

A homeostatic rule for inhibitory synapses promotes temporal sharpening and cortical reorganization

Samat B. Moldakarimov^{*†‡}, James L. McClelland^{†§¶}, and G. Bard Ermentrout^{*}

^{*}Department of Mathematics, University of Pittsburgh, Pittsburgh, PA 15260; [†]Center for the Neural Basis of Cognition, Mellon Institute, 4400 Fifth Avenue, Pittsburgh, PA 15213; and [§]Department of Psychology, Carnegie Mellon University, Pittsburgh, PA 15213

Contributed by James L. McClelland, August 31, 2006

Experience with transient stimuli leads to stronger neural responses that also rise and fall more sharply in time. This sharpening enhances the processing of transients and may be especially relevant for speech perception. We consider a learning rule for inhibitory connections that promotes this sharpening effect by adjusting these connections to maintain a target homeostatic level of activity in excitatory neurons. We analyze this rule in a recurrent network model of excitatory and inhibitory units. Strengthening inhibitory→excitatory connections along with excitatory→excitatory connections is required to obtain a sharpening effect. Using the homeostatic rule, we show that repeated presentations of a transient signal will “teach” the network to respond to the signal with both higher amplitude and shorter duration. The model also captures reorganization of receptive fields in the sensory hand area after amputation or peripheral nerve resection.

experience | interneurons | learning | computational neuroscience

Perhaps the most important proposition in the neuroscience of learning is Hebb's (1949) (1) famous postulate: when neuron A participates in firing neuron B, the strength of the connection from A to B is increased. An important consequence of Hebbian learning is that it tends to increase both the strength and the duration of the neural response elicited by a sensory stimulus. Consider a recurrent network of mutually excitatory neurons. Suppose there are initial connections of moderate strength among neurons and their near neighbors. When a transient signal is applied to a neuron, it and its neighbors will become activated, and, according to Hebb's rule, the excitatory connections between them will be strengthened. The strengthening will cause the network to respond with stronger, as well as longer, activation.

Increasing the duration of neural responses may have important benefits in many cases, but there are situations in which it may not be desirable, including cases in which it is necessary to follow a series of brief transient signals. Consider the situation confronting someone listening to fluent speech. The signal changes very rapidly, and important information for identification of many sounds is present in transients with durations <40 ms (2). If perceptual responses to speech become stronger and more robust with experience, this strengthening would reduce the network's ability to follow the sequence of changing temporal signals if the response became more temporally extended at the same time that it was strengthened.

A remedy to this problem would be to arrange neural circuitry so that the neural response not only becomes stronger, but it also becomes “sharper” (more temporally coherent or compact) with experience (3). This sharpening could have several beneficial effects. With experience, it would increase the overall temporal resolution of the system. It would produce more accurate discrimination of fine differences in timing of neural signals, which will allow the system to follow a series of signals arriving at a high rate. These benefits would co-occur with benefits at the neural level as well; for a given amount of total activation, higher temporal coherence should result in stronger postsynaptic depolarization of downstream receiving neurons, increasing their ability to excite such neurons sufficiently to induce downstream synaptic strengthening (4).

Just this sort of temporal sharpening in response to experience has been observed in an important series of experiments by Recanzone *et al.* (5–8). In these investigations, monkeys worked for reward to discriminate different frequencies of vibration applied to the skin surface of one finger. The difference between a 20-Hz standard and a higher frequency target was reduced as the monkeys improved in discriminating targets from standard stimuli in the task. Successfully trained monkeys showed an increase in the size of the population of neurons (and corresponding cortical surface) responding to each pulse of the stimulus. They also showed an increase in temporal coherence or sharpness of the response (see Fig. 1). Each monkey's behavioral performance was well predicted by the degree of coherence of its temporal response.

Inspired by these findings and considerations, we have asked the following question: how might a neural system come to sharpen, rather than broaden, its response as the strengths of excitatory connections are increased? Based on analysis and simulations, we propose a remarkably simple answer. We suggest that the brain may achieve sharpening by adaptively modifying the strength of recurrent inhibition so as to conserve average neural activity. We also suggest that the same mechanism, augmented slightly, can address key findings in the literature on cortical reorganization after amputation (9–12).

Homeostasis, the idea that the brain conserves or regulates its average activity, has been proposed and supported previously (13), and recurrent inhibition has been explored as one of the factors that may play a role (14). We add the suggestion that, within the right architecture, the homeostatic adjustment of recurrent inhibition can promote temporal sharpening, along with regulation of total neural activity.

