Co-existence of synaptic plasticity and metastable dynamics in a spiking model of cortical circuits

Xiaoyu Yang[a,b,c] and Giancarlo La Camera[b,c,1]

[a] Graduate Program in Physics and Astronomy, Stony Brook University
[b] Department of Neurobiology & Behavior, Stony Brook University
[c] Center for Neural Circuit Dynamics, Stony Brook University

Abstract

Evidence for metastable dynamics and its role in brain function is emerging at a fast pace and is changing our understanding of neural coding by putting an emphasis on hidden states of transient activity. Clustered networks of spiking neurons are capable of producing metastable dynamics, however, it is unclear how their structure may emerge in cortical circuits. Here, we demonstrate the emergence of rich metastable dynamics from a fully local synaptic plasticity rule. The metastable dynamics co-exists with ongoing plasticity; in turn, the synaptic structure is stable to ongoing dynamics and random perturbations, yet it remains sufficiently plastic to remap sensory representations to encode new sets of stimuli. Both the plasticity rule and the metastable dynamics scale well with network size, with synaptic stability increasing with the number of neurons. These results show that it is possible to generate and support metastable dynamics over meaningful hidden states using a simple but biologically plausible plasticity rule which co-exists with ongoing neural dynamics and is independent of network size.

[1]Address correspondence to Giancarlo La Camera, Department of Neurobiology and Behavior, Stony Brook University, Stony Brook, NY. Email: giancarlo.lacamera@stonybrook.edu.

Contents

1 Introduction **2**

2 Results **3**
2.1 The synaptic plasticity rule . . . 3
2.2 Formation of stable clusters with metastable dynamics . . . 4
2.3 Learning stability vs. network size . . . 6
2.4 Robustness to perturbations and remapping . . . 8

3 Discussion **9**
3.1 Model features and main results . . . 9
3.2 Comparison with previous work . . . 12
3.3 Limitations and possible extensions . . . 13
3.4 Conclusion . . . 14

4 Methods **15**

A Supplementary Figures **20**

1 Introduction

Cortical circuits express ongoing neural dynamics that is often found to be metastable. By this we mean that ensembles of neurons keep approximately constant firing rates for transient periods of time, until a subset of the neurons switch to different firing rates. A stable vector of firing rates across simultaneously recorded neurons may be thought of as a 'hidden state', of which spike trains emitted by the neurons are noisy observations. In the cortex of rodents and non-human primates, metastable states last between a few hundred of ms to a few seconds and their identity and dynamics have been linked to sensory processes, attention, expectation, navigation, decision making and behavioral accuracy (see (La Camera et al., 2019; Brinkman et al., 2022) for reviews). One way to model this type of metastable activity is to organize a neural network in clusters, or cell assemblies (Litwin-Kumar and Doiron, 2012; Mazzucato et al., 2015). The neurons in each clusters are connected by synaptic weights whose average value is larger than the average weight between neurons of different clusters. A key question is then to understand how such structure subserving metastable dynamics can emerge from, for example, experience-dependent plasticity. This problem is very similar to the problem of the stable formation of Hebb assemblies leading to persistent activity reflecting the memory of learned stimuli (Del Giudice et al., 2003; Amit and Mongillo, 2003; Litwin-Kumar and Doiron, 2014; Zenke et al., 2015). In both cases, one desires to obtain stable clusters of strongly connected neurons; the difference is that here we are interested in neural activity that, as the consequence of learning, is metastable.

The problem of the formation of stable cell assemblies has been of interest for a long time (see (Dayan and Abbott, 2005; Gerstner et al., 2014) for textbook reviews). The most recent efforts in this direction have included the combination of spike-timing-dependent plasticity (STDP) (Markram et al., 1997; Bi and Poo, 1998) and a number of homeostatic mechanisms (Watt and Desai, 2010; Turrigiano, 2017) to keep the neural activity bounded during learning. However, while most efforts so far have focused on the formation of stable neural clusters with the purpose of representing retrievable memories or the development of receptive fields, here we focus on metastable dynamics. In other words, instead of focusing

on stable neural dynamics following the presentation and removal of a stimulus, the aim of this study is to obtain neural dynamics that continuously switches among a set of hidden states, which have been stored in the network structure by training.

In pursuing this effort, we require the synaptic plasticity rule to be biologically plausible and leading to the formation of neural clusters that are stable for random perturbations. Moreover, we require that the neural clusters generate metastable dynamics that coexists with ongoing plasticity (i.e., in the absence of external stimuli). We further aim to obtain a model where new information can be accommodated, so that training with a new set of stimuli will lead to cluster rearrangement producing metastable dynamics among the new states. The requirement of biological plausibility mostly means that the plasticity rule must be local and must depend only on presynaptic and postsynaptic activity in a way that is accessible to the synapse – in particular, it must depend only on presynaptic spikes and postsynaptic variables related e.g. to membrane voltage or calcium transients (Fusi et al., 2000; Shouval et al., 2002; Clopath et al., 2010; Graupner and Brunel, 2012). We will not consider directly the role of calcium in our model. Rather, our learning rule belongs to the class of voltage-based STDP models (Fusi et al., 2000; Clopath et al., 2010). We will also allow neurons to respond to multiple stimuli, as observed in experiment (Miyashita and Chang, 1988; Miyashita, 1988; Curti et al., 2004; Jezzini et al., 2013; Fusi et al., 2016).

Here, we present a plasticity rule that fulfills all of the above requirements in a network of exponential integrate-and-fire neurons (Fourcaud-Trocmé et al., 2003). The plasticity rule is simple and depends only on presynaptic spikes arrival times, postsynaptic voltage, and postsynaptic spikes. Long-term potentiation (LTP) and long-term depression (LTD) are obtained by comparing a voltage-sensitive internal variable to an adapting threshold, reminiscent of the BCM rule (Bienenstock et al., 1982). The threshold adapts in a way to produce transient LTP among co-active neurons, followed by LTD after prolonged activation. This leads to the stable formation of neural clusters while promoting a dynamics that is automatically metastable for a large range of network sizes. Moreover, additional homeostatic mechanisms, such as inhibitory plasticity or synaptic scaling (Tetzlaff et al., 2011; Zenke and Gerstner, 2017), are not required. Our plasticity rule provides a possible explanation of the emergence of metastable dynamics observed in many brain areas (Brinkman et al., 2022) during both ongoing and evoked neural activity.

2 Results

2.1 The synaptic plasticity rule

We endowed a recurrent network of excitatory and inhibitory exponential integrate-and-fire (EIF) neurons (see Methods) with the following plasticity rule for synapses connecting excitatory neurons. Given a synapse with efficacy w_{ij} from excitatory neuron j to excitatory neuron i, a change in synaptic efficacy was triggered by the arrival of a presynaptic spike, while its polarity and strength depended on the activity of the postsynaptic neuron:

$$\dot{w}_{ij} = \left(A_{\mathrm{LTP}} e^{-\beta w_{ij}^2} [\tilde{v}_i - \theta_{\mathrm{LTP}}]_+ - A_{\mathrm{LTD}} [\theta_{\mathrm{LTD}} - \tilde{v}_i]_+ \right) s_j(t). \tag{1}$$

Here, $\dot{w} \doteq \frac{dw}{dt}$, $s_j(t) = \sum_k \delta(t - t_j^{(k)})$ is the presynaptic spike train and $\delta(t)$ is the Dirac delta function. $[x]_+ \doteq \max(x, 0)$ is the rectified linear function and $\tilde{v}_i$ is a local postsynaptic variable which dictates the polarity of plasticity: the synapse undergoes LTP when $\tilde{v}_i > \theta_{\mathrm{LTP}}$, and it undergoes LTD when $\tilde{v}_i < \theta_{\mathrm{LTD}}$, with $\theta_{\mathrm{LTD}} \leq \theta_{\mathrm{LTP}}$. The variable $\tilde{v}_i$ could represent a calcium variable or the running average

of the membrane potential. In our case, $\tilde{v}_i$ was the low-pass filter of $\Delta_T e^{\frac{V-V_T}{\Delta_T}}$, where V is the membrane potential of the EIF neuron (see Methods), and therefore it was driven most substantially during the emission of a spike. The strength of LTP was further modulated by an attenuation factor $e^{-\beta w_{ij}^2}$ that constrains the ability of strong synapses to grow excessively.

As similar learning rules of this kind (Dayan and Abbott, 2005), this rule is unstable if θ_{LTP} and θ_{LTD} are constant thresholds and the attenuation factor is missing. In such a case, stimulating a subset of neurons would result in higher firing rates and larger $\tilde{v}_i$, which in turn would lead to higher firing rates, and so on. One way to prevent this problem is to use activity-dependent thresholds as in the BCM rule (Bienenstock et al., 1982). This idea requires the LTP and LTD thresholds to be a non-linear function of the postsynaptic activity $\tilde{v}$, with faster dynamics than the dynamics of the synaptic weights. While the BCM rule uses a supralinear function of $\tilde{v}$, we chose a hyperbolic tangent function of both $\tilde{v}$ and postsynaptic spiking activity, $\tilde{s}_i$:

$$\tau_\theta \dot{\theta}_i = \theta_a \tanh[g(-\theta_i + \tilde{v}_i + \gamma \tilde{s}_i)]. \tag{2}$$

In this equation, $\tilde{s}_i$ is the low-pass filtered postsynaptic spike train, $s_i(t) = \sum_k \delta(t - t_i^{(k)})$, θ_a is a constant in units of θ_i, g is a gain factor and γ is a constant. For simplicity, we used the same threshold $\theta_i(t)$ for both θ_{LTP} and θ_{LTD}, but in principle they could be different.

The hyperbolic tangent in Eq. 2 automatically adjusts the dynamics of the threshold $\theta_i(t)$ for different postsynaptic activities. The motivation behind this specific model is that it can lead naturally to switching dynamics of neural clusters. This can be understood from the dynamics of Eq. 2, which adapts to the size of its argument $\Delta \equiv \tilde{v}_i + \gamma \tilde{s}_i$ (Fig. 1a; see Methods for details). For small arguments, the dynamics is fastest and θ_i closely follows $\tilde{v}_i$, resulting in no average change in the synaptic weights (Fig. 1b). This occurs when the postsynaptic membrane potential is characterized by subthreshold fluctuations. When the postsynaptic neuron fires a spike, its membrane potential rises significantly and θ_i is attracted to a new value with slower dynamics. As a result, θ_i will temporarily lag behind $\tilde{v}_i + \gamma \tilde{s}_i$, producing a short temporal window for LTP, followed by a longer window for LTD (Fig. 1c). An occasional spike during ongoing activity will not produce a meaningful synaptic change, however, repeated activation of the same postsynaptic neuron will produce a longer window for LTP (Fig. 1d, blue shaded area). This will occur when the postsynaptic neuron engages in recurrent excitation, and will promote the formation of clusters of co-active neurons. Prolonged activation of the same cluster, however, will turn LTP into LTD. This is due to the term $\gamma \tilde{s}_i$, so that the threshold θ_i will finally approach a value $\theta_i \approx \tilde{v}_i + \gamma \tilde{s}_i > \tilde{v}_i$, causing LTD (Fig. 1d, red shaded area). In summary, θ_i dynamics can help to form neural clusters via transient LTP, while impeding prolonged activation of the same clusters, naturally producing a switching dynamics among co-active groups of neurons.

We note that although it might seem superfluous to add the term $\gamma \tilde{s}_i$ to the equation for θ_i, it is in fact required to ensure LTD when the neuron is constantly active (i.e. $\tilde{v}_i \gg 0$). This can be either achieved by adding the offset term $\gamma \tilde{s}_i$ to $\tilde{v}_i$, or by using a supralinear function of $\tilde{v}_i$, such as $\tilde{v}_i^p$ with $p > 1$, as in the BCM rule. However, the additive solution $\tilde{v}_i + \gamma \tilde{s}_i$ has the advantage of preventing synaptic change when the postsynaptic neuron is inactive, as explained above (see also the next section). We will further comment on the comparison with the classical BCM rule in the Discussion.

