
Minimax Dynamics of Optimally Balanced Spiking Networks of Excitatory and Inhibitory Neurons

Qianyi Li

Biophysics Graduate Program
Harvard University
Cambridge, MA 02138
qianyi_li@g.harvard.edu

Cengiz Pehlevan

John A. Paulson School of Engineering and Applied Sciences
Harvard University
Cambridge, MA 02138
cpehlevan@seas.harvard.edu

Abstract

Excitation-inhibition balance is ubiquitously observed in the cortex. Recent studies suggest an intriguing link between balance on fast timescales, tight balance, and efficient information coding with spikes. We further this connection by taking a principled approach to optimal balanced networks of excitatory (E) and inhibitory (I) neurons. By deriving E-I spiking neural networks from greedy spike-based optimizations of constrained minimax objectives, we show that tight balance arises from correcting for deviations from the minimax optima. We predict specific neuron firing rates in the networks by solving minimax problems, going beyond statistical theories of balanced networks. We design minimax objectives for reconstruction of an input signal, associative memory, and storage of manifold attractors, and derive from them E-I networks that perform the computation. Overall, we present a novel normative modeling approach for spiking E-I networks, going beyond the widely-used energy-minimizing networks that violate Dale’s law. Our networks can be used to model cortical circuits and computations.

1 Introduction

While spiking neural networks empower our brains, a thorough understanding of how spikes accurately represent and transmit information is lacking. Any explanation should address a striking and widely observed property of cortical spiking networks that excitatory (E) and inhibitory (I) currents for individual neurons are balanced (detailed balance), and spikes are only generated during the brief occasions when inhibition fails to track spontaneous fluctuations in excitation (tight balance) [1–8]. In this work, we show that tight and detailed balance could arise from correcting for deviations from a minimax optimum that governs the dynamics of E-I balanced networks. Having access to an objective allows us to design spiking E-I networks that perform various biologically relevant functions.

While the theory of balanced networks goes further back [9–11], the idea that tight and detailed balance is related to the efficiency of spike coding was proposed in recent influential theoretical work [6, 12, 13]. These results are typically limited to spiking dynamics performing a greedy-minimization of a reconstruction error loss function without considering more diverse computations carried out in the cortex. Moreover, minimizing an energy function leads to symmetric connection weights [14] that violate, among other biological constraints, Dale’s law: a neuron’s influence on its post-synaptic neurons are either all excitatory or all inhibitory. This violation has been addressed by introducing separate reconstruction error cost functions for E and I neurons [12, 15]. While E neurons reconstruct an input signal, I neurons reconstruct the non-Dalian part of recurrent E interactions, assuming that I neuron dynamics equilibrates faster than E neurons. Our work extends these previous accounts in two ways.

First, we take a more principled approach and propose a common minimax dynamics that E and I spiking neurons collectively optimize, adopting an idea from previous work on networks of rate-coding (continuous outputs) E-I neurons [16]. The minimax approach, besides other benefits, provides an intriguing interpretation of the antagonism between E and I neural populations.

Second, we consider minimax objectives that perform functions beyond signal reconstruction. Energy minimizing networks with symmetric interactions have been a powerful modeling framework for neuroscience since Hopfield’s seminal contribution [14]. Phenomena such as associative memory [14], oculomotor integration [17, 18], coding of periodic variables like head direction [19–23] and grating orientation [24], and spatial navigation [25, 26] were modeled with fixed-point or manifold attractors resulting from such energy minimizing neural dynamics. Here, by moving from energy minimization to minimax optimization, we extend the reach of normative modeling to E-I spiking networks and provide derivations of circuits for each of the functions listed above.

Our technical contributions in this paper are:

- Derivation of a greedy spiking algorithm from a minimax objective. We discuss the optimality conditions leading to firing rate predictions of individual neurons and detailed balance. We argue that greedy spiking dynamics leads to tight balance, and provide necessary conditions for convergence of the dynamics.
- Applications. We design spiking networks that reconstruct signals, and exhibit fixed-point and manifold attractors, while obeying Dale’s law and remaining in tight and detailed balance. We verify our theory in simulations. These applications indicate that our approach offers a principled method for designing spiking neural network models for cortical functions.

2 Minimax dynamics of optimally balanced E-I spiking networks

In this section, we consider the spiking dynamics of an integrate-and-fire E-I network, and show how such dynamics can be derived from a minimax objective function as a greedy optimization algorithm. By analyzing the optimum of the constrained minimax objective, we observe that the network is in detailed balance and derive conditions for stability of the optimum. Finally, we derive conditions for convergence of the dynamics and demonstrate tight balance. Our derivation extends the methods of [12, 13] to minimax optima.

Our network is composed of separate E and I neuron populations of N^E and N^I neurons (Fig. 1a). For simplicity, only the E neurons receive N^0 dimensional external inputs, although our results and analysis still hold when both populations receive external inputs. The spiking dynamics is given by

$$\begin{aligned} \frac{dV_i^E}{dt} &= -\frac{V_i^E}{\tau_E} + \sum_j W_{ij}^{EE} s_j^E - \sum_j W_{ij}^{EI} s_j^I + \sum_j F_{ij} s_j^0, \\ \frac{dV_i^I}{dt} &= -\frac{V_i^I}{\tau_I} + \sum_j W_{ij}^{IE} s_j^E - \sum_j W_{ij}^{II} s_j^I. \end{aligned} \quad (1)$$

Here, V_i^E , $i = 1, \dots, N^E$ and V_i^I , $i = 1, \dots, N^I$, denote the membrane potentials for E and I neurons respectively, τ_E and τ_I are the corresponding membrane time constants, and s_j^E and s_j^I denote the spike trains of E and I neurons: e.g. $s_j^{E(I)}(t) = \sum_k \delta(t - t_{j,k})$, where $t_{j,k}$ is the time of the k -th spike of the j -th neuron. s_j^0 , $j = 1, \dots, N^0$ denotes the input signal, which is not required to be a spike train. $\mathbf{F} \in \mathbb{R}^{N^E \times N^0}$ is the feed-forward connectivity matrix and $\mathbf{W}^{EE, EI, IE, II}$ are the connectivity matrices within and between the E-I populations. We require $\mathbf{W}^{EE} \in \mathbb{R}^{N^E \times N^E}$ and $\mathbf{W}^{II} \in \mathbb{R}^{N^I \times N^I}$ to be symmetric and $\mathbf{W}^{IE} = \mathbf{W}^{EI\top}$ for our minimax objective approach. All the off-diagonal elements of the weight matrices are non-negative so that the network obeys Dale’s law. The spiking reset is incorporated into the diagonals of \mathbf{W}^{EE} and \mathbf{W}^{II} , which define how much the membrane potentials decrease after the arrival of a spike. Therefore, the diagonal elements of \mathbf{W}^{EE} are negative and \mathbf{W}^{II} are positive. The spiking thresholds for the E and I neurons are given by $T_i^E = -\frac{1}{2}W_{ii}^{EE}$ and $T_i^I = \frac{1}{2}W_{ii}^{II}$ respectively. We can obtain implicit expressions for V_i^E and

V_i^I by integrating Eq. (1):

$$\begin{aligned} V_i^E(t) &= \sum_j W_{ij}^{EE} r_j^E(t) - \sum_j W_{ij}^{EI} r_j^I(t) + \sum_j F_{ij} x_j(t), \\ V_i^I(t) &= \sum_j W_{ij}^{IE} r_j^E(t) - \sum_j W_{ij}^{II} r_j^I(t), \end{aligned} \quad (2)$$

where $r_j^{E,I}$ and x_j are defined by filtering the spike train or the input signal with an exponential kernel:

$$r_j^{E,I}(t) = \int_0^\infty e^{-(t-t')/\tau_{E,I}} s_j^{E,I}(t-t') dt', \quad x_j(t) = \int_0^\infty e^{-(t-t')/\tau_E} s_j^0(t-t') dt'. \quad (3)$$

Our intuition is that the network is in a balanced state when the dynamics reaches the optimum of an objective. Detailed balance of E and I inputs require that $V_i^{E,I}$ should fluctuate around zero for each neuron i [6]. Tight balance requires the imbalance between E and I to be short-lived [6]. As we show next, both of these goals are achieved when the dynamics Eq. (1) performs a greedy optimization of a minimax objective S , where the implicit expressions Eq. (2) correspond to the saddle-point conditions.

