Supplemental Information

Learning by the dendritic prediction of somatic spiking

Robert Urbanczik, Walter Senn

Department of Physiology, University of Bern, Bühlplatz 5, CH-3012 Bern, Switzerland

---

Contents

1 Variations on the model

While arguably an improvement over point neurons, our compartmental model still gives a highly stylized account of biological reality. One issue is that few neuronal morphologies are well described in terms of a somatic compartment and a single dendritic compartment. Hence in the next subsection we show how the model can be adapted when there is more than one dendritic compartment.

A second issue arises from the fact that synaptic input directly targeting the soma has often been found to be largely inhibitory. In a limited sense, the simulations in the main text are in line with this because in terms of conductance strength the inhibitory nudging is always considerably stronger than the excitatory nudging. This reflects the

fact that the gap between reversal potential and spiking threshold is much bigger for excitation than for inhibition. Hence for physiological values of the somatic potential, excitatory conductance is more powerful than inhibitory conductance.

Our model can however be adapted to the case where there is an even larger discrepancy between somatic excitation and inhibition. This is shown in Section 1.2 where the dendritic prediction is modified to take into account a baseline of somatic inhibition.

A further issue is our assumption that subthreshold current flows from the dendrite to the soma, but not in the reverse direction. Partial justification for this is provided by passive cable theory, since the theory posits that it is the ratio of the surface areas of the two compartments which determines the ratio between the coupling constant g_{D} for current flow from dendrite to soma and the coupling constant g_{S} for the reverse flow. For many neuronal morphologies the dendritic surface is much larger than that of the soma, suggesting that $g_{\mathrm{D}}/g_{\mathrm{S}}$ can be large. But large is not infinite, as was the case in main text, and hence we present simulation results for our model with a nonzero value of g_{S}. Finally we show how the plasticity rule can even be adapted to the case that g_{D} is not large.

1.1 Two dendritic compartments

We assume two dendritic compartments with local potentials $V_{\mathbf{w}}^{(1)}$ and $V_{\mathbf{w}}^{(2)}$. Each compartment integrates its presynaptic input just like the single compartment in the main text. To aggregate the two dendritic compartments in the soma, we replace Eq. (1) in the main text by

$$\dot{U} = -g_{\mathrm{L}}U + g_{\mathrm{D}}(V_{\mathbf{w}}^{(1)} - U) + g_{\mathrm{D}}(V_{\mathbf{w}}^{(2)} - U) + I_U^{\mathrm{som}}$$

Each compartment now makes its own prediction $V_{\mathbf{w}}^{(i)*}$ for the somatic potential and we set $V_{\mathbf{w}}^{(i)*} = \frac{2g_{\mathrm{D}}}{2g_{\mathrm{D}}+g_{\mathrm{L}}}V_{\mathbf{w}}^{(i)}$. The plasticity rule for a synapse in dendritic compartment i now uses $V_{\mathbf{w}}^{(i)*}$ as the modulation factor instead of $V_{\mathbf{w}}^{*}$ (Eq. 4 and 11, main text). Our choice for the dendritic predictions means that in addition to assuming that there is no nudging, each compartment assumes that its prediction of the somatic potential is the same as the prediction of the other compartment. While this assumption is incorrect before learning, we expect the voltage discrepancies between the compartments to decrease through learning since their overall influence on the soma is determined by the mean (or, more generally, a convex combination) of their local dendritic potentials.

Simulation results for the same task as in Fig. 1 of the main text are shown in Fig. S1. For the wiring, we assumed that the first 60% of the input neurons project to the first dendritic compartment and that the remaining input neurons project to the second dendritic compartment. Despite the asymmetry, the predictions of the two dendritic compartments converge during learning as shown by the $\mathrm{KL}(V_{\mathbf{w}}^{(1)*}, V_{\mathbf{w}}^{(2)*})$ curve in Panel S1c. Compared to the case of a single dendritic compartment, learning is slower since the dendritic predictions are less reliable. So we increased learning time

