Chapter 4

Temporal and Homeostatic Mechanisms of Synaptic Plasticity[1]

Ibrahim Ozturk[2]

David M. Halliday[3]

Abstract

This chapter provides a comprehensive examination of the biologically plausible synaptic plasticity mechanisms fundamental to learning in Spiking Neural Networks (SNNs). Because SNNs, recognized as the third generation of neural network models, encode information through the precise temporal dynamics of discrete spikes, learning is intrinsically dependent upon the continuous modulation of synaptic weights and delays. We initially review the mechanics of chemical synaptic transmission and classical Hebbian plasticity, which are operationalized biologically through Long-Term Potentiation (LTP) and Long-Term Depression (LTD).

The analysis subsequently advances to Spike-Timing Dependent Plasticity (STDP), a critical unsupervised mechanism wherein the precise relative timing of pre-synaptic and post-synaptic action potentials dictates both the magnitude and direction of synaptic weight modifications. To mitigate the inherent risk of runaway network excitation driven by these localized STDP rules, the chapter elucidates the essential regulatory role of homeostatic plasticity. Specifically, it details how activity-dependent synaptic scaling globally adjusts neuronal excitability to maintain stable, target firing rates without disrupting the relative weight distributions established by STDP.

1 This chapter is based on the PhD thesis: 'Learning spatio-temporal spike train encodings with ReSuMe, DelReSuMe, and Reward-modulated Spike-timing Dependent Plasticity in Spiking Neural Networks,' by Ibrahim Ozturk, 2017, University of York. Copyright 2017 by I. Ozturk.

2 Faculty of Engineering and Natural Sciences, Department of Electrical and Electronics Engineering, Osmaniye Korkut Ata University, ibrahimozturk@osmaniye.edu.tr, 0000-0003-3149-0527. Faculty of Technology, Department of Electronics and Communication Engineering, Karadeniz Technical University, ibrahimozturk@osmaniye.edu.tr, 0000-0003-3149-0527.

3 Department of Electronic Engineering, University of York, david.halliday@york.ac.uk, 0000-0001-9957-0983.

https://doi.org/10.58830/ozgur.pub1236.c5005

Finally, the integration of dopamine-modulated plasticity is examined, illustrating how global neuromodulatory reward signals act as reward-prediction errors to consolidate local STDP-induced changes. This mechanism bridges the gap between unsupervised Hebbian timing rules and complex, goal-oriented reinforcement learning. Collectively, these temporal and homeostatic principles establish a robust theoretical foundation for the development of more advanced learning algorithms and supervised mapping techniques in artificial networks.

1. Introduction

Spiking Neural Networks (SNNs) rely on neuron models that closely mimic biological brains, allowing many natural neurobiological principles to be directly applied to their architecture. In SNNs, information is carried through connections between a sending (pre-synaptic) node and a receiving (post-synaptic) node. Each connection typically involves a synaptic weight (which determines the impact of the signal) and a synaptic delay (which dictates the travel time of the signal).

In biological terms, the process of "learning" is achieved by regulating these connections over time, which is a phenomenon known as synaptic plasticity. If an artificial network is intended to learn via supervised, unsupervised, or reinforcement strategies, its synapses must incorporate mechanisms to dynamically facilitate or depress these connection weights.

This chapter explores these fundamental, biologically plausible synaptic plasticity approaches, which will serve as the foundation for the advanced learning algorithms proposed later in the text. The chapter is organized into the following parts:

- **Section 1.2 (The Synapse):** Introduces the biological background of neuronal junctions (synapses) and explains their basic artificial interpretation and transmission models.
- **Section 1.3 (Hebbian Plasticity):** Discusses classical Hebbian learning, an unsupervised mechanism which dictates that correlated activation between neurons strengthens their connection. It also details the biological manifestations of this rule through Long-Term Potentiation (LTP) and Long-Term Depression (LTD).
- **Section 1.4 (Spike-Timing Dependent Plasticity - STDP):** Explores STDP, a crucial extension of Hebbian plasticity where the magnitude and direction of weight changes depend heavily on the *precise relative timing* of pre-synaptic and post-synaptic spikes. This section details STDP's mathematical models, implementations, and asymmetric/symmetric variants.

