

Synaptic Scaling Balances Learning in a Spiking Model of Neocortex

Mark Rowan¹ and Samuel Neymotin²

¹ School of Computer Science, University of Birmingham, UK
m.s.rowan@cs.bham.ac.uk

² Dept. Neurobiology, Yale University School of Medicine, New Haven, USA
samuel.neymotin@yale.edu

Abstract. Learning in the brain requires complementary mechanisms: potentiation and activity-dependent homeostatic scaling. We introduce synaptic scaling to a biologically-realistic spiking model of neocortex which can learn changes in oscillatory rhythms using STDP, and show that scaling is necessary to balance both positive and negative changes in input from potentiation and atrophy. We discuss some of the issues that arise when considering synaptic scaling in such a model, and show that scaling regulates activity whilst allowing learning to remain unaltered.

1 Introduction

Spike Timing-Dependent Plasticity (STDP), a phenomenological learning rule in which synaptic potentiation and depression depend upon relative firing times [1,2], has been used to learn oscillatory rhythms in neocortical models. In an existing biologically-realistic spiking model of neocortex [3], applying excitatory to excitatory (E→E) STDP with a rhythmic training signal led to hyper-potentiation through positive feedback: strengthened synapses drove postsynaptic neurons to fire immediately, leading to further potentiation. This unbounded potentiation then pushed the network into synchronized epileptiform firing. Directly opposing E→E learning with equal excitatory to inhibitory (E→I) potentiation partially balanced this positive feedback. However, epileptiform behaviour still occurred with high-frequency signals [4].

We postulated that a homeostatic mechanism might be a solution to this problem. Neuronal homeostatic synaptic scaling is a local feedback mechanism which senses levels of activity-dependent cytosolic calcium within the cell and adjusts neuronal firing activity accordingly. This is achieved by producing alterations in excitatory AMPA receptor accumulation in response to changes in firing activity occurring over hours to days [5], leading to changes in the excitability of the neuron.

During learning, synaptic scaling plays an important role in balancing potentiation. By constantly shifting mean activation back towards a target activity level, but maintaining the learned relative distribution of presynaptic weights, global levels of activity can be regulated [6]. During periods of hypoactivity (e.g.

in degenerative disorders), synaptic scaling is also capable of raising the sensitivity of neurons via AMPA receptor upregulation, so that activity levels can be restored [5].

Previous work has demonstrated synaptic scaling with learning in a single-neuron model [6]. It has also been shown that synaptic scaling can prevent input saturation in a spiking neural network in the absence of learning [7]. In this paper, we add long-term synaptic plasticity to a spiking neural network to show that homeostatic synaptic scaling can prevent hyper-potential while preserving learned information.

2 Methods

The model was based on the anatomy of a single column of sensory neocortex [3,8,9]. It was composed of 470 neurons divided into 3 types (excitatory pyramidal cells E, fast-spiking interneurons I, and low-threshold spiking interneurons IL), distributed across the 6 layers of the neocortex. This yielded 13 neuronal populations in total, with the following numbers of cells per type: E2 (i.e. excitatory layer 2/3 cell), 150; I2 (fast spiking interneuron in layer 2/3), 25; I2L (low-threshold spiking interneuron in layer 2/3), 13; E4, 30; I4, 20; I4L, 14; E5a, 65; E5b, 17; I5, 25; I5L, 13; E6, 60; I6, 25; and I6L, 13.

The cell model was an extension of an integrate-and-fire unit with added complexity (adaptation, bursting, depolarization blockade, and voltage-sensitive NMDA conductance) in the form of rules [10], and was simulated in an event-driven fashion where cell state variables were only calculated at input events, making use of previously developed just-in-time synapses optimized for large networks supporting high-frequency synaptic events [11]. Each cell had fast inhibitory GABA_A receptors, fast excitatory AMPA receptors, and slow excitatory NMDA receptors, with each producing a voltage-step with following decay.