2.2 Formation of stable clusters with metastable dynamics

We initially tested our plasticity rule in a recurrent spiking network model comprising $N_E = 800$ excitatory and $N_I = 200$ inhibitory EIF neurons (see Methods). During training, a set of $Q = 10$ sensory stimuli

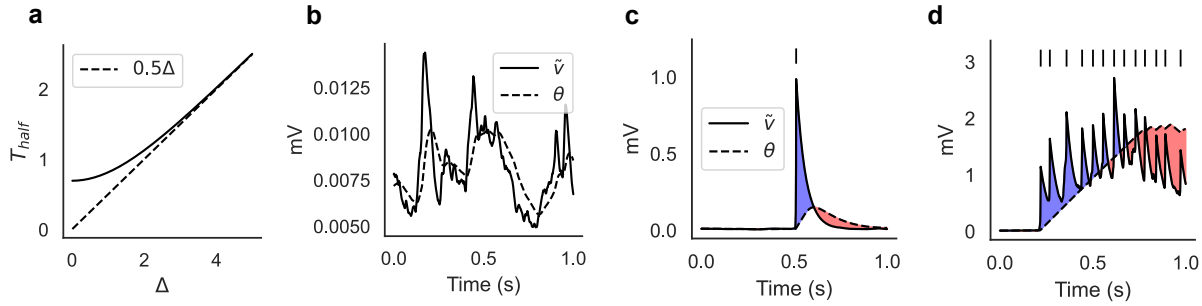

Figure 1: Illustration of the plasticity rule. (a) Time scale of $\tau_\theta \dot{\theta} = \tanh(-\theta + \Delta)$ as a function of Δ, where $\Delta \equiv \tilde{v}_i + \gamma \tilde{s}_i$ (see Eq. 2). T_{half} is defined as the time it takes for $\theta(0) = 0$ to reach half of Δ (in units of τ_θ), and for $\Delta \gtrsim 3$ it increases linearly with Δ (see Methods). (b) The learning rule ignores the subthreshold fluctuations of the membrane potential when the neuron does not fire spikes. (c) After the postsynaptic neuron fires a single spike, the synapse undergoes a short window for LTP (blue shaded area) followed by a longer window for LTD (red shaded area). (d) Repeated activation of the postsynaptic neuron will cause a transition from LTP to LTD.

were presented to the network in random order, each targeting a fixed but randomly chosen subpopulation of excitatory neurons, which we call a cluster. Although each cluster was associated to a separate stimulus, each cluster contained $f = 10\%$ of randomly selected neurons, so that the same neuron could respond to more than one stimulus as typically observed in experiments (Miyashita and Chang, 1988; Miyashita, 1988; Curti et al., 2004; Jezzini et al., 2013). Thus, the mean number of neurons in each cluster was 80, but each neuron had a probability of about 0.26 to respond to at least two stimuli, and clusters had a probability 0.61 of sharing neurons with other clusters (see Methods for details).

Stimulus presentations occurred every 2000 ms and each lasted 500 ms (however, using random inter-stimulus intervals did not alter the results). This training procedure lasted for 10 minutes and on average each sensory stimulus was presented 30 times.

As expected, after training the network exhibited metastable ongoing dynamics which was still present after 2 and 4 hours (Fig. 2a). In this figure, spikes from neurons in the same cluster have the same color, and neurons belonging to multiple clusters are duplicated. The black curves superimposed to each cluster's spikes measure the 'overlap' of the whole network's activity with the stimulus associated to that cluster (see Methods). For a given stimulus, the overlap varies between zero and one and approaches 1 when all active excitatory neurons of the network belong to the associated cluster, while the remaining neurons have negligible firing rate. The raster plots in Fig. 2 show that the neural activity after training switches between networks states characterized by specific cluster activations. These states can therefore be interpreted as memories of the stimuli, with the ongoing dynamics being akin to a random walk among these memory states. The distribution of state durations was approximately exponential with mean around $250 - 300$ ms (Fig. 2c), reminiscent of the discrete Markov processes with fast state transitions found to describe ensembles of cortical spike trains (Seidemann et al., 1996; Jones et al., 2007; Mazzucato et al., 2015; Maboudi et al., 2018; Benozzo et al., 2021; Lang et al., 2023). The metastable dynamics observed after training is the consequence of potentiated synapses inside clusters and the emergence of block structure in the synaptic matrix (Fig. 2b), a structure known to potentially produce metastable dynamics in spiking networks (Litwin-Kumar and Doiron, 2012; Deco and Hugues, 2012; Mazzucato et al., 2015). In our case, metastable dynamics is co-existent with synaptic plasticity, but it remains present also if synaptic plasticity is turned off after training (not shown).

We investigate the role of the attenuation factor $e^{-\beta w_{ij}^2}$ in Eq. 1 and the spiking term $\gamma \tilde{s}_i$ in Eq. 2. Without attenuation, training is successful, however clusters are unstable and disappear within one hour after training (Fig S1). Without the term $\gamma \tilde{s}_i$ in Eq. 2, training produces uneven clusters that are unable to sustain cluster activation within one hour after training (Fig S2).

To quantify the amount of learning, we measured the average synaptic weight among synapses connecting neurons sharing at least one sensory stimulus (dubbed 'w_1'), and the average weight among neurons that did not share any sensory stimuli ('w_0'):

$$w_{\mathbb{C}} \doteq \langle w_{ij} \rangle_{\mathbb{C}} = \frac{1}{N_{\mathbb{C}}} \sum_{(i,j) \in S_{\mathbb{C}}} w_{ij}, \tag{3}$$

where $\mathbb{C} \in \{0, 1\}$, $N_{\mathbb{C}}$ is the number of synapses of type $\mathbb{C}$, and $S_{\mathbb{C}}$ is the set of ij indices of synapses of type $\mathbb{C}$. Training significantly increased w_1 compared to w_0, thus mapping successfully the association between sensory inputs and the corresponding neural clusters activated by the stimuli. After training, the excitatory synapses were continuously modified by the plasticity rule during ongoing metastable activity, leading to the decrease of w_1 and increase of w_0 (Fig. 2d). These synaptic changes stabilize after 4 hours; after this time, plasticity is stable and coexists with metastable dynamics unfolding as a random walk among states associated with the training stimuli.

2.3 Learning stability vs. network size

Here we show that spontaneous memory decay (Fig. 2d) slows down as a function of network size. We show this by estimating the rate of change for w_0 and w_1 as a function of N, the number of neurons in the network. Scaling up the network size can be done in several ways (Renart et al., 2007; Doiron and Litwin-Kumar, 2014); we used two different scaling procedures, one in which $Q \propto N$ and one in which $Q \propto \sqrt{N}$.

In the first scenario, we scale up the number of clusters Q proportionally to $N_E = 0.8N$ while keeping constant the mean cluster size $N_E/Q = 80$ (this corresponds to a coding level $f = 1/Q \propto 1/N$, where f is the probability that a neuron is targeted by a stimulus; see Methods). As shown in Fig. 3, synaptic decay is slower in larger networks (Fig. 3a) while ongoing metastable dynamics co-exists with synaptic plasticity (Fig. 3b). In particular, nearly constant values of the mean synaptic weights are observed in networks of $20,000$ neurons (rightmost panel of Fig. 3a).

Analytical arguments imply a synaptic decay rate $\propto N^{-1}$, i.e.,

$$\frac{\Delta w_{\mathbb{C}}}{\Delta t} \propto \frac{1}{N}, \tag{4}$$

where $\Delta w_{\mathbb{C}}$ is the average change in strength over the interval Δt for synapses of type $\mathbb{C}$ (defined analogously to Eq. 3; see Methods). To confirm this prediction, we plotted $N\Delta w_0$ (Fig. 3c) and $N\Delta w_1$ (Fig. 3d) as a function of time and for increasing network size (from $N = 5,000$ to $N = 20,000$). The plots show that, after entering the ongoing metastable regime, both $|N\Delta w_0|$ and $|N\Delta w_1|$ increase linearly with time as predicted by Eq. 4. Moreover, the slopes of $N\Delta w_0$ and $N\Delta w_1$ approach a constant value that is independent of network size when N is large enough (Fig. 3c-d, right panels). The non-linear transients visible in the left panels of Fig. 3c-d are due to either a small N or not having yet reached the stable regime of synaptic decay (which takes about 4h in small networks, see Fig. 2d). For large N and 3h post training, synaptic decay is stable and $N\Delta w_{\mathbb{C}}$ is linear as predicted (Fig. 3c-d, right panels).

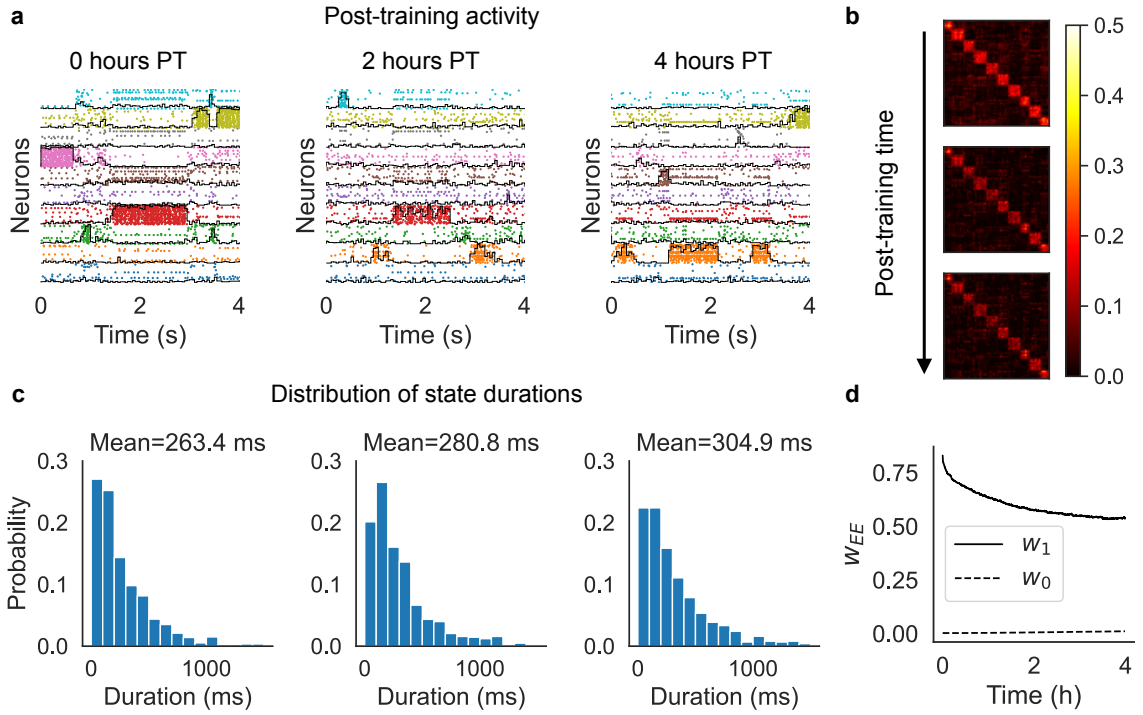

Figure 2: Cluster formation and metastable dynamics via excitatory plasticity in a network with 1,000 spiking neurons (see the text). (a) Rasterplot of excitatory neurons taken immediately after training (left), 2 hours after training (middle), and 4 hours after training (right). Neurons were ordered according to cluster membership (different colors for different clusters), and if necessary also duplicated to multiple clusters according to the sensory stimuli they respond to. *Black lines*: overlap between the network's neural activity and the stimulus associated to each cluster (see the text). (b) Synaptic matrix of the network at the same times as in (a) showing the formation of clusters (evidence from the block structure of the matrix). (c) Distribution of durations of cluster activations (see the text) for each corresponding epoch shown in (a). Samples were taken over a period of 10 minutes. (d) Averaged post-training excitatory synaptic weights as a function of time. w_1: mean weights across synapses connecting neurons sharing at least one stimulus; w_0: mean weights across synapses connecting neurons sharing no stimuli.

These results confirm that the synaptic decay rate $\propto N^{-1}$, implying slower memory decay and more stable synaptic efficacies in larger networks, despite ongoing plasticity. Moreover, the dynamics produced by the networks of Fig. 3 shows near-exponential distributions of the state durations (Fig. S3), with mean durations approaching stability 4 hours post-training and for $N > 15,000$.