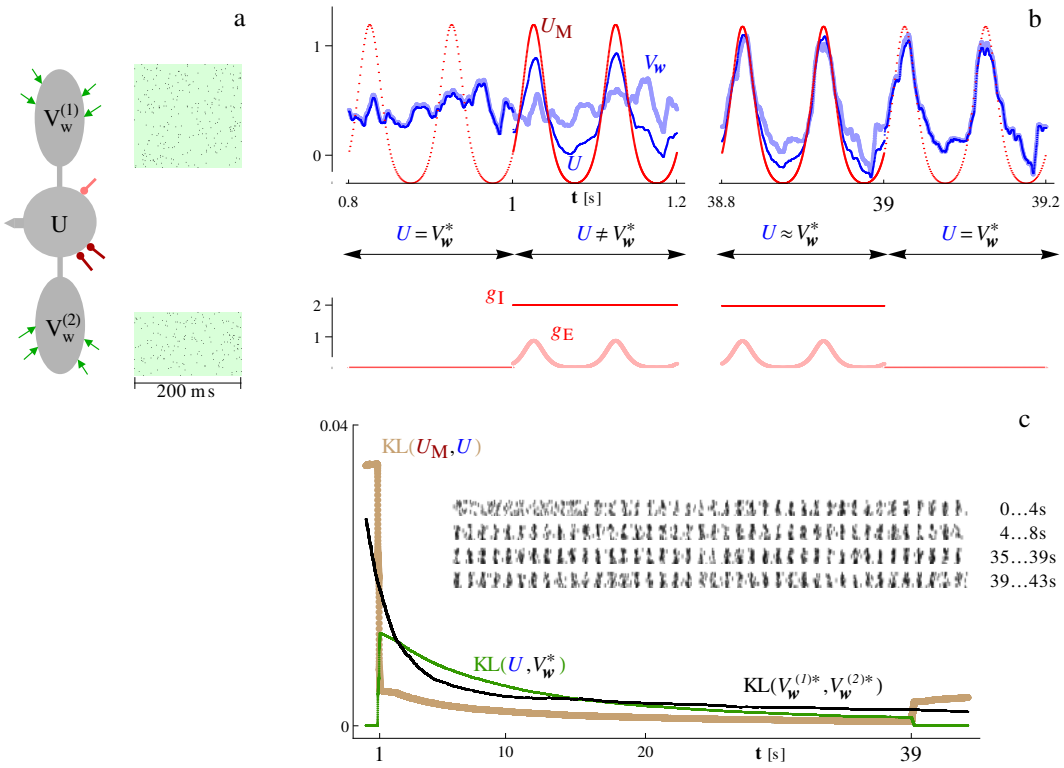

Fig. S1. Learning with two dendritic compartments. The learning scenario and the information shown is entirely analogous to Fig. 1 in the main text. In (**b**) we introduce as a purely descriptive quantity the mean $V_{\mathbf{w}}$ of $V_{\mathbf{w}}^{(1)}$ and $V_{\mathbf{w}}^{(2)}$, and we use $V_{\mathbf{w}}^{*} = \frac{2g_{\mathrm{D}}}{2g_{\mathrm{D}}+g_{\mathrm{L}}} V_{\mathbf{w}}$ in (**c**).

by a factor of 2. To account for the fact that there a two dendritic compartments, the coupling constant g_{D} was decreased from its value of $g_{\mathrm{D}} = 2$ in the main text to $g_{\mathrm{D}} = 1$. Since each of the two dendritic compartments now has roughly half as many inputs, we doubled the initial values for the strength of the dendritic synapses. All other parameters are the same as in the corresponding simulation in the main text.

1.2 Somatic baseline inhibition

Here we assume that the soma always receives inhibitory somatic input. To take this into account we consider the conductance g_{I} to be made up of two components $g_{\mathrm{I}}^{\mathrm{base}}$ and $g_{\mathrm{I}}^{\mathrm{nudge}}$, so that $g_{\mathrm{I}} = g_{\mathrm{I}}^{\mathrm{base}} + g_{\mathrm{I}}^{\mathrm{nudge}}$. A non-zero value of $g_{\mathrm{I}}^{\mathrm{base}}$ corresponds to a constant baseline of inhibition in the soma which should not drive plasticity in the dendrite. Plasticity, however, should arise when the somatic voltage is nudged due to non-zero values of $g_{\mathrm{I}}^{\mathrm{nudge}}$ and g_{E}.

From Eqs. (1,2) in the main text, we obtain that in the absence of nudging ($g_{\mathrm{I}}^{\mathrm{nudge}} =$