In addition to spikes generated by cells in the model, subthreshold Poisson-distributed spike inputs to each synapse were used to maintain activity in the model: 100–150 Hz for GABA_A, 240–360 Hz for AMPA receptors, and 40–60 Hz for NMDA receptors. These external inputs represented the inputs from other regions of the brain. To simulate additional afferent sensory inputs, low-amplitude training signals were applied to the layer 4 excitatory neurons (E4) in some simulations. STDP was implemented on AMPA synapses from E→E cells using a basic model with bilateral exponential decay (40ms maximal interspike difference, 10ms time constant) incrementing by 0.1% of baseline synaptic weight. It should be noted that STDP in this model is additive, whilst van Rossum argues that it should be multiplicative [6]. Further details of the cell model can be found in [3] and [4].

Scaling was implemented at E cell AMPA synapses by multiplying each cell i 's postsynaptic input by a scale factor w_i , representing the multiplicative accumulation of AMPA receptors at synapses. Changes in the scale factor were calculated following the formula of van Rossum et al. [6], with a_i as the cell's firing activity, a_i^{goal} as the target activity, β as the scaling strength, γ as the

“integral controller” weight, and $\frac{dw_i(t)}{dt}$ as the rate of change of the synaptic weight:

$$\frac{dw_i(t)}{dt} = \beta w_i(t)[a_i^{goal} - a_i(t)] + \gamma w_i(t) \int_0^t dt' [a_i^{goal} - a_i(t')] \quad (1)$$

The following parameter values were used: strength $\beta = 4.0 \times 10^{-8}$ /ms/Hz; integral controller weight $\gamma = 1.0 \times 10^{-10}$ /ms²/Hz; activity sensor time constant $\tau = 100 \times 10^3$ ms. Scaling was applied inversely at GABA_A synapses (i.e. by multiplying postsynaptic input by $\frac{1}{w_i}$) to enable the scaling of excitatory and inhibitory synapses in opposite directions, mimicking the effect of global growth factors such as BDNF [5,7,12,13].

Average activity level for each cell i was sensed using van Rossum’s slow-varying sensor $a_i(t)$, which increased monotonically with spike t_x at current timestep t , and decayed otherwise [6]:

$$\tau \frac{da_i(t)}{dt} = -a_i(t) + \sum_x \delta(t - t_x) \quad (2)$$

The sensor decays exponentially as it is updated at each non-firing timestep. However, the use of event-driven just-in-time synapses [4,11] meant that cell states were only updated upon each spike event rather than at every timestep, so inter-spike decay of the activity sensor could only be calculated periodically. We therefore modified the activity sensor. Here, the first term decays the sensor according to the time between spikes $t - t_x$, and the second term increments it for the new spike, with both terms updated concurrently on the occurrence of a spike at time t_x :

$$a_i(t) = a_i(t_x)e^{-\frac{1}{\tau}(t-t_x)} + \frac{1 - a_i(t_x)}{\tau} \quad (3)$$

Figure 1 shows the activity of a simulated, randomly-spiking neuron operating under the constant-timestep update policy (2), and the equivalent activity values under the periodic-update policy (3). The activity rises identically in both cases when spikes occur, but the periodic sensor does not decay until the next spike event occurs, giving the step-like appearance. The values at the spike times are correct down to round-off error at the spike times.

Instead of providing an arbitrary rate target for each cell, which would fundamentally affect network dynamics, the intrinsic dynamics of the network were used to provide set-points. Initially, with synaptic scaling off, activity sensors began at 0 Hz. They were then adjusted over 800 s of simulated time based on the activity level of the cells. Synaptic scaling was then switched on.

A time constant τ of 100 s [6] leads to a simulation timescale of several hours for synaptic scaling; far closer to the expected biological timescale than previous studies [5,14,7]. To achieve this length of simulation, the model was extended to allow periodic flushing of all spike data to disk, enabling very long runs (unlimited except for available disk space). A typical simulation of 44 hours ran in approximately real time and produced around 2 GB of spike data. The

