

Training oscillatory dynamics with spike-timing-dependent plasticity in a computer model of neocortex

Samuel A Neymotin*, Cliff C Kerr^{† ‡}, Joseph T Francis[†], and William W Lytton^{† §}

*Biomedical Engineering Program

SUNY Downstate / NYU-Poly, Brooklyn, New York 11203

Email: samn@neurosim.downstate.edu

[†]Dept. Physiology & Pharmacology

SUNY Downstate Medical Center, Brooklyn, NY 11203

[‡]Dept. Physics, U. Sydney, Sydney, Australia

[§]Dept. Neurology, Kings County Hospital Center

Brooklyn, NY 11203

Abstract—Neuronal networks are complex, adaptive systems that typically display oscillatory dynamics. The extent to which these dynamics can be shaped by training remains unknown. We explored this dynamical training in a computer model of 6-layered sensory neocortex with 470 excitatory (E) and inhibitory (I) cells. AMPA, NMDA, and GABA_A synapses were provided with Poisson input to provide baseline activation in the network. The learning rule employed spike-timing-dependent plasticity (STDP) at all AMPA synapses. We trained with a 1-16 Hz thalamic afferent signal to E4 cells (layer 4 E cells). At baseline, the power spectrum of the network activity showed oscillations with a low-amplitude peak near 6 Hz. Plasticity in the absence of a training signal (white noise input) attenuated the network response, due to the potentiation of E-to-I synapses. Plasticity coupled with an 8 Hz training signal enhanced the network's oscillations and shifted the peak to ~20 Hz. This was due to increased synaptic connection strengths between E cells caused by the near-synchronous firing of E4 cells. Plasticity coupled with a 16 Hz training signal shifted the network towards epilepsy, with high-amplitude 8 Hz oscillations and synchronous firing across all layers. The shift into epilepsy was caused by further enhancement of E-to-E synapses. In summary, our simulations demonstrate the feasibility of using plasticity and neuroprosthetic input signals to train a neuronal network's oscillatory dynamics. We predict that in order for learning in the brain to avoid transition to epilepsy, homeostatic control mechanisms must balance learning at E-to-E and E-to-I synapses.

I. INTRODUCTION

Learning in biological neuronal networks occurs by changing the strengths of the synaptic connections between neurons, allowing the system to adapt to the environment. Models of artificial neural networks have been used to solve learning problems for decades, but the learning capacity of neuronal networks composed in biologically-realistic neuronal units and layered neocortical architecture has not been as well investigated. Understanding learning in this type of highly-structured network requires investigation of ongoing dynamics and multiple different possible learning loci, as well as the investigation of particular training patterns, giving many

possible training permutations. [1], [2], [3] This differs from investigation of plasticity in artificial neural networks where rules and architectures are designed to be well-suited for optimal learning.[4], [5]

Oscillations are one of the most basic patterns of neuronal network activity, and are believed to be strongly associated with memory and cognition as well as disrupted in brain disease. [6], [7], [8], [9], [10], [11], [12], [13] We therefore decided to study the potential for development, control and production of these dynamics, with implications for neurodevelopment and dynamical homeostasis. We hypothesized that a rhythmic training signal would mold the oscillatory dynamics of a neuronal network with realistic architecture, using the biologically-plausible learning rule spike-timing-dependent plasticity (STDP) [14], [15], [16].

II. METHODS

A. Simulations

The model was based on the anatomy of a single column of sensory neocortex [17], [18], [19], [20], [21], [22], along with published models freely available at <https://senselab.med.yale.edu/modedb> [23]. It was comprised of 470 neurons divided into three types (excitatory pyramidal cells E, fast-spiking interneurons I, and low-threshold spiking interneurons IL), distributed across the six layers of the cortex. 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 and was simulated in an event-driven fashion, since cell state variables were only calculated at input events. Complexity was added using rules with various dynamical features found in real neurons: adaptation, bursting, depolarization blockade,

