Neural Computation 12.3 (2000): 519 – 529

DOI: http://dx.doi.org/10.1162/089976600300015682

Local and global gating of synaptic plasticity

Manuel A. Sánchez-Montañés[a,b], Paul F.M.J. Verschure[a] and Peter König[a]

[a] Institute of Neuroinformatics, ETH / University Zürich,

Winterthurerstr. 190, 8057 Zürich, Switzerland

[b] Grupo de Neurocomputación Biológica, ETS de Informática.

Universidad Autónoma de Madrid, 28049 Madrid, Spain

Introduction

Research on neuronal networks has been very much motivated by the ability of these systems to learn from experience (Alkon et al., 1991). In these systems information is stored in the pattern of synaptic efficacies. Among the processes influencing the modification of synaptic connections local and global mechanisms can be distinguished (Montague et al., 1996).

Examples of global signals involved in learning can be found in gating mechanisms (Abbott, 1990). In the brain this might be compared to the influence of modulatory signals, arising from subcortical structures, on cortical plasticity. Here, the cholinergic system of the basal forebrain projecting to the cerebral cortex is of particular interest: It is a necessary ingredient for the induction of cortical representations following monocular deprivation (Singer and Rauschecker, 1982). In addition, it may switch between storage and recall modes in the hippocampus (Hasselmo, 1993), and it gates the plasticity of receptive fields of neurons in the primary auditory cortex during classical conditioning (Weinberger, 1993; Bakin and Weinberger, 1996; Kilgard and Merzenich, 1998). These results support the suggestion that modulatory substances can act as a “print now” signal gating synaptic plasticity (Singer et al., 1979). As the influence of such a signal onto the cortical network is homogeneous, and is not related to the specifics of a given stimulus, it can be called a global mechanism.

Local mechanisms in learning and memory refer back to the classical work of D.O. Hebb (1949). Since then, many variations and physiological implementations have been suggested (Stent, 1973; Sejnowski, 1977; Buonomano and Merzenich, 1998). Recently Stuart and Sakmann (1994) observed that action potentials can propagate backwards into the dendritic tree . This signal provides a possible explanation for the dependence of LTP and LTD on the relative timing of pre- and postsynaptic action potentials (Gerstner et al., 1993; Markram et al., 1997). The changes in synaptic efficacy are specific to the neurons involved and may be described by rules local in space and time.

In spite of the enormous amount of work on local and global mechanisms of synaptic plasticity, their relationship has been little investigated and, implicitly, their action is assumed to be independent of each other. Here, we investigate an alternative model, where the global and local mechanisms are intimately related: The global mechanism influences the fraction of action potentials that successfully invade the dendritic tree of the stimulated neurons, a signal exploited by the local process. We demonstrate that such an integrated mechanism allows to combine positive aspects of its constituents, such as stability, specificity and flexibility, in one learning rule.

Methods

We studied local and global mechanisms of synaptic plasticity in a neuronal network consisting of excitatory and inhibitory cortical neurons, subcortical sensory input, and input from the basal forebrain. The sensory afferents target the excitatory neurons and define their receptive fields. Initially, the synaptic efficacy of these connections is randomly distributed, and updated according to the learning rule described below. The excitatory cortical neurons project to inhibitory interneurons, which in turn project back to the excitatory neurons, forming a negative feedback loop. The simulated basal forebrain afferents target the inhibitory cortical neurons and have an inhibitory effect (Freund and Gulyas, 1991; Freund and Meskenaite, 1992). All neurons are simulated as integrate-and-fire units (for quantitative details please see appendix).

