Biologically Plausible Models of Homeostasis and STDP: Stability and Learning in Spiking Neural Networks

Kristofor D. Carlson, Micah Richert, Nikil Dutt, Jeffrey L. Krichmar

***Abstract*—Spiking neural network (SNN) simulations with spike-timing dependent plasticity (STDP) often experience runaway synaptic dynamics and require some sort of regulatory mechanism to stay within a stable operating regime. Previous homeostatic models have used L_1 or L_2 normalization to scale the synaptic weights but the biophysical mechanisms underlying these processes remain undiscovered. We propose a model for homeostatic synaptic scaling that modifies synaptic weights in a multiplicative manner based on the average postsynaptic firing rate as observed in experiments. The homeostatic mechanism was implemented with STDP in conductance-based SNNs with Izhikevich-type neurons. In the first set of simulations, homeostatic synaptic scaling stabilized weight changes in STDP and prevented runaway dynamics in simple SNNs. During the second set of simulations, homeostatic synaptic scaling was found to be necessary for the unsupervised learning of V1 simple cell receptive fields in response to patterned inputs. STDP, in combination with homeostatic synaptic scaling, was shown to be mathematically equivalent to non-negative matrix factorization (NNMF) and the stability of the homeostatic update rule was proven. The homeostatic model presented here is novel, biologically plausible, and capable of unsupervised learning of patterned inputs, which has been a significant challenge for SNNs with STDP.**

I. Introduction

LEARNING plays a central role in the functioning of neural systems. Spike-timing dependent plasticity (STDP), a form of Hebbian learning where weight changes depend on the relative timing between pre and postsynaptic action potentials, has been studied extensively both experimentally [1] and theoretically [2], [3]. Computational models have used STDP to perform a variety of biologically relevant tasks that include forming sparse representations of temporal sequences [4], refining neuronal connectivity [5], and computing with neural synchrony [6]. However, STDP can exhibit runaway synaptic dynamics when a presynaptic neuron consistently excites a postsynaptic neuron to fire, increasing the synaptic strength at every subsequent pairing. This increase in synaptic strength causes the postsynaptic neuron to fire more rapidly, strengthening additional synapses and eventually resulting in a boundless increase of all presynaptic connections to the postsynaptic neuron. Neurons observed in experiments utilize a number of biophysical mechanisms to counteract this positive feedback loop and impose stability on the system.

Neurons have several homeostatic mechanisms to avoid runaway potentiation, but the two most widely studied are intrinsic plasticity and synaptic scaling [7], [8]. Intrinsic plasticity consists of the activity-dependent modification of non-synaptic ion channels, such as sodium and delayed-rectifier potassium channels, to adjust the excitability of the neuron [8]. Synaptic scaling, on the other hand, is the activity-dependent regulation of synaptic efficacy to prevent runaway increases or decreases in synaptic strength [7], [9], [10]. Although other homeostatic mechanisms such as the redistribution of synaptic efficacy [11] and balanced STDP integration windows [12], [13] may also play a role in stabilizing STDP, we focus on implementing a model describing synaptic scaling as recent theoretical work has shown synaptic scaling can stabilize a variety of different learning rules [14], including STDP.

Synaptic scaling can be experimentally observed by measuring the size of excitatory synaptic currents before and after pharmacologically increasing or decreasing the global activity of a cultured network of excitatory and inhibitory neurons [10], [15]. A decrease in the global activity leads to a scaling up of excitatory synaptic currents while an increase in the global activity leads to a scaling down of excitatory synaptic currents. The timescale over which this scaling takes place varies from hours to days [10] but has been observed in as little as four hours after induction protocols [16]. Synaptic scaling has been observed in neocortical neurons, hippocampal neurons, and spinal neurons in vitro and in vivo [10], [15], [17], [18]. Although there is experimental evidence showing presynaptic involvement in homeostatic synaptic scaling [17], [19], we focus on postsynaptic expression of synaptic scaling and acknowledge that neurons probably express synaptic scaling via multiple mechanisms. There is also evidence for two distinct forms of homeostatic synaptic plasticity, namely local synaptic scaling and global synaptic scaling. Local

Manuscript received March 1, 2013. This work was supported by the Defense Advanced Research Projects Agency (DARPA) subcontract 801888-BS.

K. D. Carlson is with the Department of Cognitive Sciences, University of California, Irvine CA 92697 USA (phone: 765-418-2594; fax: 949-824-2307; email: kdcarlso@uci.edu).

M. Richert was with the University of California, Irvine CA 92697 USA. He is now with Brain Corporation, QRC-130, 5665 Morehouse Drive, San Diego CA 92121 USA (email: richert@braincorporation.com).

N. Dutt is with the Department of Computer Science, University of California, Irvine CA 92697 USA (email: dutt@ics.uci.edu).

J. L. Krichmar is with the Department of Cognitive Sciences and the Department of Computer Science, University of California, Irvine CA 92697 USA (email: jkrichma@uci.edu).

978-1-4673-6129-3/13/$31.00 ©2013 IEEE

synaptic scaling occurs at the individual synapse level [20], [21] while global synaptic scaling induces synaptic changes across all synapses at a particular neuron [10], [16], [22]. Global synaptic scaling preserves the relative differences in strength between synapses on a particular neuron, an important feature as this is how memories are thought to be encoded in the brain [23]. We consider only postsynaptic global synaptic scaling in this study.

Multiple theoretical studies have examined the interaction between homeostatic mechanisms and STDP in SNNs [24], [25], [27]. These studies have included homeostatic features such as intrinsic plasticity, synaptic normalization, structural plasticity via the addition or removal of new synapses, and new forms of STDP that promote a balance between excitatory and inhibitory activity [24], [25], [26]. In each of these models, synaptic scaling is implemented via L_1 or L_2 normalization of the synaptic weights. This may not be biologically plausible because L_1 or L_2 normalization requires each synapse to have information about the strength of all other synapses on a particular neuron. There are currently no known biological mechanisms that would allow for this information to be shared among synapses.