The slowdown of synaptic decay rate with N was found also in an alternative scaling scenario in which $f = 1/Q$ with $Q \propto \sqrt{N}$, giving $f \propto 1/\sqrt{N}$ and $fN_E \propto \sqrt{N}$ neurons in each cluster (with $N_E = 0.8N$). In this case we have a smaller number of clusters ($\propto \sqrt{N}$), but each cluster grows in size with N. We trained several networks under this scaling and found results analogous to those of Fig. 3: training is successful, the synapses become more stable in larger networks, and mestastable dynamics is reliable across network sizes (see Fig S4). Compared with the scaling of Fig. 3, however, both w_1 and the excitatory firing rates decreased with network size in a manner that suggests a synaptic decay rate between $1/N$ and $1/\sqrt{N}$

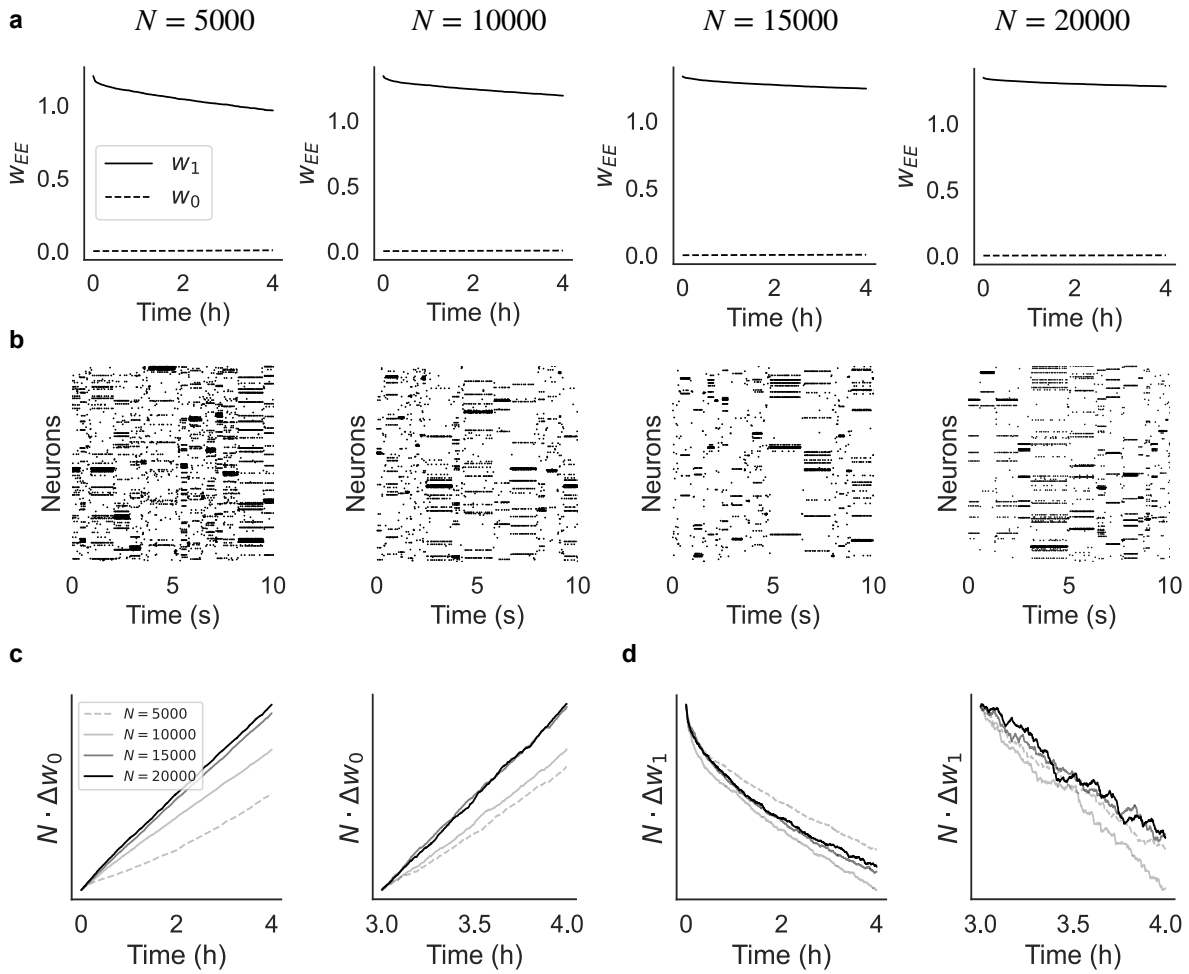

Figure 3: Effect of network size on synaptic dynamics after training. (a) Time dependence of average synaptic weights w_1 and w_0 after training for networks of different size N (same keys as Fig. 2d). In each case, the network had $N_E = 0.8N$ excitatory neurons and $Q = N_E/80 = N/100$ clusters (i.e., $Q = 50, 100, 150$, and 200 from left to right). (b) Raster plots of the network's activity 4 hours after training for the corresponding networks in (a). (c) Plots of $N\Delta w_0$ vs. time after training for the different network in (a). Observations were taken 0 to 4 hours post training (left panel) and 3 to 4 hours post training (right panel). In large networks, $|N\Delta w_0|$ increases linearly with time as predicted by Eq. 4. (d) Same as (c) for $N\Delta w_1$ vs. time. In large networks and $3h$ post-training, $N\Delta w_1$ decreases linearly with time as predicted by Eq. 4.

(Fig. S4c-d).

2.4 Robustness to perturbations and remapping

In the previous section we have shown that the synaptic weights are stable against metastable spontaneous activity. We show next that this is also true in the presence of random stimuli. To show this, we probed

the network with external sensory inputs with similar physical features (i.e. magnitude and coding level) as the training stimuli. The occurrence schedule of these external stimuli was modeled as a Poisson process with a mean inter-event interval of 10 s. We considered three different scenarios. In the first scenario, the external stimuli were randomly sampled and used only once (i.e., the subsets of target neurons were uniformly and independently sampled at each occurrence; Fig. 4a). This scenario mimics a purely noisy environment where there is no temporal correlation in the sensory inputs. In the second scenario, only half the sensory inputs were random inputs, while the other half were sampled from a finite set of 10 sensory stimuli targeting always the same neurons (Fig. 4b). This scenario mimics a combination of meaningful stimuli occurring amid some random sensory background. Finally, in the third scenario the external stimuli were all sampled from a pre-defined set of 10 sensory stimuli, mimicking a situation where the network is being retrained with new stimuli (Fig. 4c). In all cases, the network kept adjusting its synaptic weights according to our plasticity rule.

In the presence of random stimulation, the network maintained a significant memory of the learned stimuli for several hours, similar to what we found in the absence of sensory stimulation (Fig. 4a, top and middle panels). In the third scenario, the network quickly learned the new sensory inputs (Fig. 4c-bottom) while gradually forgetting the previously learned ones (Fig. 4c-middle). Finally, in the 'mixed' stimulation scenario, the network could still learn the new sensory stimuli but at a much slower rate (Fig. 4b-top, compare dashed lines in Figures 4b-top and 4c-top). A raster plot of the neural activity shows the metastable activation of clusters of neurons representing stimuli in both the first and second set (Fig. 4b, middle and bottom panels, respectively), showing that the network could accommodate the learning of new stimuli (Fig. 4b, bottom) while maintaining a trace of the previously learned ones. Note how, in between stimulus presentations, the network dynamics was metastable in all cases.

In conclusion, our plasticity rule can successfully structure a network of spiking neurons into an extensive number of clusters (cell assemblies), in a manner that spontaneously generates and maintains metastable ongoing activity. Both metastable dynamics and the learned structure are stable to random stimulus perturbations, but also flexible enough to be reshaped by new stimuli.

3 Discussion

3.1 Model features and main results

Metastable neural dynamics, defined as the repeated occupancy of a set of discrete neuronal states occurring at seemingly random transition times (Brinkman et al., 2022), has recently come to the fore as a potential mechanism mediating sensory coding, expectation, attention, navigation, decision making and behavioral accuracy (see (La Camera et al., 2019; Brinkman et al., 2022) for reviews). At the same time, model variations over a basic network of spiking neurons with clustered architecture (Miller and Katz, 2010; Deco and Hugues, 2012; Litwin-Kumar and Doiron, 2012; Mazzucato et al., 2015) have accounted for a wealth of data concerning this metastable dynamics (Mazzucato et al., 2015; Mazzucato et al., 2016; Mazzucato et al., 2019; Lang et al., 2023). One is then led to the following question: how can a cortical circuit be shaped by internal dynamics and externally-driven events so as to converge to the clustered architecture producing metastable dynamics?

The model of synaptic plasticity introduced in this work answers this question by offering a simple yet biologically plausible mechanism capable of cluster formation and metastable dynamics. The plasticity rule builds clusters of neurons with strengthened synaptic connections; after training, the neural activity switches among a number of 'states' that can be interpreted as neural representations of the stimuli

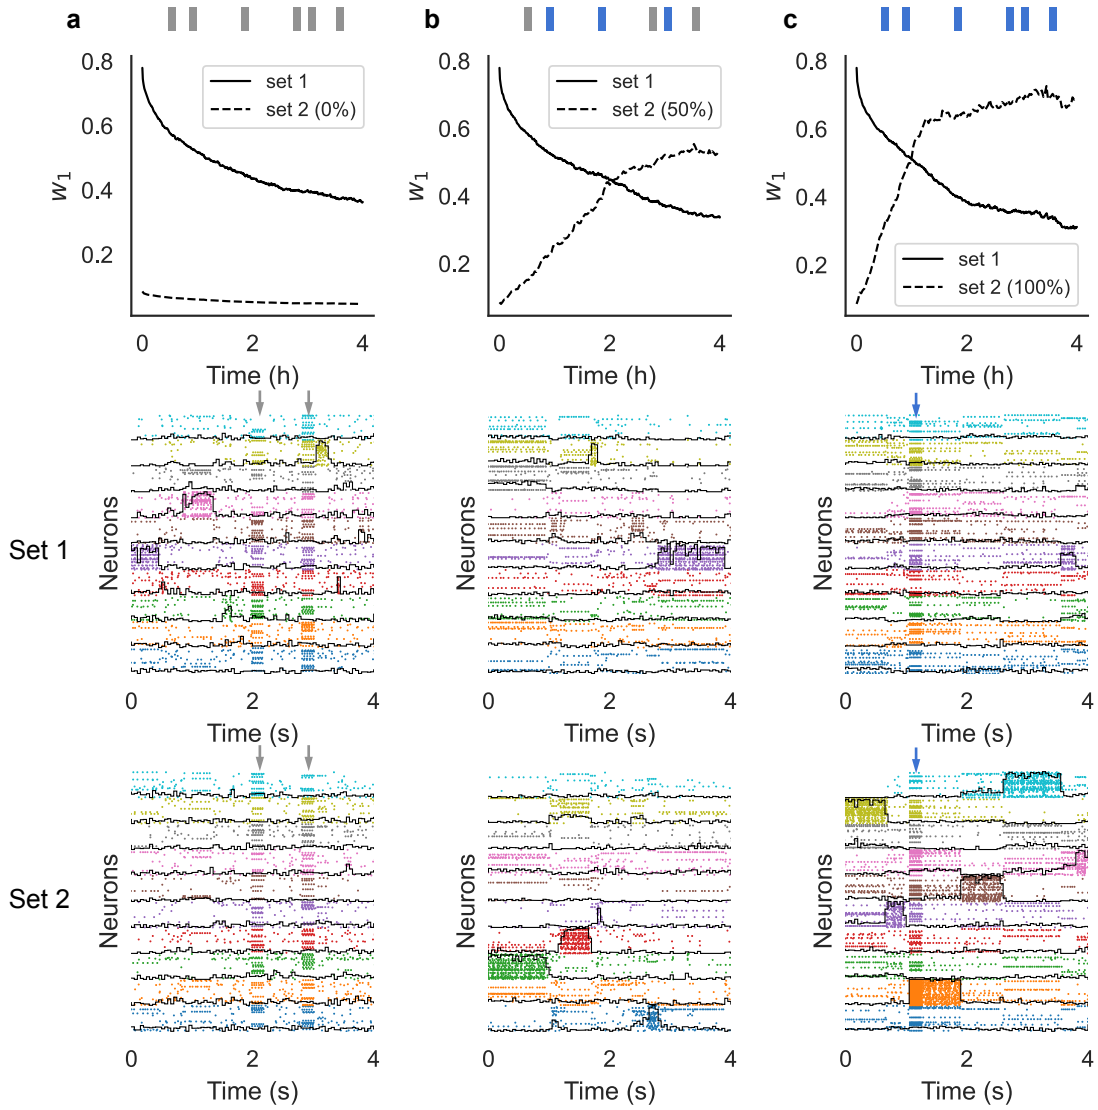

Figure 4: Sensory perturbations and remapping of the network after training. The network was trained with stimuli of Set 1 and then probed for after training, stimuli were drawn from a different set (Set 2) and delivered at random times. (a) Perturbing stimuli were randomly sampled to mimic a noisy environment. (b) 50% of perturbing stimuli were sampled randomly while the other 50% were sampled from a finite pre-defined set of reoccurring stimuli. (c) All stimuli were sampled from a finite pre-defined set of reoccurring stimuli. The top panels show w_1 (Eq. 3) after training for Set 1 (full) and Set 2 (dashed). The middle panels show raster plots of the neural activity 4 hours after training, together with the overlaps (black lines) with the stimuli of Set 1 (those used for training). The bottom panels show raster plots of the neural activity 4 hours after training, together with the overlaps with the stimuli of Set 2 (note that this set was never used in (a)). Vertical arrows indicate an external stimulation with either a random stimulus (grey arrows) or a stimulus from Set 2 (blue arrows).

used for training. Notably, the metastable dynamics generated by learning co-exists, after training, with ongoing synaptic plasticity.