The synaptic efficacy of the afferent sensory projections to the excitatory neurons evolves according to a modification of a recently proposed learning rule which utilizes a backpropagating action potential (Körding and König, 1998): First, when a “backpropagating” action potential arrives at a synapse simultaneously (i.e. within a small symmetrical temporal window) with an action potential in the presynaptic afferent fiber the efficacy of the respective synapse is increased (Gerstner et al., 1993; Markram et al., 1997; Magee et al., 1998; Bi and Poo, 1998). Furthermore, we studied the effects of using an asymmetrical temporal window, thus including a dependency of the synaptic modifications not only on the absolute delay of pre- and postsynaptic signals, but also on the temporal order of their arrival, which matches the physiological results more faithfully. Second, activation of inhibitory synapses located at the proximal dendritic tree may attenuate the retrograde propagation of the action potential in the dendritic tree (Spruston et al., 1995; Tsubokawa and Ross, 1996). In this case the efficacies of activated synapses are decreased. In addition, we implemented heterosynaptic LTD in the present simulation: synaptic efficacy is decreased in case of postsynaptic activity without coincident presynaptic activity. Thus, in this learning

rule, the changes of synaptic efficacy are crucially dependent on the temporal dynamics of the neuronal network, and in particular on the relative timing of the excitatory and inhibitory inputs to the cortical neurons.

We used two different kinds of stimulation protocols to investigate the interaction of global and local mechanisms of synaptic plasticity in analogy to experiments on auditory conditioning (Kilgard and Merzenich, 1998). In the first protocol, a pseudo-random sequence of 500 exemplars, repeatedly selected from a pool of 10 different stimuli, was presented alone. In the second protocol, one of the stimuli was paired with activity of the basal forebrain. For both protocols we analyzed a total of 40 simulations, each starting with random initial connections between sensory and excitatory populations. To control for dependencies on the details of the implementation we conducted control runs using neurons without a refractory period.

Results

As a first step, we investigated the relationship between the formation of representations of input stimuli and the frequency of their occurrence. Two of the stimuli were presented 4 times more often than the other eight (probability of occurrence of stimuli # 1 and # 2: 0.2500 vs. stimuli # 3 to # 10: 0.0675). After about 200 stimulus presentations the network had converged and stabilized. The resulting specificity of the excitatory cortical units showed a bimodal distribution. Either neurons received comparable input upon presentation of any stimulus, and thus were completely untuned (44.8%), or they responded highly specifically to only one of the 10 stimuli (50.5%). Few neurons showed intermediate degrees of specificity (4.7%). This is in good accord with our previous results (Körding and König, 1998). Next, we investigated the distribution of stimuli specificity for the whole population of neurons. Figure 1A shows that the number of neurons responding to the stimuli presented often (stimuli # 1 and 2, red and orange respectively) was similar to the size of the representation of stimuli presented rarely (stimuli # 3 to 10, yellow to violet). Compiling results from a total of 40 runs, we did not find a significant influence of presentation frequency onto the size of the representation (3.0 ± 0.4, 3.0 ± 0.6, mean ± standard deviation, fig. 1C green curve). Allowing a stronger competition for representing neurons between different stimuli, by reducing the size of the network, leads to a weak dependence of the size of their representation on presentation frequency (data not shown). Thus, the local mechanism alone leads to, firstly, the formation of stable and specific representations of stimuli, and secondly, a representation largely independent of the frequency of stimulus presentation.