- **Section 1.5 (Homeostatic Plasticity):** Examines the mechanisms necessary to maintain stable neuronal functionality and prevent runaway network activity during learning. It details how neurons regulate their own excitability through synaptic scaling to keep firing rates within a target range.
- **Section 1.6 (Dopamine Modulated Plasticity):** Describes the biological background of neuromodulators, specifically dopamine, and how they consolidate the synaptic changes proposed by local STDP rules. This provides the biological inspiration for bridging unsupervised STDP with goal-oriented Reinforcement Learning.

2. Activity-Dependent Plasticity

Artificial Neural Networks have evolved significantly, with SNNs recognized as the third generation of neural network models. Unlike previous generations that use continuous analogue signals, SNNs use discrete "spikes" to carry information, mimicking the natural computation and biological plausibility of the brain. In these networks, information is processed and encoded in the precise timing between neuron firings (Gerstner & Kistler, 2002).

The fundamental building block of this communication is the synapse. A connection between a pre-synaptic (sending) node and a post-synaptic (receiving) node carries at least two dynamics: a synaptic weight, which determines the impact of the signal on the post-synaptic potential, and a synaptic delay, which dictates the travel time of the signal. In biological terms, "learning" is the process of regulating these synaptic efficacies, which is a phenomenon known as synaptic plasticity. If a network is intended to learn via supervised, unsupervised, or reinforcement strategies, the artificial synapses must include a mechanism for weight dynamics to facilitate or depress connections over time.

3. The Synapse and Synaptic Transmission

In the mammalian nervous system, synapses are specialized junctions, which are generally chemical rather than electrical contacts that allow information to flow from the axon terminals of one neuron to the dendrites of another (see Figure 1). This transmission occurs via the diffusion of neurotransmitters (such as acetylcholine) across the synaptic cleft, which activates ion channels in the target neuron as illustrated in Figure 1.

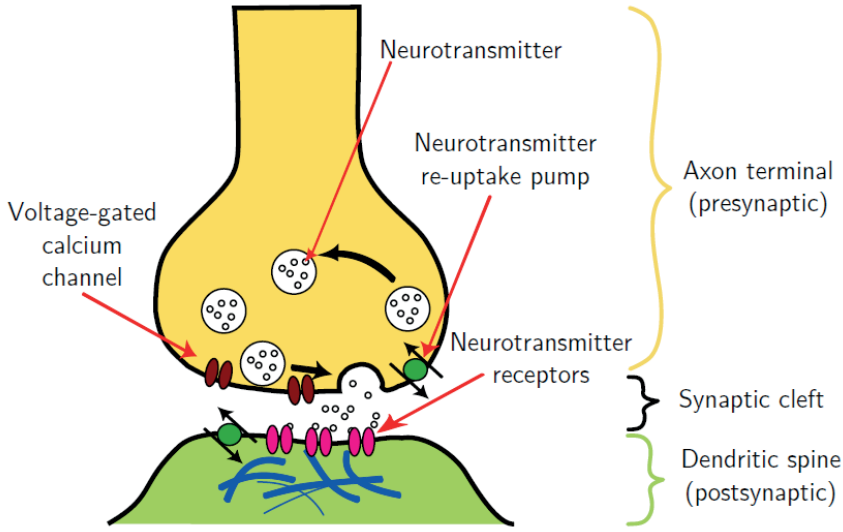

Figure 1. Biological structure of a chemical synapse, illustrating the axon terminal, neurotransmitters, and the synaptic cleft. Source: From "Learning spatio-temporal spike train encodings with ReSuMe, DelReSuMe, and Reward-modulated Spike-timing Dependent Plasticity in Spiking Neural Networks," by I. Ozturk, 2017, University of York. Copyright 2017 by I. Ozturk.

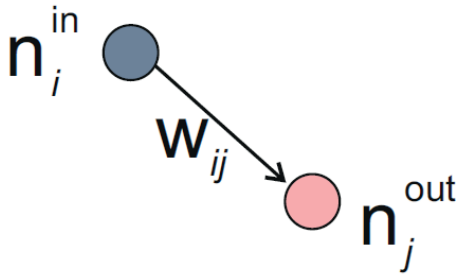

Figure 2. Artificial synapse model: bridging biological junctions to SNN abstractions via pre-synaptic/post-synaptic dynamics and synaptic weighting. Source: From "Learning spatio-temporal spike train encodings with ReSuMe, DelReSuMe, and Reward-modulated Spike-timing Dependent Plasticity in Spiking Neural Networks," by I. Ozturk, 2017, University of York. Copyright 2017 by I. Ozturk.