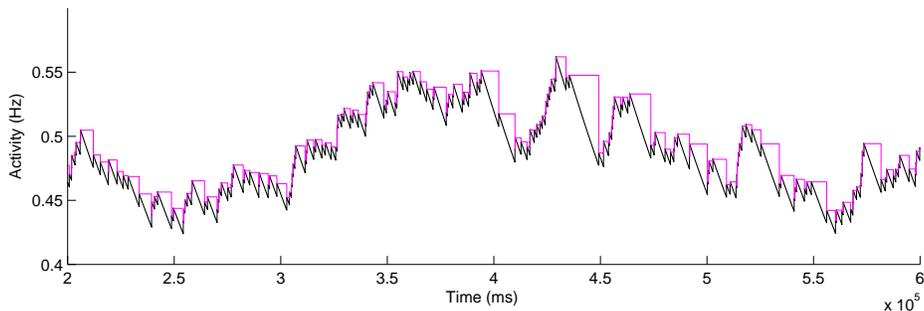


Fig. 1: Activity sensor updating at every simulation timestep (Eqn. 2; black) and at every spike for activity-driven just-in-time synapses (Eqn. 3; magenta).

model was implemented in NEURON 7.2 [15] for Linux, and is available on ModelDB at the following URL: (<https://senselab.med.yale.edu/modeldb/enterCode.asp?model=147141>).

Data analysis Simulation spike-trains were organized into multiunit activity (MUA) vectors, defined for a cell population as the number of spikes in the population over a time interval (bin). Bin sizes were set to 5 ms (200 Hz sampling rate). Analyses were performed using mean-subtracted MUA vectors, with spectra calculated by the multitaper spectral power estimation method, as implemented in the Python wrapper of the FORTRAN MTSpec library [16].

3 Simulation Results

3.1 Scaling Prolonged Activity During Deletion

In an initial experiment, we demonstrated the usefulness of synaptic scaling by altering network dynamics through gradual removal (pruning) of cells (Fig. 2). Every 1600 s, three I or E neurons were selected at random and removed from the network by setting all their synaptic weights to zero. The global external input weights were scaled down proportionally to the amount of deletion, at a quarter of the deletion rate, to prevent the external inputs from swamping internal activation and artificially raising activity. By the end of the simulation, approximately two thirds of the cells in the network had been deleted.

In the absence of scaling, average firing across E cells declined steadily as cell deletion progressed (Fig. 2 green / lower line). With scaling present, firing activity was maintained (Fig. 2 blue / upper line), with brief activity peaks caused when the inherent delay in the activity sensor led to over-compensation. These activity peaks do not correspond to deletion times, but rather to emergent instabilities in the resulting damaged network. Indeed, the network remains stable for nearly half a day following the onset of deletion after 800 s. The over-compensation can be adjusted to some degree, although not eliminated, by altering the scaling parameters β and γ (not shown).

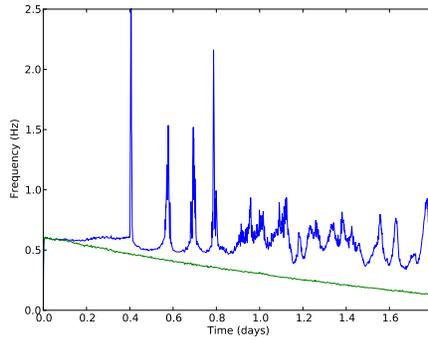


Fig. 2: E activity during pruning with (blue / upper) and without (green / lower) compensatory synaptic scaling. Run time 160,000 s (\approx 44 h).

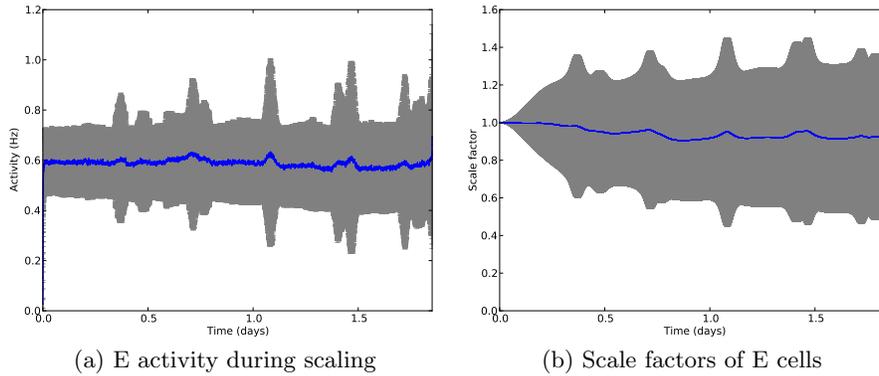


Fig. 3: Scaling does not destabilize the network (mean: blue, std: grey).