Next, we investigated the role of the global mechanism. The activation of the nucleus basalis neuron leads to a long-lasting hyperpolarization of the inhibitory neurons. This necessitates a longer integration period of the inhibitory neurons until the threshold is reached. Nevertheless, with the chosen set of parameters the inhibitory neurons do not fire in bursts and those excitatory inputs, which without nucleus basalis activation were falling into the refractory period, now contribute to the activation of the inhibitory neurons. Thus, the mean activity of the inhibitory neurons does not change. This is consistent with experimental results, which found highly complex effects of basal forebrain and direct ACh application which can not be characterized by straightforward changes of mean firing rates of cortical neurons (Sato et al., 1987; Francesconi et al., 1988; Murphy and Sillito, 1991; Jimenez-Capdeville et al., 1997). The main effect of nucleus basalis stimulation in the model is the induced delay of the inhibitory activity relative to the excitatory activity in the cortical circuit by about 3.3 milliseconds (fig. 1B). As a result, a much larger fraction of the action potentials generated by the excitatory neurons invade the dendritic tree (21 ± 16 % and 87 ± 17 % without and with nucleus basalis activation respectively). In this condition the number of neurons representing the stimulus which is paired with nucleus basalis activation (stimulus # 8, mid blue) increases about 10 fold (fig. 1A, lower panel). The representations of the other stimuli, however, remain constant. Again, compiling results from a total of 40 runs, we did not find a significant difference between the size of the representations of stimuli shown at different frequencies (3.0 ± 0.4 vs. 2.9 ± 0.6, high and low rate respectively). However, the representation of the paired stimulus was significantly increased (31.7 ± 19.8, fig. 1C, blue curve). Thus, the global mechanism biases the size of the representation towards the paired stimulus, while allowing the map to preserve a stable representation of all stimulus patterns.

As a control a modified learning rule was studied: In case the backpropagating action potential successfully invades the dendritic tree, the sign of change of synaptic efficacy is dependent on the sequence of arrival of pre- and postsynaptic action potentials. An interesting difference was found: Those cortical neurons representing a particular stimulus did not contact all of the respective input units equally, but a symmetry breaking occurred and each cortical unit connected only to a subset of afferents sufficient for its activation. This can be interpreted as an effect, where a neuron cuts down on a set of inputs sufficient for its activation. With respect to the central question of the study, however, no difference was found. Each stimulus was learned by about 3-4 neurons and the size of the representation was independent on presentation frequency. Stimulation of the basal forebrain unit lead to an huge increase

of the representation of the paired stimulus as before. Thus, in the investigated system the symmetry / asymmetry of the temporal window for coincidence detection of pre- and postsynaptic action potentials is not relevant for the size of cortical representations.

Discussion

In this study, we investigated the interaction of local and global mechanisms regulating synaptic efficacy. This relates to previous work in the auditory system of the rat and guinea pig (Weinberger, 1993; Kilgard and Merzenich, 1998). In these experiments, the animals were exposed to acoustic stimuli either with or without paired electrical stimulation of the basal forebrain. The frequent presentation of a stimulus on its own did not lead to an enlargement of its representation in primary auditory cortex. In contrast, presentation of an auditory stimulus paired with electrical stimulation of the basal forebrain induced a dramatic increase of the cortical area devoted to that particular stimulus. This matches well the behavior of the model presented here. By making the propagation of action potentials into the dendritic tree subject to attenuation by inhibitory afferents, the local learning rule implements an inhibitory mechanism (Körding and König, 1998). However, this mechanism is acting on the change of synaptic efficacy, and not directly on the neuronal activity level. Recent physiological data show that in primary visual cortex optimally activated neurons tend to fire before suboptimally activated neurons (König et al., 1995). Therefore, the optimally activated neurons can block synaptic plasticity in the suboptimally activated ones. As a consequence, this mechanism allows neurons which are optimally tuned to a stimulus to prevent the recruitment of more and more neurons for its representation. The first protocol demonstrates that the local learning rule leads to a size of neuronal representations which is independent of the stimulus presentation frequency. Nevertheless, particular stimuli, whether rarely or frequently presented, can be emphasized and allocated an increased cortical representation through the global mechanisms, as shown in the second protocol.

In our model the substantial gabaergic projection originating in the basal forebrain and terminating on cortical inhibitory neurons (Freund and Guylas, 1990; Freund and Meskenaite, 1992) increases the proportion of sucessfully backpropagating action potentials in the cortical excitatory neurons. Interestingly, the much better investigated basal forebrain cholinergic projection (Singer and Rauschecker, 1982) increases the fraction of backpropagating action potentials in cortical neurons (Tsubokawa and Ross, 1997). Thus, these two subcortical projections may act synergistically, enhancing

each other's effect. Of course these findings do not preclude additional actions of acetylcholine, e.g. the modulation of effective intracortical connectivity (Verschure and König, 1999).