Spiking dynamics is a greedy algorithm optimizing a minimax objective: Because I to E connections and E to I connections have opposite signs in Eq.s (1) and (2), a network that obeys Dale's law cannot be derived from a single minimizing objective. We instead consider the minimax optimization problem:

$$\min_{\mathbf{r}^E \geq 0} \max_{\mathbf{r}^I \geq 0} S(\mathbf{r}^E, \mathbf{r}^I), \quad S = -\frac{1}{2} \mathbf{r}^{E\top} \mathbf{W}^{EE} \mathbf{r}^E + \mathbf{r}^{E\top} \mathbf{W}^{EI} \mathbf{r}^I - \frac{1}{2} \mathbf{r}^{I\top} \mathbf{W}^{II} \mathbf{r}^I - \mathbf{x}^\top \mathbf{F}^\top \mathbf{r}^E. \quad (4)$$

We can derive from our objective function a greedy algorithm that performs the optimization, which corresponds to the spiking dynamics Eq. (1), by adopting a similar approach to [12, 13]. The details of the derivation is provided in Supplementary Information (SI) A. Here we outline the approach. We track filtered spike trains, $r_i^{E,I}$ given in Eq. (3), instantaneously. At each time, each E/I neuron makes an individual decision to spike. An E neuron spikes if the spike decreases the instantaneous value of S : $S(\mathbf{r}^E + \mathbf{e}_i, \mathbf{r}^I) < S(\mathbf{r}^E, \mathbf{r}^I)$, where \mathbf{e}_i is the i -th standard basis vector. An I neuron spikes if the spike increases the instantaneous value of S : $S(\mathbf{r}^E, \mathbf{r}^I + \mathbf{e}_i) > S(\mathbf{r}^E, \mathbf{r}^I)$. This series of spiking decisions lead to the spiking dynamics defined in Eq. (1) when $\tau_E = \tau_I = \tau$, and the dynamics in Eq.s (2) and (3) for the general case where the time constants can be different.

KKT conditions imply detailed and tight balance of active neurons: We can verify that the detailed balance is reached at the optimum of the constrained minimax problem by examining the KKT conditions [27] obtained by considering the optimizations with respect to \mathbf{r}^I and \mathbf{r}^E sequentially. These KKT conditions are satisfied for local optimal points of the minimax problem [28, 29]. We obtain, $\forall i \in \{1, \dots, N^I\}$

$$\begin{aligned} r_i^I = 0, \quad \text{and} \quad V_i^I &= \sum_j W_{ji}^{EI} r_j^E - \sum_j W_{ij}^{II} r_j^I \leq 0, \quad \text{and} \quad \lambda_i^I \geq 0, \\ \text{or} \quad r_i^I > 0, \quad \text{and} \quad V_i^I &= \sum_j W_{ji}^{IE} r_j^E - \sum_j W_{ij}^{II} r_j^I = 0, \quad \text{and} \quad \lambda_i^I = 0, \end{aligned} \quad (5)$$

and $\forall i \in \{1, \dots, N^E\}$

$$\begin{aligned} r_i^E = 0, \quad \text{and} \quad V_i^E &= \sum_j W_{ij}^{EE} r_j^E - \sum_j W_{ij}^{EI} r_j^I + \sum_j F_{ij} x_j \leq 0, \quad \text{and} \quad \lambda_i^E \geq 0, \\ \text{or} \quad r_i^E > 0, \quad \text{and} \quad V_i^E &= \sum_j W_{ij}^{EE} r_j^E - \sum_j W_{ij}^{EI} r_j^I + \sum_j F_{ij} x_j = 0, \quad \text{and} \quad \lambda_i^E = 0. \end{aligned} \quad (6)$$

From the KKT conditions, we see that for active neurons with $r_i^{E,I} > 0$, we have $V_i^{E,I} = 0$ at the saddle point, which suggests that any neuron that is active is in detailed balance.

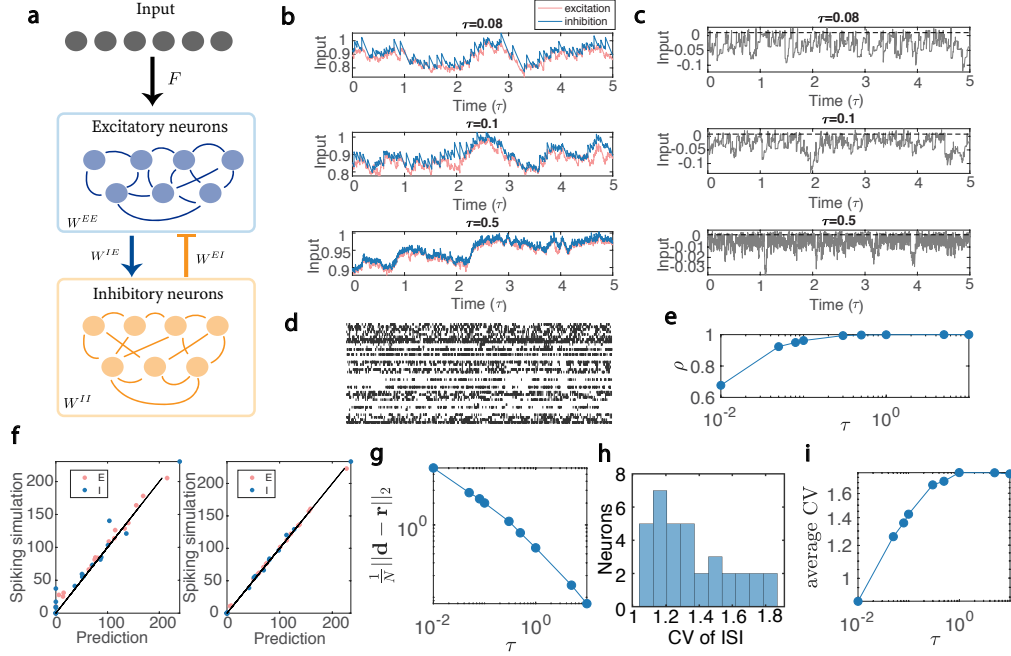


Figure 1: Optimally balanced E-I network. (a) Network of E-I neurons that obeys the Dale’s law. Weights are designed to satisfy conditions in Eq.s (7) and (8). (b) E input ($\mathbf{W}^{EE}\mathbf{r}^E + \mathbf{F}\mathbf{x}$) and I input ($\mathbf{W}^{EI}\mathbf{r}^I$) for an active E neuron with different time constants τ , normalized by the maximum value of the E input. (c) Net input ($\mathbf{W}^{EE}\mathbf{r}^E + \mathbf{F}\mathbf{x} - \mathbf{W}^{EI}\mathbf{r}^I$), normalized by the maximum value of the E input. Net input fluctuates around zero, occasionally going above the threshold (dashed line, also normalized). (d) Spiking pattern for an input with constant firing rate for 5τ , $\tau = 0.5$, corresponding to the same simulation time period as shown in the bottom panel of (b). (e) Pearson correlation coefficient between E and I input for an active neuron with different time constants τ . The network becomes more tightly balanced with increasing τ . (f) Prediction of firing rate for E and I populations. The prediction is obtained by directly solving the KKT conditions of the minimax optimization problem. Left: $\tau = 0.1$, Right: $\tau = 1$. Larger time constant results in more accurate prediction. (g) The error of firing rate prediction exhibits a power law dependence on the time constant. (h) Distribution of inter-spike interval (ISI) coefficient of variation (CV) for active neurons ($\tau = 0.08$). Most of the neurons have CV close to 1. (i) The average CV for active neurons increases as τ becomes larger, but remains close to 1, and saturates around 1.7.

Because of the greedy nature of optimization, when the network state deviates from the balanced state, the dynamics will automatically self-correct (imperfectly) by producing spikes when V_i is larger than the spiking threshold V_{th} , leading to tight balance. For the inactive (non-spiking) neurons, KKT conditions state that the voltage $V_i < 0$. This is smaller than a positive V_{th} , consistent with the neuron never spiking.

We verify our conclusions with simulations. The E and I inputs for active neurons closely track each other (Fig. 1b), and their total input fluctuates around zero (Fig. 1c). Spike raster plot of the network is shown in Fig. 1d for $\tau = 0.5$. The network is more tightly balanced with increasing time constants τ as the Pearson correlation ρ between E and I inputs for the same active neuron increases (Fig. 1e). We can predict individual neuron firing rates (defined as $f_i^{E,I}(t) = \frac{1}{\tau}r_i^{E,I}(t)$, a normalized version of $r_i^{E,I}$) in this spiking neural network by directly solving KKT conditions Eq.s (5) and (6) of the quadratic minimax problem Eq. (4). As shown in Fig. 1f, the prediction is quite accurate. We observe that the firing rate prediction becomes more accurate as τ increases (Fig. 1f&g). Similar observation was made for reconstruction error minimizing spiking networks [13]. Despite being tightly balanced, the spiking is also irregular. The coefficient of variation (CV) for active neurons are close to 1 (Fig. 1i), and the distribution of ISI is close to exponential (not shown). CV also increases as τ becomes larger, but it remains close to 1.