A more biologically accurate version of synaptic scaling has been implemented in previous works [27], [28] by using multiplicative synaptic plasticity across all synapses at a particular neuron. In these models, weights were scaled in proportion to the difference between a time-averaged postsynaptic activity signal and a target postsynaptic activity signal. The synaptic scaling model presented here is qualitatively similar to the model in [27], with non-trivial differences, but significantly different from the model presented in [28], which included a term that measured the time-averaged presynaptic activity as well.

We introduce a novel, biologically plausible mathematical description of a homeostatic synaptic scaling mechanism and show that it prevents runaway synaptic dynamics via computational experiments and mathematical analysis. The first computational experiment shows the homeostatic synaptic plasticity rule stabilizes learning with STDP while the second experiment shows a network with homeostatic plasticity and STDP that is capable of learning V1-like receptive fields from patterned visual inputs via self-organization. Additionally, the homeostatic update rule is proven to be stable and, when paired in combination with our STDP implementation, is shown to be a form of non-negative matrix factorization (NNMF). The homeostatic update rule follows experimental observations in that it adjusts the synaptic weights multiplicatively [10] and depends on both the average firing rate of the neuron and some target firing rate [29], [30].

II. Methods

To examine the computational properties of our homeostatic synaptic plasticity update rule we ran two spiking network simulations. The first simulation, which is known as the ramp test, examined the ability of STDP in a spiking network to encode graded inputs by comparing networks with and without homeostatic synaptic scaling implemented. In the second self-organizing receptive field (SORF) simulation, a network with topology that modeled the path visual input takes from the lateral geniculate nucleus (LGN) to the primary visual cortex (V1) was presented with counter-phase gratings at varying orientations. The ability of the network to self-organize and learn V1-like receptive fields was studied by again comparing networks with and without the homeostatic synaptic scaling mechanism.

The STDP update rule used in our simulations is shown in (1).

$$\frac{dw_{i,j}}{dt} = \delta + \beta\left(LTP_{i,j} + LTD_{i,j}\right) \tag{1}$$

Here, $w_{i,j}$ indicated the synaptic weight from presynaptic neuron i to postsynaptic neuron j. Additionally, δ was a bias often set to zero or a positive number to push the network towards positive weight increases for low synaptic input, while β is the learning rate. The STDP long-term potentiation (LTP) and long-term depression (LTD) magnitudes were calculated using the standard nearest neighbor STDP implementation [31]. The synaptic weight change was not proportional to the current weight value, making this STDP update rule additive in contrast to multiplicative update rule approaches. The weight changes were updated every time step (1 ms) but the weights themselves were modified once every 1 second.

To model homeostatic synaptic plasticity the STDP update rule was modified as shown in (2).

$$\frac{dw_{i,j}}{dt} = \left[\alpha \cdot w_{i,j}\left(1 - \frac{\bar{R}}{R_{\text{target}}}\right) + \beta\left(LTP_{i,j} + LTD_{i,j}\right)\right] \cdot K \tag{2}$$

Here, α was the homeostatic scaling factor, $\bar{R}$ was the average firing rate of the postsynaptic neuron, j, and R_{target} was the target firing rate for the postsynaptic neuron, j. A stability term denoted by K, damps oscillations in the weight updates and speeds up learning. K was defined as:

$$K = \frac{\bar{R}}{T \cdot \left(1 + \left|1 - \bar{R}/R_{target}\right| \cdot \gamma\right)} \tag{3}$$

In (3), γ is a tuning factor and T is the time scale over which the firing rate of the postsynaptic neuron was averaged. The STDP and homeostasis parameters used in the ramp test and SORF simulations are found in Appendix A. It is worth noting that α and β must be chosen so as to suppress runaway synaptic dynamics but still allow learning to proceed at a reasonable rate. As with the STDP rule in (1), the change in weights was updated every time step (1 ms), but the weight itself was only updated once per second.

The computational experiments consisted of SNNs with conductance based, Izhikevich-type spiking neurons [32] and were implemented on a publically available neural simulator [33]. The simulations were run using graphics

processing units (GPUs) on a 12-core computing cluster with 4 NVIDIA Tesla GPU M2090 cards, each with 6 GB of global memory and 512 cores. Excitatory Izhikevich neurons were regular spiking (RS) neurons while inhibitory neurons were fast spiking (FS) neurons, as described in [32]. Both sets of simulations used delays of 1 ms for all synaptic connections.

A. Ramp Simulation

The first network configuration consisted of 100 Poisson spiking neurons, each with a STDP-enabled plastic synaptic connection to a single RS output neuron. Each Poisson neuron had a firing rate that ranged from 0.2 Hz to 20 Hz following ascending neuron id order in increments of 0.2 Hz (e.g. Neuron 1 had a firing rate of 0.2 Hz while Neuron 100 had a firing rate of 20 Hz). The initial synaptic weights were chosen from a uniform distribution between 0.01 and 0.03. Two different network configurations were run for 1000 simulation seconds and both the synaptic weights and firing rates were monitored. The first network configuration did not include homeostasis and used plasticity update rule (1) while the second network configuration included homeostasis and used plasticity update rule (2). When homeostasis was enabled, the target firing rate of the output neuron was 35 Hz. Additional STDP and homeostatic parameters used in the ramp simulation are found in Table I of Appendix A.

B. Self-Organizing Receptive Field (SORF) Simulation

The second set of simulations examined the ability of a spiking network with STDP and homeostasis to self-organize in response to the presentation of varying angles of grayscale, counterphase grating. Two configurations were simulated, one with homeostasis enabled and one without homeostasis enabled. The topology of the network is shown in Fig. 1.