In the definition of the model several design decisions had to be made. Increasing the number of parameters in a model to make it more biologically plausible may obscure basic principles. Therefore, we decided not to address the 2-dimensional topographic arrangement of feature preferences. These seem to involve independent mechanisms, and several good reviews are available on this topic (e.g. Erwin et al., 1995).

Most notably, in the formulation of the learning rule enters only the absolute difference of timing of pre- and postsynaptic action potentials. In contrast to the experimental results (Markram et al., 1997; Bi and Poo, 1998) the sequence of their arrival is not taken into account. This simplification was made for several reasons. We did not want to create the impression that the observed effects depend on this particular feature. Indeed, in control experiments using a modified learning rule sensitive to the sequence of pre- and postsynaptic action potentials qualitatively identical results were obtained. The requirement for temporal contiguity is sufficient to generate the observed effects. This is not surprising, as the activity statistics of the input units follows a Poisson distribution, and thus did contain additional information. In a complete large scale model, including feedback of the "cortical" neurons onto units representing neurons in the thalamus new and interesting properties might emerge. These are, however, subject to current investigation and outside the scope of the present paper.

In the model the action of the basal forebrain is mediated by a subtle influence on the temporal relation of inhibitory and excitatory activity in the cortex, influencing the backpropagation of action potentials. This gives rise to two specific predictions: First, using classical electrophysiological techniques to measure the activity of single excitatory and inhibitory neurons the spike triggered average of local field potentials can be determined. This measure gives an indication of the relative timing of the recorded neurons with respect to the population activity. Therefore it is a sensitive indicator for the predicted delay of inhibitory activity during basal forebrain stimulation. Second, the fraction of retrogradely propagating action potentials (and their resulting effects on the calcium concentration) should be increased when the basal forebrain is stimulated. This can be assessed in vivo using the recently introduced two photon imaging technique (Denk et al., 1994). Thus, a test of these two specific predictions of the proposed mechanisms seems possible with presently available

techniques.

In conclusion, by allowing a global mechanism to affect synaptic plasticity via a local learning rule, a single integrated mechanism can be defined, which combines continuous learning, stability, specificity and flexibility. The global component, which was previously interpreted as a “print now” signal, might actually be best described as a “print bold” signal.

Acknowledgments

We are grateful to Konrad Körding and Mike Merzenich for valuable discussions of the previous work on the learning rule and the experimental data, Daniel Kiper for comments on a previous version of the manuscript and are happy to acknowledge the support of SPP Neuroinformatics, SNF (grant 31-51059.97, awarded to P.K.) and a FPU grant from MEC (M.A.S.-M., Spain).

References

Abbott, L.F. (1990). Learning in neural network memories. Network, 1, 105-122.

Alkon, D.L., Amaral, D.G., Bear, M.F., Black, J., Carew, T.J., Cohen, N.J., Disterhoft, J.F., Eichenbaum, H., Golski, S., Gorman, L.K., Lynch, G., Mcnaughton, B.L., Mishkin, M., Moyer, J.R., Olds, J.L., Olton, D.S., Otto, T., Squire, L.R., Staubli, U., Thompson, L.T., and Wible, C. (1991). Learning and Memory. Brain Res. Rev., 16, 193-220.

Bakin, J.S. and Weinberger, N.M. (1996). Induction of a physiological memory in the cerebral cortex by stimulation of the nucleus basalis. Proc. Natl. Acad. Sci. USA, 93, 11219-11224.

Bi, G. Q. and Poo, M. M. (1998). Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. J. Neurosci. 18, 10464-72.

Buonomano, D.V. and Merzenich, M.M. (1998). Cortical plasticity: From synapses to maps. Annu. Rev. Neurosci., 21, 149-86.

