

# SYNAPTIC PLASTICITY WITH CORRELATED FEEDBACK

## knowing how much to learn

A. D. Antrobus (antrobus@gatsby.ucl.ac.uk)

Sainsbury Wellcome Centre, 25 Howland street  
WT 4JG, London, United Kingdom

P. E. Latham (pel@gatsby.ucl.ac.uk)

Sainsbury Wellcome Centre, 25 Howland street  
WT 4JG, London, United Kingdom

### Abstract

Learning synaptic weights is difficult. Connections receive global error signals, which are low dimensional and noisy, and local signals which lack information about relative contributions to the error.

In this setting, it makes sense for connections to learn not just their ‘best guess’ of their weight, but also how confident they should be in this guess: i.e. to infer a *distribution* over their target weight. This idea was developed in (Aitchison & Latham, 2015) and (Aitchison, Pouget, & Latham, 2017). Similar concepts appear in (Hiratani & Fukai, 2018).

In the aforementioned works the update equations are discrete in time and the likelihoods used in inference are Markov. This is not how things are in biology, where signals are continuous in time and temporally correlated. Here we consider a non-Markov setting, deriving coupled ODEs which describe how the parameters of the posterior evolve as more data is observed. We use a local temporal-smoothing method to deal with the continuous feedback and discrete pre-synaptic spike events. We show that the window of smoothing can be chosen in a principled way: maximising the per-spike decrease in uncertainty.

We find our algorithm works better than the leaky delta rule with optimised learning rate. More importantly, for the simple model described, the method accurately predicts posterior variance.

**Keywords:** synaptic plasticity; bayesian inference;

probabilistic synapses

### Introduction

At a high level, our hypothesis is that connections aim to infer distributions over their target weights from sequentially observed data. To do this, they must update their estimates as new data is observed. The mathematically correct way to do this is by Bayes rule:

$$p(w^*(t)|D(0,t)) \propto p(D(t-\Delta,t)|w^*(t), D(0,t-\Delta))p(w^*(t)|D(0,t-\Delta)).$$

Here,  $w^*(t)$  is a set of target connection strengths at time  $t$  and  $D(0,t)$  is a set of observed data signals on the time interval  $(0,t)$ .

Such Bayesian approaches to synaptic learning rules have been described in (Aitchison & Latham, 2015) and (Aitchison et al., 2017). A key point in these works is that the resultant update rules for the expected target weight do not require an arbitrary learning rate,  $\eta$ , as is the case in most classical Hebbian rules. Rather the effective learning rate depends on the posterior variance, which comes from doing inference. This means more uncertain connections naturally learn faster than more certain ones. A different approach is taken in (Hiratani & Fukai, 2018) to obtain an approximation to the posterior by means of a particle filter. In this model, each of the multiple synapses which constitute a connection between two neurons is considered a particle. During learning, the strength and location of individual synapses are varied such that their cumulative effect captures the desired statistics.



In all the above mentioned works time is discretised, updates are instantaneous and likelihoods are Markov; they have the property

$$p(D(t)|w^*(t), D(0, t-1)) = p(D(t)|w^*(t)). \quad (1)$$

In biological systems, where fluctuations in potentials have decay times and target weight values may drift in time, this Markov assumption does not hold. We investigate this trickier setting below.

### The model & the problem

We consider a simple model where one post-synaptic (point) neuron receives input from  $N$  independent pre-synaptic neurons. The  $i^{\text{th}}$  pre-synaptic neuron fires Poisson spikes with rate  $\lambda_i$ . A spike on neuron  $i$  at time  $t_i^k$  causes a transient post-synaptic potential with peak amplitude  $w(t_i^k)$  and temporal profile  $\alpha(t - t_i^k)$  in the post-synaptic neuron. Let  $\{w_i^*\}_{i=1}^N$  be a set of unknown target weight functions (functions of time). This defines an error signal:

$$\delta(t) = \sum_j (w_j^*(t) - w_j(t)) S_j(t), \quad (2)$$

where  $S_j(t)$  is the smoothed spike train from neuron  $j$ .

We are agnostic as to what overall task defines the  $w_i^*$ 's. In our model context, we can think of  $w_i^*$ 's as latent variable. Our goal is to get (variational) posteriors over these functions in an online way.

To derive update equations for the distribution over target weight  $w_i^*$ , we consider the problem from the point-of-view of connection  $i$ . For connection  $i$ , the contribution to the error from other connections can be considered "noise":

$$\delta(t) = (w_i^*(t) - w_i(t)) S_i(t) + \xi_i(t), \quad (3)$$

$$\xi_i(t) := \sum_{j \neq i} (w_j^*(t) - w_j(t)) S_j(t). \quad (4)$$

Hence we will talk about learning from the point of view of a *connection*, rather than a neuron.