and voltage-sensitive NMDA conductance. [24], [25] Each cell had a membrane voltage level state variable (V_m) and a resting membrane potential (V_{RMP}), which was the baseline value for V_m in the absence of external inputs. After synaptic input events, if V_m crossed spiking threshold (V_{TH}), the cell would emit an action potential and enter an absolute refractory period, lasting τ_{refrac} ms, during which it could not fire. Refractory periods were set to prevent a maximum firing frequency from being exceeded. After an action potential, an after-hyperpolarization voltage state variable was set (V_{ahp}) and then subtracted from V_m . Then V_{ahp} decayed exponentially with time-constant (τ_{ahp}) to 0. To simulate voltage blockade a cell could not fire if V_m surpassed V_{block} . Relative-refractory period was simulated after an action potential by increasing the threshold, V_{TH} , by $W_{RR} \cdot (V_{block} - V_{TH})$ (to scale the increase in V_{TH} by a fraction of the voltage range where the cell could fire), where W_{RR} was a unitless weight parameter. V_{TH} then decayed exponentially to its baseline value with time constant τ_{RR} .

Each cell used three types of synapse: fast inhibitory GABA_A receptors, fast excitatory AMPA receptors, and slow excitatory NMDA receptors [26]. Synapses utilized our previously developed just-in-time synapses, optimized for large networks supporting high-frequency synaptic events.[26] Each cell had a voltage state-variable associated with a synapse type, V_{syn} , one for each of excitatory AMPA, NMDA, and two inhibitory GABA_A's, which simulated GABA_A at soma (fast time-constant) and GABA_A at dendrite (slower time-constant). Synaptic inputs were simulated by step-wise changes in V_{syn} and then added to the cell's V_m level. To allow for dependence on V_m , synaptic inputs changed V_{syn} by $\delta V = w_{syn} \cdot (1 - \frac{V_m}{E_{syn}})$, where w_{syn} is the synaptic weight and E_{syn} is the reversal potential, relative to V_{RMP} . E_{syn} took the following values (in mV): AMPA 65, NMDA 90, GABA_A -15. w_{syn} was positive for excitatory synapses and negative for inhibitory synapses. NMDA synapses also had an additional voltage-dependent scaling factor based on Jahr and Stevens 1990. [27], [28] For all synapses, after synaptic input events, V_{syn} decayed exponentially towards 0 with time constant τ_{syn} . τ_{syn} took the following values (in milliseconds): AMPA 20, NMDA 300, somatic GABA_A 10, dendritic GABA_A 20. Synaptic weights were constant between a given set of populations. Dendritic synapses (AMPA, NMDA, dendritic GABA_A) utilized delays chosen from a uniform distribution ranging from 3–5 ms, while somatic synapses (somatic GABA_A) had delays ranging from 1.8–2.2 ms.

Connectivity between the neurons in each of the populations is shown in Fig. 1. Connections tend to be strongest between populations within a given layer; this corresponds to the four areas of high connectivity visible along the diagonal of Fig. 1. Overall, excitatory cells have more projections than inhibitory ones, but inhibitory projections are typically stronger.

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

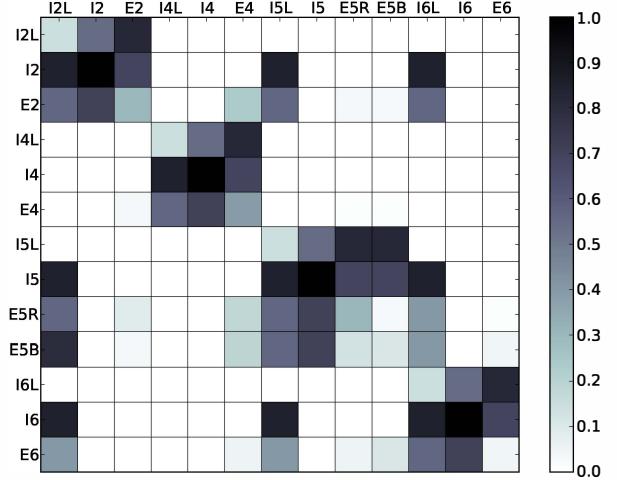


Fig. 1. Connectivity between cell types in the model. Vertical and horizontal axes represent cell types and color represents the normalized probability that a neuron from a given column will project to a neuron from a given row.

Hz for NMDA receptors. These external inputs represented the inputs from other regions of the brain. To simulate a neural prosthetic that acts by stimulating thalamic afferents, low-amplitude training signals were applied to the layer 4 excitatory neurons (E4) in some simulations.

