptic strength through R_i^{sum} imparts the flexibility for individual synapses to learn the task while maintaining strong coupling of the averaged synaptic weights.

Separation of excitatory and inhibitory neurons is an important feature of biological neural circuits that has implications to dynamical state of the network. In particular, strong recurrent connections within and between the excitatory and inhibitory neurons give rise to dynamic properties such as asynchronous spiking [15–17], balanced amplification [26], and inhibition-stabilization [19]. In recent studies, training schemes were developed to preserve the excitatory-inhibitory structure of recurrent connections. For instance, "clipping" methods skipped synaptic updates that violated the sign constraints [12, 22]; back-propagation-through-time (BPTT) was adapted to rectify the synaptic weights in order to produce zero gradients in the opposite sign regime [8]; in cases where only the read-out and/or feedback weights were trained, the strong recurrent synapses of the initial network remained unchanged while other weighs were constrained to respect the overall excitatory-inhibitory structure [9]. However, these studies focused on

respecting the excitatory-inhibitory structure without promoting the strong recurrent connections essential for generating novel network dynamics. Our study suggests that imposing additional constraints on the mean synaptic weights may lead to trained networks that inherit dynamic properties of a strongly coupled excitatory-inhibitory network.

In a recent work, Baker et al [13] considered network and stimulus configurations that deviated from the standard balanced state in that excessive inhibition present in such alternative state resulted in threshold-linear stationary rates as a function of external stimulus, as opposed to the linear relationship between positive stationary rates and external inputs in the balanced state. Using the nonlinearity present in the threshold-linear activation, they showed that nonlinear computations was possible in the inhibition-excessive excitatory-inhibitory network, pointing out the importance of dynamic regime determined by the synaptic connectivity.

Developing general purpose machine learning algorithms for training spiking neural networks is an active area of research that made significant progress in the last several years. Recent work has demonstrated that including biological mechanisms, such as spike frequency adaptation, can enhance spiking neural network's capability [27], and biologically inspired alternatives to BPTT can be applied to perform supervised and reinforcement learning tasks [28]. The progress in training algorithms sets up a fertile ground to discover and rediscover the potential role of connectivity constraints present in biological neural circuits. Spiking neural networks trained with relevant biological features and connectivity constraints may serves as a biologically interpretable alternative to artificial neural networks that can provide further insights into the inner working of biological neural circuits.

Methods

The spiking neural network consisted of N randomly connected excitatory and inhibitory theta neurons with connection probability p. The network activity was governed by a set of equations

$$\begin{aligned} \tau_m \dot{\theta}_i &= 1 - \cos\theta_i + (I_i + u_i)(1 + \cos\theta_i) \\ \tau_s \dot{u}_i &= -u_i + \sum_j W_{ij} \sum_s \delta(t - t_j^s(\theta_j)) \end{aligned} \tag{10}$$

where $\theta_i \in [0, 2\pi)$ is the phase of theta neuron i, u_i is the synaptic drive that neuron i received from other neurons in the network, I_i is the external input to neuron i, τ_s is the synaptic constant, τ_m is the neuron integration constant and $t_j^s(\theta_j)$ is the spike times of neuron j. Theta neurons emits spikes $\delta(t-t')$ at time t' when $\theta_i(t)$ crosses π.

Previous studies have shown that the recurrent synaptic connections can be trained using recursive least squares (RLS) to generate target activity patterns [22, 24, 25, 29]. Specifically, the goal was to train the recurrent connections W_{ij} such that the synaptic drive $u_i(t)$ of all neurons $i = 1, ..., N$ in the network learned to follow their own target pattern $f_i(t)$ defined on time interval $[0, T]$ when neurons were stimulated by brief external input. In this study, we included additional penalty terms to the cost function that penalized the sums of excitatory and inhibitory synaptic weights to each neuron, respectively, and then derived the RLS algorithm that included the effects of the additional penalty terms.

The learning rule described below trains the activity of one neuron in the recurrent network. However, the same scheme can be applied to all neurons synchronously (i.e., modify synaptic connections to all neuron at the same time) to update the full connectivity matrix. To simplify notation, we dropped neuron index i in the following derivation; $\mathbf{w} = (W_{i1}, ..., W_{iK})$ denotes a vector of synaptic weights that neuron i receives from presynaptic neurons, $\mathbf{r}(t) = (r_1(t), ..., r_K(t))$ denotes the synaptic filtering of spike trains of neurons presynaptic to neuron i, and $f(t), t \in [0, T]$ is the target pattern of neuron i. Neurons were labeled so that the first K_E are excitatory and the remaining $K_I = K - K_E$ are inhibitory.