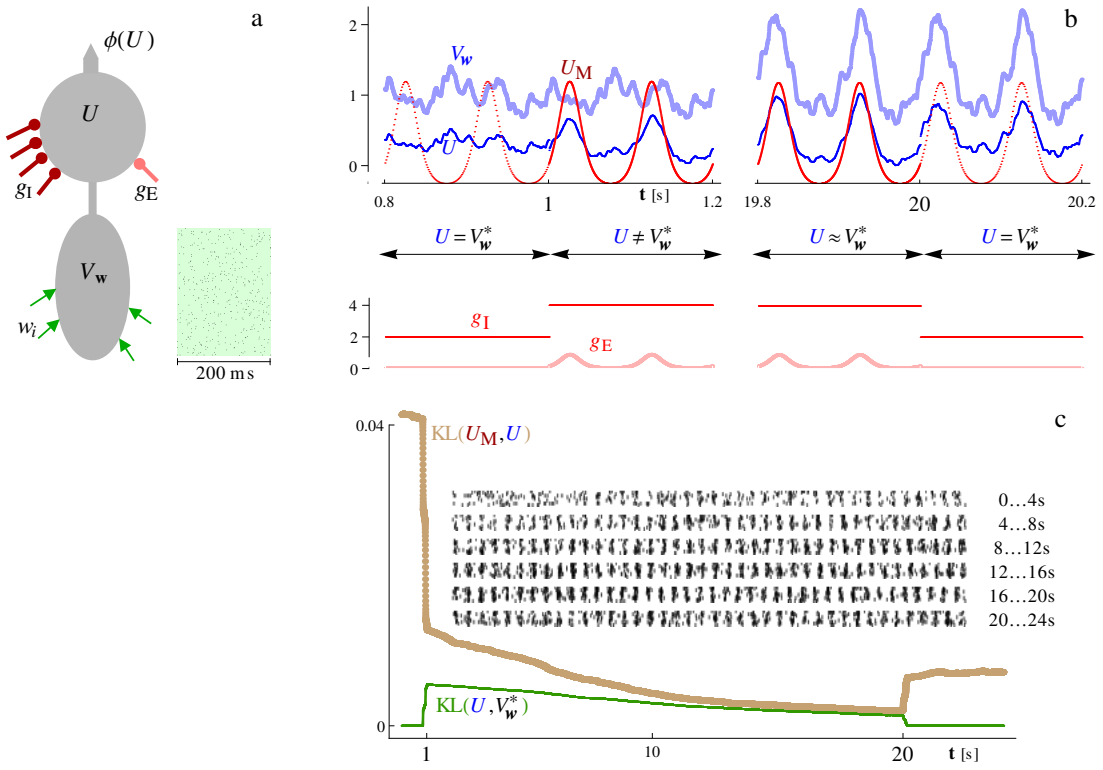

Fig. S2. Learning with a somatic baseline inhibition $g_I^{base} = 2$. The learning scenario and the information shown is entirely analogous to Fig. 1 in the main text. The calculation of the matching potential U_M is based on the expression in the main text (Eq. 3) with the following modification: Instead of using g_I for the inhibitory conductance, one has to use $g_I^{nudge} = g_I - g_I^{base}$.

$g_E = 0$) the somatic potential evolves as

$$\dot{U} = -g_L U + g_D(V_w - U) + g_I^{base}(E_I - U)\,.$$

It is now this equation which has to be used for computing the dendritic prediction V_w^* of the somatic potential. We again assume that g_D is much larger than the leak conductance g_L, but not necessarily larger than the baseline inhibition g_I^{base}. Then to an excellent approximation

$$V_w^* = \frac{g_D V_w + g_I^{base} E_I}{g_D + g_I^{base} + g_L}$$

and it is this expression which has to be used for V_w^* in the plasticity rule (Eq. 4 and 11, main text) instead of the original $V_w^* = \frac{g_D V_w}{g_D + g_I}$. The rescaling of the dendritic prediction is the only change in the model in order to accommodate a prevalence of somatic inhibition.

The simulation results in Fig. S2 highlight the strong depolarization in the dendrite now needed to counteract the somatic inhibition. Learning performance, however, is similar to the case without baseline inhibition. Due to the increased inhibition, we chose larger initial synaptic strengths in the dendrite, picking them from a Gaussian distribution with a mean and a standard deviation of 1/2. We also increased the learning rate to $\eta = 0.1$. All other parameters are the same as in the simulation for