Denk, W., Delaney, K.R., Gelperin, A., Kleinfeld, D., Strowbridge, B.W., Tank, D.W., and Yuste, R. (1994). Anatomical and functional imaging of neurons using 2-photon laser scanning microscopy. J. Neurosci. Methods, 54, 151-62.

Erwin, E., Obermayer, K., and Schulten, K. (1995). Models of orientation and ocular dominance columns in the visual cortex: A critical comparison. Neural Comput., 7, 425-468.

Francesconi, W., Müller, C.M., and Singer, W. (1988). Cholinergic mechanisms in the reticular control of transmission in the cat lateral geniculate nucleus. J. Neurophysiol., 59, 1690-1718.

Freund, T.F. and Gulyas, A.I. (1991). GABAergic interneurons containing calbindin D28K or somatostatin are major targets of GABAergic basal forebrain afferents in the rat neocortex. J. Comp. Neurol., 314, 187-99.

Freund, T.F. and Meskenaite, V. (1992). Gamma-aminobutyric acid-containing basal forebrain neurons innervate inhibitory interneurons in the neocortex. Proc. Natl. Acad. Sci. USA, 89, 738-742.

Gerstner, W., Ritz, R., and van Hemmen, J.L. (1993). Why spikes? Hebbian learning and retrieval of time-resolved excitation patterns. Biol. Cybern., 69, 503-515.

Hasselmo, M.E. (1993). Acetylcholine and learning in a cortical associative memory. Neural Comp., 5, 32-44.

Hebb, D.O. (1949). The organization of behavior. Wiley, New York.

Jimenez-Capdeville, M.E., Dykes, R.W., and Myasnikov, A.A. (1997). Differential control of cortical activity by the basal forebrain in rats: A role for both cholinergic and inhibitory influences. J. Comp. Neurol., 381, 53-67.

Kilgard, M.P. and Merzenich, M.M. (1998). Cortical map reorganization enabled by nucleus basalis activity. Science, 279, 1714-8.

König, P., Engel, A.K., Roelfsema, P.R., and Singer, W. (1995). How precise is neuronal synchronization? Neural Comput., 7, 469-85.

Körding, K.P. and König, P. (1998). A single learning rule for the formation of efficient representations and for the segmentation of multiple stimuli. Soc. Neurosci. Abst., 24, 323.16

Magee, J., Hoffman, D., Colbert, C., and Johnston, D. (1998). Electrical and calcium

signaling in dendrites of hippocampal pyramidal neurons. Annu. Rev. Physiol., 60, 327-46.

Markram, H., Lubke, J., Frotscher, M., and Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. Science, 275, 213-215.

Montague, P.R., Dayan, P., and Sejnowski, T.J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. J. Neurosci., 16, 1936-47.

Murphy, P.C. and Sillito, A.M. (1991). Cholinergic enhancement of direction selectivity in the visual cortex of the cat. Neuroscience, 40, 13-20.

Sato, H., Hata, Y., Masui, H., and Tsumoto, T. (1987). A functional role of cholinergic innervation to neurons in the cat visual cortex. J.Neurophysiol., 58, 765-780.

Sejnowski, T.J. (1977). Storing covariance with nonlinearly interacting neurons. J. Math. Biology, 4, 303-321.

Singer, W. and Rauschecker, J.P. (1982). Central core control of development plasticity in the kitten visual cortex. II: Electrical activation of mesencephalic and diencephalic projections. Exp. Brain Res., 47, 223-233.

Singer, W., von Grünau, M.W., and Rauschecker, J.P. (1979). Requirements for the disruption of binocularity in the visual cortex of strabismic kittens. Brain Res., 171, 536-540.

Spruston, N., Schiller, Y., Stuart, G., and Sakmann, B. (1995). Activity-dependent action potential invasion and calcium influx into hippocampal CA1 dendrites. Science, 268, 297-300.