The plasticity rule presented here has many desirable features. One is biological plausibility, in at least two ways: i) the plasticity rule is local, i.e., it depends only on quantities that are available to the synapse, and ii) it works in networks of spiking neurons (rather than artificial neural networks or firing rate models). Another desirable feature of our rule is that it leads to metastable dynamics with similar properties as those observed in experimental data (Abeles et al., 1995; Jones et al., 2007; Ponce-Alvarez et al., 2012; Mazzucato et al., 2015; Maboudi et al., 2018; Benozzo et al., 2021; Recanatesi et al., 2022; Lang et al., 2023), including an exponential distribution of cluster activation durations with means of few hundreds of ms (Mazzucato et al., 2015). Our model also has a minimal number of mechanisms and parameters within a class of similar models that might produce metastable dynamics – for example, a model wherein the presynaptic spike train is replaced by a low-pass filter of it; or a model that imposes an upper bound to the synaptic weights. Those mechanisms are not needed in our plasticity rule. In addition to an adaptive threshold, our model requires only two mechanisms: an attenuation term for LTP (the $e^{-\beta w_{ij}^2}$ term in Eq. 1) and a spiking term $\gamma \tilde{s}_i$ in the equation for the adaptive threshold (Eq. 2). These mechanisms require the introduction of parameters β and γ, which have a clear biological correlate. β controls the amount of LTP. Mechanisms reducing LTP but not LTD have been described *in vitro* by (Delgado et al., 2010), who showed how an increase in total excitatory and inhibitory activity (for example due to cluster activation) can rapidly reduce the amplitude of LTP but not LTD. This mechanism could be mediated through changes in calcium dynamics in dendritic spines and has been proposed as an example of 'rapid compensatory process' by (Zenke et al., 2017). The parameter γ controls the spiking contribution to the threshold's dynamics. This term could emerge from cellular mechanisms similar to those responsible for afterhyperpolarization currents (Sah, 1996; Andrade et al., 2012), which are often used in models of firing rate adaptation (Wang, 1998; Ermentrout, 1998; La Camera et al., 2004). We found no need to invoke additional compensatory mechanisms such as inhibitory plasticity or synaptic normalization often used to keep firing rates and synaptic strengths in check during training or during ongoing activity (Litwin-Kumar and Doiron, 2014; Zenke and Gerstner, 2017).

Another notable feature of our model is its behavior under network scaling. Cortical circuits are variable in size but they typically comprise large numbers of neurons (Shepherd and Grillner, 2010). It is therefore important to test whether models of plasticity maintain their properties as networks are scaled up in size. Scaling up a neural circuit is a necessary operation is several neuroscience theories (e.g., the theory of balanced networks (van Vreeswijk and Sompolinsky, 1998; Renart et al., 2007)) and it can be done in several ways. We have chosen here 2 ways, one that is most intuitive, with order N clusters but fixed cluster size, and one that is most commonly used, wherein both the size and the number of clusters grows as $\sqrt{N}$ (see (Doiron and Litwin-Kumar, 2014) for a discussion of the different scaling regimes). In both cases, we have found that training leads to the formation of stable clusters generating metastable dynamics, at least for networks up to $20,000$ neurons (larger networks becomes computationally expensive to train and simulate). Importantly, both the properties of the dynamics (Fig. S3) and the learned synaptic weights (Fig. 3 and Fig. S4) tend to stabilize as the number of neurons in the circuit grows. In the first scaling regime, we have found that the synaptic decay decreases very fast with N, i.e., $\propto 1/N$.

Finally, our model also potentially solves a conceptual problem arising in deterministic network models that try to combine attractor networks with balanced networks (Renart et al., 2007; Doiron and Litwin-Kumar, 2014). These models produce metastable dynamics due to finite size effects, so that in very large networks stable (rather than metastable) cluster activations would prevail (Doiron and Litwin-Kumar,

2014; Huang and Doiron, 2017; La Camera et al., 2019; Berlemont and Mongillo, 2022). It is not clear how to obtain metastable dynamics in these models in the limit of very large network size. In the presence of our plasticity mechanism, however, persistent activity can only be transient, because a long-lasting cluster activation always induces LTD, which will eventually inactivate the cluster. This mechanism does not depend on network size, as shown in Fig. 3, and ensures the stable coexistence of plasticity and metastable dynamics for arbitrary network size.

3.2 Comparison with previous work

Modeling the self-organization of neural circuits via synaptic plasticity has long been an important topic of research (Dayan and Abbott, 2005; Gerstner et al., 2014) and has recently been studied in a number of works bearing various similarities to our work. However, the first successful simulations of spiking networks dynamics with ongoing synaptic plasticity have appeared only about 20 years ago (Del Giudice and Mattia, 2001; Del Giudice et al., 2003; Amit and Mongillo, 2003). Previous models closest to ours are models of voltage-based STDP rules wherein an internal variable (often interpreted as postsynaptic depolarization) is compared to different thresholds for induction of LTP and LTD (Fusi et al., 2000; Clopath et al., 2010; Litwin-Kumar and Doiron, 2014; Zenke et al., 2015). Part of the motivation for some of these models was to reproduce experimental results on both STDP (Markram et al., 1997; Zhang et al., 1998) and the dependence of synaptic plasticity on postsynaptic depolarization (Artola et al., 1990; Ngezahayo et al., 2000; Sjöström et al., 2001). When applied to self-organization of cortical circuits, most previous work in this context has focused on the formation of clusters (often called cell assemblies) for the formation of stable neural activity representing memories. The goal of these works was therefore to obtain stable attractor dynamics after learning. In contrast, the goal of our work was to obtain a stable synaptic matrix that co-exists with a rich, ongoing metastable dynamics.

A notable exception among previous efforts is the model by (Litwin-Kumar and Doiron, 2014). In their model, the authors combined the voltage-based STDP rule of (Clopath et al., 2010) with inhibitory plasticity (Vogels et al., 2011) and synaptic renormalization to obtain metastable dynamics in a network of adaptive EIF neurons. Inhibitory plasticity was used to maintain a target firing rate in the excitatory neurons and prevent the formation of winner-take-all clusters during training. This problem arises due to the fact that, during training, different stimuli target different numbers of neurons creating inhomogeneities in clusters size (Doiron and Litwin-Kumar, 2014) (in our case, the initial inhomogeneity was 10%). We found that in our model, neither inhibitory plasticity nor synaptic normalization was required, presumably thanks to the presence of a dynamic threshold which keeps the synaptic weights within boundaries as in the original BCM rule (Bienenstock et al., 1982). In our model, the dynamic threshold also produces LTD after prolonged cluster activation. This, together with LTP attenuation (the $e^{-\beta w_{ij}^2}$ term in Eq. 1; see Fig. S1) helps to prevent the formation of large clusters during training, while also keeping the firing rates within an acceptable range. This confers advantages as it frees up inhibitory plasticity (and possibly other network mechanisms) to accomplish other tasks.

More recently, (Wang and Aljadeff, 2022) have proposed a plasticity rule that produces switching dynamics that resembles ours. In their model, rather than comparing a postsynaptic variable to a threshold for LTP/LTD, a Hebbian term is compared to each threshold. The Hebbian term is the low-pass filter of the product of the low-pass filters of the pre- and post-synaptic spike trains, and was interpreted as a calcium signal responsible for synaptic plasticity (Inglebert et al., 2020). The main goal of (Wang and Aljadeff, 2022) was to present a learning rule resulting in stable and unimodal weight distributions after learning. The authors show that their learning rule has intrinsic homeostatic properties without having to impose an additional homeostatic mechanism on timescales which are much shorter

than observed experimentally (Zenke et al., 2017). Perhaps due to these homeostatic properties, they reported metastable dynamics with 2 or 3 stimuli after training. However, as the aim of their work was not to investigate the potential for metastable dynamics, training was only accomplished with few stimuli, and no scaling of the network size was attempted.

We have shown that our plasticity rule allows to learn new stimuli. A similar result was obtained with the plasticity rule of (Litwin-Kumar and Doiron, 2014), which however requires inhibitory plasticity and synaptic renormalization. More recently, (Manz and Memmesheimer, 2022) have proposed a pure STDP plasticity rule that, similarly to our rule, does not require inhibitory plasticity or synaptic renormalization. However, (Manz and Memmesheimer, 2022) focus on the generation of stable clusters producing persistent activity rather than metastable dynamics. Their persistent activity has the signature of fast Poisson noise due to the stochastic nature of their model neurons, but lacks the slow fluctuations characteristic of ongoing metastable dynamics (Litwin-Kumar and Doiron, 2012). We also note that our results on remapping of sensory stimuli is different from the reorganization of clusters described in (Manz and Memmesheimer, 2022), which has been related to representational drift (Ziv et al., 2013; DeNardo et al., 2019; Rule et al., 2019; Masset et al., 2022). In our model, new clusters form due to stimulation by a new set of stimuli rather than drifting of the representations due to the noisy dynamics of plasticity as it occurs in (Manz and Memmesheimer, 2022).

Our model is reminiscent of the BCM rule (Bienenstock et al., 1982), and in fact it could be understood as a BCM rule for spiking neurons in which the non-linearity of the threshold equation is given by a hyperbolic tangent. More common implementations of BCM utilize a power function (often a quadratic function; see e.g. (Intrator and Cooper, 1992; Dayan and Abbott, 2005)). Another difference with our model is the presence, in the latter, of the additional term $\gamma \tilde{s}_i$ in the threshold dynamics, Eq. 2. In our plasticity rule, this term is required to enforce LTD when the neuron is constantly active (i.e. $\tilde{v}_i \gg 0$). One is naturally led to consider whether or not this term could be removed if we replace the hyperbolic tangent with a power of $\tilde{v}_i$, e.g.,

$$\tau_\theta \dot{\theta}_i = -\theta_i + a\tilde{v}_i^p, \tag{5}$$

where a is a constant and $p > 1$. To best match our model, p should be small, e.g. 1.1–1.2. We found that this learning rule does indeed work; however, its performance during remapping is somewhat inferior (not shown). Another problem of this learning rule is that it is very sensitive to the parameter values: for example, while $p = 1.1, a = 1.5$ gives good results, synaptic weights decay fast when $p = 1.2, a = 1.4$. This problem is somewhat ameliorated when using a squashing function of $-\theta_i + a\tilde{v}_i^p$, for example, an hyperbolic function again:

$$\tau_\theta \dot{\theta}_i = \theta_a \tanh(-\theta_i + a\tilde{v}_i^p). \tag{6}$$

In conclusion, although alternative versions of the adaptive threshold may work in principle, we found that the threshold dynamics, Eq. 2, is both more robust and better performing.