The additive noise, consisting of the signals from other connections, has temporal correlations ( $\mathbb{E}[\xi(t)\xi(t+T)] = K(T) \neq 0$ ). This means exact likelihood in equation (1) would have the form

$$\propto \exp\{D(t) \int_{-\infty}^t K^{-1}(t-\tau) D(\tau) d\tau\}. \quad (5)$$

In general, inverting this kernel is not possible and in some cases, it's not even well defined!

To overcome this problem, we consider an approximation which allows us to parameterise how much signal to extract from the feedback in a "windowed" manner after each spike. This avoids having to invert these kernels.

### Approximate Version

Consider time-averaging the error signal for a short interval after spike  $k$  arrives:

$$\bar{\delta}_\varepsilon(t_i^k) := \int_{t_i^k}^{t_i^k + \varepsilon} \delta(\tau) d\tau \quad (6)$$

$$= (w_i^*(t_i^k) - w_i(t_i^k)) \bar{S}_{i,\varepsilon}(t) + \bar{\xi}_{i,\varepsilon}(t). \quad (7)$$

Where we have assumed that the change in weights is negligible over this period,  $\varepsilon$ . We have now introduced a free model parameter,  $\varepsilon$ , which we will set in a principled way later.

For inference we assume the only global information available to all connections is the error,  $\delta$ . Each connection also receives local information: their own spike times ( $t_i^k$ ), their own implemented weights ( $w_i(t_i^k)$ ) and the PSP shape,  $\alpha(t)$ . We also assume they know the prior over target weight distributions. Hence we can write  $D(0:t) = \{\delta(0:t), w(t_i^k), t_i^k\}$  for all  $t_i^k < t$ .

Given a prior distribution  $p(w_i^*(t_i^k) | D_i(0:t_i^k))$ , how should connection  $i$  update this distribution when spike  $k$  arrives and it observes  $\{\delta(t_i^k), w(t_i^k), t_i^k\}$ ? We propose the connection integrates the error signal for the period  $\varepsilon$  after the spike. From equation (7), we have

$$p(\bar{\delta}_\varepsilon(t) | w_i^*(t_i^k), w_i(t_i^k), D(0:t_i^k)) = p(\bar{\xi}_{\varepsilon,i}(t)). \quad (8)$$

This is our likelihood distribution and it can be shown that, for our model problem, it is approximately Gaussian<sup>1</sup> with statistics

$$\mathbb{E}[\bar{\xi}_{\varepsilon,i}(t)] = 0, \quad (9)$$

$$\mathbb{V}[\bar{\xi}_{\varepsilon,i}(t)] \approx \text{Avg}_{j \neq i} [\lambda_j \sigma_j^2(t)] C[\varepsilon, \alpha], \quad (10)$$

<sup>1</sup>Omitted for brevity. See: [www.gatsby.ucl.ac.uk/~antobus/CCN2019\\_supp\\_mat.pdf](http://www.gatsby.ucl.ac.uk/~antobus/CCN2019_supp_mat.pdf)

where  $\text{Avg}_{j \neq i}[\cdot]$  indicates the empirical average over the other connections and  $C[\varepsilon, \alpha]$  is shorthand to denote a constant scale which depends only on  $\varepsilon$  and  $\alpha$ . If we assume that the prior is also Gaussian,  $p(w_i^*(t_i^k) | D_i(0 : t_i^k)) = \mathcal{N}(\mu_{i,k-1}(t_i^k), \sigma_{i,k-1}^2(t_i^k))$  (the second subscript indicates how many spikes have been seen) we can use standard results to obtain discrete increments of the mean and variance for each spike (omitted due to space constraints, see <sup>1</sup>).

**Choosing  $\varepsilon$**  : Equation (11) is the discrete change in the posterior variance after a spike arrives and is clearly always negative. It can be shown that this quantity is zero in the limits that  $\varepsilon \rightarrow 0$  and  $\varepsilon \rightarrow \infty$ . Hence we can set  $\varepsilon$  to the value for which this is most-negative, i.e. for which each spike maximally decrease the variance (see supplementary material<sup>1</sup>). Now that we have “fixed”  $\varepsilon$ , we will drop it from our notation for simplicity.

$$\Delta \sigma_i^2(t_i^k) = \left( \frac{-\sigma_{i,k-1}^4(t_i^k) \bar{S}^2(t_i^k)}{\sigma_{i,k-1}^2(t_i^k) \bar{S}^2(t_i^k) + \mathbb{V}[\bar{\xi}_i(t_i^k)]} \right) \quad (11)$$

**Drift and biophysical approximations:** So far we know how to update the posterior after each spike but what happens in between spikes? Far from a spike, the error signal is not telling the connection much about its own contribution to the error. Rather, since the connection is not receiving information, we should expect its uncertainty (i.e. variance) to increase and the posterior to regress to the original prior distribution,  $p(w_i^*)$ . This effect is captured by the drift terms in equations (15) and (16).