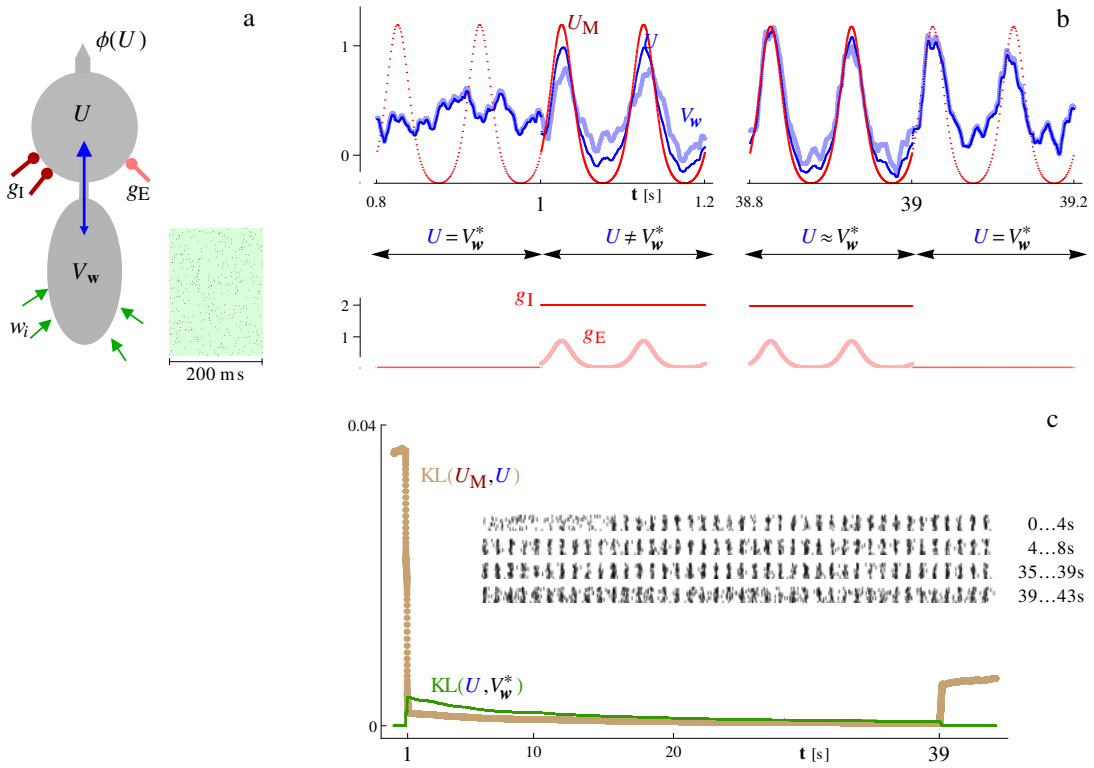

Fig. S3. Learning in the presence of subthreshold flow from soma to dendrite. The learning scenario and the information shown is entirely analogous to Fig. 1 in the main text.

Fig. 1 in the main text.

1.3 Subthreshold flow from soma to dendrite

To account for the passive propagation from the soma to the dendrite we change Eq. (7) in Methods to

$$\dot{V}_{\mathbf{w}} = -V_{\mathbf{w}}/\tau_{\mathrm{L}} + g_{\mathrm{S}}(U - V_{\mathbf{w}}) + I^{\mathrm{dnd}}/\tau_{\mathrm{L}}$$

with $g_{\mathrm{S}} = 0.2$. As the plot in the left column of Fig. S3b shows, nudging the soma now has a significant effect on the dendritic potential. The remaining results in Fig. S3 show that this slows down but does not disrupt learning. Aside of the above change to the equation for $V_{\mathbf{w}}$ and the increase in learning time, all simulation details are the same as for Fig. 1 in the main text.

1.4 Symmetrical coupling of soma and dendrite

Here we assume $g_{\mathrm{D}} = g_{\mathrm{S}} = 0.2$. The small value of g_{D} means that propagation from the dendrite to the soma results not only in increased attenuation of $V_{\mathbf{w}}$ but also in

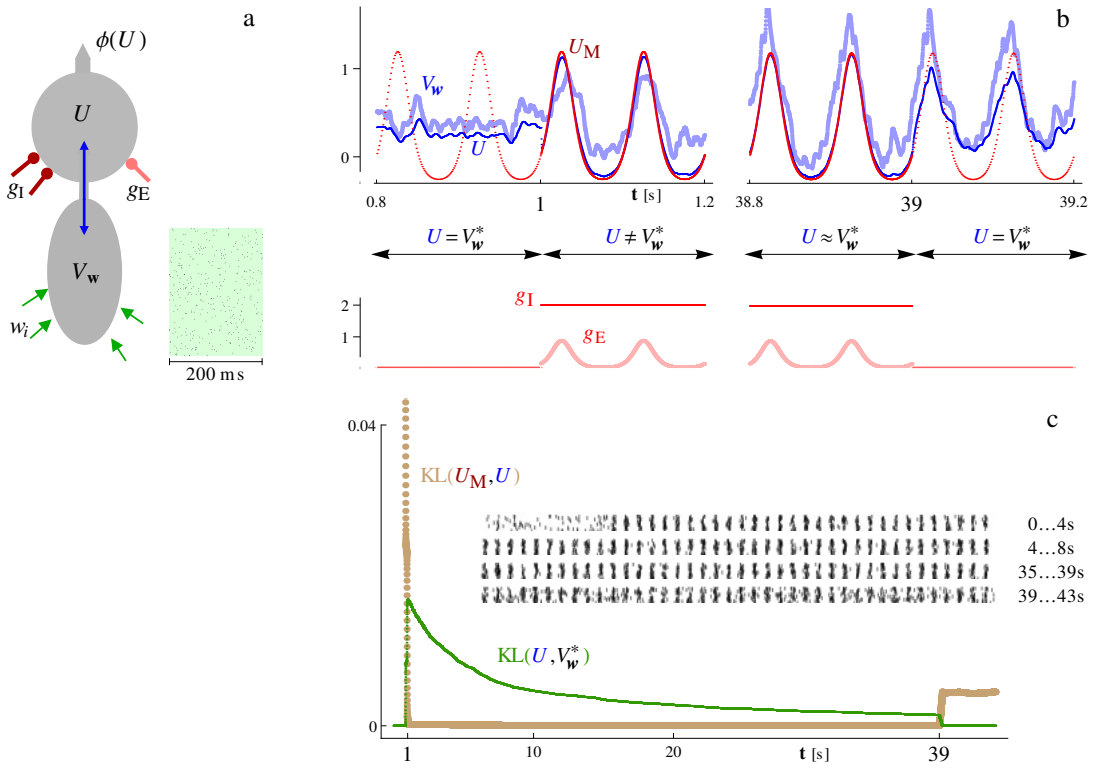

Fig. S4. Learning with a symmetrical coupling of soma and dendrite. The learning scenario and the information shown is entirely analogous to Fig. 1 in the main text.