3.2 Synaptic Scaling Did Not Disrupt Network Behavior

The model was run for 160,000 s (\approx 44 h) to examine the effects of scaling over time on network dynamics. With scaling, activity of the E cells remained steady (Fig. 3a), and scale factors remained centered around 1 (Fig. 3b). Scaling appeared to preserve stability of the network during these extremely long runs.

3.3 Unrestrained STDP Led To Hyper-Potential

We trained the network by applying a signal consisting of low-weight single spikes at 8 Hz to E4 cells for 8000 s (\approx 2.2 h) in the absence of synaptic homeostasis (Fig. 4). STDP was turned off during the final 800 s in order to test recall. We found that any training signal frequency eventually pushed the network into a state of excessive firing. This occurred even when E \rightarrow I STDP balancing was added (not shown).

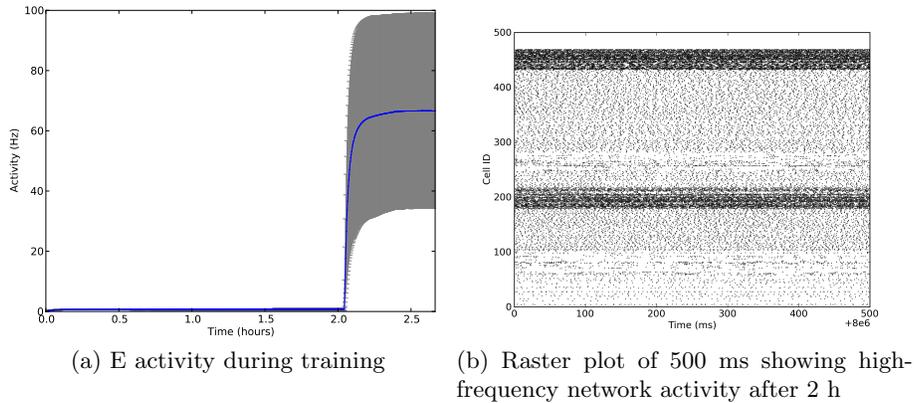


Fig. 4: Training with E→E STDP pushes network to high frequency activity.

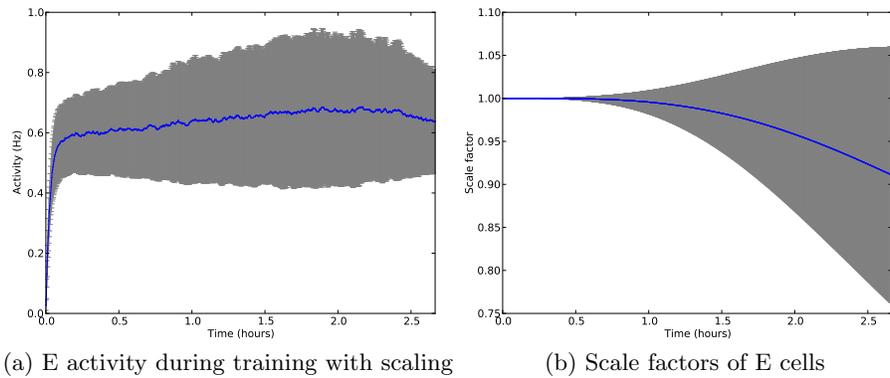


Fig. 5: Synaptic scaling maintains E activity profile during STDP.

3.4 Synaptic Scaling Prevented Overactivation

We then assessed the model with STDP, training and synaptic scaling (Fig. 5). Local E cell homeostatic scaling balanced the potentiation caused by STDP, gradually scaling down all E cells, and preventing pathological over-activation.