Using linearity in W_{ij}, the equation for synaptic drive u_i was rewritten in the form,

$$u_i(t) = \sum_j W_{ij} r_j(t) \tag{11}$$

$$\dot{r}_i = -\beta r_i + \beta \sum_s \delta(t - t_i^s(\theta_i)). \tag{12}$$

With this formulation, synaptic drive was expressed as a linear function of connectivity w_{ij}, which allowed us to derive the RLS update rule when the synaptic drive activities were trained.

For network simulations, we used a network consisting $N = 1000$ neurons for which $N_E = 500$ were excitatory and $N_I = 500$ were inhibitory. Each neuron received synaptic connections from randomly selected pN_E excitatory and pN_I inhibitory presynaptic neurons where the connection probability was $p = 0.1$. The synaptic strength between neurons were $W_{ij}^{EE} = 1.2, W_{ij}^{IE} = 2.0, W_{ij}^{EI} = -1.75, W_{ij}^{II} = -2.5$ where $W_{ij}^{\alpha\beta}$ denotes the weight of a synaptic connection from neuron $j \in \beta$ to neuron $i \in \alpha$. The external input I_i to neurons were sampled uniformly from the range $[0.1, 0.2]$. The synaptic constant was $\tau_s = 20$ ms and neuron integration constant was $\tau_m = 5$ ms. The target patterns consisted of sine waves $f_i(t) = A\sin(2\pi t/T + \phi_i)$ with amplitude $A = 0.25$, random phase $\phi_i \in [0, 2\pi)$ and period $T = 1$ sec.

1 Training a network of excitatory and inhibitory neurons

We derived an RLS training algorithm for a network consisting of excitatory and inhibitory neurons. Since maintaing strong excitatory and inhibitory synaptic strength was critical for generating dynamic features of excitatory-inhibitory network, we constrained the sums of excitatory and inhibitory synaptic weights to each neuron, respectively. The cost function was

$$C = \frac{1}{2}\sum_{t=0}^{T}(f^t - \mathbf{w}\cdot\mathbf{r}^t)^2 + \frac{\lambda}{2}\|\mathbf{w}\|_2^2 + \frac{\mu}{2}(\mathbf{w}\cdot\mathbf{1}_E)^2 + \frac{\mu}{2}(\mathbf{w}\cdot\mathbf{1}_I)^2 \tag{13}$$

where the first term is the learning error, the second is the L_2-regularization and the remaining two terms are the ROWSUM-regularization that penalizes the sums of excitatory and inhibitory synaptic connections to the neuron, respectively. Here $\|\mathbf{x}\|_2 = \sqrt{\sum_{i=1}^{K} x_i^2}$ is the standard L_2 norm, $\mathbf{1}_E = (1, .., 1, 0, ..., 0)$ and $\mathbf{1}_I = (0, ..., 0, 1, ..., 1)$ where the number of 1's in $\mathbf{1}_E$ and $\mathbf{1}_I$ is equal to K_E and K_I, respectively.

The gradient of cost function with respect to $\mathbf{w}$ is

$$\nabla_{\mathbf{w}} C = \sum_t(-f^t\mathbf{r}^t + \mathbf{r}^t(\mathbf{r}^t)'\mathbf{w}) + \lambda\mathbf{w} + \mu\mathbf{1}_E\mathbf{1}_E^T\mathbf{w} + \mu\mathbf{1}_I\mathbf{1}_I^T\mathbf{w}. \tag{14}$$

To derive the RLS algorithm, we computed the gradients at two consecutive time points

$$\mathbf{0} = \nabla_{\mathbf{w}^n} C = \sum_{t=1}^{n}(-f^t\mathbf{r}^t + \mathbf{r}^t(\mathbf{r}^t)'\mathbf{w}^n) + \lambda\mathbf{w}^n + \mu\mathbf{1}_E\mathbf{1}_E'\mathbf{w}^n + \mu\mathbf{1}_I\mathbf{1}_I'\mathbf{w}^n \tag{15}$$

and

$$\mathbf{0} = \nabla_{\mathbf{w}^{n-1}} C = \sum_{t=1}^{n-1}(-f^t\mathbf{r}^t + \mathbf{r}^t(\mathbf{r}^t)'\mathbf{w}^{n-1}) + \lambda\mathbf{w}^{n-1} + \mu\mathbf{1}_E\mathbf{1}_E'\mathbf{w}^{n-1} + \mu\mathbf{1}_I\mathbf{1}_I'\mathbf{w}^{n-1} \tag{16}$$

and subtracted Eq. 16 from Eq. 15 to obtain

$$\begin{aligned} \mathbf{w}^n &= \mathbf{w}^{n-1} + e^n P^n \mathbf{r}^n \\ e^n &= f^n - \mathbf{w}^{n-1}\cdot\mathbf{r}^n \end{aligned} \tag{17}$$

where

$$P^n = \left[\sum_{t=1}^{n}\mathbf{r}^t(\mathbf{r}^t)' + \lambda I + \mu\mathbf{1}_E\mathbf{1}_E' + \mu\mathbf{1}_I\mathbf{1}_I'\right]^{-1} \quad \text{for} \quad n \geq 1 \tag{18}$$

with the initial value

$$P^0 = [\lambda I + \mu\mathbf{1}_E\mathbf{1}_E' + \mu\mathbf{1}_I\mathbf{1}_I']^{-1}. \tag{19}$$