Further optimality conditions: KKT conditions are necessary but not sufficient. The second order sufficient conditions [27–29] of the minimax problem (by considering maximization and minimization sequentially, details in SI B.1) give us

$$\hat{\mathbf{W}}^{II} \succcurlyeq 0, \quad \hat{\mathbf{W}}^{EI} \hat{\mathbf{W}}^{II^{-1}} \hat{\mathbf{W}}^{IE} - \hat{\mathbf{W}}^{EE} \succcurlyeq 0, \quad (7)$$

where $\hat{\mathbf{W}}^{II}$ is the principal submatrix of \mathbf{W}^{II} whose columns and rows correspond to nonzero elements of \mathbf{r}^I , $\hat{\mathbf{W}}^{EE}$ is the principal submatrix of \mathbf{W}^{EE} whose columns and rows correspond to nonzero elements of \mathbf{r}^E , and $\hat{\mathbf{W}}^{EI}$ is a submatrix of \mathbf{W}^{EI} with its rows corresponding to nonzero elements of \mathbf{r}^E and its columns corresponding to nonzero elements of \mathbf{r}^I .

Convergence conditions: The dynamics of approaching a saddle point of our minimax objective can be quite different than that of a minimum. The conditions discussed above describe an optimum of the objective, however, unlike in a minimization problem, where the objective itself act as Lyapunov function of the dynamics, these optimality conditions do not guarantee that the spiking dynamics converges to the optimum. The convergence of spiking dynamics is challenging to prove. Instead, we characterize the convergence conditions for a rate dynamics optimizing the minimax objective Eq. (4), hypothesizing that similar conditions would hold for the spiking dynamics. This thinking is also motivated by previous work that proved the convergence of a certain type of spiking dynamics to the fixed point of a rate dynamics, both minimizing the reconstruction error cost function [30, 31]. By constructing an energy (or Lyapunov) function [16, 32] on the rate dynamics, we derive the sufficient convergence conditions below (see SI B.2 for details):

1. The following eigenvalue condition guarantees the convergence of any bounded trajectory that does not exhibit a change in the set of active/inactive neurons through its evolution:

$$\lambda_{\min}(\hat{\mathbf{W}}^{II}) > \lambda_{\max}(\hat{\mathbf{W}}^{EE}). \quad (8)$$

2. Replacing the positive semidefiniteness in the second order sufficient conditions for the optimality of the objective (Eq. (7)) with strict positive definiteness guarantees the boundedness of trajectories.

We empirically observed that these conditions are also valid for the convergence of spiking dynamics.

3 Applications

Having access to an objective function allows us to design balanced E-I networks to perform specific functions. Such normative approach has been common in neuroscience, using energy minimizing dynamics to derive efficient coding circuits [12, 15, 33–37]. Energy functions for various types of attractors have been thoroughly studied: fixed point attractor networks are commonly used as models for associative memory [14], ring attractor networks are used to model head-direction systems [19–23] and orientation selectivity [24] and grid attractor models are used to model grid cell responses [25, 26]. Here, we revisit all these systems and show how balanced E-I circuits performing the same tasks can be obtained by designing the minimax objectives and weights.

3.1 Input reconstruction

Consider the typical objective for a signal reconstruction problem: minimization of the mean squared error with an l_2 -regularizer on the response,

$$\arg \min_{\mathbf{r}^E} \|\mathbf{x} - \mathbf{F}^\top \mathbf{r}^E\|_2^2 + \frac{\lambda}{2} \|\mathbf{r}^E\|_2^2 = \arg \min_{\mathbf{r}^E} -\mathbf{x}^\top \mathbf{F}^\top \mathbf{r}^E + \frac{1}{2} \mathbf{r}^{E\top} \mathbf{F} \mathbf{F}^\top \mathbf{r}^E + \frac{\lambda}{2} \mathbf{r}^{E\top} \mathbf{r}^E. \quad (9)$$

Our goal is to transform this problem to the one in Eq. (4), allowing a mapping to an E-I network. To achieve our goal, we first perform a matrix factorization $\mathbf{F} \mathbf{F}^\top = \mathbf{U} \mathbf{\Sigma} \mathbf{U}^\top$ where $\mathbf{U} \in \mathbb{R}^{N^E \times N^I}$, $\mathbf{\Sigma} \in \mathbb{R}_{\geq 0}^{N^I \times N^I}$. We emphasize that the elements of $\mathbf{\Sigma}$ are non-negative, $\mathbf{\Sigma}$ is symmetric and can be, but does not have to be, diagonal. For the factorization to be plausible $N^I \geq \text{rank}(\mathbf{F})$. We will show examples of two different ways of performing the factorization, 1) by simply choosing $\mathbf{\Sigma} = \mathbf{I}$ and $\mathbf{U} = \mathbf{F}$, and 2) performing a singular value decomposition (SVD) on $\mathbf{F} \mathbf{F}^\top$. We also separate out the positive and negative parts in \mathbf{U} , and denote them by $[\mathbf{U}]_+ = \mathbf{U}_+$, and $[\mathbf{U}]_- = \mathbf{U}_-$. With all this set

up, the reconstruction error cost function can be written as:

$$\min_{\mathbf{r}^E} -\mathbf{x}^\top \mathbf{F}^\top \mathbf{r}^E + \mathbf{r}^{E\top} \left(\mathbf{U}_+ \Sigma \mathbf{U}_-^\top + \mathbf{U}_- \Sigma \mathbf{U}_+^\top + \frac{\lambda}{2} \mathbf{I} \right) \mathbf{r}^E + \frac{1}{2} \mathbf{r}^{E\top} (\mathbf{U}_+ - \mathbf{U}_-) \Sigma (\mathbf{U}_+ - \mathbf{U}_-)^\top \mathbf{r}^E. \quad (10)$$

Next, using a ‘‘variable substitution’’ trick commonly used in statistical mechanics and the similarity matching framework [38], we transform the problem into a minimax problem and introduce the I neurons:

$$\min_{\mathbf{r}^E} \max_{\mathbf{r}^I} -\mathbf{x}^\top \mathbf{F}^\top \mathbf{r}^E + \mathbf{r}^{E\top} \left(\mathbf{U}_+ \Sigma \mathbf{U}_-^\top + \mathbf{U}_- \Sigma \mathbf{U}_+^\top + \frac{\lambda}{2} \mathbf{I} \right) \mathbf{r}^E + \mathbf{r}^{I\top} \Sigma (\mathbf{U}_+ - \mathbf{U}_-) \mathbf{r}^E - \frac{1}{2} \mathbf{r}^{I\top} \Sigma \mathbf{r}^I. \quad (11)$$

The equivalence of Eq. (11) to Eq. (10) can be seen by performing the \mathbf{r}^I maximization explicitly. With this formulation, we have $\mathbf{W}^{EE} = 2(\mathbf{U}_+ \Sigma \mathbf{U}_-^\top + \mathbf{U}_- \Sigma \mathbf{U}_+^\top + \frac{\lambda}{2} \mathbf{I})$, $\mathbf{W}^{IE} = \Sigma (\mathbf{U}_+ - \mathbf{U}_-)$, and $\mathbf{W}^{II} = \Sigma$. Next, we present simulations of this network.

Reconstruction of synthetic data First, we test our network on a synthetic dataset with randomly generated N^0 dimensional i.i.d. uniformly distributed input in $[0, r_{max}]^{N^0}$. We choose $\Sigma = \mathbf{I}$, i.e. there are no recurrent connections between I neurons. To be able to reconstruct all non-negative inputs, we choose \mathbf{F} such that $\mathbf{F}^+ = \mathbf{F}(\mathbf{F}^\top \mathbf{F})^{-1}$ is non-negative. We then design our weights correspondingly to satisfy the constraints described in Eq.s (7) and (8). In Fig. 2a, we see that the reconstruction is remarkably accurate. In this example dimensionality is expanded, i.e. $N^E > N^0$, but $N^I = \text{rank}(\mathbf{F}) = N^0$ to satisfy the factorization rank constraints.

We also simulate time varying inputs. Here, the network computes a leaky integration of the input signal dependent on the timescale of the network. We simulate step inputs given by $s(t)$, as shown in Fig. 2b. The target line is obtained by directly computing $\frac{dx}{dt} = -\frac{x}{\tau} + s(t)$, where $\tau = \tau_E = \tau_I$. The prediction line is given by decoding the firing rate of the spiking network, i.e. computing $\mathbf{F}^\top \mathbf{r}^E(t)$.