Depending on the synapse type, the post-synaptic potential (PSP) can be modified in two ways:

- **Excitatory Post-Synaptic Potential (EPSP):** Increases the electric polarization of the membrane, moving the neuron closer to its firing threshold.

- **Inhibitory Post-Synaptic Potential (IPSP):** Changes the charge negatively, lowering the membrane potential further from the firing threshold.

In SNN models, the input current $I_j(t)$ to a post-synaptic neuron j is commonly expressed as a weighted sum of pre-synaptic currents: $I_j(t) = \sum_{i=1}^{N_{\text{synapses}}} w_{ij}(t) \cdot f\left(t - t_f^i\right)$ where w_{ij} is the synaptic strength, t_f^i is the firing time of the pre-synaptic neuron, and f is a synaptic current function (often modeled as a Dirac delta or alpha function).

4. Hebbian Plasticity

In 1949, Donald O. Hebb proposed a groundbreaking unsupervised mechanism for synaptic plasticity. Hebb postulated that when an axon of cell A persistently takes part in firing cell B, a growth process or metabolic change occurs that increases A's efficiency in firing B (Hebb, 1949). In the context of artificial neural networks, simultaneous activation of pre- and post-synaptic cells increases the weight between them, while separate activation decreases it. The basic mathematical formulation for Hebbian learning is: $\Delta w_{ij} = \eta x_i x_j$ where Δw_{ij} is the change in synaptic weight, η is a small, positive learning rate, and x_i, x_j are the activations of the respective neurons.

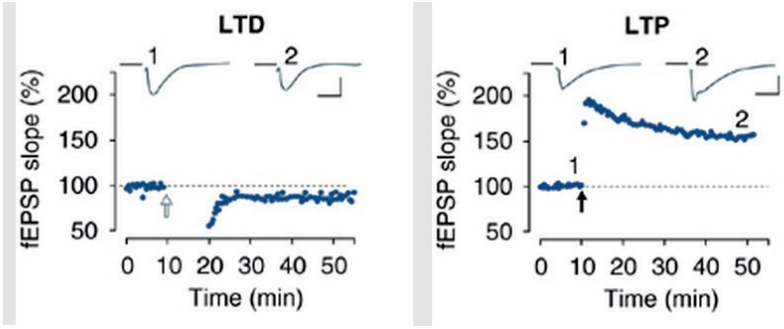

Figure 3. NMDAR-dependent synaptic plasticity in the hippocampal CA1 region. Sample experimental traces illustrating the induction of Long-Term Potentiation (LTP) and Long-Term Depression (LTD) as biological correlates for Hebbian and anti-Hebbian learning rules.Source: From "Learning spatio-temporal spike train encodings with ReSuMe, DelReSuMe, and Reward-modulated Spike-timing Dependent Plasticity in Spiking Neural Networks," by I. Ozturk, 2017, University of York. Copyright 2017 by I. Ozturk.

Decades later, biological experiments confirmed Long-Term Potentiation (LTP), a long-lasting increase in synaptic strength following rapid, successive stimulation (Bliss & Gardner-Medwin, 1973; Citri & Malenkas, 2008). Conversely, Long-Term Depression (LTD) was discovered as an activity-

dependent decrease in synaptic strength. Because some forms of LTD do not require pre-synaptic activity to be depressed, this reverse process is often defined as "anti-Hebbian" learning (Lynch et al., 1977) as sample experiments demonstrating these biological correlates for LTP and LTD are shown in Figure 3.

5. Spike-Timing Dependent Plasticity (STDP)

While traditional Hebbian plasticity relies heavily on rate-based coding, biological observations have demonstrated that the *precise relative firing times* of pre- and post-synaptic neurons dictate synaptic modification (Bi & Poo, 1998). This mechanism is known as Spike-Timing Dependent Plasticity (STDP).

There are two main shapes of the STDP form, illustrated as symmetric in Figure 4 and asymmetric in Figure 5. Symmetric STDP relies solely on the time difference between presynaptic and postsynaptic firing times, whereas asymmetric STDP also depends heavily on the specific temporal order of those firing times.