To update P^n iteratively, we used the Woodbury matrix identity

$$(A + UCV)^{-1} = A^{-1} - A^{-1}U(C^{-1} + VA^{-1}U)^{-1}VA^{-1} \tag{20}$$

where A is invertible and N-by-N, U is N-by-T, C is invertible and T-by-T and V is T-by-N matrices. Then P^n can be calculated iteratively

$$P^n = P^{n-1} - \frac{P^{n-1}\mathbf{r}^n(\mathbf{r}^n)'P^{n-1}}{1 + (\mathbf{r}^n)'P^{n-1}\mathbf{r}^n}.$$

2 Translation invariance of RLS algorithm

Here we show that the RLS algorithm is invariant to translation in the sense that adding arbitrary constants to the quadratic penalty terms does not change the synaptic update rule derived from the RLS. Translation invariance arises because of the recursive nature of RLS, i.e., the gradients at two time points are compared to derive the synaptic update rule, hence any constant term present in the gradient is removed when the difference of two gradients are computed. This justifies the form of penalty terms in the cost function Eq. 13 since including additional constants to the penalty terms, $\|\mathbf{w}\|_2$, $(\mathbf{w}\cdot\mathbf{1}_E)^2$ and $(\mathbf{w}\cdot\mathbf{1}_I)^2$, does not affect the synaptic update rule (Eq. 17).

To demonstrate translation invariance, we shifted w_i in the penalty terms by arbitrary number m_i and rewrote the cost function in the following form:

$$C = \frac{1}{2}\sum_{t=0}^{n}(f^t - \mathbf{w}\cdot\mathbf{r}^t)^2 + \frac{\lambda}{2}\|\mathbf{w}-\mathbf{m}\|_2^2 + \frac{\mu}{2}\big((\mathbf{w}-\mathbf{m})\cdot\mathbf{1}_E\big)^2 + \frac{\mu}{2}\big((\mathbf{w}-\mathbf{m})\cdot\mathbf{1}_I\big)^2. \tag{21}$$

Although this cost function is different from the one considered in the previous section, we now show that the synaptic update rule derived from this alternative form is identical to the previously derived Eq. 17. Differentiating the cost function Eq. 21 with respect to $\mathbf{w}^n$ and $\mathbf{w}^{n-1}$ yields

$$\begin{aligned}\mathbf{0} = \nabla_{\mathbf{w}^n} C = &\sum_{t=0}^{n}(-f^t\mathbf{r}^t + \mathbf{r}^t(\mathbf{r}^t)'\mathbf{w}^n) \\ &+ \lambda(\mathbf{w}^n - \mathbf{m}) + \mu\mathbf{1}_E\mathbf{1}_E'(\mathbf{w}^n - \mathbf{m}) + \mu\mathbf{1}_I\mathbf{1}_I'(\mathbf{w}^n - \mathbf{m})\end{aligned} \tag{22}$$

and

$$\begin{aligned}\mathbf{0} = \nabla_{\mathbf{w}^{n-1}} C = &\sum_{t=0}^{n-1}(-f^t\mathbf{r}^t + \mathbf{r}^t(\mathbf{r}^t)'\mathbf{w}^{n-1}) \\ &+ \lambda(\mathbf{w}^{n-1} - \mathbf{m}) + \mu\mathbf{1}_E\mathbf{1}_E'(\mathbf{w}^{n-1} - \mathbf{m}) + \mu\mathbf{1}_I\mathbf{1}_I'(\mathbf{w}^{n-1} - \mathbf{m}).\end{aligned} \tag{23}$$

Since the RLS algorithm updates the synaptic weights by computing the differences of gradients at time points n and $n-1$, the $\mathbf{m}$ terms drop out in the final synaptic update rule when we subtract Eq. 23 from Eq. 22. Therefore, the synaptic update rule we obtain from the alternative cost function (Eq. 21) is identical to one (Eq. 17) derived previously from the original cost function (Eq. 13).