A pair of 32 by 32 Poisson spike generator arrays received a 32 by 32 gray scale pixel presentation stimulus with one-to-one connectivity. Neurons in one Poisson group, the OnPoiss group, had linear spiking responses to the input with a maximum firing rate for white pixels and no firing rate response to pixels less bright than the mean gray value, while neurons in the OffPoiss group had a maximum firing rate for black pixels and no firing rate response to pixels more bright than the mean gray value. The maximum response of the On(Off)Poiss groups was chosen to be 28 Hz. Each of the On(Off)Poiss spike groups had one-to-one excitatory connections with On(Off)Buffer groups consisting of RS neurons. When spike trains generated by Poisson neurons were filtered through RS neurons, which had a refractory period, the resulting spike trains were more biologically realistic. Neurons in the On(Off)Buffer groups had plastic excitatory synapses with full connectivity to four Exc group RS neurons. To induce competition, the Exc group neurons had plastic excitatory synapses with full connectivity to an inhibitory (Inh) group of four FS neurons, which in turn had inhibitory feedback with full connectivity back to the Exc group neurons. The On(Off)Poiss→On(Off)Buffer group weights were randomly initialized in the range of 0.200 to 0.600, the On(Off)Buffer→Exc group weights were randomly initialized in the range of 0.004 to 0.015, the Exc→Inh group weights were randomly initialized in the range of 0.116 to 0.463, and the Inh→Exc group weights were randomly initialized in the range of 0.065 to 0.259. These weight ranges, excluding the On(Off)Poiss→On(Off)Buffer weight ranges, were found using an optimization technique called evolutionary algorithms (EA), which is discussed next.

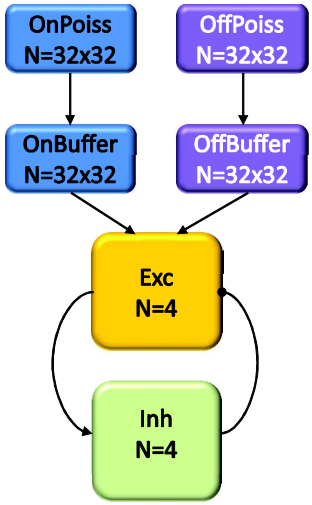

Fig. 1. Topology of the spiking network used in the SORF simulation. Lines with arrowhead endpoints indicate excitatory connections while lines with circle endpoints indicate inhibitory connections. Arrow notations (→) in text denote a connection from one group (left-hand side) to another group (right-hand side). Weights from the On(Off)Buffer→Exc groups were plastic as were the weights from the Exc→Inh groups. Different STDP parameter values were used for the On(Off)Buffer→Exc synapses than for the Exc→Inh synapses and are shown in Table II of Appendix A.

Hand tuning neural networks is a difficult task, for instance obtaining persistent firing activity in recurrent networks requires precise tuning of the connection weights [34]. Additionally, a large range of biologically plausible parameter values often exist [35] and it is unclear which combination of these parameters will give the desired network behavior. A variety of automated parameter tuning techniques have been developed for use in parameter tuning neural networks [36]; in this work evolutionary algorithms were used to tune 14 network parameters. An EA computational library called Evolving Objects [37] was modified and integrated into the spiking neural network simulator to find suitable values for the synaptic weights, the homeostatic parameters, the STDP parameters, and the maximum firing rate for the On(Off) Poisson spiking neurons. The EA ranked parameter fitness using firing rate responses of the four output neurons to the training and testing tasks. Once suitable parameters were found, they were used for the entirety of the SORF simulation.

The simulations first consisted of a training phase, where 40 counterphase grating orientations (from $\pi/40$ to π) were randomly presented for 2 seconds with 500 ms of 1 Hz random input noise between pattern presentations for

approximately 100 minutes. During the testing phase, all plasticity mechanisms were disabled and the network was then presented with all 40 test orientations while the firing rates of the output neurons and the synaptic weights were observed. The training and testing protocols were done for two different network configurations, one with homeostasis and one without homeostasis. Additional STDP and homeostatic parameters used in the SORF simulations are found in Table II of Appendix A.

III. Results

A. Ramp Simulation

The results for the ramp simulation are shown in Fig. 2 and Fig. 3. In Fig. 2, the synaptic strength of each synapse is shown as a function of the Poisson spiking neuron ID at simulation time t=1000 seconds. For networks with homeostatic synaptic scaling (blue x's), the synaptic weights roughly track the Poisson neuron spike rates with lower weights for neurons with lower firing rates and higher weights for neurons with higher firing rates. For networks without homeostatic synaptic scaling (red points), the synaptic weights saturate completely, showing runaway potentiation. Note that at the end of the ramp test simulation without homeostasis, even the synapse experiencing the lowest firing rate saturates.

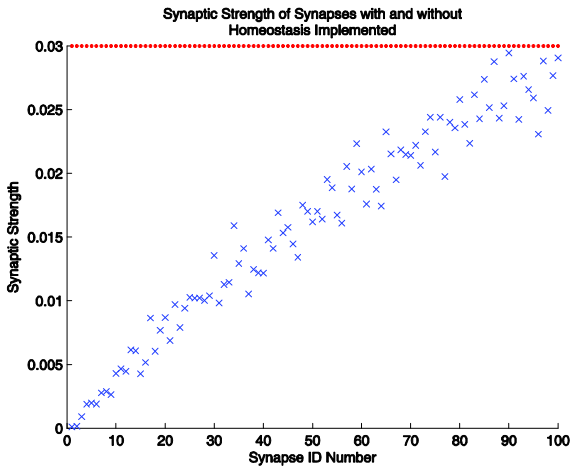

Fig. 2. Plot of synaptic strength vs. synapse ID number at simulation time t=1000 seconds. The blue x's represent synapses that underwent STDP learning *with* homeostatic synaptic scaling enabled while the red points represent synapses that underwent STDP learning *without* homeostatic synaptic scaling.

Fig. 3 shows the firing rate of the output neuron versus time with homeostasis enabled (blue line) and without homeostasis enabled (red line). The output firing rate stays near its target rate of 35 Hz with homeostasis but reaches 55 Hz in the case where homeostasis is excluded. This simple simulation shows that the homeostasis update rule can stabilize both synaptic dynamics and firing rates while at the same time allowing STDP learning to occur. It may be possible to produce stable learning via balanced STDP in the ramp simulation with carefully chosen STDP parameters [13], [38]. However balanced STDP may still require STDP integration windows to change over time or synaptic scaling to rescue synapses too weak to cause postsynaptic activations, both of which are homeostatic processes.