3.3 Limitations and possible extensions

Our model has a minimal number of ingredients and did not use widespread properties of cortical neurons, such as external noise, firing rate adaptation or short-term plasticity. One reason is model economy, i.e., the desire to use only minimal, essential ingredients. Another reason is that external noise, firing rate adaptation and short-term plasticity all facilitate metastable dynamics (Moreno-Bote et al., 2007; Giugliano et al., 2008; Deco and Hugues, 2012; Cao et al., 2016; Jercog et al., 2017; Setareh et al., 2017;

Ballintyn et al., 2019), leading to a possible confound on the role of long-term modifications on generating and maintaining ongoing metastable dynamics.

Although synaptic plasticity depends on pre- and postsynaptic calcium (Sjöström and Nelson, 2002; Shouval et al., 2002; Graupner and Brunel, 2012), we do not consider directly the role of calcium in our model. Rather, our learning rule belongs to the class of voltage-based STDP models (Clopath et al., 2010). These models aim to capture the dependence of LTP and LTD on the membrane potential (Ngezahayo et al., 2000; Sjöström et al., 2001; Wang and Maffei, 2014) while producing STDP curves as an emergent phenomenon. Although voltage-based STDP can capture in vitro results where the membrane potential at the electrode is well defined, in a real neuron with an extended geometry the membrane potential is not generally uniform along the dendritic tree. At the time of a postsynaptic spike, one might assume that a back-propagating action potential imposes a uniform voltage on the dendritic tree, at least in proximal dendrites (Stuart and Sakmann, 1994; Nuriya et al., 2006), although faithful propagation depends on many factors including the order of the somatic action potential within a train and the location of the spines (Spruston et al., 1995; Waters et al., 2005). At the moment of a presynaptic spike, however, the voltage at different locations will be different and probably reflect random subthreshold fluctuations. The plasticity rule in such a case would be at the whim of random fluctuations in membrane potential, unrelated to the timing correlation between pre and postsynaptic spikes. We argue that this does not pose a problem in our model because when our postsynaptic variable $\tilde{v}_i$ follows random subthreshold fluctuations, our plasticity rule produces no net average synaptic change (Fig. 1b). On the other hand, when the postsynaptic neuron emits a spike, $\tilde{v}_i$ follows closely the transient exponential increase in membrane potential (Fig. 1c), modeling the more uniform effect of the backpropagating action potential along the dendrites. It is even tempting to speculate that, during a prolonged period of active firing, the amplifying buildup of $\gamma\tilde{s}_i$ could reflect the facilitation of dendritic calcium spikes by trains of backpropagating action potentials (Larkum et al., 1999).

Our model did not include inhibitory plasticity (Maffei, 2011; Vogels et al., 2013). This was because adding inhibitory plasticity to our plasticity rule did not qualitatively change our results (not shown), and therefore was deemed unnecessary. A benefit of including inhibitory plasticity in our model could be to autonomously set the level of inhibitory activity best suited for metastable dynamics (Zenke et al., 2015). This useful function of inhibitory plasticity contrasts with its homeostatic role in keeping the postsynaptic firing rates close to a desired target value (Vogels et al., 2011; Litwin-Kumar and Doiron, 2014), a role that in our model is fulfilled by the adaptive threshold for LTP and LTD.

3.4 Conclusion

We have proposed a biologically plausible model of synaptic plasticity that produces stable cell assemblies in a deterministic spiking network of excitatory and inhibitory neurons. The ensuing clustered architecture of the network produces ongoing metastable dynamics that co-exists with synaptic plasticity. The metastable dynamics supports seemingly random switching among hidden states representing the stimuli used for training, and has several characteristic traits of the metastable dynamics observed in brain regions of rodents and primates. Our model could also provide an explanation for how a clustered network of spiking neurons can generate metastable dynamics for large network sizes in the absence of external noise.

Acknowledgments

We wish to thank Drs. Arianna Maffei, Yonatan Aljadeff and Paul Adams for a critical reading of the manuscript and very useful discussions. This work was partially supported by a U01 grant from the NIH/NINDS Brain Initiative (1UF1NS115779) to G.L.C.

4 Methods

Spiking network model. Here we define the 'basic' network, from which all other networks were obtained by scaling up the number of neurons N and the number of stimuli Q (see 'Network scaling' below). The basic network comprised $N_E = 800$ excitatory and $N_I = 200$ inhibitory exponential integrate-and-fire (EIF) neurons. The ratio of excitatory to inhibitory neurons was $N_E/N_I = 4$, in agreement with previous studies and experimental observations (Braitenberg and Schüz, 1991; Amit and Brunel, 1997; Shepherd and Grillner, 2010). Before training, neurons were randomly connected with fixed probability (0.2 among excitatory neurons and 0.5 in all the other cases) and constant synaptic efficacies ($w_{EE} = 0.005$, $w_{EI} = -0.34$, $w_{IE} = 0.54$ and $w_{II} = -0.46$ mV). Only connected excitatory neurons underwent plasticity.

We modeled the impact of $Q = 10$ random stimuli as follows. Each neuron had a probability ('coding level') $f = 0.1$ to be targeted by a stimulus, meaning that the neuron would receive external input during the presentation of that stimulus (we call such neurons 'responsive'). This resulted in a typical number of $fN_E \pm \sqrt{f(1-f)N_E} = 80 \pm 8.4$ neurons targeted by each stimulus and a mean fraction $(1-f)^Q \approx 0.35$ of non-responsive neurons in the basic network (in keeping with similar numbers found in the experimental literature, see e.g. (Perin et al., 2011)). We call a 'cluster' a subpopulation of neurons targeted by a given stimulus during training. There were a mean fraction $1-(1-f)^Q \approx 0.65$ of neurons in (overlapping) clusters. A randomly picked neuron had a probability $P_{\geq 2} = 1-(1-f)^Q - Qf(1-f)^{Q-1} \approx 0.26$ to respond to at least two stimuli, which results in about $P_{ov} = 1-(1-f)^{Q-1} \approx 0.61$ probability for a *responsive* neuron to respond to multiple stimuli (and therefore belong to multiple clusters). Parameter values are summarized in Table 1.

Network scaling. For the analysis of Fig. 3, N, N_E and Q were increased with $N_E = 0.8N$, $Q = N/100$ and coding level $f = 1/Q$, resulting in a constant mean cluster size of $N_Q = fN_E = 80$ neurons in all cases. For the analysis of Fig. S4, we used $f = 1/Q$ with $Q = \sqrt{N/10}$ stimuli, resulting in $N_Q = 0.8\sqrt{10N}$ neurons in each cluster (compatible with $QN_Q = 0.8N = N_E$ excitatory neurons in total). The factors in the second scaling were chosen so as to agree with the 'basic' network for $N = 1,000$ ($f = 0.1, Q = 10, N_Q = 80$).

Single neuron dynamics. The membrane potential of the i-th EIF neuron followed the dynamical equation

$$\frac{dV_i}{dt} = -\frac{V_i - V_L}{\tau_m} + \frac{\Delta_T}{\tau_m} e^{\frac{V_i - V_T}{\Delta_T}} + h_{iE}^{syn} + h_{iI}^{syn} + h_i^{st} + h_i^{ext},$$

where $V_L = 0\,\text{mV}$ is the 'leak' potential (practically equal to the resting potential in this model), $\Delta_T = 1\,\text{mV}$ is a 'sharpness' parameter related to spike width and upstroke velocity, and $V_T = 20\,\text{mV}$ is a

EIF neurons		
V_L	0	leak potential (mV)
V_T	20	soft spiking threshold (mV)
V_{peak}	25	peak potential (mV)
V_r	0	reset potential (mV)
τ_m	15(E), 10(I)	membrane time constant (ms)
τ_{syn}	3(E), 2(I)	synaptic time constant (ms)
Network (pre-training)		
N	from 1,000 to 20,000	total number of neurons
N_E/N_I	4	ratio of excitatory to inhibitory neurons
$p_{\alpha\beta}$	$p_{EE} = 0.2$, $p_{EI} = p_{IE} = p_{II} = 0.5$	connection probability
w_{EE}	0.005	$E \to E$ synaptic efficacy (mV)
w_{EI}	-0.34	$I \to E$ synaptic efficacy (mV)
w_{IE}	0.54	$E \to I$ synaptic efficacy (mV)
w_{II}	-0.46	$I \to I$ synaptic efficacy (mV)
h^{st}	0.5	stimulus current ($\mathrm{mV} \cdot \mathrm{ms}^{-1}$)
h^{ext}	1.50(E), 2.19(I)	external current ($\mathrm{mV} \cdot \mathrm{ms}^{-1}$)
Q	10 (basic network)	number of stimuli
f	0.1 (basic network)	coding level
Plasticity rule		
τ_v	50	time constant of $\tilde{v}_i$ (ms)
τ_θ	500	time constant of θ_i (ms)
τ_s	1000	time constant of $\tilde{s}_i$ (ms)
θ_a	1	"amplitude" in Eq. 2 (mV)
g	5	gain factor in Eq. 2 (mV^{-1})
γ	50	contribution of spikes in Eq. 2 (mVms)
β	0.1	LTP attenuation factor in Eq. 1 (mV^{-2})
A_{LTP}	0.005	LTP strength
A_{LTD}	0.015	LTD strength

Table 1: Model parameters. See the text for details.

reference potential somewhat related to the spike threshold (it would be the spike threshold in the limit $\Delta_T \to 0$). A spike was said to be emitted when $V_i \geq V_{peak} = 25\,\mathrm{mV}$, after which the membrane potential was reset to $V_r = 0\,\mathrm{mV}$. The membrane time constant τ_m was 15 ms for excitatory neurons and 10 ms for inhibitory neurons. The term h_i^{st} represents a constant stimulus input current which was present only during stimulus presentation. h_i^{ext} represents a constant external current, presumably coming from more distant regions or other brain areas. The term $h_{i\alpha}^{syn}(t)$, where $\alpha \in \{E, I\}$, is the recurrent synaptic input coming from excitatory and inhibitory neurons of the network, respectively. The synaptic input to neuron i obeyed the equation

$$\tau_{syn,\alpha}\, \dot{h}_{i\alpha}^{syn}(t) = -h_{i\alpha}^{syn}(t) + \sum_{j\in\alpha} w_{ij} \sum_k \delta(t - t_j^{(k)}),$$

where $t_j^{(k)}$ denotes the arriving time of the k-th spike of neuron j, w_{ij} is the synapse weight from neuron j to neuron i, and $\tau_{syn,\alpha}$ the synaptic time constant of the presynaptic inputs, equal to $\tau_{syn,E} = 3\,\mathrm{ms}$ for excitatory inputs and $\tau_{syn,I} = 2\,\mathrm{ms}$ for inhibitory inputs (see Table 1).

Plasticity rule and training protocol. In our model, only the synapses between excitatory neurons were plastic and they were subject to the following plasticity rule:

$$\dot{w}_{ij} = \left(A_{\mathrm{LTP}} e^{-\beta w_{ij}^2} [\tilde{v}_i(t) - \theta_i(t)]_+ - A_{\mathrm{LTD}} [\theta_i(t) - \tilde{v}_i(t)]_+ \right) s_j(t), \tag{7}$$

where $s_j(t) = \sum_k \delta(t - t_j^{(k)})$ is the presynaptic spike train and $[x]_+ := \max(x, 0)$ is the rectified linear function. The dynamical threshold $\theta_i(t)$ evolved according to Eq. 2 of the main text:

$$\tau_\theta \dot{\theta}_i = \theta_a \tanh[g(-\theta_i + \tilde{v}_i + \gamma \tilde{s}_i)],$$

where the notation $\tilde{x}$ indicates a variable obeying the equation $\tau_x \dot{\tilde{x}} = -\tilde{x} + x(t)$, i.e., $\tilde{x}$ is a low-pass filter of the variable $x(t)$. Specifically, $\tilde{v}_i$ is the exponential voltage term $\Delta_T e^{\frac{V_i - V_T}{\Delta_T}}$ low-pass filtered with time constant $\tau_v = 50\,\mathrm{ms}$, while $\tilde{s}_i$ is the postsynaptic spike train low-pass filtered with time constant $\tau_s = 1\,\mathrm{s}$. τ_θ, θ_a and γ were constant (see Table 1).