STDP was implemented on AMPA synapses using a basic model with bilateral exponential decay (40 ms maximal interspike difference, 10 ms time constant) and hard cutoffs for maximal weights. We applied STDP at AMPA synapses from E → E and E → I cells. In addition, we set maximal post-STDP weights at E → I and E → I AMPA synapses to be 15 and 10 times the weights at baseline, respectively. STDP increments were balanced in favor of inhibition, with E → E and E → I increments set at 5 and 10% of baseline synaptic weight, respectively.

The model was implemented in NEURON 7.2 [29] for Linux, and is available on ModelDB (<https://senselab.med.yale.edu/modeldb>). One minute of simulated time took approximately one minute of CPU time on an Intel XEON 2.27 GHz CPU.

B. 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). Organization of the simulations and analysis of output data was done using Neural Query System.[30] 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.

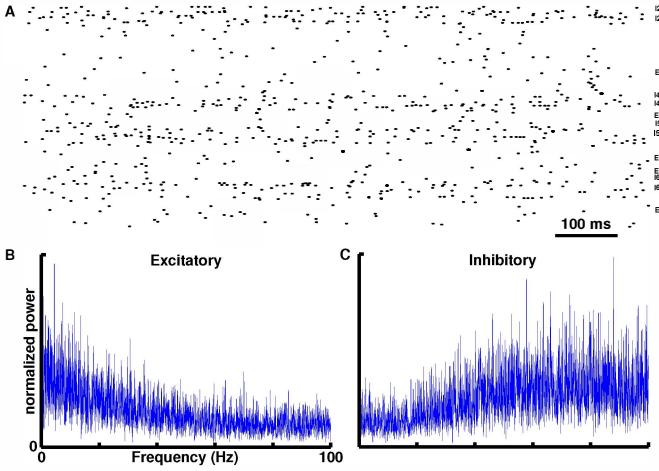


Fig. 2. Baseline activity of network. A) raster. B) Excitatory and C) Inhibitory multiunit activity (MUA) power spectra. Power is normalized; overall I power is higher. Note that we have used unsmoothed power spectra to provide greater detail

III. RESULTS

Simulations were run for 500 s of simulated time. They were organized into three periods: baseline (100 s), training (300 s), and recall (100 s). At baseline, each synapse received Poisson inputs and inputs from other cells. During training, STDP was turned on at AMPA synapses. Rhythmic training signals were applied to E4 cells to simulate neuroprosthetic input. During recall, STDP and all training signals were turned off, to assess the network's response to training.

A. Baseline network shows weak oscillations

Fig. 2A shows the baseline activity of the network. Average firing rates of excitatory (E) and inhibitory (I) cells were between 0.7–1 Hz and 3.3–4.2 Hz, respectively, as shown in Fig. 3. Periods of weak and strong correlations between firing times of cells were visible as vertical stripe patterns in the raster, which alternated with more diffuse firing. Although driving inputs to cells were Poisson-distributed, with a flat white-noise profile, the interactions between cells in the network produced a non-white spectrum, evidenced by the peak (defined as maximum spectral power) near 6 Hz in the E MUA power spectrum (Fig. 2B). The E MUA spectrum also had attenuation of higher-frequency oscillations, and had similar overall appearance to physiological recordings.[20] The I MUA power spectrum (Fig. 2C) showed little power from 0–40 Hz and had activity spread across a broad range of frequencies in the gamma band (40–100 Hz). Note that individual E and I cells fired at rates significantly lower than the population rhythms, demonstrating that the oscillation frequencies emerge primarily from network-level interactions.

B. Balanced E,I STDP flattens power spectra

Applying STDP to only E→E connections led to epileptiform activity characterized by recurrent coordinated activation of a large proportion of the cells in the network (not shown).

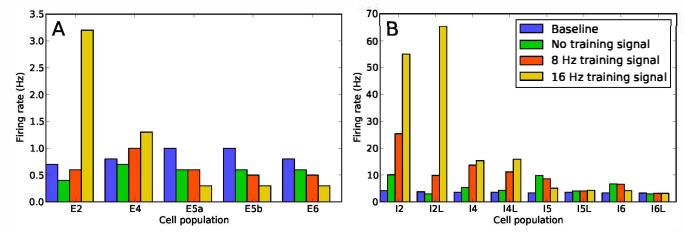


Fig. 3. Average neuronal firing rates in each population in the model, for the baseline case with no plasticity (blue) and training cases with plasticity and different input signals (green, red, yellow). (A) Excitatory populations. (B) Inhibitory populations.