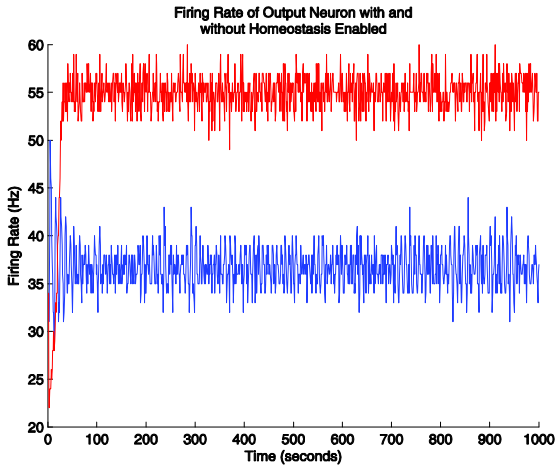

Fig. 3. Plot of the output neuron firing rate as a function of simulation time for STDP learning with (blue line) and without (red line) homeostatic synaptic scaling.

B. SORF Simulation

The ability of homeostatic synaptic scaling to stabilize STDP for a more complex network topology with patterned stimulus is examined next. Experiments show simple cells in visual cortical area V1 have Gaussian tuning curves [39] in response to visual bar stimuli [40]. Fig. 4 shows the combined synaptic weights for the On(Off)Buffer group neurons and the four Exc neurons in a representative simulation trial. In (a), we see the synaptic weights for a network with randomly initialized rates before training occurred while in (b) we see the synaptic weights for a network trained without homeostasis. While the network trained without homeostasis in (b) does exhibit some pattern formation, the weights do not resemble Gabor functions. In (c) we see synaptic weights from a network trained with homeostasis, indicating the network has successfully self-organized to form Gabor-like spatial receptive fields. It should be noted that networks using STDP learning without homeostasis (b) never resulted in tuning curves, even with extensive parameter optimization using EAs. All SNNs in the SORF simulations required both STDP learning and homeostasis for the development of self-organizing V1 receptive fields. It is possible that SNNs with different connection topologies or neuron types may be able to self-organize V1 receptive fields with only STDP learning.

Fig. 5 shows the firing rate of the output neurons to all 40 orientations of the counterphase grating for networks with homeostasis. The blue points show the firing rate of the output neuron for a particular orientation of stimulus presentation while the dashed black line shows an idealized Gaussian tuning curve. The neuronal responses from networks trained with homeostasis qualitatively match the idealized tuning curve while those networks without homeostasis had untuned neuronal responses.

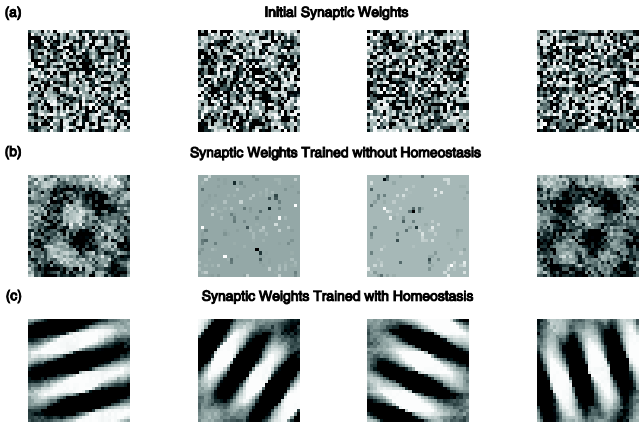

Fig. 4. Visualization of the synaptic strengths between the On(Off)Buffer groups and the 4 Exc neurons. White areas denote higher weight values (with a maximum value of 0.06) while dark areas denote low weight values (with a minimum value of 0). **(a)** The top row shows the initial synaptic weights before training occurred. **(b)** The middle row shows the synaptic weights trained on a network using STDP learning *without* homeostatic synaptic scaling. **(c)** The bottom row shows the synaptic weights trained on a network using STDP learning *with* homeostatic synaptic scaling.

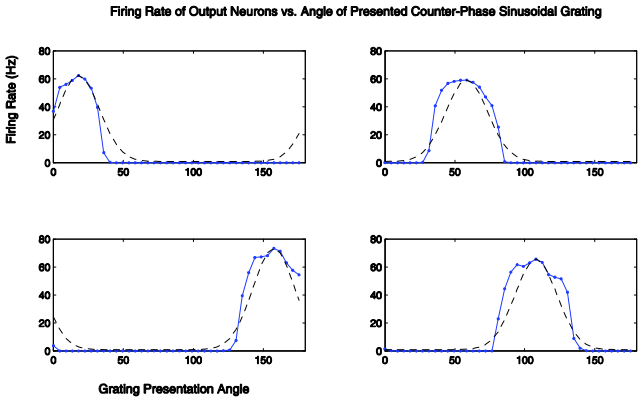

Fig. 5. Plot of the firing rate of the four Exc neurons vs. the angle of presented counter-phase grating for a network trained using STDP with homeostasis enabled. The blue points represent the firing rate for a particular neuron at a given presentation stimulus angle. The dashed black line represents an ideal Gaussian tuning curve firing rate response.