3 Analytical upper bound on synaptic weight updates

In this section, we derived analytical estimates on how the training and network parameters constrained the trained synaptic weights. The sums of excitatory and inhibitory synaptic weights to a neuron were defined, respectively, as

$$z_\alpha^n = \sum_{i\in\alpha} w_i^n = \mathbf{w}^n\cdot\mathbf{1}_\alpha \quad \text{for} \quad \alpha = E, I. \tag{24}$$

We rearranged the RLS learning rule (Eq. 17) to find an expression for changes in the sums of excitatory and inhibitory synaptic weights:

$$\delta z_\alpha^n = z_\alpha^n - z_\alpha^{n-1} = e^n P^n\mathbf{r}^n\cdot\mathbf{1}_\alpha = e^n\boldsymbol{\sigma}_\alpha^n\cdot\mathbf{r}^n \tag{25}$$

where

$$\boldsymbol{\sigma}_\alpha^n = P^n\mathbf{1}_\alpha \quad \text{for} \quad \alpha = E, I. \tag{26}$$

Here we give a brief overview of the following analysis. We first obtained analytical expressions for the initial value $\sigma_\alpha^0 = (\lambda + \mu K_\alpha)^{-1}$ and $\boldsymbol{\sigma}_\alpha^n, n \geq 1$ to be used in the subsequent analysis. We found that the initial value σ_α^0 is of order $O((\mu K_\alpha)^{-1})$ and used it as the main quantity to bound the rest of $\boldsymbol{\sigma}_\alpha^n, n \geq 1$ (Section 3.1). To show that $\boldsymbol{\sigma}_\alpha^n$ is uniformly bounded in n, we decomposed the spiking activity as a sum of spiking rate and noise, and then obtained an approximation $\widehat{\boldsymbol{\sigma}}_\alpha^d$ of $\boldsymbol{\sigma}_\alpha^n$ which decreases monotonically and thus has a uniform upper bound $\widehat{\boldsymbol{\sigma}}_\alpha^{d_0}$ for some $d \gg d_0$ (Section 3.2). We numerically

verified that $\boldsymbol{\sigma}_\alpha^n \approx \widehat{\boldsymbol{\sigma}}_\alpha^d$, and they both decrease monotonically which led to the uniform upper bound on $\|\boldsymbol{\sigma}_\alpha^n\|_2 \approx \|\widehat{\boldsymbol{\sigma}}_\alpha^d\|_2 \leq C_\alpha \sigma_\alpha^0$ and mean synaptic weights $|\overline{\delta z_\alpha^n}| = K_\alpha^{-1}\delta z_\alpha^n \sim O\big((\mu K_\alpha)^{-1}\big)$ (Section 3.3).

The main result deduced from the analysis is that the changes in mean synaptic weights $|\overline{\delta z_\alpha^n}|$ is constrained by the product μK_α. By either increasing the μ-pentalty or the number of synpases K_α modified by the ROWSUM-regularization, the changes in mean synaptic weights can be constrained tightly.

3.1 Analytical expressions for $\boldsymbol{\sigma}_\alpha^n$

Here we derived expressions for $\boldsymbol{\sigma}_\alpha^n$ to be used in the following analysis. To find the initial $\boldsymbol{\sigma}_\alpha^0$ for $\alpha = E, I$, we substituted $A = \lambda I_K$, $U = [\mathbf{1_E}, \mathbf{1_I}]$, $V = U'$ and $C = I_2$ to the Woodbury matrix identity (Eq. 20) to obtain

$$P^0 = \frac{1}{\lambda} I_K - \frac{\mu}{\lambda(\lambda + \mu K_E)} \mathbf{1}_E \mathbf{1}_E' - \frac{\mu}{\lambda(\lambda + \mu K_I)} \mathbf{1}_I \mathbf{1}_I'. \tag{27}$$

Here the subscript in I_K denote the dimensionality of $K \times K$ identity matrix. Then,

$$\begin{aligned} &\boldsymbol{\sigma}_\alpha^0 = \sigma_\alpha^0 \mathbf{1}_\alpha \\ \text{where} \quad &\sigma_\alpha^0 = \frac{1}{\lambda + \mu K_\alpha} \quad \text{for} \quad \alpha = E, I. \end{aligned} \tag{28}$$

We can see immediately that the initial value σ_α^0 is of order $O(1/(\mu K_\alpha))$. To obtain estimates for $\boldsymbol{\sigma}_\alpha^n$ with $n \geq 1$, we used the Woodbury matrix identity (Eq. 20) to express P^n in the form

$$P^n = P^0 - P^0 R^n [I + (R^n)' P^0 R^n]^{-1} (R^n)' P^0 \tag{29}$$

where $R^n = [\mathbf{r}^1, ..., \mathbf{r}^n]$ and obtained

$$\boldsymbol{\sigma}_\alpha^n = [I - P^0 R^n [I + (R^n)' P^0 R^n]^{-1} (R^n)'] \boldsymbol{\sigma}_\alpha^0. \tag{30}$$