We present our argument in the following series of steps. First, we review some relevant physiology of the mammalian somatosensory systems thought to be the sites of the plastic changes under consideration. We then consider the effect of increasing the recurrent excitation in a simplified preliminary model of this system. We show that, to achieve temporal sharpening when the strength of the excitation is increased, the strength of recurrent inhibition must also be increased. We suggest a simple homeostatic mechanism for regulating the inhibition based on conserving average neural activity in the population of excitatory units “seen” by the inhibitory neuron, and we review data from systems-level experiments consistent with this homeostatic approach. We implement this mechanism in a model addressing the strengthening and sharpening of

Author contributions: S.B.M., J.L.M., and G.B.E. designed research; S.B.M. performed research; G.B.E. contributed new reagents/analytic tools; S.B.M. analyzed data; S.B.M., J.L.M., and G.B.E. wrote the paper; and G.B.E. and S.B.M. developed and wrote up the mathematical analysis in supporting information.

The authors declare no conflict of interest.

Abbreviation: Ln, layer n.

[‡]Present address: Computational Neurobiology Laboratory, The Salk Institute for Biological Studies, 10010 North Torrey Pines Road, La Jolla, CA 92037.

[¶]To whom correspondence should be sent at the present address: Department of Psychology, Jordan Hall, Stanford University, Stanford, CA 94305. E-mail: jlm@psych.stanford.edu.

© 2006 by The National Academy of Sciences of the USA

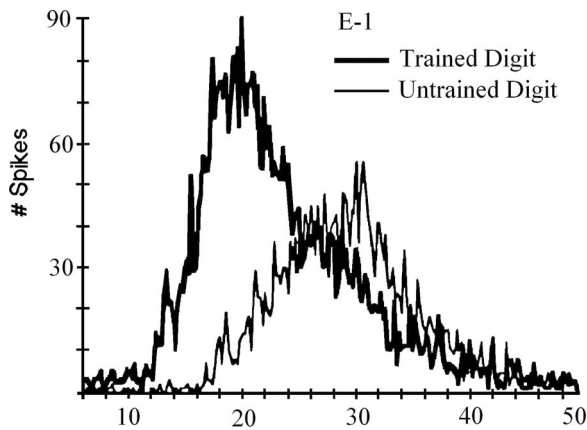


Fig. 1. Temporal sharpening effect. [Reproduced with permission from ref. 8 (Copyright 1992, The American Physiological Society).]

neural responses as revealed in the Recanzone *et al.* (5–8) experiments. Finally, we use the same model system to model cortical reorganization after peripheral alteration, adding a mechanism inspired by a proposal of Bienenstock *et al.* (15) for the regulation of homeostasis itself that may promote reorganization.

Sites of Cortical Plasticity

Sensory maps in cortex are regulated by experience, both during development and in adulthood. Whereas early studies focused on thalamocortical synapses as the primary site of plasticity, recent studies have identified intracortical excitatory synapses from layer 4 (L4) to L2/3 neurons and excitatory horizontal synapses across columns within L2/3 as sites of plasticity as well (9). It has been shown (16–22) that, in neonatal animals (<4 days old), thalamocortical synapses are likely to be a principal site of plasticity. In older animals, receptive field plasticity is most rapid in L2/3 and occurs only later, if at all, in L4. Based on these experimental findings, we examine a model in which the main sites of plasticity are recurrent horizontal connections among excitatory and inhibitory cells treated as a simplified model of somatosensory cortex layer L2/3.

Simplified Preliminary Model for Analysis

Excitatory neurons in layer L2/3 excite through direct excitatory projections and inhibit via inhibitory interneurons. Here, we consider a simplified model that relies on one excitatory and one inhibitory unit (Fig. 2A) to model the mean activities of these excitatory and inhibitory populations. To model recurrent excitation, we include a self-excitatory connection from the excitatory unit back to itself. Inhibition is modeled as an excitatory connection from the excitatory unit to the inhibitory unit, plus an inhibitory connection from the inhibitory unit back to the excitatory unit. Dynamics of unit activities are given by

$$\tau_e \frac{dE}{dt} = -E + F_e(I_{\text{exte}} + \alpha E - \beta I)$$

$$\tau_i \frac{dI}{dt} = -I + F_i(I_{\text{exti}} + \gamma E).$$

E is the firing rate of the excitatory unit and I is the firing rate of the inhibitory unit. I_{exte} and I_{exti} are external inputs to the excitatory and inhibitory units, F_e and F_i are gain functions of the excitatory and inhibitory units, τ_e and τ_i are constants governing the response rate of each type of unit, and α , β , γ are parameters.