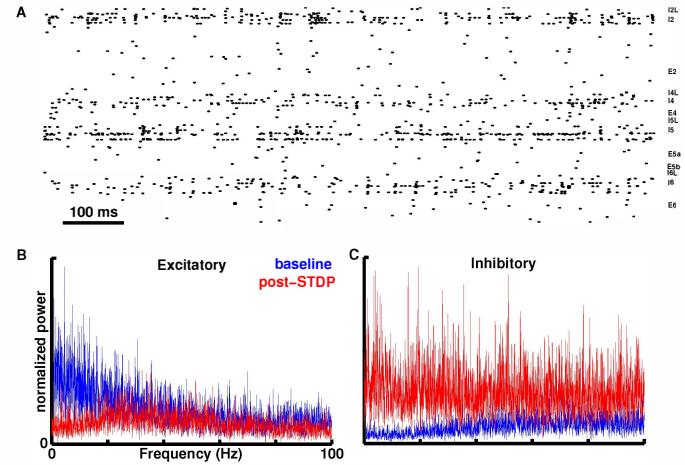


Fig. 4. Network activity after 300 s of STDP with no training signal. A) raster. B) Excitatory and C) Inhibitory MUA power spectra during baseline (blue) and after STDP (red).

We therefore applied STDP to AMPA synapses at both E→E and from E→I projections. After applying STDP for 300 s with only white noise (no training signal), network activity was substantially altered (Fig. 4). As shown in Fig. 3, I cell average firing rates increased for most populations (NB: high-rate rows in Fig. 4A correspond to I cells). As a result of higher I cell firing, average E cell firing rates decreased for all populations (Fig. 3). The most marked decrease was in infragranular layers. Although there was sparse E cell firing, there was still synchrony present between different cells, as seen in the occasional vertical stripe in the E2 population. Fig. 4C confirms dominant I dynamics, as seen with the large, broad-band increase in I MUA power over baseline. This increase in I MUA power attenuated a broad range of low-frequency oscillations, including the peak near 6 Hz in the E MUA spectral power (Fig. 4B) that was evident at baseline. The new E MUA spectral peak was around 30 Hz. Interestingly, the reduction in E activity via inhibition tended to increase the E population's peak frequency.

C. Balanced STDP with training signal enhances oscillations

We next applied an 8 Hz training signal to E4 cells to simulate neuroprosthetic stimulation of thalamic afferents. Each E4 cell was activated every 125 ms, with a jitter of ±

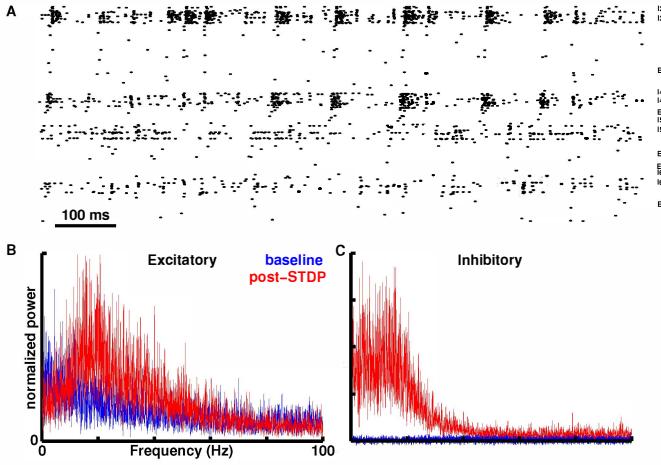


Fig. 5. Network activity after 300 s of STDP with an 8 Hz training signal applied to E4 cells. A) raster. B) Excitatory and C) Inhibitory MUA power spectra during baseline (blue) and after training (red). Note that the vertical axes in B and C use normalized scales.

0.5 ms to avoid coincident (epileptogenic) activation of large numbers of cells. Network activity gradually altered during the 300 seconds of STDP with training signal application. Resultant network dynamics after turning off the training signal and STDP are shown in Fig. 5. Note that the activity spectrum has been permanently altered, and no longer depends on the periodic drive of the training signal (Fig. 5B,C).