The time-scale of θ dynamics is roughly a linear function of its argument (Fig. 1a). Specifically, consider the equation $\tau_\theta \dot{\theta} = \tanh(-\theta + \Delta)$, when a change in its fixed point Δ occurs at $t = 0$ ($\theta = 0$ when $t < 0$). The time it takes for θ to reach half of Δ can be computed to be $T_{\mathrm{half}}(\Delta) = \tau_\theta \ln \frac{\sinh(\Delta)}{\sinh(\Delta/2)}$, implying a fixed time constant $T_{\mathrm{half}} \approx \tau_\theta \ln 2$ when $\Delta \ll 1$ and $T_{\mathrm{half}} \approx \tau_\theta \Delta/2$ when $\Delta > 3$ (Fig. 1a).

In the basic network, initial training was performed by randomly presenting one stimulus out of Q every 2000 ms for a duration of 10 minutes (presentation at random times did not affect the results). Each stimulus lasted 500 ms and and was presented an average of 30 times during the training period. The training time was linearly scaled in larger networks (i.e., 20 minutes for a network with $2Q$ clusters, and so on), resulting in the same mean number of presentations per stimulus in all networks. Perturbing stimuli after training (Fig. 4), whether reoccurring or not, were presented for 200 ms at random times characterized by an exponential inter-event distribution with an average of 10 s.

Quantification of neural and synaptic activity. A cluster is defined as a subpopulation of excitatory neurons targeted by a given stimulus during training. An active cluster (see below) is interpreted as an internal representation of the corresponding stimulus, and therefore a memory of that stimulus in the absence of external stimulation (active clusters could also represent decision variables and other abstract variables, see e.g. (La Camera et al., 2019)). Cluster activity was measured by the normalized firing rate of its neurons, specifically, for the q-th cluster,

$$m_q(t) = \frac{\sum_{i \in E} \eta_{q,i} \nu_i(t)}{\sum_{i \in E} \nu_i(t)}, \quad q = 1 \ldots Q, \tag{8}$$

where $\nu_i(t)$ is the firing rate of neuron i at time t and $\eta_{q,i} = 1$ if neuron i is a target of stimulus q, otherwise it is zero. We call m_q the 'overlap' between the network activity and stimulus q. By definition, $0 \leq m_q \leq 1$: $m_q = 1$ implies that only neurons responding to stimulus q have non-zero firing rates; $m_q = 0$ means that the neurons targeted by stimulus q are all silent. A memory for stimulus q was said to be active if $m_q > 0.5$. Since $\sum_q m_q \approx 1$, this guarantees that at any given time there is at most one active memory state ($\sum_q m_q$ is not exactly 1 because neurons can be targeted by multiple stimuli).

To quantify the amount of learning, we defined the average synaptic weight among synapses connecting neurons sharing at least one sensory stimulus (dubbed 'w_1'), and the average weight among neurons that did not share any sensory stimuli ('w_0') according to Eq. 3,

$$w_{\mathbb{C}} \doteq \frac{1}{N_{\mathbb{C}}} \sum_{(i,j) \in S_{\mathbb{C}}} w_{ij}, \tag{9}$$

where $\mathbb{C} \in \{0, 1\}$, $N_{\mathbb{C}}$ is the number of synapses of type $\mathbb{C}$, and $S_{\mathbb{C}}$ is the set of ij indices of synapses of type $\mathbb{C}$. The mean synaptic *change* for synapses of type $\mathbb{C}$, $\Delta w_{\mathbb{C}}$, was defined analogously, i.e., by replacing w_{ij} with Δw_{ij} in Eq. 9:

$$\Delta w_{\mathbb{C}} \doteq \frac{1}{N_{\mathbb{C}}} \sum_{(i,j)\in S_{\mathbb{C}}} \Delta w_{ij}. \tag{10}$$

Scaling laws for synaptic decay rate. Here we show that, in the limit of large N, the synaptic decay rate is proportional to excitatory firing rates and the coding level f. Under the scaling used in Fig. 3, i.e., $Q \propto N_E \propto N$ and $f \sim 1/Q$ as $N \to \infty$, this leads to a synaptic decay rate proportional to $1/N$.

According to the plasticity rule, Eq. 7, the rate of change in synaptic weight w_{ij} is proportional to the product of a post-synaptic function of membrane potential, ϕ_i, and the presynaptic spike train $s_j(t) = \sum_k \delta(t - t_j^{(k)})$. The total synaptic weight change for synapses of a given class $\mathbb{C}$ (see Eq. 9) during an interval Δt, is therefore

$$\sum_{(i,j)\in S_{\mathbb{C}}} \Delta w_{ij} \propto \int_t^{t+\Delta t} \sum_i^{N_E} \phi_i \sum_{j:(i,j)\in S_{\mathbb{C}}}^{N_E} s_j(u)\, du, \tag{11}$$

where i, j in the right hand side are such that $i, j \in \mathbb{C}$. The postsynaptic term is the sum of two contributions, one coming from active neurons $\phi_i^{(a)}$ and one coming from inactive neurons $\phi_i^{(in)}$, where active neurons are those firing in an active cluster. Of these two contributions, only $\phi_i^{(a)}$ has non-zero mean during learning, because occasional pre- and post-synaptic spikes in inactive neurons produce no average synaptic change, as shown in Fig. 1c.

We further assume that we can separate the pre- and postsynaptic terms in the integral (on account that the latter term is much slower than the former), obtaining

$$\sum_{(i,j)\in S_{\mathbb{C}}} \Delta w_{ij} \propto \sum_i^{N_E} \phi_i^{(a)} \int_t^{t+\Delta t} \sum_{j:(i,j)\in S_{\mathbb{C}}}^{N_E} s_j(u)\, du. \tag{12}$$

Assuming that only a finite number of clusters is active at any given time, the mean number of active postsynaptic neurons connected by a synapse of type $\mathbb{C}$ is proportional to $fN_E \propto fN$ (the mean number of neurons in each cluster), giving $\sum_i^{N_E} \phi_i^{(a)} \propto fN$.

Assuming Poisson spike trains of average rate $\nu_{\mathbb{C}}$, the presynaptic term scales as

$$\int_t^{t+\Delta t} \sum_{j:(i,j)\in S_{\mathbb{C}}}^{N_E} s_j(u)\, du \;\propto\; p_{\mathbb{C}} N \nu_{\mathbb{C}} \Delta t, \tag{13}$$

where $p_{\mathbb{C}}$ is the probability that a synapse is of type $\mathbb{C}$. Putting all together, we get

$$\sum_{(i,j)\in S_{\mathbb{C}}} \Delta w_{ij} \propto p_{\mathbb{C}} f N^2 \nu_{\mathbb{C}} \Delta t. \tag{14}$$

From this equation and Eq. 10 with $N_{\mathbb{C}} \approx \langle N_{\mathbb{C}} \rangle = p_{\mathbb{C}} N_E^2 \sim p_{\mathbb{C}} N^2$, the rate of change of $w_{\mathbb{C}}$ in an interval Δt is then

$$\frac{\Delta w_{\mathbb{C}}}{\Delta t} = \frac{1}{N_{\mathbb{C}} \Delta t} \sum_{(i,j)\in S_{\mathbb{C}}} \Delta w_{ij} \propto \frac{p_{\mathbb{C}} f N^2 \nu_{\mathbb{C}} \Delta t}{p_{\mathbb{C}} N^2 \Delta t} \propto \nu_{\mathbb{C}} f, \tag{15}$$

where $\nu_{\mathbb{C}}$ is the average firing rate of the excitatory presynaptic neurons connected by a synapse of type $\mathbb{C}$.

When $f \sim 1/N$ (Fig. 3), the firing rates in the active clusters remain nearly constant with N (not shown), leading to Eq. 4 of the main text.

When $f \sim 1/\sqrt{N}$ (Fig. S4), both w_1 (Fig. S4a) and the firing rates (not shown) scale with N, in a manner that suggests a synaptic decay rate between $1/N$ and $1/\sqrt{N}$ (Fig. S4c-d).

A Supplementary Figures

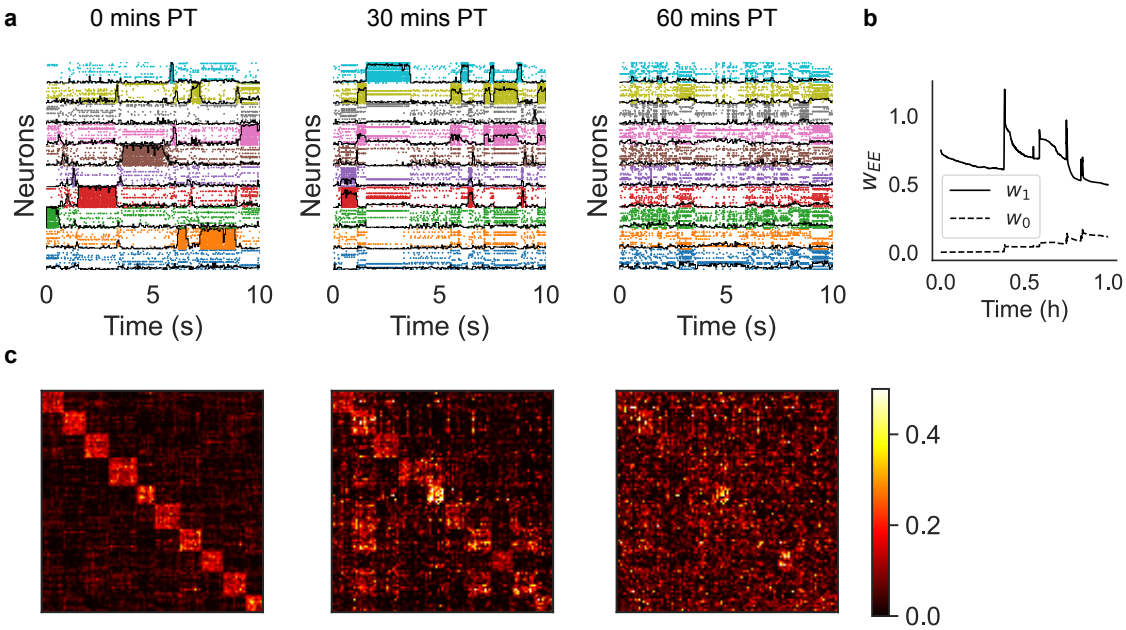

Figure S1: Results from training the basic network of Fig. 2 with $\beta = 0$ in Eq. 1. (a) Rasterplot of excitatory neurons taken immediately after training (left), 30 minutes after training (middle), and 60 minutes after training (right). Same keys as Fig. 2 of the main text. (b) Averaged post-training excitatory synaptic weights as a function of time. w_1: mean weights across synapses connecting neurons sharing at least one stimulus; w_0: mean weights across synapses connecting neurons sharing no stimuli. (c) Synaptic matrix of the network at the same times as in (a) showing the formation of clusters from the block structure of the matrix.

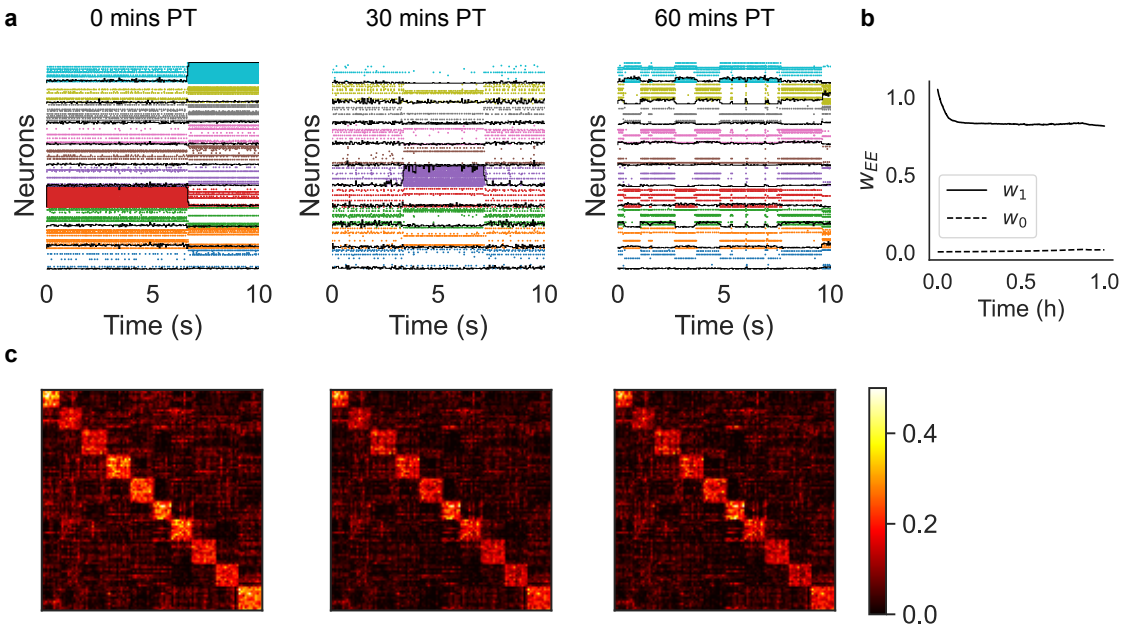

Figure S2: Results from training the basic network of Fig. 2 with $\gamma = 0$ in Eq. 2. Same keys as Fig. S1.

