

Normalisation of Weights and Firing Rates in Spiking Neural Networks with Spike-Timing-Dependent Plasticity

Katarzyna Kozdon¹, Peter Bentley^{1,2}

¹University College London, London, United Kingdom

²Braintree Ltd, London, United Kingdom

k.kozdon@cs.ucl.ac.uk

Abstract

Maintaining the ability to fire sparsely is crucial for information encoding in neural networks. Additionally, spiking homeostasis is vital for spiking neural networks with changing numbers of weights and neurons. We discuss a range of network stabilisation approaches, inspired by homeostatic synaptic plasticity mechanisms reported in the brain. These include weight scaling, and weight change as a function of the network's spiking activity. We tested normalisation of the sum of weights for all neurons, and by neuron type. We examined how this approach affects firing rate and performance on clustering of time-series data in the form of moving geometric shapes. We found that neuron type-specific normalisation is a promising approach for preventing weight drift in spiking neural networks, thus enabling longer training cycles. It can be adapted for networks with architectural plasticity.

Introduction

Learning systems need to maintain a balance between plasticity and stability. In the brain, mechanisms such as long-term potentiation (LTP) and long-term depression (LTD) are vital for learning, but on their own they would lead to a powerful destabilisation of neuronal activity (Abbott and Nelson (2000); Miller and Mackay (1994)). Strengthening of the synapses would lead to the increased likelihood of postsynaptic firing, whereas weakening of the synapses would lead to their silencing and eventually a cascade of further downstream changes in the same direction.

Similar problems are encountered in artificial intelligence. In unsupervised spiking neural networks (SNN) with spike-timing-dependent plasticity (STDP), weight drift is a known problem which impedes learning by leading to either weight silencing or saturation (Miller and Mackay (1994); Abbott and Nelson (2000); Oja (1982)). In these networks, weight update is a function of the spike timing in a pair of connected neurons. Unlike in non-spiking supervised neural networks (NNN), the weight update is not aimed to minimise the loss function. Techniques such as L1 (Tibshirani (1997)) and L2 (Nigam et al. (1999)) regularisation used to prevent excessive weight growth and parameter overfitting in NNN are not directly applicable to SNN, because of the

complex relationship between the weights and spiking activity.

Additional difficulties arise when creating networks with plastic architecture. What value should be assigned to new weights, associated either with new synapses of existing neurons, or new neurons? The selected weight will not only affect the capacity of the new connection to learn, but will also cause downstream effects which are hard to predict in a complex systems, and may not be desirable. Thus, in order to utilise SNN for data processing, we need to identify methods preventing weight drift, silencing and saturation of spiking. Ideally, these methods should be applicable to SNN with changing numbers of neurons and synapses. These normalisation mechanisms cannot simply counteract the changes as it would effectively lead to erasing memories and undoing learning (Fox and Stryker (2017)).

One of the classical ways to tame Hebbian learning-induced bistability is the Bienenstock, Cooper and Munro model (Bienenstock et al. (1982)), according to which correlated spiking leads to LTP if the postsynaptic rate is higher than threshold, and to LTD if the firing rate is lower. The threshold changes as a function of average postsynaptic firing rate. However, there is no direct evidence that this mechanism exists in the brain.

Oja rule (Oja (1982)) is another classical approach to stabilising networks with STDP. It extends Hebbian plasticity with scaling down synaptic efficacy as a function of the square of the firing rate.

Overall, models suggest that weight normalisation mechanisms are crucial for preserving the functionality of SNN. However, little research exploring and comparing different weight normalisation methods in applied SNN, and in SNN with plastic architecture, has been published.

The brain has developed homeostatic mechanisms for preserving its functionality while changing during development and learning. These mechanisms act both at the global (Turrigiano (2008)) and local (Yu and Goda (2009)) level. Four main forms have been distinguished (Fox and Stryker (2017)):

- synaptic scaling (SS),

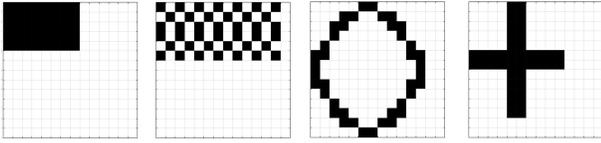


Figure 1: Examples of bitmaps used as input data: a square, grid, ellipse and cross, each composed of 40 black pixels. One of the shapes was placed at a random location within the visual field (20 x 25 pixels) and moved up, down, left or right.