Reconstruction of natural image patches with varying number of I neurons. Next, we reconstruct natural image patches. Using the sparse coding algorithm of [33], we learn a dictionary for 13×13 natural image patches, concatenated to the rows of the matrix \mathbf{F} . We build our network using an SVD, $\mathbf{F}\mathbf{F}^\top = \mathbf{U}\Sigma\mathbf{U}^\top$, and keep only the largest N^I singular values. For this example we allow $N^I < \text{rank}(\mathbf{F})$, leading to signal loss. We observe that the reconstruction becomes more accurate as we increase the number of I neurons (Fig. 2c&d) because larger ratio of the variance of the dictionary we learned from natural image patches is captured by the I neurons (Fig. 2e).

3.2 Attractor networks

Attractor networks are widely used as models of neural and cognitive phenomena [14, 19–22, 24, 26, 39, 40]. Frequently such attractor networks minimize energy functions and violate Dale’s law. Here, we show how minimax objectives can be used to design optimally balanced E-I attractor networks for popular manifold attractors used in neuroscience.

Fixed point attractors. Fixed point attractors have been used to model associative memory [14]. We can store attractors in our network as discrete fixed points with properly designed weights, allowing different fixed points to have overlapping sets of active neurons. In Fig. 3, for a given set of fixed points, the network weights are obtained by performing a nonlinear constrained optimization problem. We minimize the mean-squared-error of the membrane potentials $V_i^{E,I}$ for the active neurons at the fixed points, subject to the constraints given by the KKT conditions for the inactive neurons (Eq.s (5) and (6)), the second order sufficient (Eq. (7)) condition, and the condition for convergence of rate dynamics (Eq. (8)) (see SI C.1). We show in Fig. 3 that with the optimized weights, our network can store different attractors (Fig. 3a&b) while remaining tightly balanced (Fig. 3c).

Ring attractor. Ring attractors have been used to model the head direction system [19–23] and orientation selectivity in primary visual cortex [24]. Its dynamics can be derived from energy functions. To store a ring attractor in our E-I network, we design the weight matrices such that the

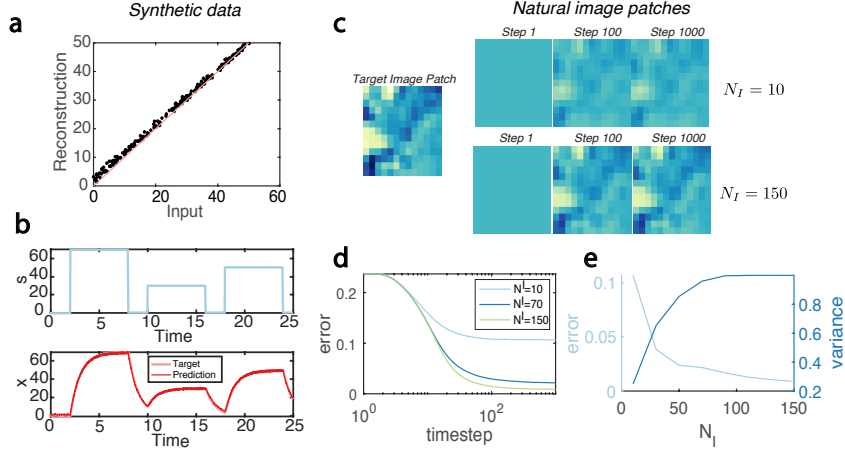


Figure 2: Reconstruction of input signal: (a) Our network with $N^E = 60$, $N^I = 10$ achieves input reconstruction accurately. The input is 20 different 10-dimensional i.i.d. signals sampled from a uniform distribution with $r_{max} = 50$. (b) Reconstruction of filtered step inputs can be achieved with our network. $\tau = 1$. (c) Reconstruction of a target image patch at simulation steps 1, 100 and 1000 for $N^E = 400$ and $N^I = 10$ and $N^I = 150$. (d) Reconstruction error $\frac{1}{N^I} \|\mathbf{F}^T \mathbf{r}^E - \mathbf{x}\|^2$ measured as a function of time, for the $N^I = 10, 70, 150$. (e) Left: The error after convergence measured as a function of N^I . Right: The ratio of variance that is accounted for by the N^I neurons.

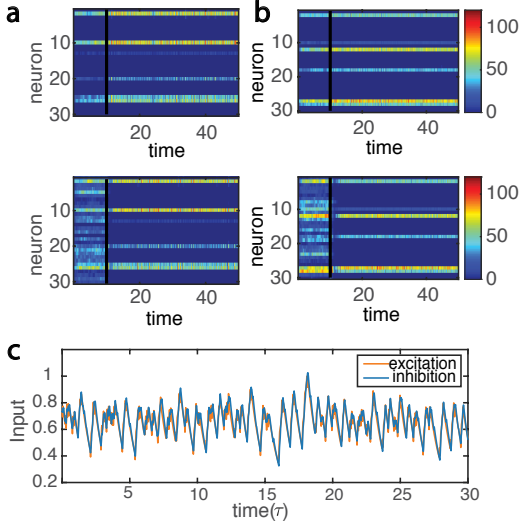


Figure 3: Balanced E-I network trained to store 2 attractor states with $N^E = N^I = 15$: (a),(b) Two different attractor states with overlapping sets of active neurons in the same spiking E-I network, the network is initialized with different inputs that stop after 10s, the network correctly converges to the more adjacent fixed point afterwards. The simulation timescales are $\tau_E = 0.5$, $\tau_I = 0.4$. Top: noiseless case, Bottom: initialized with noisy inputs. Color indicates the firing rate. (c) The E input and I input for an active neuron (normalized by maximum E input) closely track each other. The network remains tightly balanced in the attractor states.

effective energy function for the E neurons matches the energy function in a standard ring attractor network [20]:

$$\mathcal{L} = -\frac{1}{2N^E} \sum_{ij} r_i^E (w_0 + w_1 \cos(\theta_i - \theta_j)) r_j^E - \sum_i (h_0 + h_1 \cos(\theta_0 - \theta_i)) r_i^E + \sum_i \int_0^{r_i^E} f^{-1}(r_i^E), \quad (12)$$

where w_0, w_1, h_0, h_1 are scalar parameters that control the system's behavior (see SI C.2), and $f(x) = [x]_+$ represents ReLU nonlinearity (see SI B.2). Each neuron i in the E population is assigned a distinct preferred angle $\theta_i \in \left\{0, \frac{2\pi}{N^E}, \dots, \frac{2\pi(N^E-1)}{N^E}\right\}$.

We assume that the inhibitory neurons are all active ($r_i^I(t) > 0$) at the optima for simplicity, which can be achieved by carefully designing the weights (eg. $\mathbf{W}^{II^{-1}} \mathbf{W}^{IE} > 0$). We obtain an effective energy function for E neurons by maximizing over \mathbf{r}^I in Eq. (4) and plugging in the optimum value:

$$L_{eff} = -\frac{1}{2} \mathbf{r}^{E\top} (\mathbf{W}^{EE} - \mathbf{W}^{EI} \mathbf{W}^{II^{-1}} \mathbf{W}^{IE}) \mathbf{r}^E - \mathbf{x}^\top \mathbf{F}^\top \mathbf{r}^E. \quad (13)$$

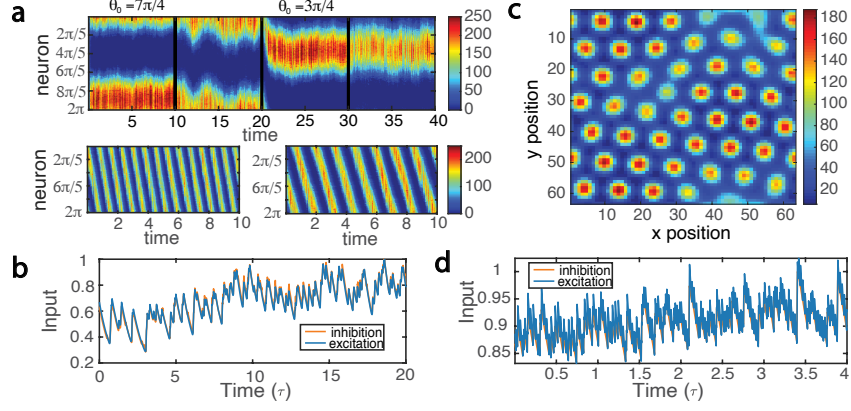


Figure 4: Ring attractor and grid attractor in E-I balanced network: (a) Top: During 0-10s and 20-30s, the network is initialized with inhomogeneous input with two different θ_0 's, this inhomogeneity is turned off during 10-20s and 30-40s. The network is able to pick up inhomogeneous response even when the input to each neuron is identical, the attractor state the network ends up in depends on initialization. Bottom: Spiking network with the same parameters as in corresponding panel in (a), except that h_1 is zero throughout the simulation (Left: $\gamma = 0.15$, Right: $\gamma = 0.08$). Larger γ corresponds to higher angular frequency. (b) The network remains tightly balanced in the attractor states. The E and I inputs for 20τ ($\tau = 0.1$) are both normalized by the maximum of the E input, they closely track each other. (c) Spiking model for grid attractors. Each pixel represents a neuron. Color indicates firing rate of individual neurons as defined in Eq. (3). (d) E and I input for one of the active neurons in (c) for a duration of $4\tau_E$, both normalized by the maximum of the E input. The network is tightly balanced.