3.2 Asymptotic approximation of $\boldsymbol{\sigma}_\alpha^n$

To show that $\boldsymbol{\sigma}_\alpha^n$ can be bounded uniformly in n, it suffices to show that the right hand side of Eq. 30 does not depend on n. To this end, we analyzed its asymptotic behavior as the training iteration n grew arbitrarily large. We used the fact that (1) same target pattern was learned repeatedly during training and (2) fast synaptic weight updates (e.g., every 10 ms) by the RLS algorithm kept the spiking activity close to the target pattern. Then, the actual spiking activity $\mathbf{r}^t$ was approximated with a firing rate function $\bar{\mathbf{r}}^t$ induced by the target patterns where $\bar{\mathbf{r}}^t$ is a periodic function with a period equal to the target length (m-steps) and repeated d-times to cover the entire training duration, i.e., $n = d \cdot m$.

To make it clear that $\bar{\mathbf{r}}$ is periodic in time, we used a time index $\bar{t}$ that denotes the time point within the target pattern that have period m. Specifically, the spiking activity is expressed as

$$\mathbf{r}^t = \bar{\mathbf{r}}^{\bar{t}} + \boldsymbol{\varepsilon}_s^{\bar{t}} \quad \text{for} \quad t = 1, ..., n \tag{31}$$

where $t = sm + \bar{t}$ and $s = 1, ..., d$ indicate the training iteration number. We introduced a noise term $\boldsymbol{\varepsilon}_s^{\bar{t}}$ whose i^{th} component is the fluctuation of neuron i's spiking activity around the firing rate $\bar{\mathbf{r}}^{\bar{t}}$ at the sth training session. For the following analysis, we assumed a simple noise model that $\boldsymbol{\varepsilon}_s^{\bar{t}}$ is a multivariate random variable at each $\bar{t}$ with no correlations between neurons within and across training sessions

$$\text{Cov}[\boldsymbol{\varepsilon}_s^{\bar{t}}, \boldsymbol{\varepsilon}_s^{\bar{t}}] = \text{E}[\boldsymbol{\varepsilon}_s^{\bar{t}} (\boldsymbol{\varepsilon}_s^{\bar{t}})'] = \alpha I_K \quad \text{for} \quad 1 \leq s \leq d \tag{32}$$

$$\text{Cov}[\boldsymbol{\varepsilon}_s^{\bar{t}}, \boldsymbol{\varepsilon}_q^{\bar{t}}] = \text{E}[\boldsymbol{\varepsilon}_s^{\bar{t}} (\boldsymbol{\varepsilon}_q^{\bar{t}})'] = 0 \quad \text{for} \quad s \neq q \text{ and } 1 \leq s, q \leq d. \tag{33}$$

To derive an alternative expression for $\boldsymbol{\sigma}_\alpha^n$ amenable to asymptotic analysis, we expressed the corre-

lation of spiking activities in terms of the firing rates and fluctuations around them:

$$\sum_{t=1}^{n} \mathbf{r}^t(\mathbf{r}^t)' = \sum_{\bar{t}=1}^{m}\sum_{s=1}^{d}(\bar{\mathbf{r}}^{\bar{t}} + \boldsymbol{\varepsilon}_s^{\bar{t}})(\bar{\mathbf{r}}^{\bar{t}} + \boldsymbol{\varepsilon}_s^{\bar{t}})' \tag{34}$$

$$= \sum_{\bar{t}=1}^{m}\left(d\,\bar{\mathbf{r}}^{\bar{t}}(\bar{\mathbf{r}}^{\bar{t}})' + d\,\bar{\mathbf{r}}^{\bar{t}}(\boldsymbol{\varepsilon}^{\bar{t}})' + d\,\boldsymbol{\varepsilon}^{\bar{t}}(\bar{\mathbf{r}}^{\bar{t}})' + \sum_{s=1}^{d}\boldsymbol{\varepsilon}_s^{\bar{t}}(\boldsymbol{\varepsilon}_s^{\bar{t}})'\right) \tag{35}$$

$$\approx d\sum_{\bar{t}=1}^{m}\left(\bar{\mathbf{r}}^{\bar{t}}(\bar{\mathbf{r}}^{\bar{t}})' + \bar{\mathbf{r}}^{\bar{t}}(\boldsymbol{\varepsilon}^{\bar{t}})' + \boldsymbol{\varepsilon}^{\bar{t}}(\bar{\mathbf{r}}^{\bar{t}})' + d\,\boldsymbol{\varepsilon}^{\bar{t}}(\boldsymbol{\varepsilon}^{\bar{t}})'\right) \tag{36}$$

$$\approx d\sum_{\bar{t}=1}^{m}(\bar{\mathbf{r}}^{\bar{t}} + \sqrt{d}\,\bar{\boldsymbol{\varepsilon}}^{\bar{t}})(\bar{\mathbf{r}}^{\bar{t}} + \sqrt{d}\,\bar{\boldsymbol{\varepsilon}}^{\bar{t}})'. \tag{37}$$