Stent, G.S. (1973). A physiological mechanism for Hebb's postulate of learning. Proc. Natl. Acad. Sci. USA, 70, 997-1001.

Stuart, G.J. and Sakmann, B. (1994). Active propagation of somatic action potentials into neocortical pyramidal cell dendrites. Nature, 367, 69-72.

Tsubokawa, H. and Ross, W.N. (1996). IPSPs modulate spike backpropagation and associated Ca^{2+}_i changes in the dendrites of hippocampal CA1 pyramidal neurons. J. Neurophysiol., 76, 2896-906.

Tsubokawa, H. and Ross, W.N. (1997). Muscarinic modulation of spike backpropagation in the apical dendrites of hippocampal CA1 pyramidal neurons. J. Neurosci., 17, 5782-91.

Verschure, P.F.M.J. and König, P. (1999). On the role of biophysical properties of cortical neurons in binding and segmentation of visual scenes. Neural Comp., 11, 1113-1138

Weinberger, N.M. (1993). Learning induced changes of auditory receptive fields. Cur. Opin. Neurobiol., 3, 570-577.

Appendix

The network includes an input layer, a layer of excitatory cells, a layer of inhibitory cells, and one unit representing the activity in the nucleus basalis.

The input layer consists of 100 neurons, and projects in an all-to-all manner to the set of excitatory cortical neurons. This connectivity is initially random, with weights homogeneously distributed between 0.3 and 0.7 in units of the postsynaptic threshold. Hence, the receptive fields of the excitatory neurons are initially unspecific. The input neurons are activated for 30 ms and generate spikes following a Poisson process. For physiological realism we implemented a refractory period for these neurons (5 ms). As a control simulations without such a refractory period were performed. The spatial distribution of their firing rate follows a Gaussian shape with a peak of 100 Hz, and dispersion of 3 units.

The excitatory neurons are modeled as integrate and fire neurons with a time constant of 9.5 ms. The synapses of the projection between input and excitatory neurons are modified according to the learning rule described below. During all the simulations the weights are kept in the 0-1 range. The set of excitatory units projects in a one-to-one manner to the inhibitory neurons. This projection is subject to a transmission delay of 2 ms.

The inhibitory units project all-to-all back to the excitatory neurons. Here as well, a transmission delay of 2 ms is taken into account. The dynamics of the inhibitory neurons are identical to those of the excitatory neurons described above. In addition to excitatory input, they receive inhibitory afferents from the basal forebrain.

The learning rule is a modification of one proposed in (Körding and König, 1998). The change of synaptic efficacy is contingent on the relative timing of pre- and postsynaptic action potentials. If these two signals are coincident on a time scale $\tau_0 = 10\text{ms}$ (Markram et al., 1997), the synaptic efficacy is increased (equation 1, $\alpha_{LTP} > 0$). In case the action potential is attenuated on its way into the dendritic tree by inhibitory input within 3 ms after its initiation, the synaptic efficacy is decreased (equation 2, $\alpha_{LTD} < 0$). The relative timing is evaluated as before. Furthermore, the introduction of heterosynaptic depression leads to a weakening of inactive synapses in the case of postsynaptic activity (equation 3, $\alpha_{heteroLTD} < 0$):

$$(1)\quad \Delta\omega = \alpha_{LTP}\frac{\tau_0}{\tau_0 + |t_{post} - t_{pre}|}$$

$$(2)\quad \Delta\omega = \alpha_{LTD}\frac{\tau_0}{\tau_0 + |t_{post} - t_{pre}|}$$

$$(3)\quad \Delta\omega = \alpha_{heteroLTD}$$

with $\tau_0 = 10\text{ms}$ and t_{post} and t_{pre} timing of the pre - and postsynaptic action potential.