a noticeable time lag. Hence in this section we no longer use the attenuation approximation $V_{\mathbf{w}}^* \approx \frac{g_D V_{\mathbf{w}}}{g_D + g_I}$ but obtain V_w^* by low pass filtering $V_{\mathbf{w}}$. So V_w^* is the solution of

$$\dot{U} = -g_L U + g_D(f - U) \tag{S1}$$

for $f(t) = V_{\mathbf{w}}(t)$. Similarly, we need to take the lag into account for the presynaptic term, so we let PSP_i^* be the solution of (S1) for $f(t) = \mathrm{PSP}_i(t)$. The modified plasticity rule then reads

$$\mathrm{PI}_i(t) = \Big(S(t) - \phi(V_{\mathbf{w}}^*(t))\Big)\, h\left(V_{\mathbf{w}}^*(t)\right) \mathrm{PSP}_i^*(t)\,.$$

The results with this rule in Fig. S4 show that learning is not disrupted by the symmetrical coupling, even if it is a bit slower than for Fig. 1 in the main text. The learning rate was $\eta = 0.15$ and a mean value of 0.3 was used for initializing the dendritic weights. All other simulation parameters where as in the preceding subsection.

2 Mathematical Analysis

In this section we return to the basic model considered in the main text, but assume for simplicity that there is no refractory period, i.e. somatic spiking is an inhomogeneous Poisson process with rate $\phi(U(t))$.

2.1 The plasticity rule as a gradient ascent procedure

We first combine the two equations defining the somatic potential (Eq. 1&2, main text) into the single equation

$$\dot{U} = -g_{\text{tot}}U + g_{\text{D}}V_{\mathbf{w}} + g_{\text{E}}E_{\text{E}} + g_{\text{I}}E_{\text{I}} \quad \text{with } g_{\text{tot}} = g_{\text{L}} + g_{\text{D}} + g_{\text{E}} + g_{\text{I}}\,. \tag{S2}$$

We now consider the limit that the coupling g_{D} of the soma to dendrite is strong. In particular we assume that g_{D} is much larger than g_{L}, but make no assumption about the magnitude of g_{D} relative to the nudging conductances g_{E} and g_{I}. Expanding the solution of (S2) in powers of $1/g_{\text{D}}$ we obtain

$$U = (1-\lambda)U_{\text{M}} + \lambda V^*_{\mathbf{w}} + \mathcal{O}(\tfrac{1}{g_{\text{D}}}) \quad \text{with } \lambda = \tfrac{g_{\text{D}}+g_{\text{L}}}{g_{\text{tot}}}\,. \tag{S3}$$

So to leading order the somatic potential is a convex combination of the matching potential (U_{M}, Eq. 3 main text) and the dendritic prediction $V^*_{\mathbf{w}} = \frac{g_{\text{D}}}{g_{\text{D}}+g_{\text{L}}}V_{\mathbf{w}}$. Note that the mixing factor λ will usually be time varying, when the nudging conductances and thus g_{tot} depend on time.

We now introduce the class of objective functions on which we want to do gradient ascent, namely

$$c_\lambda(u,v) = \int_0^v \mathrm{d}\nu\; h(\nu)\left(\phi\left((1-\lambda)u + \lambda\nu\right) - \phi(\nu)\right)\,. \tag{S4}$$

In the context of our model, λ is going to be the above mixing factor, ϕ the firing rate and h the weight function introduced in Eq. (4) of the main text. But for the moment just assume that $0 < \lambda < 1$, that ϕ is monotonically increasing and that h is positive. For the partial derivative w.r.t. to v of the cost function (S4) we have

$$\tfrac{\partial}{\partial v}c_\lambda(u,v) = h(v)\left(\phi\left((1-\lambda)u + \lambda v\right) - \phi(v)\right) \tag{S5}$$

and this derivative is positive for $v < u$ but negative for $v > u$. Hence $c_\lambda(u,v)$ as function of v has a single maximum at $v = u$ and is thus a suitable objective function for the purpose of achieving $v = u$ by gradient ascent in v.