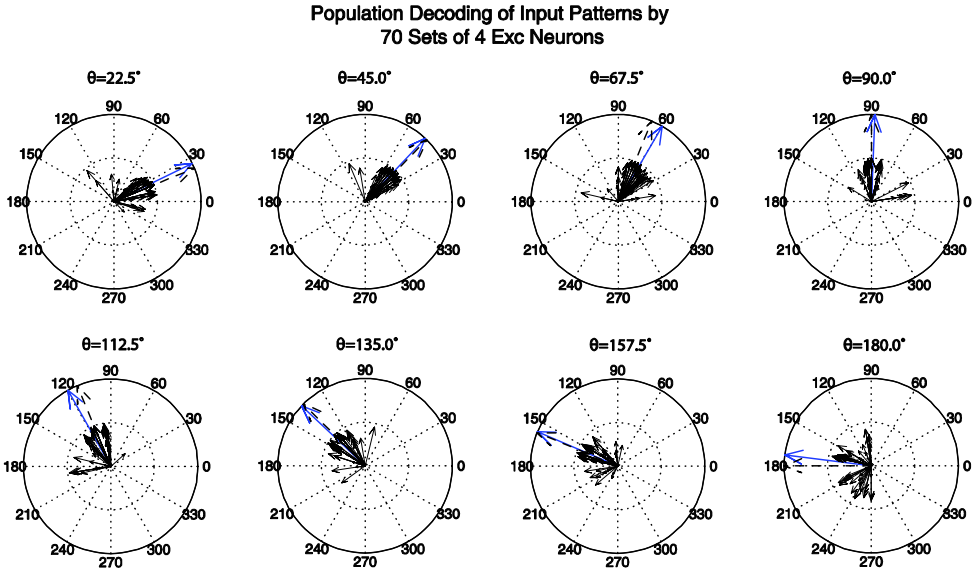

Fig. 6. Population decoding for eight different presentation angles (θ) calculated using the firing rate of the four Exc neurons. The stimulus responses were summed over 70 individual network runs presented with the same input patterns in random order. Each network had four Exc neurons with unique tuning curves because each network was presented with a different random order of input patterns. The blue arrow is the decoded population resultant vector (normalized), the dashed black arrow is the correct idealized vector (normalized), and the smaller, solid black arrows are the component vectors, divided by a factor of 2 for display purposes.

Fig. 6 shows the population decoding of 8 test angles with 70 sets of 4 Exc neurons using a common population decoding algorithm [41]. When the network is presented with a grating orientation, the firing rate and preferred stimulus orientation are combined to create a vector component for every neuron (shown with solid black arrows). These components are then summed over all neurons (280 in this case) to create a population vector (shown with a blue arrow) and compared with the grating presentation vector (shown with a dashed arrow). The mean error of the decoded population vectors was 3.8° with a standard deviation of 2.9°.

The results found in Fig. 4, Fig. 5, and Fig. 6 were tested for repeatability and found to be very robust. Over 70 runs of the simulation were performed both with and without homeostasis, each with a different random stimulus presentation order. In runs where homeostasis was included, both V1-like receptive fields and Gaussian tuning curves always formed, but rarely with the same preferred orientations. For networks with homeostasis, the preferred tuning curve orientations shown in Fig. 5 were randomly distributed uniformly from 0 to π when sampled over many runs. SORF simulations that did not include homeostasis had untuned neural responses and were unable decode presentation angles in the population decoding task.

IV. Discussion

The model of homeostatic synaptic scaling presented here is a phenomenological description of many biophysical processes working together at the cellular and molecular levels. Neurons most likely integrate calcium signals, perhaps to measure spiking activity, which activate gene expression in the soma leading to subsequent changes at all synapses for a particular neuron [7], [29]. Synaptic weights are scaled multiplicatively, preserving the relative weight differences between synapses located on the same neuron [10], [22]. The homeostatic update rule described here integrates both of these observations to provide a biologically plausible mechanism that stabilizes STDP in spiking neural networks.

It may be possible to implement the BCM learning rule using the homeostasis framework presented here if additional modifications are included. The BCM rule attains stability with a sliding LTP/LTD threshold that depends on average postsynaptic activity [42]. Homeostatic intrinsic excitability involves changing the excitability of the target neuron through the modification of non-synaptic currents, and may provide a mechanism for implementing a BCM-style threshold. This could be incorporated into the current homeostasis framework by adding non-synaptic ion channels subject to activity-dependent conductance changes or by modifying Izhikevich parameters of the neuron as a function of neuronal firing rates. There is also evidence that activity-dependent regulation of NMDAR subunit ratios NR2A and NR2B may play a role in homeostatic regulation [43]. Altering the NR2A/B ratio can lead to changes in calcium influx through the NMDAR which, in turn, could affect the BCM threshold as well. It is possible to model this phenomenon by including a description of the NR2A/B ratio in an NMDA conductance and allowing the ratio to change via a synaptic scaling mechanism.

Our homeostasis rule has mathematical properties that allow it be compared to other well-known algorithms. Note that the homeostatic synaptic scaling rule is proven to be stable, and when combined with the STDP update rule, can be shown to be a form of NNMF (see Appendix B and C for proofs). Furthermore, Oja's rule is analogous to our STDP/homeostasis rule in the following manner. Recall that Oja's rule results from modifying Hebb's rule in an effort to stabilize the learning dynamics just as our homeostasis/STDP rule results from applying homeostatic synaptic scaling to the STDP rule. In addition, both Oja's rule and our homeostasis/STDP rule can be proven to be stable [44]. Finally, Oja's rule has been shown to perform principal component analysis (PCA) [44] while our homeostasis/STDP rule has been shown to perform NNMF. PCA and NNMF both perform the same basic mathematical operation: the factorization of a data matrix into the inner product of two matrices, where one matrix contains basis vectors that optimally describe the data as a linear approximation and the other matrix contains the weighting components for this linear approximation. These algorithms differ only in the additional constraints they place on the system.

V. Conclusion

We presented a biologically plausible homeostatic synaptic scaling model capable of stabilizing STDP learning that follows two important experimental observations concerning synaptic scaling. First, synaptic scaling is driven by somatic calcium signaling, which may be tracking the average firing rate of the neuron. Second, synaptic scaling is multiplicative and preserves the difference in weight strengths between synapses on the same postsynaptic neuron.

Simulations showed the homeostatic synaptic scaling model stabilized the firing rate and synaptic weights of a simple SNN with STDP. Additional simulations demonstrated that homeostasis was necessary for SNNs with STDP and homeostasis to successfully self-organize into V1-like receptive fields and Gaussian tuning curves in response to randomly oriented counterphase grating stimuli. The homeostatic synaptic scaling model was proven to be stable and, when used in conjunction with STDP, was shown to be a form of NNMF. The model succeeds in introducing a biologically plausible synaptic scaling mechanism to the SNN domain and can be used as a stepping stone for building models that more closely integrate STDP and homeostasis.