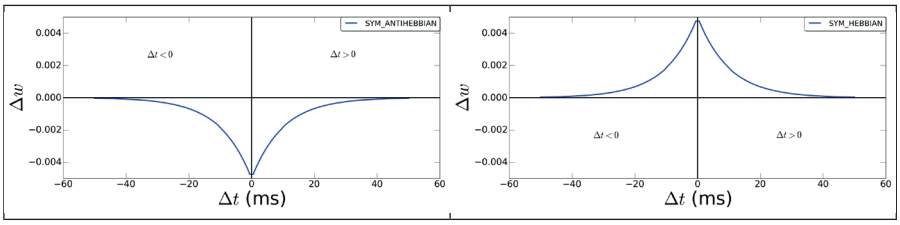

Figure 4. Symmetric-STDP. The left hand plot shows symmetric anti-Hebbian-STDP window. The right hand plot shows symmetric Hebbian-STDP window.Source: From "Learning spatio-temporal spike train encodings with ReSuMe, DelReSuMe, and Reward-modulated Spike-timing Dependent Plasticity in Spiking Neural Networks," by I. Ozturk, 2017, University of York. Copyright 2017 by I. Ozturk.

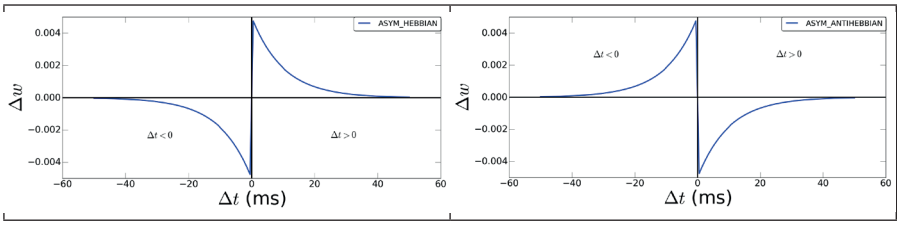

Figure 5. Asymmetric-STDP. Regardless of the order of pre- and post-synaptic firing times, STDP function W (Δt) modify the weights in the same manner for Δt < 0 and Δt > 0. The left hand plot shows asymmetric anti-Hebbian-STDP window. The right hand plot shows asymmetric Hebbian-STDP window.Source: From "Learning spatio-temporal spike train encodings with ReSuMe, DelReSuMe, and Reward-modulated Spike-timing Dependent Plasticity in Spiking Neural Networks," by I. Ozturk, 2017, University of York. Copyright 2017 by I. Ozturk.

a. The Asymmetric Learning Window

STDP operates within a critical temporal window, typically 20-25 ms (Song et al., 2000). The exact interval, defined as $\Delta t = t_j^{\text{post}} - t_i^{\text{pre}}$, determines the magnitude and direction of the weight change.

- **Potentiation (LTP):** If a post-synaptic neuron fires shortly after a pre-synaptic spike ($\Delta t \geq 0$), the synapse is strengthened.
- **Depression (LTD):** If the pre-synaptic spike arrives after the post-synaptic neuron has fired ($\Delta t < 0$), the synapse is weakened.

The standard STDP modification function, $W\left(\Delta t\right)$, is phenomenologically modeled using exponential decay constants (τ_{pre} and τ_{post}):

$$W\left(\Delta t\right) = \begin{cases} +F_{\text{pre}}\left(w\right) exp\left(-\Delta t / \tau_{\text{pre}}\right) \text{ if } \Delta t \geq 0 \\ -F_{\text{post}}\left(w\right) exp\left(+\Delta t / \tau_{\text{post}}\right) \text{ if } \Delta t < 0 \end{cases}$$

where $F_{pre}\left(w\right)$ and $F_{post}\left(w\right)$ control the amplitude of the learning window based on the current weight (Gerstner & Kistler, 2002).

b. Weight Dependence

The dependence of synaptic changes on the current weight (w_{ij}) is modulated using a non-negative exponent μ.

- **Additive STDP ($\mu = 0$):** Modifies weights independently of their current strength. In this model, $F_{pre}\left(w_{ij}\right) = A_{pre}$ and $F_{post}\left(w_{ij}\right) = A_{post}$.
- **Multiplicative STDP ($\mu = 1$):** Scales the weight change linearly based on the current weight, creating different equilibrium dynamics. Additive STDP is most commonly utilized for reliable network training.

c. Implementation via Multi-Spike Interactions

Simulating STDP by directly summing all pairs of spikes mathematically is computationally inefficient and physiologically implausible, as a biological neuron cannot remember all its individual firing times. Instead, models implement continuous pre-synaptic and post-synaptic traces (a_{pre} and a_{post}). Each spike leaves a trace that decays exponentially over time (Masquelier et al., 2008).