3.5 Synaptic Scaling Preserved Learning

Synaptic scaling served to maintain cell firing near the target rate, here the baseline rate. However, it was possible that the scaling-down of activity would simply reverse the potentiation caused by STDP, resulting in a loss of learned information. In order to determine whether scaling allowed the learning of oscillations to persist, the power spectra of the E cells were obtained at various points during

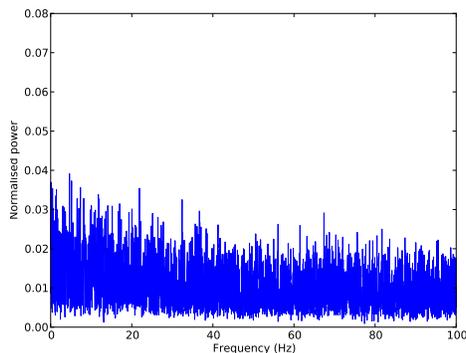


Fig. 6: Baseline power spectrum of E cells (frequency (Hz) vs normalized power).

the learning process using the multitaper spectral analysis method, with spikes sorted into 5 ms bins (Fig. 6) [16]. These plots show unsmoothed normalised power of the E cells within the network at each of a range of frequencies from 0-100 Hz.

STDP was applied at E→E synapses for 8000 s (\approx 2.2 h) with an 8 Hz sensory signal (Fig. 7d). In one simulation, synaptic scaling was also switched on for E cells. Power spectra were obtained for the period from 5600-6400 s, shortly after the middle of training (Figs. 7a and 7b), and again during the recall period at the end of learning (Figs. 7c and 7d).

In both simulations, it can be seen that STDP has caused a shift in the power spectra, with an increase in the amplitude of oscillations at low frequencies from 0-10 Hz and a decrease above 10 Hz (Figs. 7a and 7b). This demonstrates that the network has learned from the training signal. Shortly after 7400 s (2 h), the network without synaptic scaling transitioned to high-frequency activity, without retention of the 8 Hz training signal (Fig. 7c; note different scale). However, in the network with synaptic scaling turned on, lower frequency activity was maintained, with a peak near the 8 Hz that was imposed during training (Fig. 7d). Synaptic scaling therefore prevented over-activation and preserved learning.

4 Discussion

This research has introduced homeostatic synaptic scaling with dynamically-obtained target activity rates to a realistic spiking model of neocortex which learned oscillatory frequencies via STDP. We demonstrated that scaling is necessary for upregulation of neural activity during decline in input. This might have implications for neurodegenerative brain disorders, in which cortical activation might be expected to decrease. Peaks of activity were observed during deletion due to periodic over-compensation by the scaling mechanism. Experimental observations demonstrating hyperactivity in cells near beta-amyloid plaques in Alzheimer’s disease, and the increased incidence of seizures in Alzheimer’s