In terms of $c_\lambda(u,v)$ the instantaneous cost function for our plasticity rule is

$$c(g_{\text{E}}, g_{\text{I}}; V^*_{\mathbf{w}}) = c_\lambda(U_{\text{M}}, V^*_{\mathbf{w}}) \tag{S6}$$

with λ given by Eq. (S3) and U_{M} given in the main text (Eq. 3). In view of (S5) the partial derivative with respect to the strength of the i-th synapse is

$$\frac{\partial}{\partial w_i} c(g_{\mathrm{E}}, g_{\mathrm{I}}; V^*_{\mathbf{w}}) = h(V^*_{\mathbf{w}}) \left(\phi\left((1-\lambda) U_{\mathrm{M}} + \lambda V^*_{\mathbf{w}} \right) - \phi(V^*_{\mathbf{w}}) \right) \frac{\partial}{\partial w_i} V^*_{\mathbf{w}} \,.$$

Using (S3) this can be rewritten as

$$\frac{\partial}{\partial w_i} c(g_{\mathrm{E}}, g_{\mathrm{I}}; V^*_{\mathbf{w}}) = \left(\phi\left(U\right) - \phi(V^*_{\mathbf{w}}) \right) h(V^*_{\mathbf{w}}) \frac{\partial}{\partial w_i} V_{\mathbf{w}} + \mathcal{O}(\tfrac{1}{g_{\mathrm{D}}}) \tag{S7}$$

Upto the $\frac{1}{g_{\mathrm{D}}}$ correction, the above right hand side is equal to the expectation value of our plasticity rule (Eq. 4, main text). Hence the proposed plasticity rule implements a stochastic gradient ascent procedure on the temporal average of (S6), driving the weight vector $\mathbf{w}$ towards a value satisfying $V^*_{\mathbf{w}} = U_{\mathrm{M}}$.

2.2 Mean field analysis of unsupervised learning

Considering a network of N compartmental neurons with somato-somatic connection matrices A (excitation) and B (inhibition), we derive a self-consistency relation for the vector $\mathbf{R}_{\mathrm{M}}$ of the neuronal firing rates after learning. Our starting point is the equation for the matching potential in a single neuron (3, main text), which for convenience we state in the vector form

$$\mathbf{U}_{\mathrm{M}}(t) = \frac{\mathbf{g}_{\mathrm{E}}(t) E_{\mathrm{E}} + \mathbf{g}_{\mathrm{I}}(t) E_{\mathrm{I}}}{\mathbf{g}_{\mathrm{E}}(t) + \mathbf{g}_{\mathrm{I}}(t)} \,. \tag{S8}$$

Here $\mathbf{U}_{\mathrm{M}}, \mathbf{g}_{\mathrm{E}}$ and $\mathbf{g}_{\mathrm{I}}$ denote the vectors of matching potential, excitatory nudging and inhibitory nudging for the N neurons. The division in (S8) is component-wise.

If N and/or the firing rates are large enough, we can neglect the temporal fluctuations due to the Poisson spiking. Then the rate vector $\mathbf{R}_{\mathrm{M}}$ determines all of the quantities in the above equation. In particular

$$\begin{aligned} \mathbf{R}_{\mathrm{M}} &= \phi(\mathbf{U}_{\mathrm{M}}) \\ \mathbf{g}_{\mathrm{E}} &= \alpha\, A\, \mathbf{R}_{\mathrm{M}} \\ \mathbf{g}_{\mathrm{I}} &= \alpha\, B \mathbf{R}_{\mathrm{M}} \end{aligned}$$

where the scaling factor α accounts for the synaptic release kernel. For simplicity, we have assumed the kernel to be the same for excitation and inhibition. Plugging this into Eq. (S8) yields the desired self-consistency relation:

$$\mathbf{R}_{\mathrm{M}} = \phi\left(\frac{(E_{\mathrm{E}}\, A + E_{\mathrm{I}}\, B)\mathbf{R}_{\mathrm{M}}}{(\, A + \, B)\mathbf{R}_{\mathrm{M}}} \right) . \tag{S9}$$

In general this equation will have many solutions. For instance, if A and B have topographic structure and a specific vector $\mathbf{R}_{\mathrm{M}}$ satisfies (S9), a topographically shifted

copy of the vector will also satisfy the equation. One important property of the learning dynamics can however be read of directly from (S9), because the equation is invariant under common rescalings of A and B, i.e under the transformation $A \to \gamma A$, $B \to \gamma B$. So the set of possible outcomes of learning is determined only by the balance of somatic excitation and inhibition but not by the absolute strength of the nudging.

3 Simulation Details

The differential equations were integrated using Euler's method with a time step of 0.2. Initial values of the weights of dendritic synapses were picked from a Gaussian distribution with mean μ and standard deviation 2μ, the value of μ being task dependent (see below).