- adjustment of the levels of inhibition and excitation,
- basal firing rate homeostasis,
- plasticity of the neuronal membrane properties.

Relevant to our project, it has been hypothesised that SS may be needed ”to ensure that firing rates do not become saturated during developmental changes in the number and strength of synaptic inputs, as well as stabilising synaptic strengths during Hebbian modification and facilitating competition between synapses” (Turrigiano et al. (1998)). Another attractive property of SS is stabilisation of neuronal activity without disrupting information storage and processing related to differences in synaptic strengths.

Model overview and experimental setup

The model set-up is summarised in table 1, and is a modification of Kozdon and Bentley (2018). Briefly, we used the exponential leaky integrate-and-fire model (Barranca et al. (2014); Fourcaud-Trocme et al. (2003)). Our network was composed of three feed-forward, fully-connected layers with 500, 50 and 10 neurons. 80% of the neurons were excitatory and 20% were inhibitory.

Our input data were bitmaps representing four different shapes (Fig. 1) moving in one of four directions at the speed of one pixel per frame. Each bitmap had 500 pixels, and each iteration of the network activity corresponded to a new frame.

We used evolutionary approach described in detail in Kozdon and Bentley (2018). Clustering inputs by movement direction was used as a measure of fitness. A population of networks was initialised with random starting weights and learning hyperparameters. After training and testing, the top third of the networks were treated as non-sexually reproducing parents, each giving rise to three children: a child-clone, a child-clone which undergoes another training cycle, and a child with a mutation in one of the four learning parameters and exposed to another cycle of training.

We examined the spiking activity in each layer, and performance of the networks.

Table 1: Model hyperparameters.

Neural networks		
input neurons	500	
hidden neurons	50	
output neurons	10	
Training and testing set-up		
frames per training input	10	
inputs per training cycle	20	
frames per testing input	50	
inputs per testing cycle	80	
time step	0.02 s	
Electrophysiological parameters		
	<i>excitatory</i>	<i>inhibitory</i>
resting potential	-63.70	-59.3
reset potential	-68.7	-64.3
firing threshold	-45.80	-33.3
resistance	67.70	133.3
τ	16.70	36.5
Evolution constraints		
	<i>min</i>	<i>max</i>
LTP	0.001	0.1
inhLTP	0.001	0.1
LTD	0.001	0.1
discharge	0.06	5.0

Experimental setup: global weight normalisation

Normalisation was performed after each training epoch (10 frames of the input data). Based on our preliminary experiments, we selected the normalised sum of weights w_{norm} to be tied to the number of all neurons (n_n) in the network and equal:

$$w_{norm} = n_n \cdot 100 \quad (1)$$

We tested six set-ups:

1. No normalisation. The weights were capped at 4 as in the previous paper (Kozdon and Bentley (2018)) (“control_capped”).
2. No normalisation and weight-capping (“control”).
3. Normalisation with weight-capping (“norm_capped”).
4. Normalisation without weight-capping (“norm”).
5. Separate normalisation for excitatory and inhibitory neurons, with weight-capping (“norm_capped_ie”).
6. Separate normalisation for excitatory and inhibitory neurons, without weight-capping. A proportion of w_{norm} was assigned to each neuronal population; 80% was assigned to excitatory and 20% to inhibitory synapses (“norm_ie”).

We analysed a population of 12 networks over 20 generations. 10 repeats were performed. Significance was calculated using the Kruskal-Wallis test.

Results

Global weight normalisation

Effects of normalisation were different depending on the presence of the weight cap 2).

In networks without normalisation but with weight-capping (“control_capped”), spiking in the hidden and output layers (Fig. 2 a and b) decreased with time. Networks without normalisation and weight-capping (“control”) demonstrated a wide range of spiking patterns, and ranged from fully saturated to fully silent. Both of these unnormalised approaches lead to activity patterns which are undesirable for information processing.

Networks with weight caps and normalisation (“norm_capped” and “norm_capped_ie”) demonstrated the most stable spiking behaviour, and were resistant to silencing while avoiding saturation during the examined period.

Networks with normalisation and capping (“norm_capped”), and both set-ups with neuron type-specific weight normalisation (“norm_capped_ie” and “norm_ie”) were significantly more precise than the controls (Fig. 3). Their performance also improved with time; normalisation with weight-capping had a lower precision than normalisation without the cap, but a steeper learning rate.