The noise approximation in Eq. 36 is equivalent to following relationship

$$d\,\bar{\boldsymbol{\varepsilon}}^{\bar{t}}(\bar{\boldsymbol{\varepsilon}}^{\bar{t}})' = d^{-1}\sum_{s=1}^{d}\boldsymbol{\varepsilon}_s^{\bar{t}}(\boldsymbol{\varepsilon}_s^{\bar{t}})' + d^{-1}\sum_{s=1}^{d}\sum_{q\neq s}\boldsymbol{\varepsilon}_s^{\bar{t}}(\boldsymbol{\varepsilon}_q^{\bar{t}})' \approx d^{-1}\sum_{s=1}^{d}\boldsymbol{\varepsilon}_s^{\bar{t}}(\boldsymbol{\varepsilon}_s^{\bar{t}})' \tag{38}$$

where

$$\bar{\boldsymbol{\varepsilon}}^{\bar{t}} = d^{-1}\sum_{s=1}^{d}\boldsymbol{\varepsilon}_s^{\bar{t}} \tag{39}$$

and the cross-training correlations in Eq. 38 was ignored by the noise assumption given in Eq. 33. In Eqs. 36 and 37, the dominant term is $d\,\boldsymbol{\varepsilon}^{\bar{t}}(\boldsymbol{\varepsilon}^{\bar{t}})'$, so $\bar{\mathbf{r}}^{\bar{t}}(\boldsymbol{\varepsilon}^{\bar{t}})'$ and $\boldsymbol{\varepsilon}^{\bar{t}}(\bar{\mathbf{r}}^{\bar{t}})'$ in Eq. 36 were replaced by $\sqrt{d}\,\bar{\mathbf{r}}^{\bar{t}}(\boldsymbol{\varepsilon}^{\bar{t}})'$ and $\sqrt{d}\,\boldsymbol{\varepsilon}^{\bar{t}}(\bar{\mathbf{r}}^{\bar{t}})'$ in Eq. 37.

Next, we substituted the approximated $\sum_t \mathbf{r}^t(\mathbf{r})'$ (Eq. 37) into the definition of P^n (Eq. 29) and suppressed $O(\sqrt{d})$-terms to obtain

$$\widehat{\boldsymbol{\sigma}}_\alpha^d = [I_K - P^0\bar{R}(d^{-1}I_m + \bar{R}'P^0\bar{R} + d\,\mathcal{E}'P^0\mathcal{E})^{-1}\bar{R}']\boldsymbol{\sigma}_\alpha^0 \tag{40}$$

where

$$\bar{R} = [\bar{\mathbf{r}}^1, ..., \bar{\mathbf{r}}^m] \quad \text{and} \quad \bar{\mathcal{E}} = [\bar{\boldsymbol{\varepsilon}}^1, ..., \bar{\boldsymbol{\varepsilon}}^m]. \tag{41}$$

Note that for uncorrelated noise, the term $\mathcal{E}'P^0\mathcal{E}$ yields an identity matrix which effectively serves as a regularizer that makes it possible to invert $\bar{R}'P^0\bar{R} + d\,\mathcal{E}'P^0\mathcal{E}$.

We verified numerically that the actual $\boldsymbol{\sigma}_\alpha^n$ can be well approximated by $\widehat{\boldsymbol{\sigma}}_\alpha^d$ (Fig. 5b). Moreover, since $\|\widehat{\boldsymbol{\sigma}}_\alpha^d\|_2$ decreases monotonically with d, there exists a constant d_0 such that

$$\|\widehat{\boldsymbol{\sigma}}_\alpha^d\|_2 \leq \|\widehat{\boldsymbol{\sigma}}_\alpha^{d_0}\|_2 \quad \text{for} \quad d \gg d_0 \gg 1. \tag{42}$$