Due to the synchronous inputs of the training signal arriving in layer 4, E4 activity approximately doubled during training (0.8→1.5 Hz). This shifted the balance for activity coming out of layer 4, which was now dominated by excitation (Fig. 3 red bars). Training then altered the synaptic strengths from these driven E4 cells to other cell populations. Enhanced excitation could then spread from these populations throughout the network. On average, E4→E4 and E4→E2 connections were strengthened to 1.4 and 1.5 times baseline strength, respectively. Firing rates of E2 cells also increased after training, while other E populations decreased slightly, decreases that were caused by the increased firing of I cells. In addition to firing rate changes, synchrony between E cells increased.

Fig. 5B shows substantial enhancements of the E MUA spectral power over baseline. The peak spectral power shifted from the theta range at ~6 Hz to a broad range over 10-20 Hz. There are additional broad increases in the gamma band between 30-100 Hz, including sharp increases at 40 and 50 Hz. These increases appeared to reflect synchronizing effects due to the enhanced inhibition in the network[11], [20]. The I MUA power spectrum shows larger increases in spectral power (Fig. 5C). Instead of the broad increase seen in Fig. 4C, the enhanced activation of E cell activity helped focus the I MUA spectral peak in the low-frequency band. However, the I MUA spectral power also increased substantially over baseline across full range.

Training at other frequencies from 6-10 Hz resulted in output frequencies comparable to Fig. 5. Lower frequencies

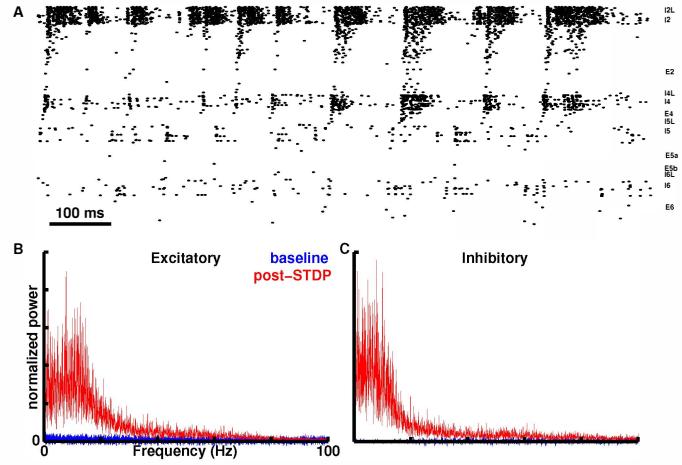


Fig. 6. Network activity after 300 s of STDP with a 16 Hz training signal applied to E4 cells. A) raster. B) Excitatory and C) Inhibitory MUA power spectra during baseline (blue) and after training (red). Note that the vertical axes in B and C use normalized scales.

were not as effective, giving results comparable to those of the control (Fig. 4).

D. Overtraining induces epileptiform oscillations

Increasing the rate of the training signal above 10 Hz further pushed the dynamic balance towards excitation. Following training with a 16 Hz signal, high amplitude (epileptiform) activity was seen, with a sharper spectral band (Fig. 6). In the raster plot an ~8 Hz seizure is evident. The extra drive to E4 cells spread throughout the network and produced pathological excitation for the E and I cells. Average E population firing rates increased substantially, including a greater than four-fold increase in E2 rates (Fig. 3). This produced added excitation that also increased I4 and I2 cell firing rates, in some cases by more than a factor of 10 due to prolonged bursting. The added excitation and inhibition in the network created synchronous population spikes that spread mainly between the E and I cells of layers 2 and 4. Almost all E and I cells in layers 2 and 4 were active during the population spikes, but the population spikes also included E and I cells in layers 5 and 6. E cells were largely silent between population spikes. I cell activity was not confined solely to the population bursts. The E2 activity during population spikes ranged from narrow (10 ms) to wide (75 ms) while the I2 activity typically lasted longer (approximately 100 ms). E4 participation in the spikes was typically shorter. The longer duration of layer 2 activity in the bursts is likely due to enhanced recurrent connectivity in the E2 cells.

The population spikes activated at ~ 8 Hz are seen in both the E and I MUA spectra (Fig. 6B,C) as sharp peaks. Although increases were seen in the low frequency bands, there were also large increases in the full frequency range.