3.1 KL divergence for firing rates

To evaluate learning performance (e.g. in Fig. 1c), we want to asses how close spike trains produced by a neuron with somatic potential $U(t)$ are to the spike trains which would be produced if the potential where $U_{\mathrm{M}}(t)$. Neglecting refractoriness, we assume that the potentials give rise to Poisson spike trains with rates $\phi(U(t))$ and, respectively, $\phi(U_{\mathrm{M}}(t))$. In a short time bin of duration δ the firing probabilities then are $q\delta$ and $p\delta$ for

$$q = \phi(U(t)) \quad \text{and} \quad p = \phi(U_{\mathrm{M}}(t))\,. \tag{S10}$$

The firings in the time bin are Bernoulli random variables and the textbook definition for their KL-divergence yields:

$$\mathrm{kl}(U_{\mathrm{M}}(t), U(t)) = p\delta \log \tfrac{p\delta}{q\delta} + (1 - p\delta) \log \tfrac{1-p\delta}{1-q\delta}\,.$$

Expanding for small δ this becomes

$$\mathrm{kl}(U_{\mathrm{M}}(t), U(t)) = \Big(p \log \tfrac{p}{q} + q - p\Big)\,\delta + \mathcal{O}(\delta^2) \tag{S11}$$

Assuming the stimulus is presented from $t = 0$ to $t = T$, averaging over the stimulus duration and using (S10,S11), we set:

$$\mathrm{KL}(U_{\mathrm{M}}, U) = \tfrac{1}{T} \int_0^T \mathrm{d}t\ \phi(U_{\mathrm{M}}(t)) \log \tfrac{\phi(U_{\mathrm{M}}(t))}{\phi(U(t))} + \phi(U(t)) - \phi(U_{\mathrm{M}}(t))\,. \tag{S12}$$

The analogous formula is used for $\mathrm{KL}(U, V^*_{\mathbf{w}})$, when comparing the actual somatic potential to the dendritic prediction $V^*_{\mathbf{w}}$.

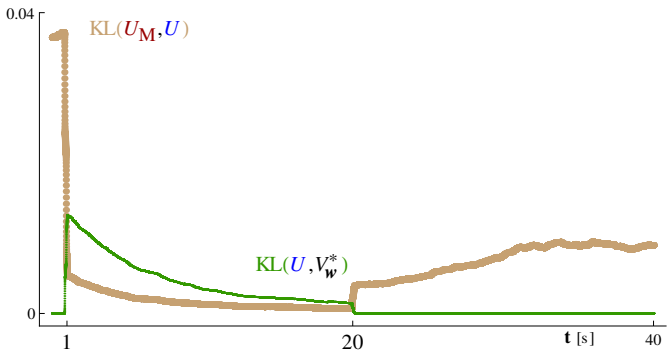

Fig. S5. Evaluation of stability in the absence of nudging. The plot is analogous to Fig. 1c of the main text. But the time frame for which the simulation was run after learning, with ongoing plasticity in the absence of any direct synaptic input to the soma, was extended from 4s in the main text to the 20s shown above.

3.2 Details for Figure 1, main text

The dendritic input pattern shown in Fig. 1a is made up of 200 homogeneous Poisson spike trains, each with a mean rate of 10 HZ. The mean value of a dendritic weight was $\mu = 0.2$ before learning, and a learning rate of $\eta = 0.07$ was used.

We also evaluated in more detail the effects of the synaptic diffusion occurring in the absence of nudging as a consequence of the stochastic nature of the plasticity rule (Fig. S5). While the diffusion is detrimental, performance is still much better than before learning even after 20s of plasticity in the absence of nudging

3.3 Details for Figure 2, main text

Learning time was divided in epochs with a duration of 500 ± 100 ms (Gaussian distribution). For each epoch one of the nudging patterns in Fig. 2a was picked at random and applied during the entire epoch. Total learning time elapsed between Fig. 2b and Fig. 2c was 500 s (simulated biological time). The mean value of a dendritic weight was $\mu = 0.1$ before learning, and a learning rate of $\eta = 0.01$ was used.

We quantitatively evaluated network performance using a protocol similar to the one shown in Fig. 2b and 2c. For each trial a randomly chosen pattern was nudged for 50 ms, followed by a 100 ms period without any nudging before the onset of the next trial. During the 100 ms without nudging we evaluated the persistence of the pattern which had last been nudged by the average value of $\mathrm{KL}(U_\mathrm{M}, U)$. The average is over the neurons making up the last pattern as well as over the 100 ms duration. For each neuron and time point the appropriate value of U_M for this pattern, obtained from the g_E values in Fig. 2a, was used in computing $\mathrm{KL}(U_\mathrm{M}, U)$. Before learning, the average

value of this discrepancy measure was 0.065 ± 0.001. During the 500 s of learning considered in the main text this measure decreased five-fold to 0.013 ± 0.002.

3.4 Details for Figure 3, main text

The inhibitory neurons in the network were point neurons with a soma modeled like the soma of the compartmental neurons (Eq. 1, main text; Eq. 10, Methods), but for the fact that $g_{\mathrm{D}} = 0$ in the case of the inhibitory neurons. Interaction between compartmental and inhibitory neurons was conductance based and modeled just like the interaction between the somata of compartmental neurons (Eq. 9, Methods).