To match this effective energy function with Eq. (12), we design weights such that $(\mathbf{W}^{EE} - \mathbf{W}^{EI}\mathbf{W}^{II^{-1}}\mathbf{W}^{IE})_{ij} = \frac{1}{N_E}(w_0 + w_1 \cos(\theta_i - \theta_j)) - \delta_{ij}$ (δ_{ij} term is due to the nonlinearity, similar to SI B.2). For simplicity, we choose \mathbf{W}^{II} to be an identity matrix and \mathbf{W}^{EI} to be a uniform matrix whose elements are all equivalent. The amplitude of the two matrices are tuned for the constraints to be satisfied. The input to the i -th neuron is given by $h_0 + h_1 \cos(\theta_0 - \theta_i)$. θ_0 reflects the inhomogeneity of the input: if $h_1 > 0$, neurons with their preferred angle θ_i closer to θ_0 receive larger inputs. As the standard ring attractor model, this network exhibits different behaviors in different parameter regimes [41] (see SI C.2). In a certain regime the network self-organizes into a bump attractor state even with homogeneous input (Fig. 4a). It also exhibits tight balance (Fig. 4b).

We can introduce anisotropy to our weights to obtain further functions while loosing the objective function interpretation. By following [41, 42], when we design the weights as $(\mathbf{W}^{EE} - \mathbf{W}^{EI}(\mathbf{W}^{II})^{-1}\mathbf{W}^{IE})_{ij} = \frac{1}{N_E}(w_0 + w_1 \cos(\theta_i - \theta_j) - w_1\gamma \sin(\theta_i - \theta_j)) - \delta_{ij}$ (\mathbf{W}^{II} and \mathbf{W}^{EI} are chosen as above), the network produces traveling waves, where γ controls the angular frequency (Fig. 4a). In this example, the network stays on a limit cycle and does not reach an equilibrium, and therefore is not balanced (SI C.2).

Grid attractor. We next discuss a grid attractor, which is used to model the grid cells of the entorhinal cortex [26]. Here, neurons are arranged on a square sheet with periodic boundary conditions (a torus). Each neuron is assigned a positional index \mathbf{x}_i . We design the weight matrices such that the effective energy function for the E neurons resembles that of the grid attractor :

$$L = -\frac{1}{2} \sum_{ij} r_i^E W_0(\mathbf{x}_i - \mathbf{x}_j) r_j^E + \sum_i A_i r_i^E + \sum_i \int_0^{r_i^E} f^{-1}(r_i^E), \quad (14)$$

where $W_0(\mathbf{x}) = ae^{-\gamma|\mathbf{x}|^2} - e^{-\beta|\mathbf{x}|^2}$ [26]. We design the weights such that $(\mathbf{W}^{EE} - \mathbf{W}^{EI}(\mathbf{W}^{II})^{-1}\mathbf{W}^{IE})_{ij} = W_0(\mathbf{x}_i - \mathbf{x}_j) - \delta_{ij}$ (\mathbf{W}^{II} and \mathbf{W}^{EI} are chosen same as the ring attractor, for detailed parameters see SI C.3). Neurons receive homogeneous input, i.e. $A_i = A, \forall i$. We see that although the input is homogeneous, the spiking neural network self-organizes into a grid-like

pattern reminiscent of grid cells on the entorhinal cortex (Fig. 4c). Furthermore, the active neurons in the network remain tightly balanced with the E-I currents cancelling each other precisely (Fig. 4d).

4 Conclusion

Our work provides a novel normative framework for modeling of biologically realistic spiking neural networks. While energy minimization has been used widely in computational neuroscience [12, 14, 15, 19–26, 33–37], we extend normative modeling to minimax optimization in optimally balanced E-I spiking networks. Compared to existing work on minimax dynamics of rate neurons [16], we made several major contributions. We extended the framework to spiking neurons and by inspecting the KKT conditions, we proved that the active neurons are in detailed balance at the saddle point. Simulations confirmed tight balance. We also described conditions on the weights for the rate dynamics (gradient descent-ascent) to converge. Furthermore, by relating our effective objective to the well-studied energy functions for different types of computational goals, we designed novel E-I circuit architectures for various applications from minimax objectives, where the saddle points can be either discrete or forming manifolds, similar to the minima in traditional minimization objective formulations. These applications could be potentially useful for the neuroscience community. Finally, minimax objectives can be a useful strategy when faced with uncertainty and previous works have demonstrated that minimax objectives can be used to solve problems such as source localization [43] and sensorimotor control [44].

Broader Impact

This work introduces a principled approach to designing spiking neural networks of excitatory (E) and inhibitory (I) neurons to perform various computations. As any spiking neural network model, networks derived from our approach could be applied in the field of neuromorphic computing, where information is transmitted through spikes instead of rate, and is thus more energy efficient. Previous work have shown that E-I balanced networks can serve as fast responding modules [9, 10], and our approach could be applied to designing E-I balanced modules that potentially speed up solving optimization problems in general neuromorphic computing systems.

Furthermore, we provided several conditions for the regular functioning of our spiking networks. These conditions could potentially have implications on understanding neural connectivity in the cortex, and how pathological activities in the brain may arise from disrupted synaptic interactions.

Acknowledgments and Disclosure of Funding

We thank Sreya Vemuri for help in the initial stages of this project. This work was supported by funds from the Intel Corporation.

References

- [1] Michael Wehr and Anthony M Zador. Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. *Nature*, 426(6965):442–446, 2003.
- [2] Yousheng Shu, Andrea Hasenstaub, and David A McCormick. Turning on and off recurrent balanced cortical activity. *Nature*, 423(6937):288–293, 2003.
- [3] Bilal Haider, Alvaro Duque, Andrea R Hasenstaub, and David A McCormick. Neocortical network activity in vivo is generated through a dynamic balance of excitation and inhibition. *Journal of Neuroscience*, 26(17):4535–4545, 2006.
- [4] Michael Okun and Ilan Lampl. Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nature neuroscience*, 11(5):535, 2008.
- [5] Bassam V Atallah and Massimo Scanziani. Instantaneous modulation of gamma oscillation frequency by balancing excitation with inhibition. *Neuron*, 62(4):566–577, 2009.
- [6] Sophie Denève and Christian K Machens. Efficient codes and balanced networks. *Nature neuroscience*, 19(3):375, 2016.