Appendix

A. STDP and Homeostasis Parameters

Table I shows STDP and homeostasis parameters for the ramp simulations while Table II shows the aforementioned parameters for the SORF simulations. Two sets of STDP update rules were used. The STDP update rule that corresponds to exc→exc synaptic connections, denoted by W_{exc}, is described in (4) and (5). The STDP update rule that corresponds to exc→inh connections, denoted by W_{inh}, is described in (6) and (7).

TABLE I
STDP and Homeostasis Parameters in the Ramp Simulation

Parameter	Value	Parameter	Value
δ	0	A_+	2.0e-4
β	1	A_-	6.6e-5
T	5 seconds	τ_+	20 ms
γ	50	τ_-	60 ms
R_{target}	35 Hz		

Here, A_+, and A_- refer to the LTP and LTD magnitudes shown in STDP equations (4) and (5) where τ_+ and τ_- are the decay constants.

$$W_{\text{Exc.}}(t) = A_+ \exp(-t/\tau_+) \text{ for } t > 0 \quad (4)$$

$$W_{\text{Exc.}}(t) = -A_- \exp(t/\tau_-) \text{ for } t < 0 \quad (5)$$

$$W_{\text{Inh.}}(t) = -A_- \exp(-t/\tau_+) \text{ for } t > 0 \quad (6)$$

$$W_{\text{Inh.}}(t) = A_+ \exp(t/\tau_-) \text{ for } t < 0 \quad (7)$$

TABLE II
STDP AND HOMEOSTASIS PARAMETERS IN THE SORF SIMULATION

Parameter	Value	Parameter	Value
α	0.1	A_- Exc	3.0e-5
β	1	τ_+ Exc	60 ms
T	10 seconds	τ_- Exc	90 ms
γ	50	A_+ Inh	1.3e-5
R_{target} Exc	10 Hz	A_- Inh	4.1e-5
R_{target} Inh	75 Hz	τ_+ Inh	51 ms
A_+ Exc	4.5e-5	τ_- Inh	78 ms

The SORF simulation had STDP rules for both exc→exc synapses described by (4) and (5) and also exc→inh synapses described by (6) and (7). Parameters followed by 'Exc' refer to STDP at exc→exc synapses while those followed by 'Inh' refer to STDP at exc→inh synapses.

B. Proof of Stability

We follow a similar stability analysis as described in Joshi and Treisch [45]. The stability term K is omitted in our analysis for simplicity as it only serves to improve performance of our learning rule. We assume a single synapse, linear neuron and map STDP to a rate-based learning paradigm by assuming independent pre and postsynaptic Poisson firing statistics [38]. This approach yields rate-based Hebbian learning: $v_{post}(t) = w(t)v_{pre}(t)$, where v is voltage and can be considered proportional to firing rate. Our homeostasis learning rule then becomes:

$$\frac{dw}{dt} = \alpha\left(1 - \frac{\overline{v_{post}(t)}}{v_{target}}\right)w(t) + \beta v_{pre}(t)v_{post}(t)$$

$$\overline{v_{post}(t)} = \frac{\int_0^N v_{post}(t-s)\,ds}{T}$$

For this equation, there are two fixed points: w = 0 (unstable) and when

$$\beta v_{pre}(t)v_{post}(t) = \alpha\left(1 - \frac{\overline{v_{post}}(t)}{v_{target}}\right)w(t)$$

To study the stability of this learning rule, we assume v_{pre} is constant and since we also assume the neuron is linear and $w \neq 0$:

$$\beta v_{pre}^2 = \alpha\left(1 - \frac{\overline{v_{post}(t)}}{v_{target}}\right)$$

$$\overline{v_{post}}(t) = \frac{\beta}{\alpha}v_{pre}^2 v_{target} + v_{target}$$

If T=1 then $\overline{v_{post}}(t) = v_{post}$(t).

If $\overline{v_{post}}(t) > \frac{\beta}{\alpha}v_{pre}^2 v_{target} + v_{target}$ then $\frac{dw}{dt} < 0$ and consequently $\frac{dv_{post}}{dt} < 0$. Conversely, for $\overline{v_{post}}(t) < \frac{\beta}{\alpha}v_{pre}^2 v_{target} + v_{target}$, $\frac{dv_{post}}{dt} > 0$. Therefore, this is a stable fixed point.

From simulations, it can be shown that the system is critically damped when $\alpha \propto \frac{1}{N}$, $\beta = \frac{\alpha}{N\overline{v_{pre}}^2}$ and is under-damped (exhibits oscillations) when $\beta > \frac{\alpha}{N\overline{v_{pre}}^2}$. For the multi-dimensional input case, the system is critically damped when $\beta = \frac{\alpha}{N\sum_{i-1}^{m}\overline{v_{pre_i}}^2}$, where m is the number of inputs.

C. NNMF Proof

1) Theorem 1:

Non-negative Matrix Factorization update rule:

$$W \leftarrow W + \alpha(VH^T - WHH^T)$$

$$V_{n\times m} = W_{n\times r}H_{r\times m}, \text{ minimize } \|V - WH\|$$

For r=1, H is the activation of a single neuron over time and H=WTV. VHT is the Hebbian learning rule, which if 0-centered is qualitatively equivalent to the Spike-Timing Dependent Plasticity, STDP, rule (LTP and LTD combined).
Since r=1, $-WHH^T = -W\sum_{i=1}^{m}h_i^2$
By Theorem 2, $-W\sum_{i=1}^{m}h_i^2 \propto -W\sum_{i=1}^{m}h_i + \varepsilon$
$-W\sum_{i=1}^{m}h_i + \varepsilon \propto -W\bar{h}$, where $\bar{h}$ is the average activation
Assuming convergence, such that $\bar{h}$ is near C, the target average activation, and $1 - \bar{h}/C$ is zero centered then, the NNMF rule can be rewritten:

$$dW = \alpha\left(1 - \frac{\bar{h}}{C}\right)W + \beta STDP,$$

which is the STDP learning rule with homeostasis.
This also presumes that the noise term averages out over many presentations.