IV. DISCUSSION

Our baseline neocortical network demonstrated oscillations in a physiological frequency range, but these oscillations had

relatively low amplitudes. We utilized STDP and training signals to both alter the amplitude and shift the peak frequency of oscillations produced by the network, demonstrating the feasibility of training a neuronal network's oscillatory dynamics. We found that in order to avoid epileptic dynamics, a balance was needed between learning at E→E and E→I AMPA synapses. Depending on this balance, the network could be shifted between the extremes of silence and seizures. We predict that in order for learning in the brain to avoid transition to epilepsy, homeostatic metaplasticity control mechanisms exist to balance these two types of learning.

An open research question is what kinds of problems can be solved using neuronal networks with complex, multi-layered architectures and multiple cell types. In our neocortical network with a moderate number of cells, the input-to-output transfer functions were not easy to decipher. Further complications arose from the positive and negative feedback loops caused by interactions between excitatory and inhibitory populations. Future research will need to ascertain what the limits are of learning in realistic neuronal networks, and the role each circuit element contributes to the information processing. The manifold dynamical interplay arising from complex wiring and intrinsic neuronal dynamics is then made still more complex as synaptic strengths constantly change. In light of this, it is perhaps surprising that memories, which are thought to be stored primarily as connection strengths between neurons, can be so stable.

We predict that a neural prosthetic could potentially restore lost function to a damaged neocortical network, using an appropriate signal as a stimulus to the network. This could be given electrically subcortically or could potentially be applied noninvasively by repetitive external stimulation through the sensory pathway. Not surprisingly, excessive stimulation could lead to epilepsy, something sometimes seen with either electrical stimulation or, in responsive individuals, with strobe light stimulation.[31], [32] Such over-stimulation could cause irreversible damage by allowing the brain to "learn" epileptic dynamics. Embedding of patterned activity for retraining after brain damage would need to be carefully controlled to avoid this possibility. Similarly, while benefits of neural-prosthetic repair could be enhanced through use of neuromodulatory drugs, or activation of intrinsic neuromodulators such as acetylcholine or dopamine that promote learning, these drugs would also carry risk depending on the relative influence on E→E and E→I synapses.

ACKNOWLEDGMENT

Research supported by DARPA grant N66001-10-C-2008.

The authors would like to thank Larry Eberle and Amy Delman (SUNY Downstate) for Neurosim lab support and administration; Michael Hines and Ted Carnevale (Yale) for ongoing NEURON support; and Tom Morse (Yale) for ModelDB support.