For a control simulation we used a modified learning rule, where the sign of synaptic modification is dependent on the sequence of pre- and postsynaptic action potentials:

$$(1)\quad \Delta\omega = \alpha_{LTP}\frac{sign(t_{post} - t_{pre})\tau_0}{\tau_0 + |t_{post} - t_{pre}|}$$

with parameters as before.

For each stimulus presentation, every synapse is updated only once, even if there are several presynaptic spikes generated by the presynaptic neuron.

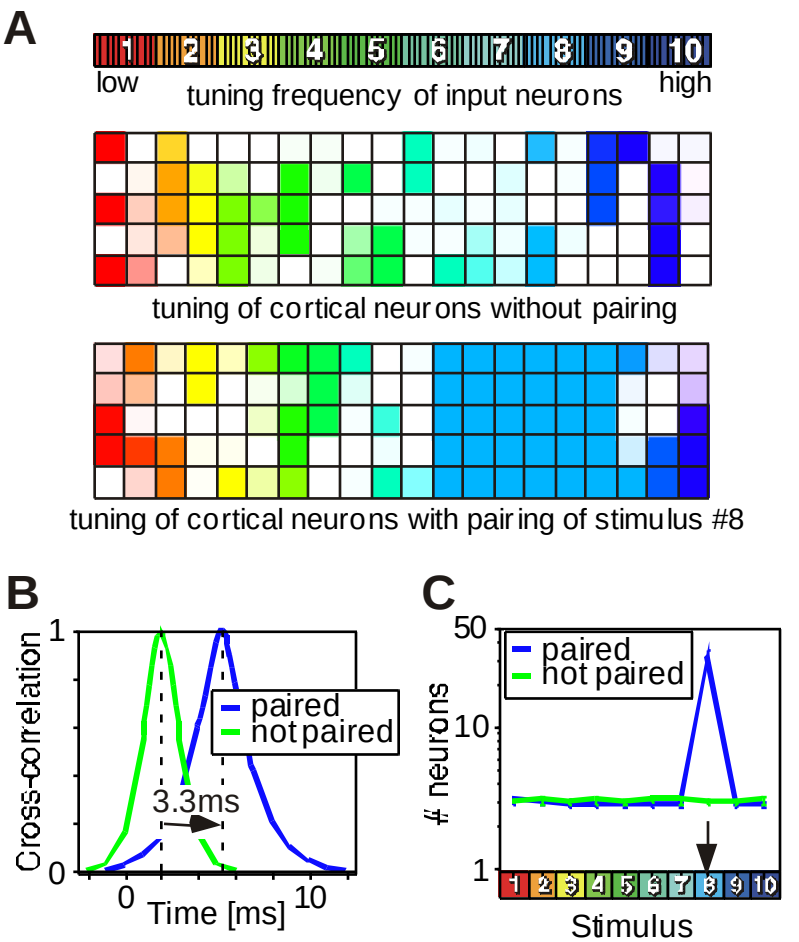

Figure 1

(A) The tuning of the 100 input neurons is shown in the upper panel using a false color code. Low frequencies correspond to reddish colors, high frequencies to violet, and intermediate frequencies follow a rainbow pattern. The distribution of the preferred frequency of the cortical excitatory units after training, without nucleus basalis activity, is shown in the middle panel using the same color code: Hue is indicating the preferred stimulus frequency, while saturation is indicating the tuning strength. Each square corresponds to one neuron and, for better visibility, they are arranged in five rows in order of increasing preferred stimulus frequency. Thus, the graph might be compared to a top-view onto the cortex as used by (Kilgard and Merzenich, 1998). In the lower panel the tuning of cortical excitatory neurons is shown, while one stimulus (8, mid-blue) was paired with nucleus basalis activity. (B) Temporal relation of the activity of cortical inhibitory neurons in relation to activity of excitatory neurons with (blue) and without (green) simultaneous nucleus basalis activity. (C) Size of the representation of all 10 stimuli with (blue), or without nucleus basalis activity (green) paired with stimulus # 8. The error bars indicate standard error of the mean.