Discussion

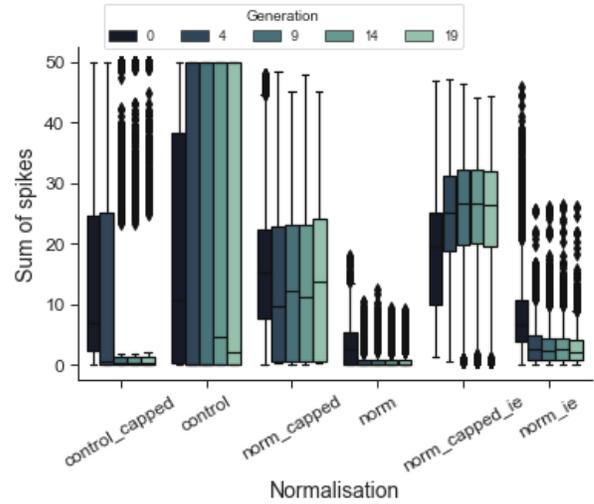
Our results demonstrate that even simple weight normalisation approaches can reduce saturation and silencing in SNN. These effects translated into improved performance on the unsupervised clustering tasks.

The simplest approach to network stabilisation is a global normalisation of the sum of weights, irrespective of the neuron type. The sum of weights is a function of the number of neurons, which gives it a potential to be suitable for networks with a changing number of neurons.

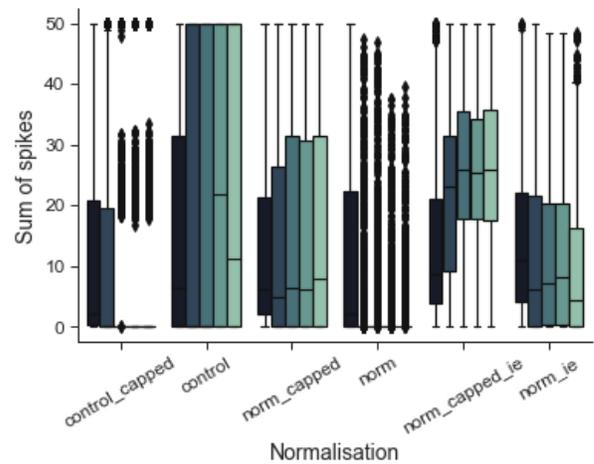
We also tested a neuron type-dependent modification of this approach, and normalised the sum of excitatory and inhibitory weights separately, assigning a proportion of the sum of weights to each population. This approach aimed to preserve the balance between the two. However, as the activity of each neuron depends on the strengths and types of all incoming weights, this is an indirect and global way of balancing excitation and inhibition, and does not normalise the exact balance between these two. This approach could be further extended to networks with multiple neuronal populations, even if the numbers of neurons fluctuate.

We are currently testing other approaches:

- Normalisation of the spiking activity level:
 - of the output neurons.
 - of all neurons.



(a) Hidden layer.



(b) Output layer.

Figure 2: Analysis of spiking activity. Median number of spikes per neuron per 50 iterations. Bars indicated standard deviation, and dots outliers.

- Activity-based adjustment of the level of inhibition and excitation.

When stabilising SNN, we primarily care about preserving their ability to encode information using spikes. Weight normalisation is an indirect approach to achieve this; a more direct approach would be to adjust the weights based on the spiking activity of the network. Spiking of the output layer is our direct readout. It is desired for the output neurons to fire sparsely (both activation of all neurons and total silence are to be avoided) so that multiple classes of input data can be encoded. Thus, we could focus solely on adjusting weights as a function of spiking in the output layer. This approach benefits from direct focus on stabilising our readout. However, spiking of the output layer may not provide enough

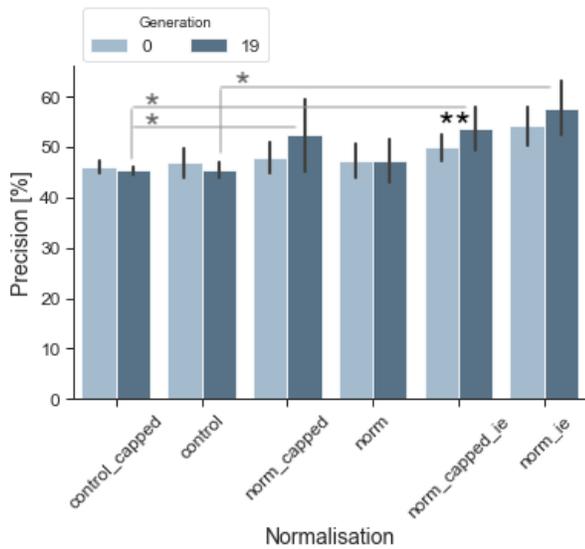


Figure 3: Fitness of the networks. Double asterisk indicates significant difference between generation 0 and 19. Single asterisks indicate significant precision differences between normalised networks and their respective controls in generation 19. Bars indicate standard deviation.