REFERENCES

- [1] J Arthur and K Boahen. Learning in silicon: Timing is everything. *Advances in neural information processing systems*, 18:75, 2006.
- [2] K Boahen. Neurogrid: emulating a million neurons in the cortex. In *Conf Proc IEEE Eng Med Biol Soc*, 2006.
- [3] H Jaeger and H Haas. Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication. *Science*, 304:78–80, 2004.
- [4] GE Hinton and RR Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006.
- [5] JJ Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proc Nat Acad Sci*, 79(8):2554–2558, 1982.
- [6] G Buzsaki and A Draguhn. Neuronal oscillations in cortical networks. *Science*, 304(5679):1926–1929, 2004.
- [7] C Gray and W Singer. Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proc Nat Acad Sci*, 86:1698–1702, 1989.
- [8] N Kopell, GB Ermentrout, MA Whittington, and RD Traub. Gamma rhythms and beta rhythms have different synchronization properties. *Proc Natl Acad Sci U S A*, 97:1867–1872, 2000.
- [9] J Lisman and G Buzsaki. A neural coding scheme formed by the combined function of gamma and theta oscillations. *Schizophr Bull*, 34(5):974–980, 2008.
- [10] C Von der Malsburg and W Schneider. A neural cocktail-party processor. *Biol Cybern*, 54:29–40, 1986.
- [11] SA Neymotin, MT Lazarewicz, M Sherif, D Contreras, LH Finkel, and Lytton WW. Ketamine disrupts theta modulation of gamma in a computer model of hippocampus. *J Neurosci*, 31(32):11733–11743, 2011.
- [12] KM Spencer, PG Nestor, MA Niznikiewicz, DF Salisbury, ME Shenton, and RW Carley. Abnormal neural synchrony in schizophrenia. *J Neurosci*, 23:7407–7411, 2003.
- [13] PJ Uhlhaas, DE Linden, W Singer, C Haenschel, M Lindner, K Maurer, and E Rodriguez. Dysfunctional long-range coordination of neural activity during gestalt perception in schizophrenia. *J Neurosci*, 26:8168–8175, 2006.
- [14] Y Dan and M Poo. Spike timing-dependent plasticity of neural circuits. *Neuron*, 44(1):23–30, 2004.
- [15] PD Roberts and CC Bell. Spike timing dependent synaptic plasticity in biological systems. *Biol Cybern*, 87(5):392–403, 2002.
- [16] S Song, KD Miller, and LF Abbott. Competitive hebbian learning through spike-timing-dependent synaptic plasticity. *Nat Neurosci*, 3:919–926, 2000.
- [17] T Binzegger, RJ Douglas, and KA Martin. A quantitative map of the circuit of cat primary visual cortex. *J Neurosci*, 24:8441–8453, 2004.
- [18] CC Kerr, SA Neymotin, GL Chadderdon, CT Fietkiewicz, JT Francis, and WW Lytton. Electrostimulation as a prosthesis for repair of information flow in a computer model of neocortex. *IEEE Transactions on Neural Systems & Rehabilitation Engineering*, Submitted March 28, 2011.
- [19] S Lefort, C Tomm, JC Floyd-Sarria, and CC Petersen. The excitatory neuronal network of the C2 barrel column in mouse primary somatosensory cortex. *Neuron*, 61:301–316, 2009.
- [20] SA Neymotin, H Lee, E Park, AA Fenton, and WW Lytton. Emergence of physiological oscillation frequencies in a computer model of neocortex. *Front Comput Neurosci*, 5:19, 2011b.
- [21] SA Neymotin, KM Jacobs, AA Fenton, and WW Lytton. Synaptic information transfer in computer models of neocortical columns. *J Comput Neurosci*, 30(1):69–84, 2011a.
- [22] RD Traub, D Contreras, MO Cunningham, H Murray, FE LeBeau, A Roopun, A Bibbig, WB Wilent, MJ Higley, and MA Whittington. Single-column thalamocortical network model exhibiting gamma oscillations, sleep spindles, and epileptogenic bursts. *J Neurophysiol*, 93:2194–2232, 2005.
- [23] BE Peterson, MD Healy, PM Nadkarni, PL Miller, and GM Shepherd. ModelDB: an environment for running and storing computational models and their results applied to neuroscience. *J Am Med Inform Assoc.*, 3(6):389–398, 1996.
- [24] WW Lytton and M Stewart. A rule-based firing model for neural networks. *Int J Bioelectromagnetism*, 7:47–50, 2005.
- [25] WW Lytton and M Stewart. Rule-based firing for network simulations. *Neurocomputing*, 69(10–12):1160–1164, 2006.
- [26] WW Lytton, A Omurtag, SA Neymotin, and ML Hines. Just-in-time connectivity for large spiking networks. *Neural Comput*, 20(11):2745–2756, 2008.

- [27] CE Jahr and CF Stevens. A quantitative description of NMDA receptor-channel kinetic behavior. *J Neurosci*, 10(6):1830, 1990.
- [28] CE Jahr and CF Stevens. Voltage dependence of NMDA-activated macroscopic conductances predicted by single-channel kinetics. *J Neurosci*, 10(9):3178–3182, 1990.
- [29] NT Carnevale and ML Hines. *The NEURON Book*. Cambridge University Press, New York, 2006.
- [30] WW Lytton. Neural query system: data-mining from within the NEURON simulator. *Neuroinformatics*, 4:163–176, 2006.
- [31] S Neymotin, DJ Uhlrich, KA Manning, and WW Lytton. Data mining of time-domain features from neural extracellular field data. *Studies in Computational Intelligence*, 151:119–140, 2008.
- [32] DJ Uhlrich, KA Manning, ML Laughlin, and WW Lytton. Photic-induced sensitization: acquisition of an augmenting spike-wave response in the adult rat through repeated strobe exposure. *Journal of Neurophysiology*, 94:3925–3937, 2005.