- [7] Peter Rupprecht and Rainer W Friedrich. Precise synaptic balance in the zebrafish homolog of olfactory cortex. *Neuron*, 100(3):669–683, 2018.
- [8] Cindy Poo and Jeffry S Isaacson. Odor representations in olfactory cortex: “sparse” coding, global inhibition, and oscillations. *Neuron*, 62(6):850–861, 2009.
- [9] Carl van Vreeswijk and Haim Sompolinsky. Chaotic balanced state in a model of cortical circuits. *Neural computation*, 10(6):1321–1371, 1998.
- [10] Carl Van Vreeswijk and Haim Sompolinsky. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, 274(5293):1724–1726, 1996.
- [11] Michael N Shadlen and William T Newsome. Noise, neural codes and cortical organization. *Current opinion in neurobiology*, 4(4):569–579, 1994.
- [12] Martin Boerlin, Christian K Machens, and Sophie Denève. Predictive coding of dynamical variables in balanced spiking networks. *PLoS computational biology*, 9(11), 2013.
- [13] David G Barrett, Sophie Denève, and Christian K Machens. Firing rate predictions in optimal balanced networks. In *Advances in Neural Information Processing Systems*, pages 1538–1546, 2013.
- [14] 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.
- [15] David GT Barrett, Sophie Deneve, and Christian K Machens. Optimal compensation for neuron loss. *Elife*, 5:e12454, 2016.
- [16] H Sebastian Seung, Tom J Richardson, Jeffrey C Lagarias, and John J Hopfield. Minimax and hamiltonian dynamics of excitatory-inhibitory networks. In *Advances in neural information processing systems*, pages 329–335, 1998.
- [17] H Sebastian Seung. How the brain keeps the eyes still. *Proceedings of the National Academy of Sciences*, 93(23):13339–13344, 1996.
- [18] Mark S Goldman, A Compte, and Xiao-Jing Wang. Neural integrator models. In *Encyclopedia of neuroscience*, pages 165–178. Elsevier Ltd, 2010.
- [19] William E Skaggs, James J Knierim, Hemant S Kudrimoti, and Bruce L McNaughton. A model of the neural basis of the rat’s sense of direction. In *Advances in neural information processing systems*, pages 173–180, 1995.
- [20] Kechen Zhang. Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory. *Journal of Neuroscience*, 16(6):2112–2126, 1996.
- [21] Sung Soo Kim, Hervé Rouault, Shaul Druckmann, and Vivek Jayaraman. Ring attractor dynamics in the drosophila central brain. *Science*, 356(6340):849–853, 2017.
- [22] Xiaohui Xie, Richard HR Hahnloser, and H Sebastian Seung. Double-ring network model of the head-direction system. *Physical Review E*, 66(4):041902, 2002.
- [23] Rishidev Chaudhuri, Berk Gercek, Biraj Pandey, Adrien Peyrache, and Ila Fiete. The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nature neuroscience*, 22(9):1512–1520, 2019.
- [24] Rani Ben-Yishai, R Lev Bar-Or, and Haim Sompolinsky. Theory of orientation tuning in visual cortex. *Proceedings of the National Academy of Sciences*, 92(9):3844–3848, 1995.
- [25] Himanshu Mhatre, Anatoli Gorchetchnikov, and Stephen Grossberg. Grid cell hexagonal patterns formed by fast self-organized learning within entorhinal cortex. *Hippocampus*, 22(2):320–334, 2012.
- [26] Yoram Burak and Ila R Fiete. Accurate path integration in continuous attractor network models of grid cells. *PLoS computational biology*, 5(2), 2009.
- [27] H. W. Kuhn and A. W. Tucker. Nonlinear programming. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 481–492, Berkeley, Calif., 1951. University of California Press.
- [28] Kyle L Luther, Runzhe Yang, and H Sebastian Seung. Unsupervised learning by a “softened” correlation game: duality and convergence. In *2019 53rd Asilomar Conference on Signals, Systems, and Computers*, pages 876–883. IEEE, 2019.
- [29] Chi Jin, Praneeth Netrapalli, and Michael I Jordan. What is local optimality in nonconvex-nonconcave minimax optimization? *arXiv preprint arXiv:1902.00618*, 2019.
- [30] Ping Tak Peter Tang, Tsung-Han Lin, and Mike Davies. Sparse coding by spiking neural networks: Convergence theory and computational results. *arXiv preprint arXiv:1705.05475*, 2017.
- [31] Chi-Ning Chou, Kai-Min Chung, and Chi-Jen Lu. On the algorithmic power of spiking neural networks. *arXiv preprint arXiv:1803.10375*, 2018.

- [32] H Sebastian Seung. Convergence of gradient descent-ascent analyzed as a newtonian dynamical system with dissipation. *arXiv preprint arXiv:1903.02536*, 2019.
- [33] Bruno A Olshausen and David J Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607–609, 1996.
- [34] Anthony J Bell and Terrence J Sejnowski. The “independent components” of natural scenes are edge filters. *Vision research*, 37(23):3327–3338, 1997.
- [35] Surya Ganguli and Haim Sompolinsky. Compressed sensing, sparsity, and dimensionality in neuronal information processing and data analysis. *Annual review of neuroscience*, 35:485–508, 2012.
- [36] Christopher J Rozell, Don H Johnson, Richard G Baraniuk, and Bruno A Olshausen. Sparse coding via thresholding and local competition in neural circuits. *Neural computation*, 20(10):2526–2563, 2008.
- [37] Wieland Brendel, Ralph Bourdoukan, Pietro Vertech, Christian K Machens, and Sophie Denève. Learning to represent signals spike by spike. *PLoS computational biology*, 16(3):e1007692, 2020.
- [38] Cengiz Pehlevan and Dmitri B Chklovskii. Neuroscience-inspired online unsupervised learning algorithms: Artificial neural networks. *IEEE Signal Processing Magazine*, 36(6):88–96, 2019.
- [39] H Sebastian Seung, Daniel D Lee, Ben Y Reis, and David W Tank. Stability of the memory of eye position in a recurrent network of conductance-based model neurons. *Neuron*, 26(1):259–271, 2000.
- [40] A David Redish, Adam N Elga, and David S Touretzky. A coupled attractor model of the rodent head direction system. *Network: Computation in Neural Systems*, 7(4):671–685, 1996.
- [41] David Hansel and Haim Sompolinsky. Modeling feature selectivity in local cortical circuits, 1998.
- [42] Chris Eliasmith. A unified approach to building and controlling spiking attractor networks. *Neural computation*, 17(6):1276–1314, 2005.
- [43] Praveen Venkatesh and Pulkit Grover. Lower bounds on the minimax risk for the source localization problem. In *2017 IEEE International Symposium on Information Theory (ISIT)*, pages 3080–3084. IEEE, 2017.
- [44] Yuki Ueyama. Mini-max feedback control as a computational theory of sensorimotor control in the presence of structural uncertainty. *Frontiers in computational neuroscience*, 8:119, 2014.
- [45] Peter Dayan and Laurence F Abbott. Theoretical neuroscience: computational and mathematical modeling of neural systems. 2001.
- [46] Garth P McCormick. Second order conditions for constrained minima. In *Traces and Emergence of Nonlinear Programming*, pages 259–270. Springer, 2014.
- [47] MATLAB. *version 9.6.0 (R2019a)*. The MathWorks Inc., Natick, Massachusetts, 2019.

Supplementary Information (SI)

A Spiking dynamics as a greedy optimization algorithm on the minimax objective

We first show the derivation of the spiking dynamics for I neurons. Because we are maximizing over \mathbf{r}^I , the neuron fires a spike when it increases the objective. The firing condition for neuron k correspond to $S(\mathbf{r}^E, \mathbf{r}^I + \mathbf{e}_k) > S(\mathbf{r}^E, \mathbf{r}^I)$, where \mathbf{e}_k denotes the standard basis vector. By plugging in Eq. (4), we have

$$\begin{aligned} & \sum_{ij} r_i^E W_{ij}^{EI} r_j^I - \frac{1}{2} \sum_{ij} r_i^I W_{ij}^{II} r_j^I < \sum_{ij} (r_i^E W_{ij}^{EI} (r_j^I + \delta_{jk}) - \frac{1}{2} (r_i^I + \delta_{ik}) W_{ij}^{II} (r_j^I + \delta_{jk})) \\ \implies & \sum_i W_{ik}^{EI} r_i^E - \sum_j r_j^I W_{jk}^{II} > \frac{1}{2} W_{kk}^{II}, \end{aligned} \quad (\text{SI.1})$$

where δ_{jk} denotes the j -th element of the standard basis vector \mathbf{e}_k . We define the membrane potential and the firing threshold of the I neurons.

$$\begin{aligned} V_k^I & \equiv \sum_i W_{ik}^{EI} r_i^E - \sum_i W_{ik}^{II} r_i^I \\ T_k^I & \equiv \frac{1}{2} W_{kk}^{II}, \end{aligned} \quad (\text{SI.2})$$

Next, we derive the dynamics of the membrane potential. For $\tau_E = \tau_I$ and using Eq. (3) for the definition of $r_i^{I/E}$, we arrive at:

$$\frac{dV_k^I}{dt} = -\frac{1}{\tau} V_k^I - \sum_i W_{ik}^{II} s_i^I + \sum_i W_{ik}^{EI} s_i^E. \quad (\text{SI.3})$$

We recognize this as the standard integrate-and-fire spiking dynamics with threshold $T_k^I = \frac{1}{2} W_{kk}^{II}$ [45].