Somatosensory cortex produces a transient response to the onset of a stimulus with little or no sustained component (23, 24). It has been suggested (23, 24) that this effect is achieved in part by making

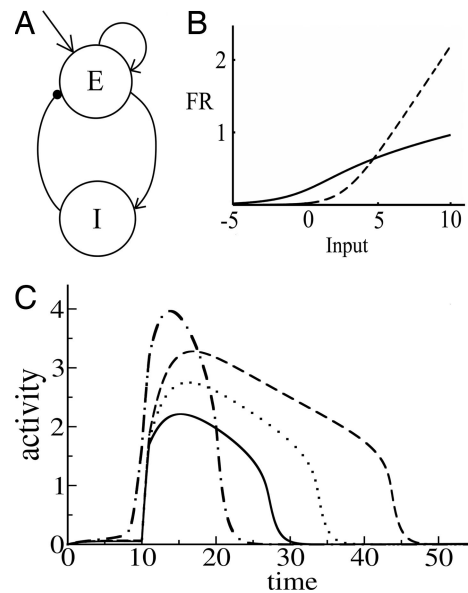


Fig. 2. Simplified network: structure, gain functions, and dynamics. (A) The two-unit network. E and I are excitatory and inhibitory units. (B) Gain functions. The solid line represents the excitatory gain function and the dashed line is the inhibitory gain function; $\epsilon_e = 0.2$; $\epsilon_i = 0.2$; $th_e = 1$; $th_i = 3$. (C) The activity level of the E cell over time as a response to a transient stimulus. Solid, dotted, and dashed lines represent responses with increasing self-excitatory strength ($\alpha = 3, 3.5$, and 4) and fixed inhibitory-to-excitatory strength ($\beta = 1.2$). Dot-dashed line represents response with both excitatory and inhibitory connections increased relative to the solid curve baseline ($\alpha = 5$; $\beta = 3$). Other parameters were fixed with these values: $\tau_e = 1$; $\tau_i = 20$; $\epsilon_e = 0.2$; $\epsilon_i = 0.2$; $th_e = 1$; $th_i = 2$; $l_{\text{exte}} = 5$; $\gamma = 1$.

the inhibitory unit respond more slowly than the excitatory unit. Accordingly, we set τ_i much larger than τ_e . To ensure that the response even to a sustained stimulus is transient, it has further been suggested that the inhibitory gain function must be much steeper than the excitatory gain function. Indeed, in the cortex, pyramidal neurons and interneurons show different $f-I$ curves (25). Regular-spiking (RS) excitatory pyramidal neurons show a continuous (type 1) $f-I$ curve with negative acceleration at high input levels. Fast-spiking inhibitory interneurons have a discontinuous (type 2) $f-I$ relationship, such that they do not fire below a critical input level but fire at much higher frequency than RS neurons at high input levels. We summarize these facts with the following gain functions (Fig. 2B) for the E and I units

$$F_e(x) = \sqrt{\epsilon_e \cdot \ln(1 + \exp((x - th_e)/\epsilon_e))} \quad [1]$$

$$F_i(x) = \epsilon_i \cdot \ln(1 + \exp((x - th_i)/\epsilon_i)). \quad [2]$$

Here th_e and th_i are thresholds and ϵ_e and ϵ_i are parameters describing the smoothness of the gain functions.

Given these features of the model, if we apply a strong stimulus to both units simultaneously, the excitatory unit will respond more quickly, but the activity of the inhibitory unit will eventually be much higher than the activity of the excitatory unit and so will come to suppress the excitatory unit, producing a transient response both to transient and sustained stimuli. Within this regime, we can now investigate how changes in connections may give rise to temporal sharpening of the transient response along with experience-dependent strengthening.

Model Analysis. We implemented and analyzed the model using the XPPAUT package (26). To understand the effects of the strength of the excitatory and inhibitory connections, we observed the behavior of the network to changes in the strength of the connec-




Fig. 3. Response duration in the two-unit model as a function of the strength of the excitatory connection (parameter α) and the strength of the inhibitory connection (parameter β); $\tau_e = 1$, $\tau_i = 10$, $\gamma = 2.5$.