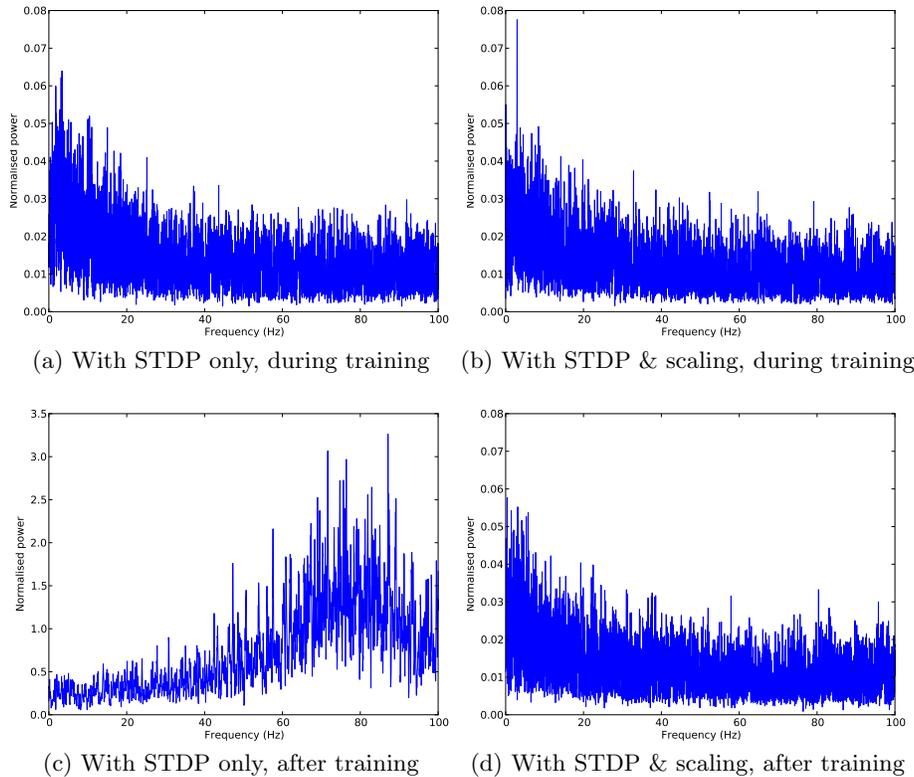


Fig. 7: Power spectra during (top)/after (bottom) STDP, with (R)/without (L) scaling.

patients, suggests these activity peaks may have a biological basis [17,14,18]. Additionally, synaptic scaling may play a significant role in the progression of Alzheimer's disease [19,20,21], and further understanding of this mechanism and its relationship to learning and disease pathology may be crucial to finding better treatments.

We also showed that scaling does not negatively affect the network at baseline, but that it is stable. We demonstrated that $E \rightarrow E$ scaling is sufficient to balance the hyper-potential caused by unrestrained STDP. Potentiation strengthens the co-incident connections between neurons in a positive feedback cycle, eventually leading to hyper-potential, but scaling acts to shift the mean activation constantly back towards the target activity. At the same time, the relative (learned) distribution between postsynaptic weights remains unaltered by scaling, and we subsequently demonstrated this principle by showing that learning of an 8 Hz oscillatory signal is not erased by scaling.

This model investigated training and scaling at E→E synapses between E cells. While there is some evidence of STDP in I cells [22], I cells do not appear to perform scaling, but rather: “homeostatic regulation of inhibition is a noncell-autonomous process that either requires changes in both pre-and postsynaptic activity simultaneously or is triggered by global changes in network activity” (Turrigiano et al., 2011 [13]). In our model, directly enabling synaptic scaling in I cells was found to lead to dramatic instabilities in the network dynamics (even when operating the network at baseline, i.e. without STDP or a sensory signal), which is consistent with Turrigiano’s observations. Rather, the network appears to be most stable when I cells are allowed to adjust their activity passively according to the changing output from neighboring E cells, thus requiring only one dimension for the E/I balance rather than needing a second simultaneously active dimension for scaling.

STDP was implemented using an incremental step of 0.1% of baseline synaptic weight, which may seem low. Increasing this step size, however, meant that short bursts of high-frequency activity were seen during learning, as the activity sensors could not respond quickly enough to cause sufficient compensatory scaling (although the network did soon scale back to previous firing rates). However, 8000 s (2 h) of sustained training may also be very long compared to biological learning from hippocampal backprojections, which is known to include periods of recall and consolidation between periods of learning [23]. This would make an interesting avenue for future research.

Acknowledgements. Research funded by EPSRC, DARPA grant N66001-10-C-2008, NIH grant R01MH086638. The authors would like to thank John Bullinaria (Birmingham) and William Lytton (SUNY Downstate) for their helpful comments; Michael Hines and Ted Carnevale (Yale) for NEURON simulator support; Tom Morse (Yale) for ModelDB support; and the anonymous reviewers for their constructive feedback.