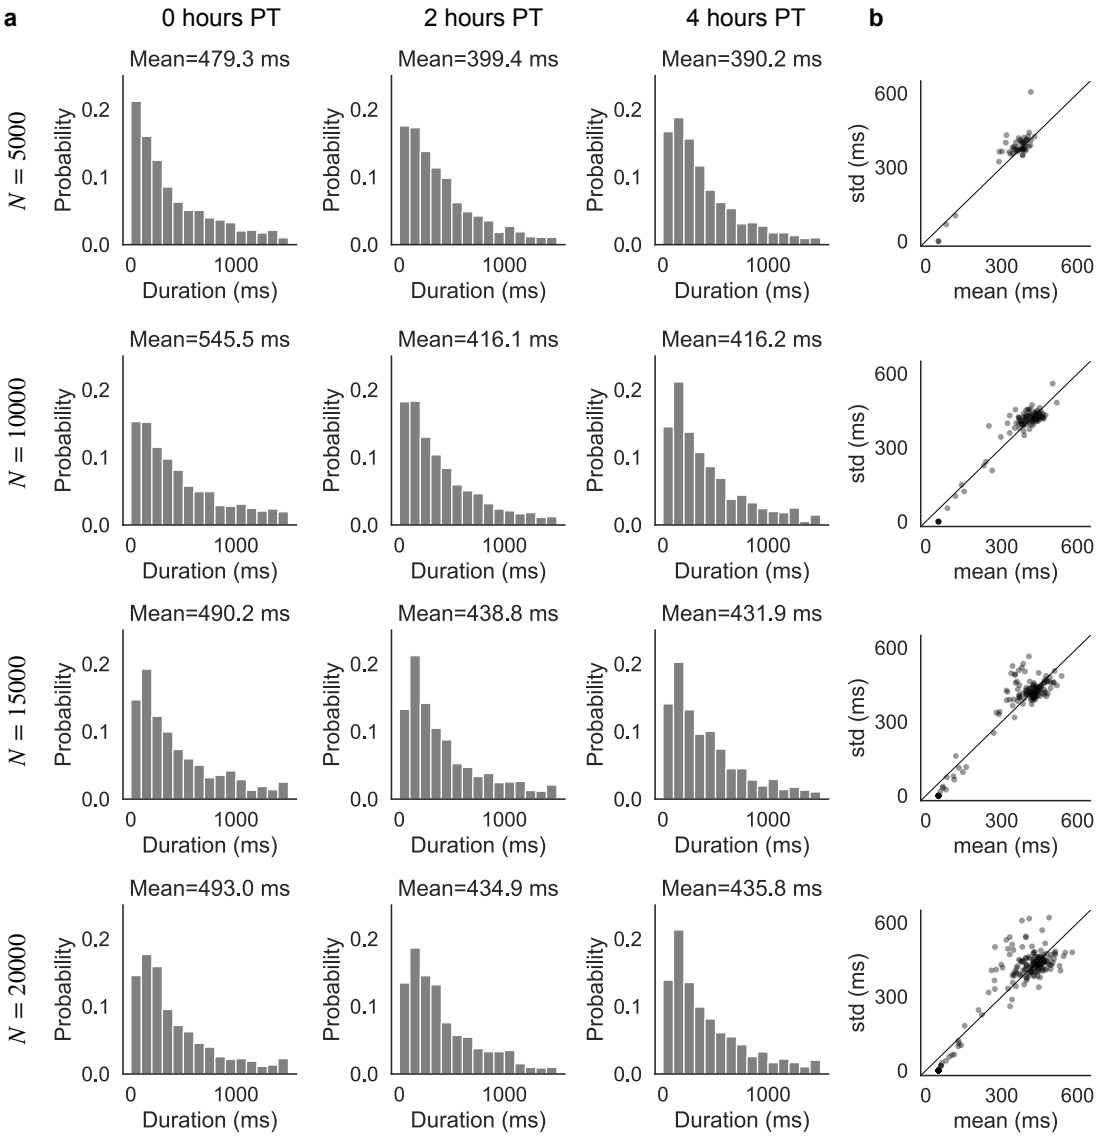

Figure S3: Distributions of durations of cluster activations vs. network size for the networks of Fig. 3. (a) Distributions of durations right after training (left column), 2 hours post-training (middle) and 4 hours post-training (right) for different network sizes. Means tend to decrease with post-training time and increase with network size, approaching stability 4 hours post-training and for $N \geq 15,000$. (b) Scatterplots of standard deviation vs. mean of durations for the corresponding networks in (a), superimposed to the identity line. Each circle corresponds to a cluster. For the majority of the clusters, the standard deviations are approximately equal to the means as expected for an exponential distribution.

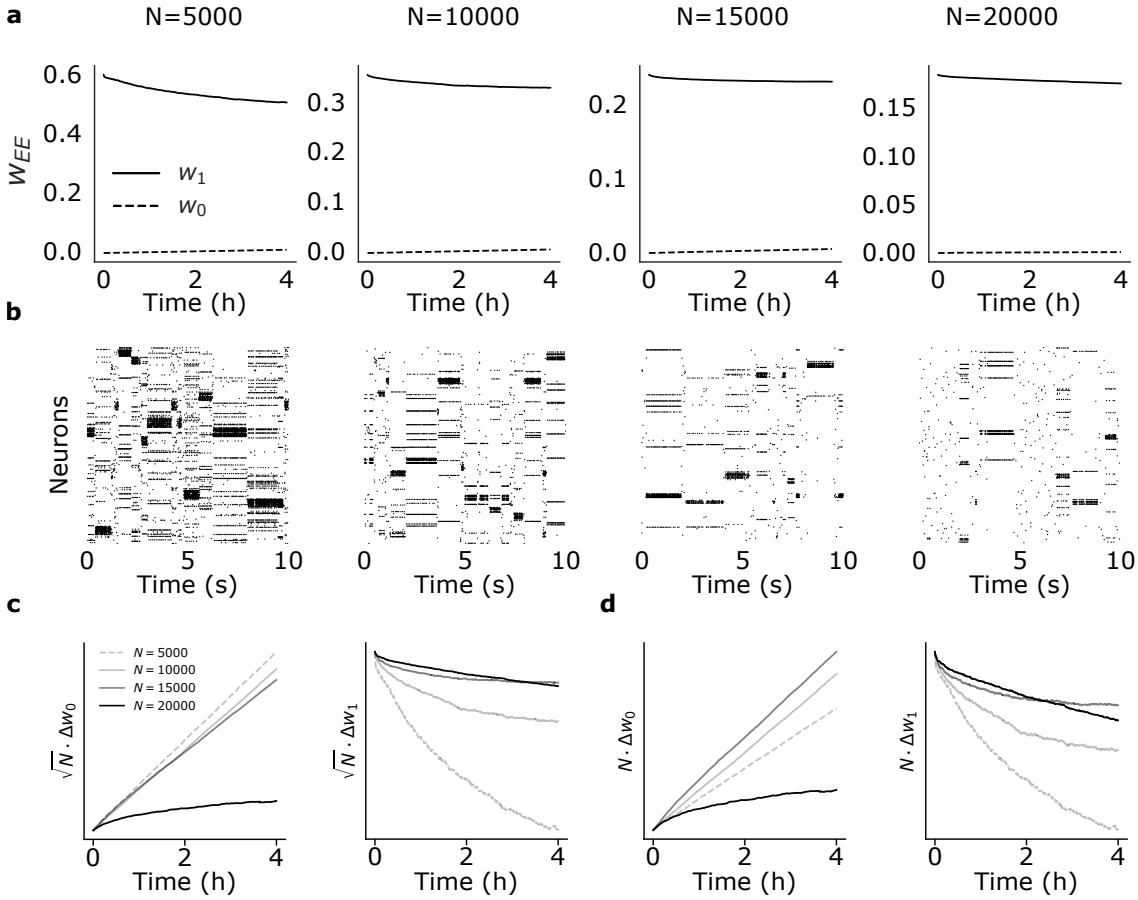

Figure S4: Effect of network size on synaptic dynamics after training. Same as Fig. 3 of the main text but for a different scaling, namely $f \propto 1/\sqrt{N}$, $Q \propto \sqrt{N}$ (see the main text). (a) Time dependence of average synaptic weights w_1 and w_0 after training for networks of different size N, with $N_E = 0.8N$ excitatory neurons. Scaling laws were $f = 1/Q$ with $Q = \sqrt{N/10}$, giving $N_Q = fN_E = 0.8\sqrt{10N}$ neurons in each cluster (where Q and N_Q were rounded to the nearest integer). From left to right, $Q = 22, 32, 39$ and 45. (b) Raster plots of the network's activity 4 hours after training for the corresponding networks in (a). (c) Plots of $\sqrt{N}\Delta w_0$ and $\sqrt{N}\Delta w_1$ vs. time after training for the different network in (a). Observations were taken 0 to 4 hours post training. (d) Same as (c) for $N\Delta w_{\mathbb{C}}$ vs. time. Panels c-d are inconclusive but suggest a synaptic decay rate between $1/N$ and $1/\sqrt{N}$.

References

Abeles M, Bergman H, Gat I, Meilijson I, Seidemann E, Tishby N, Vaadia E (1995) Cortical activity flips among quasi-stationary states. *Proc. Natl. Acad. Sci. U. S. A.* 92: 8616–8620.

Amit DJ, Brunel N (1997) Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb. Cortex* 7: 237–252.

Amit DJ, Mongillo G (2003) Spike-driven synaptic dynamics generating working memory states. *Neural Comput* 15: 565–96.

Andrade R, Foehring RC, Tzingounis AV (2012) The calcium-activated slow AHP: cutting through the gordian knot. *Front Cell Neurosci* 6: 47.

Artola A, Bröcher S, Singer W (1990) Different voltage-dependent thresholds for inducing long-term depression and long-term potentiation in slices of rat visual cortex. *Nature* 347: 69–72.

Ballintyn B, Shlaer B, Miller P (2019) Spatiotemporal discrimination in attractor networks with short-term synaptic plasticity. *J Comput Neurosci* 46: 279–297.

Benozzo D, La Camera G, Genovesio A (2021) Slower prefrontal metastable dynamics during deliberation predicts error trials in a distance discrimination task. *Cell Rep* 35: 108934.

Berlemont K, Mongillo G (2022) Glassy phase in dynamically-balanced neuronal networks. *bioRxiv* 2022.03.14.484348.

Bi GQ, Poo MM (1998) Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J Neurosci* 18: 10464–72.

Bienenstock EL, Cooper LN, Munro PW (1982) Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J Neurosci* 2: 32–48.

Braitenberg V, Schüz A (1991) *Anatomy of the cortex.* Berlin: Springer-Verlag.

Brinkman BAW, Yan H, Maffei A, Park IM, Fontanini A, Wang J, La Camera G (2022) Metastable dynamics of neural circuits and networks. *Appl Phys Rev* 9: 011313.

Cao R, Pastukhov A, Mattia M, Braun J (2016) Collective activity of many bistable assemblies reproduces characteristic dynamics of multistable perception. *J Neurosci* 36: 6957–72.

Clopath C, Büsing L, Vasilaki E, Gerstner W (2010) Connectivity reflects coding: a model of voltage-based STDP with homeostasis. *Nat Neurosci* 13: 344–52.