tions (Fig. 2C). First, we recorded the network's response to a transient stimulus with baseline choices of the strengths for the $E-E$ and $E-I$ connections (Fig. 2C, solid line). As we increase the $E-E$ connection strength (α in the equations), the amplitude of activity of the E unit increases in response to a transient external stimulus (Fig. 2C, dotted line). However, as Fig. 2C shows, the response duration also increases. Further increases just amplify the effect (Fig. 2C, dashed line). Thus, strengthening of the self-excitatory loop alone results in a stronger but also longer lasting response. If we increase the inhibitory connection (parameter β) sufficiently along with the change in the excitatory connection, we can observe a response with higher amplitude and at the same time with shorter duration compared with baseline (Fig. 2C, dot-dash line). In general, our simulations indicate that, even with a positively accelerating inhibitory gain function and a negatively accelerating excitatory gain function, sharpening does not occur when only the excitatory connections are allowed to change. Sharpening along with strengthening can occur if the strength of the inhibitory loop is increased along with strengthening of the excitatory connection.

It was surprising at first that sharpening was not an automatic consequence of strengthening in this circuit, given the positive acceleration of the inhibitory gain function and the negative acceleration of the excitatory gain function. In supporting information, which is published on the PNAS web site, we investigate this issue and explore how changing parameters α and β can influence the amplitude and duration of the neural response. With the chosen gain functions and inhibition slower than excitation, the analysis shows that sharpening cannot occur with a change in the $E-E$ connection alone. The result can be proven for the case of inhibition much slower than excitation but also holds when τ_e and τ_i are not very different and holds for other choices of the gain functions as long as the excitatory gain function is negatively accelerating. Neurally realistic excitatory gain functions have this property, which is itself a basic requirement for stability. We have also shown that parameter α controls the amplitude of the response and its duration (increasing α increases both the amplitude and the duration of the response) and parameter β controls mainly the duration (increasing β shortens the duration). Finally (Fig. 3), we have shown that, for any increase in α , it is always possible to find a value of β such that the resulting response time will be shorter than the original response time before the change in α .

All the adjustments we have discussed so far have been performed manually. Now we ask: is there a simple policy whereby a network could adjust its inhibitory connections to produce temporal sharpening when the strength of the response becomes stronger? The key observation is the following: if the network adjusts its inhibitory connections to ensure that the time-averaged activity of the excitatory unit remains constant, then a higher-amplitude




Fig. 4. Somatosensory learning model: structure and dynamics. (Upper) Network of excitatory (upper row) and inhibitory (lower row) units used to model learning in somatosensory cortex layers 2/3. After training, several units are involved in the response to a transient stimulus applied to the trained unit, as indicated approximately by shading. (Lower) Activity levels of the trained unit as a function of time. Different curves represent activity change during training. Parameters were as follows: $\tau_e = 3$; $\tau_i = 5$; $\varepsilon_e = 0.2$; $\varepsilon_i = 0.15$; $th_e = 0.1$; $th_i = 0.2$; $l_{exte} = 5$; $\tau_{syne} = 300$; $\tau_{syni} = 150$; $g_{ee} = 5$; $g_{ei} = 3.5$; $g_{ie} = 1$; $w_{ij}^e = 1$. The standard deviation of the Gaussian weight function $r(i, j)$ is 3 for both excitatory-to-excitatory and excitatory-to-inhibitory connections.

response will necessarily be briefer than the original low-amplitude response. We examine the implications of this observation through the use of a simple homeostatic learning rule for inhibitory connections in the somatosensory learning model described below.

Somatosensory Learning Model. The model described in this section will be used to show how the homeostatic learning rule suggested above can produce sharpening along with strengthening as observed in the experiments of Recanzone *et al.* (5–8). It also addresses the expansion of the cortical area activated by the stimulus as a result of successful training in the Recanzone experiments (5–8). Finally, it addresses the reorganization of somatosensory receptive fields after amputation, a pattern for which others (9–12) have proposed a role for adaptive modification of inhibition as well as excitation.

Our model is based on a recurrent network of excitatory and inhibitory units intended as a simplified representation of cortical neurons in layer L2/3 of mammalian somatosensory cortex (Fig. 4 Upper). For each excitatory unit there is a corresponding inhibitory unit. There are connections from excitatory to excitatory units, from excitatory to inhibitory units, and from inhibitory to excitatory units. Excitatory connections to both excitatory and inhibitory units are scaled by a Gaussian function of distance. Each inhibitory unit projects only to the single corresponding excitatory unit. Both excitatory and inhibitory connections onto excitatory units are plastic. Connections to inhibitory units from excitatory units are fixed.

As in the preliminary model, excitatory and inhibitory units are classical Wilson–Cowan rate units (27) with dynamics of activities described by equations

$$\tau_e \frac{du_i}{dt} = -u_i + F_e(I_i^{exte} + EE_i - EI_i)$$