- When a pre-synaptic spike arrives, a_{pre} is incremented by A_{pre}, and the weight is decreased based on the current value of the post-synaptic trace a_{post}.

- When a post-synaptic spike arrives, a_{post} is incremented by A_{post}, and the weight is increased based on the current value of the pre-synaptic trace a_{pre}.

d. Weight Bounds

To prevent runaway excitation, where highly active neurons cause their synaptic weights to increase uncontrollably, synaptic weights must be bounded. Hard boundaries are applied to clip weights so they cannot exceed a maximum value (w_{max}) or fall below a minimum value (w_{min}).

e. Homeostatic Plasticity: The Balancing Act

While STDP drives learning through local timing mechanisms, the network must maintain stable overall functionality and firing rates. To achieve this, neocortical circuits employ homeostatic plasticity, an activity-dependent mechanism where neurons detect changes in their own firing rates and regulate their own excitability (Turrigiano, 2008).

6. Synaptic Scaling

Synaptic scaling globally scales a neuron's synaptic strengths multiplicatively to maintain target neuronal firing rates without destroying the relative weight distribution established by STDP. A simplified mathematical representation of this is:

$$\frac{dw_{ij}(t)}{dt} = \beta w_{ij}(t)\left[N_{\text{des}} - N_{\text{act}}\right]$$

where β is a scaling factor, N_{des} is the desired number of post-synaptic spikes, and N_{act} is the actual number of spikes (van Rossum et al., 2000).

However, continuous up-and-down scaling can cause instability, conflicting with the primary STDP learning mechanism. To solve this, scaling is often restricted to a permissible range ($N_{\text{min}}^{\text{des}} < N_{\text{act}} < N_{\text{max}}^{\text{des}}$):

$$\frac{dw_{ij}(t)}{dt} = \begin{cases} \beta w_{ij}(t)\left[N_{\text{max}}^{\text{des}} - N_{\text{act}}\right] & \text{if } N_{\text{act}} > N_{\text{max}}^{\text{des}} \\ \beta w_{ij}(t)\left[N_{\text{min}}^{\text{des}} - N_{\text{act}}\right] & \text{if } N_{\text{act}} < N_{\text{min}}^{\text{des}} \end{cases}$$

If the neuronal activity remains within this range, homeostatic scaling is inactive, and the local STDP mechanism is entirely responsible for learning. If

activity falls outside the boundaries, the scaling mechanism acts as a regularizer, adjusting all of a neuron's synapses globally to restore stability.

7. Dopamine-Modulated Plasticity

Beyond local (STDP) and global (Homeostatic) mechanisms, learning in biological systems is heavily influenced by neuromodulators like dopamine, which is closely associated with reward-driven reinforcement learning. Experimental evidence shows that subcortical dopamine signals act as a reward-prediction error (Schultz, 1998), functionally similar to Temporal Difference (TD) learning. By using a neuromodulator to consolidate the synaptic changes proposed by STDP as illustrated by the schematic of the reward-modulated STDP learning rule in Figure 6, networks can be trained through a hybrid semi-supervised approach, bridging the gap between unsupervised Hebbian mechanisms and complex, goal-oriented behavioral learning (Fremaux et al., 2010; Ozturk & Halliday, 2016). Practical applications of these reward-driven algorithms extend beyond theoretical models to training agents in specific environmental challenges, such as spatial navigation. For instance, utilizing knowledge-based reinforcement learning in a maze task demonstrates how reward mechanisms can significantly optimize an agent's trial-and-error exploration to successfully reach a target destination (Ozturk & Halliday, 2014).