3.3 Upper bound on the size of synaptic updates

By its definition in Eq. 40, $\widehat{\boldsymbol{\sigma}}_\alpha^{d_0}$ can be expressed in terms of the initial value σ_α^0

$$\widehat{\boldsymbol{\sigma}}_\alpha^{d_0} = \sigma_\alpha^0\mathbf{v}_\alpha \tag{43}$$

$$\text{where} \quad \mathbf{v}_\alpha = [I_K - P^0\bar{R}(d_0^{-1}I_m + \bar{R}'P^0\bar{R} + d_0\,\mathcal{E}'P^0\mathcal{E})^{-1}\bar{R}']\mathbf{1}_\alpha. \tag{44}$$

Next we made two observations regarding the upper bound on $\mathbf{v}_\alpha$. First, it can be bounded by a constant C_α that does not depend on d or n since $P^0 = P^0(\mu, \lambda, K)$ is a function of the penalty terms μ, λ and the number of synapses K, and $\bar{R}$ and $\bar{\mathcal{E}}$ are functions of spiking rate $\bar{\mathbf{r}}$ and noise $\bar{\boldsymbol{\varepsilon}}$, respectively, thus

$$\|\mathbf{v}_\alpha\|_2 \leq C_\alpha. \tag{45}$$

Second, the constant C_α does not depend on μ and K_α if they become sufficiently large. This can be seen from the fact that P^0 converged elementwise as μ or K_α become arbitrarily large. Specifically,

$$P^0 \approx \frac{1}{\lambda} I_K - \frac{1}{\lambda K_E} \mathbf{1}_E \mathbf{1}_E' - \frac{1}{\lambda K_I} \mathbf{1}_I \mathbf{1}_I' \quad \text{if} \quad \mu \gg 1 \tag{46}$$

$$P^0 \approx \frac{1}{\lambda} I_K \quad \text{if} \quad K_E, K_I \gg 1. \tag{47}$$

These observations led to the uniform upper bound

$$\|\widehat{\boldsymbol{\sigma}}_\alpha^d\|_2 \leq \|\widehat{\boldsymbol{\sigma}}_\alpha^{d_0}\|_2 \leq C_\alpha \sigma_\alpha^0 \quad \text{for} \quad d \gg d_0 \gg 1, \tag{48}$$

where C_α does not depend on d, μ and K_α.

Applying such upper bound to the synaptic learning rule (Eq. 25), we showed that the update size of the sums of excitatory and inhibitory synaptic weights, respectively, at each update are constrained by

$$|\delta z_\alpha^n| \approx |e^n \widehat{\boldsymbol{\sigma}}_\alpha^d \cdot \boldsymbol{r}^n| \tag{49}$$

$$\leq |e^n| \cdot \|\widehat{\boldsymbol{\sigma}}_\alpha^{d_0}\|_2 \cdot \|\mathbf{r}^n\|_2 \tag{50}$$

$$\leq |e^n| \cdot C_\alpha \sigma_\alpha^0 \cdot K r_{max} \tag{51}$$

$$\approx |e^n| B_\alpha \mu^{-1} \quad \text{for all} \quad n \tag{52}$$

where $B_\alpha = C_\alpha r_{max} K / K_\alpha$. We used $\boldsymbol{\sigma}_\alpha^n \approx \widehat{\boldsymbol{\sigma}}_\alpha^d$ to approximate $|\delta z_\alpha^n|$ in the first line, the uniform upper bound on $\widehat{\boldsymbol{\sigma}}_\alpha^d$ (Eq. 48) in the second and third lines, and $\sigma_\alpha^0 = O(1/(\mu K_\alpha))$ (Eq. 28) in the last line.

This estimate showed that the upper bound on the sum of synaptic weights (1) is inversely proportional to μ and (2) does not depend on the number of synaptic connections K_α modified under the ROWSUM-regularization as long as the proportion $q_\alpha = K_\alpha / K$ remains constant. Finally, we concluded that the upper bound on the average synaptic weights was inversely proportional to the product μK_α:

$$|\overline{\delta z_\alpha^n}| = K_\alpha^{-1} |\delta z_\alpha^n| \leq |e^n| B_\alpha (\mu K_\alpha)^{-1}. \tag{53}$$

Acknowledgments

This work was supported by the Intramural Program of the NIH, NIDDK.

References

[1] Valerio Mante, David Sussillo, Krishna V Shenoy, and William T Newsome. Context-dependent computation by recurrent dynamics in prefrontal cortex. *nature*, 503(7474):78–84, 2013.

[2] 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.

[3] Warasinee Chaisangmongkon, Sruthi K Swaminathan, David J Freedman, and Xiao-Jing Wang. Computing by robust transience: how the fronto-parietal network performs sequential, category-based decisions. *Neuron*, 93(6):1504–1517, 2017.

[4] Jing Wang, Devika Narain, Eghbal A Hosseini, and Mehrdad Jazayeri. Flexible timing by temporal scaling of cortical responses. *Nature neuroscience*, 21(1):102–110, 2018.