For the topographic connectivity both the 40×40 compartmental neurons and the 20×20 were placed on square grids within the unit square. To compensate for the relatively small number of neurons we are able to simulate, periodic boundary conditions were used. For implementing these, we define the cyclic distance of two points p and q in the unit square as

$$d_{\mathrm{c}}(p, q) = \sqrt{|p_1 - q_1|_{\mathrm{c}}^2 + |p_2 - q_2|_{\mathrm{c}}^2}$$

with $|.|_{\mathrm{c}}$ defined as

$$|\delta|_{\mathrm{c}} = \min\{\, \mathrm{mod}(\delta, 1),\, 1 - \mathrm{mod}(\delta, 1) \,\}.$$

The definition of $|.|_{\mathrm{c}}$ reflects the fact that the maximal distance between two points on the unit line is 0.5, if one thinks of the endpoints of the line as being one and the same point.

Now, the wiring of the network is given by the following rules:

- If *compartmental* neuron a lies at distance d_{c} from a second *compartmental* neuron b, the probability of there being an excitatory connection from a targeting the soma of b is $p = e^{-10 d_{\mathrm{c}}}$. For such a connection, the excitatory synapse conveys conductance with strength $w^{\mathrm{E}} = 0.06$.
- If *compartmental* neuron a lies at distance d_{c} from *inhibitory* neuron i, the probability of there being an excitatory connection from a targeting i is $p = e^{-8 d_{\mathrm{c}}}$. These connections are facilitating, see Eqs. (S13,S14) below, with a baseline conductance of $w^{\mathrm{E}}_{\mathrm{base}} = 0.008$
- If *inhibitory* neuron i lies at distance d_{c} from *compartmental* neuron b, the probability of there being a inhibitory connection from i targeting the soma of b is $p = \frac{3}{3+(9 d_{\mathrm{c}})^2}$. For such a connection, the inhibitory synapse conveys conductance with strength $w^{\mathrm{I}} = 0.3$.
- All of the 100 input neurons project onto the dendrites of all of the compartmental neurons.

The above wiring implements long range inhibition because the distribution of connections from inhibitory neurons to compartmental neurons is scale free. Nevertheless, this distribution peaks at short distances, so there is also substantial local inhibition, competing with the short range excitation. If there is local excitation only, all input

patterns tend to get mapped to one and the same area by learning. The reason is that once the lengths of the dendritic weight vectors of the neurons in one area of the map start to increase, this becomes self reinforcing until all input patterns are mapped to this area. The standard remedy for this is, to use an update rule which does not change the length of a neurons weight vector (Ref. 15, main text). We did not adopt this solution, since we are not aware of a mechanism which might coordinate updates in all the synapses of a neuron so that total synaptic strength remains unchanged. Instead, we found that introducing a competition between short range excitation and inhibition, provides a remedy. In the conductance based formulation, the effective strength of inhibition increases with increasing somatic potential. This dampens the self reinforcing growth of the weight vectors in an area of the map giving other areas a chance to catch up.

Robustness of map formation can be further improved by shifting the balance of excitatory and inhibitory nudging towards inhibition when there is prolonged high activity in one area of the map. The reason is that the self reinforcing growth of weight vectors decreases the stimulus selectivity of a neuron's response, and neurons which fire in response to many of the stimuli will show prolonged times of high activity. Shifting the balance can be achieved, for instance, by short term depression in the excitatory somato-somatic connection. Other mechanisms could be short term facilitation in synapses targeting the inhibitory neurons, or even a small after-depolarizing current in these neurons.

For the simulations, we implemented facilation in the afferents of the inhibitory neurons. So the actual conductance w^{E} of a synapse targeting such a neuron can be higher than its baseline value $w^{\mathrm{E}}_{\mathrm{base}}$ given above. It is obtained as

$$w^{\mathrm{E}} = w^{\mathrm{E}}_{\mathrm{base}}(1+f) \tag{S13}$$

where the facilitation variable f increases with the presynaptic firing rate. In particular,

$$\tau_{\mathrm{f}}\dot{f} = -f + F\sum_{s\in X^{\mathrm{E}}}\delta(t-s) \tag{S14}$$

Here τ_{f} is the time constant for facilitation (we used 200 ms), F determines the facilitation strength (we used $F = 40$), and X^{E} denotes the presynaptic spike train of the synapse. The parameter choices mean that a sustained presynaptic rate of 0.1 kHz leads to a value of f which is approximately 4, i.e to a roughly five-fold increase in synaptic strength.

Learning time was divided into epochs of 100 ms. For each epoch one of the six input patterns in Fig. 3b was chosen at random, and during the epoch the input neurons emitted independent Poisson spike trains with mean frequencies given by the chosen pattern. Total learning time elapsed between the top and bottom row of Fig. 3c was 1600 s (simulated biological time). The mean value of a dendritic weight was $\mu = 0.3$ before learning, and a learning rate of $\eta = 0.015$ was used.