Equation (11) and its analogy for  $\mu$  give us discrete, per-spike updates. So, ignoring drift for now, we could write the value of  $\mu_{i,k}(t)$  for  $t \gg t_i^k$  as

$$\mu_{i,k}(t) = \mu_{i,k}(0) + \sum_{l=1}^k \Delta \mu_i(t_i^l) \quad (12)$$

$$\approx \mu_{i,k}(0) + \sum_{l=1}^k \Delta \mu_i(t_i^l) \sigma(t - t_i^l) \quad (13)$$

$$\frac{d\mu_{i,k}(t)}{dt} \approx \sum_{l=1}^k \Delta \mu_i(t_i^l) \frac{d\sigma(t - t_i^l)}{dt} \quad (14)$$

where  $\sigma(t)$  is a sigmoidal function satisfying  $0 \leq \sigma(t) < 1$ . The integral of the (normalised) PSP

shape,  $\alpha$ , is just such a sigmoid, so we can “smooth” our updates out in time by multiplying them by the PSPs.

In the end, we obtain equations (15) and (16) as our final learning rules. Equation (15) is made of two parts. The first two terms are the drift component: the prior pulls the mean back towards the prior mean,  $\mu_{i,0}$ . The Third term looks Hebbian: a product of a learning rate, the error and the PSP. We can see that each neuron’s learning rate depends on the ratio of its variance and the variance of the smoothed noise. This means connections with high uncertainty (large variance) learn faster. The learning rule for the precision also has a drift component, pulling it up to the prior variance  $\sigma_0^2$  (which is typically large) and a data-driven component, which is always negative but approaches zero for highly confident connections ( $\sigma^2(t) \rightarrow 0$ ).

$$\begin{aligned} \frac{d}{dt} \mu_i(t) = & \\ \tau_w^{-1} (\mu_{i,0} - \mu_i(t)) + \frac{\sigma_i^2(t) \bar{\alpha} / |\alpha|}{\bar{\alpha}^2 \sigma_i^2(t) + \mathbb{V}[\bar{\xi}_i(t)]} \delta(t) S_i(t) & \end{aligned} \quad (15)$$

$$\begin{aligned} \frac{d}{dt} \sigma_i^2(t) = & \\ 2\tau_w^{-1} \left( 1 - \frac{\sigma_0^2}{\sigma_i^2(t)} \right) - \frac{\sigma_i^4(t) \bar{\alpha}^2 / |\alpha|}{\bar{\alpha}^2 \sigma_i^2(t) + \mathbb{V}[\bar{\xi}_i(t)]} S_i(t). & \end{aligned} \quad (16)$$

## Results

A natural benchmark is testing if this algorithm against the leaky delta rule (LDR: equation (17)). The learning rate ( $\eta$ ) is numerically optimised in mean-square error of the output for each set of simulation parameters.

$$\frac{d}{dt} \mu_i(t) = \tau_w^{-1} (\mu_{i,0} - \mu_i(t)) + \eta \delta(t) S_i(t). \quad (17)$$

This rule only tracks a point-estimate for the expected weight, not the variance. For our simulations we used a log-normal distribution of presynaptic firing rates, to mimic distributions commonly found in biology (Shafi et al., 2007).

Due to space limits, we present just the take-home point: Figure 1 shows that the Bayesian rule does significantly better than the optimised Hebbian rule in mean squared error of the output (a), as well as achieving a smaller error between the posterior mean and the ground-truth target values (b).

If the Bayesian rule is doing a reasonable job of inference, we find that true target value to be within  $1 - \sigma(t)$  of  $\mu(t)$  about 68% of the time. In the simulation presented here, we find that our means are within this interval  $72.26 \pm 3.64\%$  of the time.

### References

- Aitchison, L., & Latham, P. E. (2015). Synaptic sampling: A connection between psp variability and uncertainty explains neurophysiological observations. *arXiv, ID:1505.04544*.
- Aitchison, L., Pouget, A., & Latham, P. E. (2017). Probabilistic synapses. *arXiv, ID:1410.1029*.
- Hiratani, N., & Fukai, T. (2018). Redundancy in synaptic connections enables neurons to learn optimally. *Proceedings of the National Academy of Sciences of the United States of America, 115*, E6871–E6879.
- Shafi, M., Zhou, Y., Quintana, J., Chow, C., Fuster, J., & Bodner, M. (2007). Variability in neuronal activity in primate cortex during working memory tasks. *Neuroscience, 146*, 1082 - 1108.




Figure 1: **(a)** Mean square error in output in units of the inference algorithm mean square error. The leaky-delta rule is approximately 50% worse than the inference method. Both methods are much better than the control, achieving approximately  $10^{-2}$  the control error. **(b)** Square-difference between the posterior means and the target weights, averaged over the 100 neurons. Shaded area is one standard deviation computed over neurons. The inference method achieves a lower average and much tighter variance.