[5] Ashok Litwin-Kumar and Srinivas C Turaga. Constraining computational models using electron microscopy wiring diagrams. *Current Opinion in Neurobiology*, 58:94–100, October 2019.

[6] Claire Eschbach, Akira Fushiki, Michael Winding, Casey M. Schneider-Mizell, Mei Shao, Rebecca Arruda, Katharina Eichler, Javier Valdes-Aleman, Tomoko Ohyama, Andreas S. Thum, Bertram Gerber, Richard D. Fetter, James W. Truman, Ashok Litwin-Kumar, Albert Cardona, and Marta Zlatic. Recurrent architecture for adaptive regulation of learning in the insect brain. *Nature Neuroscience*, 23(4):544–555, April 2020. Number: 4 Publisher: Nature Publishing Group.

[7] Aref Arzan Zarin, Brandon Mark, Albert Cardona, Ashok Litwin-Kumar, and Chris Q Doe. A multilayer circuit architecture for the generation of distinct locomotor behaviors in Drosophila. *eLife*, 8:e51781, December 2019. Publisher: eLife Sciences Publications, Ltd.

[8] H Francis Song, Guangyu R Yang, and Xiao-Jing Wang. Training excitatory-inhibitory recurrent neural networks for cognitive tasks: a simple and flexible framework. *PLoS computational biology*, 12(2), 2016.

[9] Wilten Nicola and Claudia Clopath. Supervised learning in spiking neural networks with force training. *Nature communications*, 8(1):2208, 2017.

[10] Brian DePasquale, Mark M Churchland, and LF Abbott. Using firing-rate dynamics to train recurrent networks of spiking model neurons. *arXiv preprint arXiv:1601.07620*, 2016.

[11] Robert Kim, Yinghao Li, and Terrence J Sejnowski. Simple framework for constructing functional spiking recurrent neural networks. *Proceedings of the National Academy of Sciences*, 116(45):22811–22820, 2019.

[12] Alessandro Ingrosso and LF Abbott. Training dynamically balanced excitatory-inhibitory networks. *PloS one*, 14(8), 2019.

[13] Cody Baker, Vicky Zhu, and Robert Rosenbaum. Nonlinear stimulus representations in neural circuits with approximate excitatory-inhibitory balance. *BioRxiv*, page 841684, 2019.

[14] Sompolinsky, Crisanti, and Sommers. Chaos in random neural networks. *Physical review letters*, 61(3):259–262, jul 1988.

[15] Carl van Vreeswijk and Haim Sompolinsky. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, 274(5293):1724, 1996.

[16] 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.

[17] 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.

[18] Srdjan Ostojic. Two types of asynchronous activity in networks of excitatory and inhibitory spiking neurons. *Nature neuroscience*, 17(4):594–600, 2014.

[19] Misha V Tsodyks, William E Skaggs, Terrence J Sejnowski, and Bruce L McNaughton. Paradoxical effects of external modulation of inhibitory interneurons. *Journal of neuroscience*, 17(11):4382–4388, 1997.

[20] Simon Haykin. *Adaptive Filter Theory (3rd Ed.)*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1996.

[21] David Sussillo and Larry Abbott. Generating Coherent Patterns of Activity from Chaotic Neural Networks. *Neuron*, 63(4):544–557, 2009.

[22] Christopher M Kim and Carson C Chow. Learning recurrent dynamics in spiking networks. *eLife*, 7:e37124, 2018.

[23] A Sanzeni, Bradley Akitake, Hannah C Goldbach, Caitlin E Leedy, Nicolas Brunel, and Mark H Histed. Inhibition stabilization is a widespread property of cortical networks. *bioRxiv*, page 656710, 2019.

[24] Kanaka Rajan, Christopher D. Harvey, and David W. Tank. Recurrent Network Models of Sequence Generation and Memory. *Neuron*, 90(1):128–142, 2016.

[25] 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.

[26] Brendan K Murphy and Kenneth D Miller. Balanced amplification: a new mechanism of selective amplification of neural activity patterns. *Neuron*, 61(4):635–648, 2009.

[27] Guillaume Bellec, Darjan Salaj, Anand Subramoney, Robert Legenstein, and Wolfgang Maass. Long short-term memory and learning-to-learn in networks of spiking neurons. In *Advances in Neural Information Processing Systems*, pages 787–797, 2018.

[28] Guillaume Bellec, Franz Scherr, Anand Subramoney, Elias Hajek, Darjan Salaj, Robert Legenstein, and Wolfgang Maass. A solution to the learning dilemma for recurrent networks of spiking neurons. *bioRxiv*, page 738385, 2019.

[29] 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.