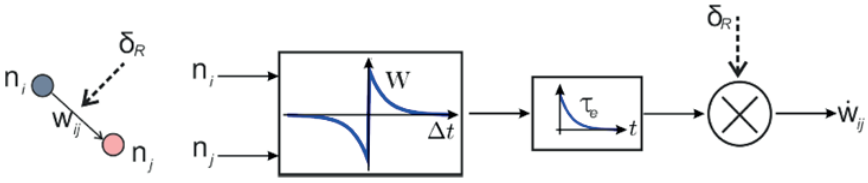

Figure 6. Schematic of the reward-modulated STDP learning rule. Block diagram illustrating the interaction between the local STDP coincidence window and a global Temporal Difference (TD) reward signal (δ) in governing synaptic weight modulation (Δw).Source: From "Learning spatio-temporal spike train encodings with ReSuMe, DelReSuMe, and Reward-modulated Spike-timing Dependent Plasticity in Spiking Neural Networks," by I. Ozturk, 2017, University of York. Copyright 2017 by I. Ozturk.

8. Conclusion

This chapter has presented the fundamental mechanisms of synaptic inputs and transmission, emphasizing that learning in Spiking Neural Networks (SNNs) is a highly complex process because information is intrinsically encoded in the time domain. To understand how artificial networks can mimic this

natural computation, various synaptic plasticities were detailed alongside their biological backgrounds.

We initially explored classical Hebbian plasticity before advancing to Spike-Timing Dependent Plasticity. STDP captures the essence of unsupervised learning in the nervous system by relying on the precise relative timing of pre-synaptic and post-synaptic spikes, and its additive form serves as the foundational learning method for training SNNs.

Beyond local STDP rules, the chapter discussed homeostatic plasticity, specifically synaptic activity-dependent scaling which is as a necessary regulatory process. This mechanism plays a crucial role in maintaining an optimal and stable level of firing activity across the network, ensuring that neurons do not become over-activated or entirely silenced during the learning process. Finally, the introduction of dopamine-modulated plasticity established a biological basis for linking these local, unsupervised timing mechanisms with global, reward-driven reinforcement signals.

Together, these temporal and homeostatic principles provide a robust, biologically plausible foundation. These concepts will serve as the core building blocks for developing the more advanced Reinforcement Learning algorithms and supervised mapping techniques explored in subsequent chapters.

References

Bi, G. Q., & Poo, M. M. (1998). Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. Journal of Neuroscience, 18(24), 10464–10472.

Bliss, T. V., & Gardner-Medwin, A. R. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the unanaesthetized rabbit following stimulation of the perforant path. Journal of Physiology, 232(2), 357–374.

Citri, A., & Malenkas, R. C. (2008). Synaptic plasticity: Multiple forms, functions, and mechanisms. Neuropsychopharmacology, 33, 18–41.

Fremaux, N., Sprekeler, H., & Gerstner, W. (2010). Functional requirements for reward-modulated spike-timing-dependent plasticity. The Journal of Neuroscience, 30(40), 13326–13337.

Gerstner, W., & Kistler, W. M. (2002). Spiking neuron models: Single neurons, populations, plasticity. Cambridge University Press.

Hebb, D. O. (1949). The organization of behavior: A neuropsychological theory. Wiley and Sons.

Lynch, G. S., Dunwiddie, T., & Gribkoff, V. (1977). Heterosynaptic depression: A postsynaptic correlate of long-term potentiation. Nature, 266(5604), 737–739.

Masquelier, T., Guyonneau, R., & Thorpe, S. J. (2008). Spike timing dependent plasticity finds the start of repeating patterns in continuous spike trains. PLoS ONE, 3(1), Article e1377.

Ozturk, I., & Halliday, D. M. (2014). Knowledge based reinforcement learning in maze navigation task [Poster presentation]. YDS 2014.

Ozturk, I., & Halliday, D. M. (2016, December). Mapping spatio-temporally encoded patterns by reward-modulated STDP in spiking neurons. In Proceedings of the International Conference on Computational Intelligence (IEEE SSCI 2016). Athens, Greece.

Schultz, W. (1998). Predictive reward signal of dopamine neurons. Journal of Neurophysiology, 80(1), 1–27.

Song, S., Miller, K., & Abbott, L. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. Nature Neuroscience, 3(9), 919–926.

Turrigiano, G. G. (2008). The self-tuning neuron: Synaptic scaling of excitatory synapses. Cell, 135(3), 422–435.

Van Rossum, M. C., Bi, G. Q., & Turrigiano, G. G. (2000). Stable Hebbian learning from spike timing-dependent plasticity. Journal of Neuroscience, 20(23), 8812–8821.