For the E neurons, we proceed similarly. The firing condition for a neuron k is $S(\mathbf{r}^E + \mathbf{e}_i, \mathbf{r}^I) < S(\mathbf{r}^E, \mathbf{r}^I)$, where \mathbf{e}_i is the i -th standard basis vector. Eq. (4) implies

$$\begin{aligned} & \sum_{ij} -\frac{1}{2} r_i^E W_{ij}^{EE} r_j^E + r_i^E W_{ij}^{EI} r_j^I - \sum_{ij} x_i F_{ij} r_j^E \\ & > \sum_{ij} (-\frac{1}{2} (r_i^E + \delta_{ik}) W_{ij}^{EE} (r_j^E + \delta_{jk}) + (r_i^E + \delta_{ik}) W_{ij}^{EI} r_j^I - \sum_{ij} x_i F_{ij} (r_j^E + \delta_{jk})) \\ \implies & \sum_i (W_{ki}^{EE} r_i^E - W_{ki}^{EI} r_i^I + F_{ik} x_i) > -\frac{1}{2} W_{kk}^{EE}. \end{aligned} \quad (\text{SI.4})$$

Defining the membrane potential and the firing threshold

$$\begin{aligned} V_k^E & \equiv \sum_i W_{ki}^{EE} r_i^E - \sum_i W_{ki}^{EI} r_i^I + \sum_i F_{ik} x_i \\ T_k^E & \equiv -\frac{1}{2} W_{kk}^{EE}, \end{aligned} \quad (\text{SI.5})$$

we can obtain dynamics of the membrane potential for $\tau_E = \tau_I = \tau$:

$$\frac{dV_k^E}{dt} = -\frac{1}{\tau} V_k^E + \sum_i W_{ki}^{EE} s_i^E - \sum_i W_{ki}^{EI} s_i^I + \sum_i F_{ik} s_i^0, \quad (\text{SI.6})$$

with spiking threshold $T_k^E = -\frac{1}{2} W_{kk}^{EE}$.

B Convergence of the dynamics

B.1 Second order sufficient condition for optimality

We cite a theorem from [46].

Theorem 1 *The solution x^*, λ^* obeying the KKT conditions is a constrained local minimum if for the Lagrangian*

$$L(x, \lambda) = f(x) + \sum_{i=1}^m \lambda_i g_i(x), \quad (\text{SI.7})$$

we have

$$\mathbf{s}^T \nabla_{xx}^2 L(x^*, \lambda^*) \mathbf{s} \geq 0, \quad (\text{SI.8})$$

where $\mathbf{s} \neq 0$ is a vector satisfying

$$\nabla_x g_i(x^*)^T \mathbf{s} = 0, \quad (\text{SI.9})$$

where only those active inequality constraints $g_i(x)$ corresponding to strict complementarity (i.e. where $\lambda_i > 0$) are applied.

We apply Thm. 1 to our minimax objective, for the maximization problem with \mathbf{s} as the activities of inhibitory neurons that are active, and $\nabla_{xx} L(x^*, \lambda^*)$ gives $\hat{\mathbf{W}}^{II}$. $\hat{\mathbf{W}}^{II}$ is the submatrix where strict complementarity is applied (namely when $\mathbf{r}^I > 0$), and therefore the second order sufficient condition for optimality is $\hat{\mathbf{W}}^{II} \succcurlyeq 0$. Plugging in the optimal solution for \mathbf{r}^I in the maximization problem and similarly deriving the condition for optimality for the minimization over \mathbf{r}^E according to Thm. 1, we obtain $\hat{\mathbf{W}}^{EI} \hat{\mathbf{W}}^{II^{-1}} \hat{\mathbf{W}}^{IE} - \hat{\mathbf{W}}^{EE} \succcurlyeq 0$.

B.2 Rate dynamics and convergence

We can prove the convergence of a rate dynamics derived from the same minimax objective, for trajectories that do not include switching between active and inactive neurons, i.e., active neurons at initialization remain active, and silent neurons remain silent throughout the trajectory. In practice we observe convergence even when a transition between active and inactive states occurs during the dynamics for some cases.

We use a slightly modified objective function of the form

$$S = -\frac{1}{2} \mathbf{r}^{E^T} \mathbf{W}^{EE^*} \mathbf{r}^E + \mathbf{r}^{E^T} \mathbf{W}^{EI} \mathbf{r}^I - \frac{1}{2} \mathbf{r}^{I^T} \mathbf{W}^{II^*} \mathbf{r}^I - \mathbf{x}^T \mathbf{F} \mathbf{r}^E + \sum_i H(r_i^E) - \sum_i G(r_i^I). \quad (\text{SI.10})$$

The last two terms are related to nonlinear neural activations. For ReLU neurons with thresholds θ_E and θ_I , $H(r_i^E) = \int_0^{r_i^E} f^{-1}(x) dx = \frac{1}{2}(r_i^E + \theta_i^E)^2$, $G(r_i^I) = \int_0^{r_i^I} f^{-1}(x) dx = \frac{1}{2}(r_i^I + \theta_i^I)^2$. Function $H(x)$ and $G(x)$ are only defined for $x \geq 0$. If we choose $\mathbf{W}^{EE^*} = \mathbf{W}^{EE} + \mathbf{I}$, $\mathbf{W}^{II^*} = \mathbf{W}^{II} - \mathbf{I}$, $\theta_E = \theta_I = 0$, then this objective reduces to Eq. (4). For the modifications induced by nonzero thresholds, one can simply change the thresholds of E neurons from $T_i^E = -\frac{1}{2} W_{ii}^{EE}$ to $T_i^E = -\frac{1}{2} W_{ii}^{EE} + \theta_i^E$, and that of the I neurons from $T_i^I = \frac{1}{2} W_{ii}^{II}$ to $T_i^I = \frac{1}{2} W_{ii}^{II} + \theta_i^I$.

The rate dynamics that optimizes this objective is

$$\begin{aligned} \tau_I \frac{du_i^I}{dt} &= -u_i^I + \sum_j W_{ij}^{IE} r_j^E - \sum_j (W_{ij}^{II} - \delta_{ij}) r_j^I \\ \tau_E \frac{du_i^E}{dt} &= -u_i^E + \sum_j (W_{ij}^{EE} + \delta_{ij}) r_j^E - \sum_j W_{ij}^{EI} r_j^I + \sum_j F_{ij}^T x_j \\ r_j^{E(I)} &= f(u_j^{E(I)}) = [u_j^{E(I)} - \theta_j^{E(I)}]_+. \end{aligned} \quad (\text{SI.11})$$

To investigate convergence, we can construct an energy function of the system [16] assuming $\tau_E = \tau_I = \tau$:

$$L = \frac{1}{2} |\dot{\mathbf{r}}^E|^2 + \frac{1}{2} |\dot{\mathbf{r}}^I|^2 + \gamma S, \quad \gamma \in \mathbb{R}. \quad (\text{SI.12})$$

Next we show that the energy function is decreasing. Except for a set of measure zero ($\forall i, u_i^{E/I} = \theta_i^{E/I}$), the ReLU function is differentiable and we have

$$\dot{r}_i^{E/I} = \frac{dr_i^{E/I}}{du_i^{E/I}} \dot{u}_i^{E/I}, \quad \ddot{r}_i^{E/I} = \frac{d^2 r_i^{E/I}}{du_i^{E/I^2}} \dot{u}_i^{E/I^2} + \frac{dr_i^{E/I}}{du_i^{E/I}} \ddot{u}_i^{E/I} = \frac{dr_i^{E/I}}{du_i^{E/I}} \ddot{u}_i^{E/I}. \quad (\text{SI.13})$$

We compute the time derivative of the energy function:

$$\begin{aligned} \dot{L} &= (\dot{\mathbf{r}}^E, \dot{\mathbf{r}}^I)^T (\dot{\mathbf{r}}^E, \dot{\mathbf{r}}^I) + \gamma \dot{S} \\ &= (\dot{\mathbf{u}}^E, \dot{\mathbf{u}}^I)^T \begin{bmatrix} \frac{d\mathbf{r}^E}{d\mathbf{u}^E} & 0 \\ 0 & \frac{d\mathbf{r}^I}{d\mathbf{u}^I} \end{bmatrix} \begin{bmatrix} -\mathbf{I} + \mathbf{W}^{EE^*} \frac{d\mathbf{r}^E}{d\mathbf{u}^E} & -\mathbf{W}^{EI} \frac{d\mathbf{r}^I}{d\mathbf{u}^I} \\ \mathbf{W}^{IE} \frac{d\mathbf{r}^E}{d\mathbf{u}^E} & -\mathbf{I} - \mathbf{W}^{II^*} \frac{d\mathbf{r}^I}{d\mathbf{u}^I} \end{bmatrix} (\dot{\mathbf{u}}^E, \dot{\mathbf{u}}^I) \\ &\quad + \gamma (-\dot{\mathbf{u}}^E, \dot{\mathbf{u}}^I)^T \begin{bmatrix} \frac{d\mathbf{r}^E}{d\mathbf{u}^E} & 0 \\ 0 & \frac{d\mathbf{r}^I}{d\mathbf{u}^I} \end{bmatrix} (\dot{\mathbf{u}}^E, \dot{\mathbf{u}}^I) \\ &= (\dot{\mathbf{u}}^E, \dot{\mathbf{u}}^I)^T \begin{bmatrix} \frac{d\mathbf{r}^E}{d\mathbf{u}^E} \mathbf{W}^{EE^*} \frac{d\mathbf{r}^E}{d\mathbf{u}^E} - \frac{d\mathbf{r}^E}{d\mathbf{u}^E} - \gamma \frac{d\mathbf{r}^E}{d\mathbf{u}^E} & 0 \\ 0 & \frac{d\mathbf{r}^I}{d\mathbf{u}^I} \mathbf{W}^{II^*} \frac{d\mathbf{r}^I}{d\mathbf{u}^I} - \frac{d\mathbf{r}^I}{d\mathbf{u}^I} + \gamma \frac{d\mathbf{r}^I}{d\mathbf{u}^I} \end{bmatrix} (\dot{\mathbf{u}}^E, \dot{\mathbf{u}}^I) \\ \dot{L} &= \dot{\mathbf{u}}^{E^T} (\hat{\mathbf{W}}^{EE} - \gamma \mathbf{I}) \dot{\mathbf{u}}^E - \dot{\mathbf{u}}^{I^T} (\hat{\mathbf{W}}^{II} - \gamma \mathbf{I}) \dot{\mathbf{u}}^I. \end{aligned} \quad (\text{SI.14})$$