2) Corollary 1:

A real valued approximation of STDP is $STDP \cong (V - \bar{V})H$
Therefore, using an update rule of $dW_i = (V_i - \bar{V})h_i + (1 - \frac{\bar{h}}{C})W_i$, where $\bar{h}$ is defined as the running average over a long period of time, performs NNMF and is also a rate-based approximation to our STDP with homeostasis learning rule.

3) Theorem 2:

Since h is non-negative and a squaring non-linearity is strictly monotonic, $\sum_{i=1}^{m}h_i^2$ must be correlated with $\sum_{i=1}^{m}h_i$ and for large m, both sums must be Gaussian distributed. Since both are Gaussian, the interaction between the two is fully defined by their covariance, i.e. there are no higher order interactions. Since the covariance must be strictly positive due to the correlation both sums must be linearly related.
In the limit of large m, $\sum_{i=1}^{m}h_i^2 = \alpha\sum_{i=1}^{m}h_i + \varepsilon$, where ε is a Gaussian noise term.

ACKNOWLEDGMENT

We thank Jayram M. Nageswaran for his work developing the custom spiking neural network simulator and with the integration of the Evolving Objects libraries into the simulator. We also thank Michael Avery and Michael Beyeler for valuable feedback and discussion on this project.

REFERENCES

[1] N. Caporale and Y. Dan, "Spike timing-dependent plasticity: a Hebbian learning rule," *Annu. Rev. Neurosci.*, vol. 31, pp. 25–46, 2008.

[2] K. Carlson and N. Giordano, "Interplay of the magnitude and time-course of postsynaptic Ca2+ concentration in producing spike timing-dependent plasticity," *J. Comput. Neurosci.*, vol. 30, no. 3, pp. 747–758, 2011.

[3] H. Z. Shouval, M. F. Bear, and L. N. Cooper, "A unified model of NMDA receptor-dependent bidirectional synaptic plasticity," *Proc. Natl. Acad. Sci. USA*, vol. 99, no. 16, pp. 10831–10836, Aug. 2002.

[4] S. Byrnes, A. N. Burkitt, D. B. Grayden, and H. Meffin, "Learning a sparse code for temporal sequences using STDP and sequence compression," *Neural Comput.*, vol. 23, no. 10, pp. 2567–2598, Oct. 2011.

[5] C. Clopath, L. Büsing, E. Vasilaki, and W. Gerstner, "Connectivity reflects coding: a model of voltage-based STDP with homeostasis," *Nat. Neurosci.*, vol. 13, no. 3, pp. 344–352, Mar. 2010.

[6] R. Brette, "Computing with neural synchrony," *PLoS Comput. Biol.*, vol. 8, no. 6, p. e1002561, Jun. 2012.

[7] A. J. Watt and N. S. Desai, "Homeostatic plasticity and STDP: keeping a neuron's cool in a fluctuating world," *Front. Synaptic Neurosci.*, vol. 2, Jun. 2010.

[8] N. S. Desai, "Homeostatic plasticity in the CNS: synaptic and intrinsic forms," *J. Physiol.-Paris*, vol. 97, no. 4–6, pp. 391–402, Nov. 2003.

[9] G. Turrigiano, "Homeostatic synaptic plasticity: local and global mechanisms for stabilizing neuronal function," *Cold Spring Harbor Perspect. Biol.*, vol. 4, no. 1, Jan. 2012.

[10] G. G. Turrigiano, K. R. Leslie, N. S. Desai, L. C. Rutherford, and S. B. Nelson, "Activity-dependent scaling of quantal amplitude in neocortical neurons," *Nature*, vol. 391, no. 6670, pp. 892–896, Feb. 1998.

[11] H. Markram and M. Tsodyks, "Redistribution of synaptic efficacy between neocortical pyramidal neurons," *Nature*, vol. 382, no. 6594, pp. 807–810, Aug. 1996.

[12] L. F. Abbott and S. B. Nelson, "Synaptic plasticity: taming the beast," *Nat. Neurosci.*, vol. 3, no. 11, pp. 1178–1183, Nov. 2000.

[13] S. Song, K. D. Miller, and L. F. Abbott, "Competitive hebbian learning through spike-timing-dependent synaptic plasticity," *Nat. Neurosci.*, vol. 3, no. 9, pp. 919–926, Sep. 2000.

[14] C. Tetzlaff, C. Kolodziejski, M. Timme, and F. Woergoetter, "Synaptic scaling in combination with many generic plasticity mechanisms stabilizes circuit connectivity," *Front. Comput. Neurosci.*, vol. 5, Nov. 2011.

[15] R. J. O'Brien, S. Kamboj, M. D. Ehlers, K. R. Rosen, G. D. Fischbach, and R. L. Huganir, "Activity-dependent modulation of synaptic AMPA receptor accumulation," *Neuron*, vol. 21, no. 5, pp. 1067–1078, Nov. 1998.

[16] K. Ibata, Q. Sun, and G. G. Turrigiano, "Rapid synaptic scaling induced by changes in postsynaptic firing," *Neuron*, vol. 57, no. 6, pp. 819–826, Mar. 2008.

[17] J. Burrone, M. O'Byrne, and V. N. Murthy, "Multiple forms of synaptic plasticity triggered by selective suppression of activity in individual neurons," *Nature*, vol. 420, no. 6914, pp. 414–418, Nov. 2002.

[18] N. S. Desai, R. H. Cudmore, S. B. Nelson, and G. G. Turrigiano, "Critical periods for experience-dependent synaptic scaling in visual cortex," *Nat. Neurosci.*, vol. 5, no. 8, pp. 783–789, Jun. 2002.

[19] C. J. Wierenga, M. F. Walsh, and G. G. Turrigiano, "Temporal regulation of the expression locus of homeostatic plasticity," *J. Neurophysiol.*, vol. 96, no. 4, pp. 2127–2133, Oct. 2006.