resolution to fine-tune weights of the whole network. Alternatively, as the activity of the output layer is a function of the activity of the preceding layers, we could change the weights as a function of spiking of the whole network. This approach is less direct, but the breadth of observed behaviours may allow us to fine-tune the activity of the network.

Weight update can be multiplicative or additive. In the first case, weights are changed proportionally to their size. This method affects distribution of the weights. In additive weight update, a set amount Δw is added or subtracted from each weight. Weights should not be allowed to drop below 0 to avoid switching between excitation and inhibition. This method allows to activate silent synapses.

Selective scaling of inhibitory and excitatory neurons as the function of spiking activity of the whole network or the output layer is a possible modification of this approach. If the number of spikes exceeds the permitted range, excitatory weights can be decreased and/or inhibitory increased; if spiking is below the permitted range, the reverse is executed. This approach can be adapted to minimise the weights without silencing of the network, and thus can be a form of activity-dependent mechanism reducing overfitting.

Conclusions

Spiking neural networks with STDP require a mechanism for preventing weight drift. Our results show that even simple normalisation of the sum of weights, tied to the number of neurons provides stabilisation. Interestingly, the results are improved when normalisation was neuron type-specific,

maintaining the sum of inhibitory and excitatory weights separately. This method is beneficial for maintaining trainable weights in networks with static architecture; it can also be implemented in networks with fluctuating number of neurons and weights, and with multiple neuronal populations, thus enabling their training and stabilisation of their activity levels.

Acknowledgements

This work was supported by Biotechnology and Biological Sciences Research Council and the London Interdisciplinary Biosciences Consortium.

References

- Abbott, L. F. and Nelson, S. B. (2000). Synaptic plasticity: Taming the beast. *Nature Neuroscience*, 3(11s):1178–1183.
- Barranca, V. J., Johnson, D. C., Moyher, J. L., Sauppe, J. P., Shkarayev, M. S., Kovačič, G., and Cai, D. (2014). Dynamics of the exponential integrate-and-fire model with slow currents and adaptation. *Journal of computational neuroscience*, 37(1):161–80.
- Bienenstock, E. L., Cooper, L. N., and Munro, P. W. (1982). Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 2(1):32–48.
- Fourcaud-Trocme, N., Hansel, D., Vreeswijk, C. V., and Brunel, N. (2003). How spike mechanisms determine response to fluctuating inputs. *The Journal of Neuroscience*, 23(37):11628–11640.
- Fox, K. and Stryker, M. (2017). Integrating Hebbian and homeostatic plasticity: Introduction. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372(1715).
- Kozdon, K. and Bentley, P. (2018). The Evolution of Training Parameters for Spiking Neural Networks with Hebbian Learning. *The 2018 Conference on Artificial Life*, pages 276–283.
- Miller, K. D. and Mackay, D. J. C. (1994). The Role of Constraints in Hebbian Learning. *Neural computation*, 6:100–126.
- Nigam, K., Lafferty, J., and Mccallum, A. (1999). Using Maximum Entropy for Text Classification. *IJCAI-99 Workshop on Machine Learning for Information Filtering*, pages p. 61–67.
- Oja, E. (1982). Simplified neuron model as a principal component analyzer. *Journal of Mathematical Biology*, 15(3):267–273.
- Tibshirani, R. (1997). The LASSO method for variable selection in the Cox model. *Statistics in Medicine*, 16(4):385–395.
- Turrigiano, G. G. (2008). The Self-Tuning Neuron: Synaptic Scaling of Excitatory Synapses.
- Turrigiano, G. G., Leslie, K. R., Desai, N. S., Rutherford, L. C., and Nelson, S. B. (1998). Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature*, 391(6670):892–896.
- Yu, L. M. and Goda, Y. (2009). Dendritic signalling and homeostatic adaptation.