If $\hat{\mathbf{W}}^{II} - \gamma \mathbf{I}$ is positive definite and $\hat{\mathbf{W}}^{EE} - \gamma \mathbf{I}$ is negative definite, then any bounded trajectory that does not cross $u_i = 0$ boundaries (and cause changes in the set of active/inactive neurons) are convergent. If $\min \lambda(\hat{\mathbf{W}}^{II}) > \max \lambda(\hat{\mathbf{W}}^{EE})$, then there exist γ such that $\min \lambda(\hat{\mathbf{W}}^{II}) > \gamma > \max \lambda(\hat{\mathbf{W}}^{EE})$, and the condition is satisfied.

Now we proved that any bounded trajectories that do not cross the $u_i = 0$ boundaries are convergent, we need to further prove that the trajectories are bounded. We applying Thm. 2 in [32].

Theorem 2 *Given a twice differentiable objective S , suppose that $\lambda_{inf}(S_{\mathbf{x}\mathbf{x}}) > \lambda_{sup}(S_{\mathbf{y}\mathbf{y}})$. If either*

1. $\lambda_{inf}(S_{\mathbf{x}\mathbf{x}} > 0)$ and $-V(\mathbf{y}) = -\min_{\mathbf{x}} S(\mathbf{x}, \mathbf{y})$ is radially unbounded, or
2. $\lambda_{inf}(S_{\mathbf{y}\mathbf{y}} < 0)$ and $U(\mathbf{y}) = \max_{\mathbf{x}} S(\mathbf{x}, \mathbf{y})$ is radially unbounded

is also satisfied, then any trajectory of gradient descent-ascent is bounded.

We see that the condition for boundedness is the same as the condition for second order optimality as discussed in SI. B.1.

C Attractor networks

C.1 Optimization problem for designing weights in fixed point attractor networks

We design the network weights in order to store fixed point attractors by solving an optimization problem. For a given attractor state m , we denote the set of active neurons as \mathcal{A}_m and the set of silent neurons as \mathcal{I}_m . We formulate the following optimization problem.

$$\begin{aligned}
& \min \sum_m \sum_{i \in \mathcal{A}_m} \|V_i^m\|^2, \quad \text{subject to} \\
& \exists \kappa > 0, V_i^m \leq -\kappa, \forall i \in \mathcal{I}_m, \forall m && \text{(constraint 1)} \\
& \hat{\mathbf{W}}_m^{EI} \hat{\mathbf{W}}_m^{II^{-1}} \hat{\mathbf{W}}_m^{IE} - \hat{\mathbf{W}}_m^{EE} \succcurlyeq 0, \forall m && \text{(constraint 2)} \\
& \exists \gamma_m \in \mathbb{R}, \min \lambda(\hat{\mathbf{W}}_m^{II}) > \gamma_m > \max \lambda(\hat{\mathbf{W}}_m^{EE}), \forall m && \text{(constraint 3)} \\
& \mathbf{W}^{EI/II/IE} \geq 0 && \text{(constraint 4)} \\
& W_{ij}^{EE} \geq 0, \forall i \neq j; W_{ii}^{EE} \leq 0, \forall i && \text{(constraint 5)} \\
& \mathbf{W}^{EE} = \mathbf{W}^{EE^T}; \mathbf{W}^{II} = \mathbf{W}^{II^T}; \mathbf{W}^{EI} = \mathbf{W}^{IE^T} && \text{(constraint 6).} \quad \text{(SI.15)}
\end{aligned}$$

Here $\hat{\mathbf{W}}_m^{EI/EE/IE/II}$ denote the submatrices with rows and columns corresponding to the active neurons in attractor state m . The expression for V_i^m is given by plugging in \mathbf{r}^E and \mathbf{r}^I in the attractor state m into Eq. (2) for E and I neurons respectively. Constraint 1 corresponds to the KKT condition on the inactive neurons. Constraints 2 and 3 guarantee convergence of rate dynamics when the trajectory does not cross the $u_i = 0$ boundaries. Constraints 4,5,6 are imposed for symmetry and nonnegativity of the matrices.

This is a nontrivial optimization problem with nonlinear constraints. We used the Sequential Quadratic Programming (SQP) algorithm for MATLAB function *fmincon* [47] to solve this optimization problem. We start with different initial values and only accept solutions that satisfy $\min \sum_m \sum_{i \in \mathcal{A}_m} \|V_i^m\|^2 = 0$.

C.2 Different parameter regimes in ring attractor network

The energy function of a standard ring attractor model is given by

$$L = -\frac{1}{2N_E} \sum_{ij} r_i^E (w_0 + w_1 \cos(\theta_i - \theta_j)) r_j^E - \sum_i (h_0 + h_1 \cos(\theta_0 - \theta_i)) r_i^E + \sum_i \int_0^{r_i^E} f^{-1}(r_i^E) \quad \text{(SI.16)}$$

Here $h_0 + h_1 \cos(\theta_0 - \theta_i)$ is the input term, and $f(x) = [x]_+$ is ReLU nonlinearity. This network has different parameter regimes for different behaviors [41]. When $w_0 \geq 1$, the network is unstable and the firing rate goes to infinity. When $w_1 < 2$, we have signal amplification, the network has larger response amplitude than its input. Finally, when $w_1 \geq 2$, there is symmetry breaking, the network can pick up inhomogeneous response even when the input is homogeneous ($h_1 = 0$). The stable states of the network lie on a ring, and which steady state the network reaches depends on the initial conditions. The same parameter regimes and behaviors apply for our E-I network.

The parameters used in the simulations for Fig. 4a&b are $w_0 = 0.5$, $w_1 = 2.7$, $h_0 = 10$. For the top panel, $\gamma = 0$. For 0-10s, $h_1 = 5$ and $\theta_0 = \frac{7}{4}\pi$. For 20-30s $h_1 = 5$ and $\theta_0 = \frac{3}{4}\pi$. $h_1 = 0$ otherwise, and the input is thus homogeneous.




Figure SI.1: The E and I inputs for 30τ , they are not balanced for network with travelling wave with the above parameters except $h_1 = 0$ throughout the simulation and $\gamma = 0.08$.

For the travelling wave, we show that the E and I inputs are not balanced, because the system does not reach equilibrium, as shown in Fig. SI.1.

C.3 Grid attractor

The energy function is given by

$$L = -\frac{1}{2} \sum_{ij} r_i^E W_0(\mathbf{x}_i - \mathbf{x}_j) r_j^E - \sum_i A_i r_i^E + \sum_i \int_0^{r_i^E} f^{-1}(r_i^E). \quad (\text{SI.17})$$

Here A_i is the input to each neuron, in our simulation, we assume homogeneous input $A_i = A, \forall i$. $f(x) = [x]_+$ represents ReLU nonlinearity. Function $W_0(\mathbf{x})$ is given by $W_0(\mathbf{x}) = ae^{-\gamma|\mathbf{x}|^2} - e^{-\beta|\mathbf{x}|^2}$. For simulation of the spiking network with grid attractor, the parameters we used are $N_E = N_I = 63^2$, neurons are arranged in a 63×63 square sheet, $\tau_E = 0.5, \tau_I = 0.2, \tau = 0.5, \lambda = 6, \beta = 3\lambda^2, \alpha = 1.1, \gamma = 1.2\beta, A = 2$.

For code used to reproduce results in this paper, see <https://github.com/Pehlevan-Group/BalancedEIMinimax>.