References

1. Dan, Y., Poo, M.: Spike timing-dependent plasticity of neural circuits. *Neuron* **44**(1) (2004) 23–30
2. Zhang, L., Tao, H., Holt, C., Harris, W., Poo, M.: A critical window for cooperation and competition among developing retinotectal synapses. *Nature* **395**(6697) (1998) 37–44
3. Neymotin, S., Lee, H., Park, E., Fenton, A., Lytton, W.: Emergence of physiological oscillation frequencies in a computer model of neocortex. *Front Comput Neurosci* **5** (2011)
4. Neymotin, S., Kerr, C., Francis, J., Lytton, W.: Training oscillatory dynamics with spike-timing-dependent plasticity in a computer model of neocortex. In: *Signal Processing in Medicine and Biology Symposium (SPMB)*, IEEE (2011) 1–6
5. Turrigiano, G.: The self-tuning neuron: synaptic scaling of excitatory synapses. *Cell* **135**(3) (2008) 422–435

6. Van Rossum, M., Bi, G., Turrigiano, G.: Stable Hebbian learning from spike timing-dependent plasticity. *J Neurosci* **20**(23) (2000) 8812
7. Chandler, B., Grossberg, S.: Joining distributed pattern processing and homeostatic plasticity in recurrent on-center off-surround shunting networks: Noise, saturation, short-term memory, synaptic scaling, and BDNF. *Neural Networks* (2012)
8. Binzegger, T., Douglas, R., Martin, K.: A quantitative map of the circuit of cat primary visual cortex. *The Journal of Neuroscience* **24**(39) (2004) 8441–8453
9. Lefort, S., Tómm, C., Floyd Sarria, J., Petersen, C.: The excitatory neuronal network of the C2 barrel column in mouse primary somatosensory cortex. *Neuron* **61**(2) (2009) 301
10. Lytton, W., Stewart, M.: Rule-based firing for network simulations. *Neurocomputing* **69**(10) (2006) 1160–1164
11. Lytton, W., Omurtag, A., Neymotin, S., Hines, M.: Just-in-time connectivity for large spiking networks. *Neural Comput* **20**(11) (2008) 2745–2756
12. Rutherford, L., Nelson, S., Turrigiano, G.: BDNF has opposite effects on the quantal amplitude of pyramidal neuron and interneuron excitatory synapses. *Neuron* **21**(3) (1998) 521–530
13. Turrigiano, G.: Too many cooks? intrinsic and synaptic homeostatic mechanisms in cortical circuit refinement. *Annu Rev Neurosci* **34** (2011) 89–103
14. Fröhlich, F., Bazhenov, M., Sejnowski, T.: Pathological effect of homeostatic synaptic scaling on network dynamics in diseases of the cortex. *The Journal of Neuroscience* **28**(7) (2008) 1709–1720
15. Carnevale, N., Hines, M.: *The NEURON Book*. Cambridge University Press, New York (2006)
16. Prieto, G., Parker, R., Vernon III, F.: A Fortran 90 library for multitaper spectrum analysis. *Computers & Geosciences* **35**(8) (2009) 1701–1710
17. Busche, M., Eichhoff, G., Adelsberger, H., Abramowski, D., Wiederhold, K., Haass, C., Staufenbiel, M., Konnerth, A., Garaschuk, O.: Clusters of hyperactive neurons near amyloid plaques in a mouse model of Alzheimer’s disease. *Science Signalling* **321**(5896) (2008) 1686
18. Trasande, C., Ramirez, J.: Activity deprivation leads to seizures in hippocampal slice cultures: is epilepsy the consequence of homeostatic plasticity? *J Clin Neurophysiol* **24**(2) (2007) 154–164
19. Small, D.H.: Network dysfunction in Alzheimer’s disease: does synaptic scaling drive disease progression? *Trends Mol Med* **14**(3) (2008) 103 – 108
20. Rowan, M.: Information-selectivity of beta-amyloid pathology in an associative memory model. *Front Comput Neurosci* **6**(2) (January 2012)
21. Rowan, M.: Effects of Compensation, Connectivity and Tau in a Computational Model of Alzheimer’s Disease. In: *Proc. IJCNN, IEEE* (2011) 543–550
22. Lamsa, K., Kullmann, D., Woodin, M.: Spike-timing dependent plasticity in inhibitory circuits. *Frontiers in Synaptic Neuroscience* **2** (2010)
23. McClelland, J., McNaughton, B., O’Reilly, R.: Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychol Rev* **102**(3) (1995) 419–457