Curti E, Mongillo G, La Camera G, Amit DJ (2004) Mean-Field and capacity in realistic networks of spiking neurons storing sparsely coded random memories. *Neural Computation* 16: 2597–2637.

Dayan P, Abbott LF (2005) *Theoretical neuroscience: computational and mathematical modeling of neural systems.* Massachusetts Institute of Technology Press, Cambridge, Mass.

Deco G, Hugues E (2012) Neural network mechanisms underlying stimulus driven variability reduction. *PLoS Comput Biol* 8: e1002395.

Del Giudice P, Mattia M (2001) Long and short-term synaptic plasticity and the formation of working memory: A case study. *Neurocomputing* 38-40: 1175–1180.

Del Giudice P, Fusi S, Mattia M (2003) Modelling the formation of working memory with networks of integrate-and-fire neurons connected by plastic synapses. *J Physiol Paris* 97: 659–81.

Delgado JY, Gómez-González JF, Desai NS (2010) Pyramidal neuron conductance state gates spike-timing-dependent plasticity. *J Neurosci* 30: 15713–25.

DeNardo LA, Liu CD, Allen WE, Adams EL, Friedmann D, Fu L, Guenthner CJ, Tessier-Lavigne M, Luo L (2019) Temporal evolution of cortical ensembles promoting remote memory retrieval. *Nat Neurosci* 22: 460–469.

Doiron B, Litwin-Kumar A (2014) Balanced neural architecture and the idling brain. *Front Comput Neurosci* 8: 56.

Ermentrout B (1998) Linearization of F-I curves by adaptation. *Neural Comput* 10: 1721–9.

Fourcaud-Trocmé N, Hansel H, van Vreeswijk C, Brunel N (2003) How spike generation mechanisms determine the neuronal response to fluctuating inputs. *J. Neurosci.* 23: 11628–11640.

Fusi S, Annunziato M, Badoni D, Salamon A, Amit DJ (2000) Spike-driven synaptic plasticity: theory, simulation, vlsi implementation. *Neural Comput* 12: 2227–58.

Fusi S, Miller EK, Rigotti M (2016) Why neurons mix: high dimensionality for higher cognition. *Curr Opin Neurobiol* 37: 66–74.

Gerstner W, Kistler WM, Naud R, Paninski L (2014) *Neuronal dynamics: from single neurons to networks and models of cognition.* Cambridge University Press, UK.

Giugliano M, La Camera G, Fusi S, Senn W (2008) The response function of cortical neurons: theory and experiment. II. Time-varying and spatially distributed inputs. *Biol. Cybern.* 99: 279–301.

Graupner M, Brunel N (2012) Calcium-based plasticity model explains sensitivity of synaptic changes to spike pattern, rate, and dendritic location. *Proc Natl Acad Sci U S A* 109: 3991–6.

Huang C, Doiron B (2017) Once upon a (slow) time in the land of recurrent neuronal networks.... *Curr Opin Neurobiol* 46: 31–38.

Inglebert Y, Aljadeff J, Brunel N, Debanne D (2020) Synaptic plasticity rules with physiological calcium levels. *Proc Natl Acad Sci U S A* 117: 33639–33648.

Intrator N, Cooper L (1992) Objective function formulation of the BCM theory of visual cortical plasticity: Statistical connections, stability conditions. *Neural Networks* 5: 3–17.

Jercog D, Roxin A, Barthó P, Luczak A, Compte A, de la Rocha J (2017) Up-down cortical dynamics reflect state transitions in a bistable network. *Elife* 6.

Jezzini A, Mazzucato L, La Camera G, Fontanini A (2013) Processing of hedonic and chemosensory features of taste in medial prefrontal and insular networks. *J Neurosci* 33: 18966–78.

Jones LM, Fontanini A, Sadacca BF, Miller P, Katz DB (2007) Natural stimuli evoke dynamic sequences of states in sensory cortical ensembles. *Proc Natl Acad Sci U S A* 104: 18772–7.

La Camera G, Rauch A, Senn W, Lüscher HR, Fusi S (2004) Minimal models of adapted neuronal response to *in vivo*-like input currents. *Neural Computation* 16: 2101–2124.

La Camera G, Fontanini A, Mazzucato L (2019) Cortical computations via metastable activity. *Curr Opin Neurobiol* 58: 37–45.

Lang L, La Camera G, Fontanini A (2023) Temporal progression along discrete coding states during decision-making in the mouse gustatory cortex. *PLoS Comput Biol* 19: e1010865.

Larkum ME, Zhu JJ, Sakmann B (1999) A new cellular mechanism for coupling inputs arriving at different cortical layers. *Nature* 398: 338–41.

Litwin-Kumar A, Doiron B (2012) Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nat Neurosci* 15: 1498–505.

Litwin-Kumar A, Doiron B (2014) Formation and maintenance of neuronal assemblies through synaptic plasticity. *Nat Commun* 5: 5319.

Maboudi K, Ackermann E, de Jong LW, Pfeiffer BE, Foster D, Diba K, Kemere C (2018) Uncovering temporal structure in hippocampal output patterns. *Elife* 7:e34467.

Maffei A (2011) The many forms and functions of long term plasticity at GABAergic synapses. *Neural Plasticity* 2011: 1–9.

Manz P, Memmesheimer RM (2022) Purely STDP-based assembly dynamics: stability, learning, overlaps, drift and aging. *bioRxiv* 10.1101/2022.06.20.496825.

Markram H, Lübke J, Frotscher M, Sakmann B (1997) Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275: 213–5.

Masset P, Qin S, Zavatone-Veth JA (2022) Drifting neuronal representations: Bug or feature? *Biol Cybern* 116: 253–266.

Mazzucato L, Fontanini A, La Camera G (2015) Dynamics of multistable states during ongoing and evoked cortical activity. *J Neurosci* 35: 8214–31.

Mazzucato L, La Camera G, Fontanini A (2019) Expectation-induced modulation of metastable activity underlies faster coding of sensory stimuli. *Nat Neurosci* 22: 787–796.

Mazzucato L, Fontanini A, La Camera G (2016) Stimuli reduce the dimensionality of cortical activity. *Front. Syst. Neurosci.* 10: 11.

Miller P, Katz DB (2010) Stochastic transitions between neural states in taste processing and decision-making. *J Neurosci* 30: 2559–70.

Miyashita Y (1988) Neural correlate of visual associative long-term memory in the primate temporal cortex. *Nature* 335: 817–820.

Miyashita Y, Chang H (1988) Neural correlate of pictorial short-term memory in the primate temporal cortex. *Nature* 331: 68–70.

Moreno-Bote R, Rinzel J, Rubin N (2007) Noise-induced alternations in an attractor network model of perceptual bistability. *J. Neurophysiol.* 98: 1125–1139.

Ngezahayo A, Schachner M, Artola A (2000) Synaptic activity modulates the induction of bidirectional synaptic changes in adult mouse hippocampus. *J Neurosci* 20: 2451–8.

Nuriya M, Jiang J, Nemet B, Eisenthal KB, Yuste R (2006) Imaging membrane potential in dendritic spines. *Proc Natl Acad Sci U S A* 103: 786–90.

Perin R, Berger TK, Markram H (2011) A synaptic organizing principle for cortical neuronal groups. *Proc Natl Acad Sci U S A* 108: 5419–24.

Ponce-Alvarez A, Nácher V, Luna R, Riehle A, Romo R (2012) Dynamics of cortical neuronal ensembles transit from decision making to storage for later report. *J. Neurosci.* 32: 11956–11969.

Recanatesi S, Pereira-Obilinovic U, Murakami M, Mainen Z, Mazzucato L (2022) Metastable attractors explain the variable timing of stable behavioral action sequences. *Neuron* 110: 139–153.e9.

Renart A, Moreno-Bote R, Wang XJ, Parga N (2007) Mean-driven and fluctuation-driven persistent activity in recurrent networks. *Neural Comput* 19: 1–46.

Rule ME, O'Leary T, Harvey CD (2019) Causes and consequences of representational drift. *Curr Opin Neurobiol* 58: 141–147.

Sah P (1996) Ca^{2+}-activated K^+ currents in neurons: types, physiological roles and modulation. *Trends Neurosci.* 19: 150–154.

Seidemann E, Meilijson I, Abeles M, Bergman H, Vaadia E (1996) Simultaneously recorded single units in the frontal cortex go through sequences of discrete and stable states in monkeys performing a delayed localization task. *J. Neurosci.* 16: 752–768.

Setareh H, Deger M, Petersen CCH, Gerstner W (2017) Cortical dynamics in presence of assemblies of densely connected weight-hub neurons. *Front Comput Neurosci* 11: 52.

Shepherd GM, Grillner S (2010) *Handbook of brain microcircuits.* Oxford University Press, New York.

Shouval HZ, Bear MF, Cooper LN (2002) A unified model of NMDA receptor-dependent bidirectional synaptic plasticity. *Proc Natl Acad Sci U S A* 99: 10831–6.

Sjöström PJ, Turrigiano GG, Nelson SB (2001) Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron* 32: 1149–64.

Sjöström PJ, Nelson SB (2002) Spike timing, calcium signals and synaptic plasticity. *Curr Opin Neurobiol* 12: 305–14.

Spruston N, Schiller Y, Stuart G, Sakmann B (1995) Activity-dependent action potential invasion and calcium influx into hippocampal ca1 dendrites. *Science* 268: 297–300.

Stuart GJ, Sakmann B (1994) Active propagation of somatic action potentials into neocortical pyramidal cell dendrites. *Nature* 367: 69–72.

Tetzlaff C, Kolodziejski C, Timme M, Wörgötter F (2011) Synaptic scaling in combination with many generic plasticity mechanisms stabilizes circuit connectivity. *Front Comput Neurosci* 5: 47.

Turrigiano GG (2017) The dialectic of Hebb and homeostasis. *Philosophical Transactions of the Royal Society B: Biological Sciences* 372: 20160258.

van Vreeswijk C, Sompolinsky H (1998) Chaotic balanced state in a model of cortical circuits. *Neural Comput* 10: 1321–71.

Vogels TP, Froemke RC, Doyon N, Gilson M, Haas JS, Liu R, Maffei A, Miller P, Wierenga CJ, Woodin MA, Zenke F, Sprekeler H (2013) Inhibitory synaptic plasticity: spike timing-dependence and putative network function. *Front Neural Circuits* 7: 119.

Vogels TP, Sprekeler H, Zenke F, Clopath C, Gerstner W (2011) Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science* 334: 1569–73.

Wang B, Aljadeff J (2022) Multiplicative shot-noise: A new route to stability of plastic networks. *Phys Rev Lett* 129: 068101.

Wang L, Maffei A (2014) Inhibitory plasticity dictates the sign of plasticity at excitatory synapses. *J Neurosci* 34: 1083–93.

Wang XJ (1998) Calcium coding and adaptive temporal computation in cortical pyramidal neurons. *J Neurophysiol* 79: 1549–66.

Waters J, Schaefer A, Sakmann B (2005) Backpropagating action potentials in neurones: measurement, mechanisms and potential functions. *Prog Biophys Mol Biol* 87: 145–70.

Watt AJ, Desai NS (2010) Homeostatic plasticity and STDP: Keeping a neuron's cool in a fluctuating world. *Front Synaptic Neurosci* 2: 5.

Zenke F, Agnes EJ, Gerstner W (2015) Diverse synaptic plasticity mechanisms orchestrated to form and retrieve memories in spiking neural networks. *Nat Commun* 6: 6922.

Zenke F, Gerstner W (2017) Hebbian plasticity requires compensatory processes on multiple timescales. *Philos Trans R Soc Lond B Biol Sci* 372.

Zenke F, Gerstner W, Ganguli S (2017) The temporal paradox of hebbian learning and homeostatic plasticity. *Curr Opin Neurobiol* 43: 166–176.

Zhang LI, Tao HW, Holt CE, Harris WA, Poo M (1998) A critical window for cooperation and competition among developing retinotectal synapses. *Nature* 395: 37–44.

Ziv Y, Burns LD, Cocker ED, Hamel EO, Ghosh KK, Kitch LJ, El Gamal A, Schnitzer MJ (2013) Long-term dynamics of CA1 hippocampal place codes. *Nat Neurosci* 16: 264–6.