[20] T. Branco, K. Staras, K. J. Darcy, and Y. Goda, "Local dendritic activity sets release probability at hippocampal synapses," *Neuron*, vol. 59, no. 3, pp. 475–485, Aug. 2008.

[21] M. A. Sutton, H. T. Ito, P. Cressy, C. Kempf, J. C. Woo, and E. M. Schuman, "Miniature neurotransmission stabilizes synaptic function via tonic suppression of local dendritic protein synthesis," *Cell*, vol. 125, no. 4, pp. 785–799, May 2006.

[22] T. C. Thiagarajan, M. Lindskog, and R. W. Tsien, "Adaptation to synaptic inactivity in hippocampal neurons," *Neuron*, vol. 47, no. 5, pp. 725–737, Sep. 2005.

[23] R. C. Malenka and R. A. Nicoll, "Long-term potentiation - a decade of progress?," *Science*, vol. 285, no. 5435, pp. 1870–1874, Sep. 1999.

[24] A. Lazar, G. Pipa, and J. Triesch, "SORN: a self-organizing recurrent neural network," *Front. Comput. Neurosci.*, vol. 3, Oct. 2009.

[25] P. Zheng, C. Dimitrakakis, and J. Triesch, "Network self-organization explains the statistics and dynamics of synaptic connection strengths in cortex," *PLoS Comput. Biol.*, vol. 9, no. 1, p. e1002848, Jan. 2013.

[26] J. Triesch, "A gradient rule for the plasticity of a neuron's intrinsic excitability," in *Artificial Neural Networks: Biological Inspirations - ICANN 2005, Pt 1, Proceedings*, vol. 3696, W. Duch, J. Kacprzyk, and S. Zadrozny, Eds. Berlin: Springer-Verlag Berlin, 2005, pp. 65–70.

[27] M. C. W. van Rossum, G. Q. Bi, and G. G. Turrigiano, "Stable hebbian learning from spike timing-dependent plasticity," *J. Neurosci.*, vol. 20, no. 23, pp. 8812–8821, Dec. 2000.

[28] D. V. Buonomano, "A learning rule for the emergence of stable dynamics and timing in recurrent networks," *J. Neurophysiol.*, vol. 94, no. 4, pp. 2275–2283, Oct. 2005.

[29] G. G. Turrigiano, "The self-tuning neuron: synaptic scaling of excitatory synapses," *Cell*, vol. 135, no. 3, pp. 422–435, Oct. 2008.

[30] G. G. Turrigiano and S. B. Nelson, "Homeostatic plasticity in the developing nervous system," *Nat. Rev. Neurosci.*, vol. 5, no. 2, pp. 97–107, Feb. 2004.

[31] E. M. Izhikevich and N. S. Desai, "Relating STDP to BCM," *Neural Comput.*, vol. 15, no. 7, pp. 1511–1523, Jul. 2003.

[32] E. M. Izhikevich, "Simple model of spiking neurons," *IEEE Trans. Neural Netw.*, vol. 14, no. 6, pp. 1569 – 1572, Nov. 2003.

[33] J. M. Nageswaran, N. Dutt, and J. L. Krichmar, "An efficient simulation environment for modeling large-scale cortical processing," *Front. Neuroinform.*, vol. 5, p. 19, 2011.

[34] H. S. Seung, D. D. Lee, B. Y. Reis, and D. W. Tank, "Stability of the memory of eye position in a recurrent network of conductance-based model neurons," *Neuron*, vol. 26, no. 1, pp. 259–271, Apr. 2000.

[35] E. Marder and A. L. Taylor, "Multiple models to capture the variability in biological neurons and networks," *Nat. Neurosci.*, vol. 14, no. 2, pp. 133–138, Feb. 2011.

[36] W. Van Geit, E. De Schutter, and P. Achard, "Automated neuron model optimization techniques: a review," *Biol. Cybern.*, vol. 99, no. 4, pp. 241–251, 2008.

[37] M. Keijzer, J. J. Merelo, G. Romero, and M. Schoenauer, "Evolving objects: a general purpose evolutionary computation library," in *Artficial Evolution*, vol. 2310, P. Collet, C. Fonlupt, J. K. Hao, E. Lutton, and M. Schoenauer, Eds. Berlin: Springer-Verlag Berlin, 2002, pp. 231–242.

[38] R. Kempter, W. Gerstner, and J. L. Van Hemmen, "Hebbian learning and spiking neurons," *Phys. Rev. E*, vol. 59, no. 4, p. 4498, 1999.

[39] J. Jones and L. Palmer, "An evaluation of the two-dimensional gabor filter model of simple receptive-fields in cat striate cortex," *J. Neurophysiol.*, vol. 58, no. 6, pp. 1233–1258, Dec. 1987.

[40] G. Henry, B. Dreher, and P. Bishop, "Orientation specificity of cells in cat striate cortex," *J. Neurophysiol.*, vol. 37, no. 6, pp. 1394–1409, 1974.

[41] P. Dayan and L. F. Abbott, *Theoretical neuroscience*, vol. 31. MIT press Cambridge, MA, 2001.

[42] E. Bienenstock, L. Cooper, and P. Munro, "Theory for the development of neuron selectivity - orientation specificity and binocular interaction in visual-cortex," *J. Neurosci.*, vol. 2, no. 1, pp. 32–48, 1982.

[43] K. K. A. Cho, L. Khibnik, B. D. Philpot, and M. F. Bear, "The ratio of NR2A/B NMDA receptor subunits determines the qualities of ocular dominance plasticity in visual cortex," *Proc. Natl. Acad. Sci. USA*, vol. 106, no. 13, pp. 5377–5382, Mar. 2009.

[44] E. Oja, "A simplified neuron model as a principal component analyzer," *J. Math. Biol.*, vol. 15, no. 3, pp. 267–273, 1982.

[45] P. Joshi and J. Triesch, "A globally asymptotically stable plasticity rule for firing rate homeostasis," in *Artificial Neural Networks - ICANN 2008*, vol. 5164, V. Kůrková, R. Neruda, and J. Koutník, Eds. Springer Berlin Heidelberg, 2008, pp. 567–576.