Predictive plasticity in dendrites: from a computational principle to experimental data

Dominik Spicher, Claudia Clopath, Walter Senn

Summary

Plasticity of excitatory cortical synapses is thought to be the main mediator of mammalian learning. Due to the selective advantage that the ability to adapt to a changing environment grants, it is reasonable to assume that the processes governing these changes in connection strength are in some sense optimal. Yet, this has been difficult to reconcile with an "embarrassment of riches" of the LTP and LTD phenomenology. Here, we present an attempt at bridging this gap by showing that a mathematically derived model can exhibit some of these experimentally observed effects, while still retaining functional capabilities under diverse learning paradigms.

Our work is based on a published plasticity model [3] that postulates that learning is driven by an intraneuronal prediction error where the weights of "student inputs" onto a dendritic compartment change in order to reproduce voltage changes imposed on a somatic compartment by "teacher inputs." We show here that this two-compartment model of a pyramidal neuron can be extended in some simple ways, such as using conductance-based instead of current-based inputs, bringing it closer to the biophysics of pyramidal neurons. This allows us to reproduce a diverse set of experimental observations on cortical plasticity, such as different characteristics of the spike-timing dependence of plasticity.

Additionally, we show within a simple setup of a pattern recognition task that the extended model, while being less analytically tractable, can still perform well under unsupervised and reinforcement learning paradigms. Therefore, a single learning rule derived from the optimization of a well-defined cost function can be brought into correspondence with a large body of experimental evidence on synaptic plasticity, while still providing a diverse set of relevant functionality.

Methods

The principal idea of the learning rule is that synaptic plasticity in dendrites serves to minimize a prediction error between dendritic and somatic voltages. Briefly (see [3] for the full model), we conceptualize a two-compartment neuron with somatic and dendritic voltages U and V respectively. We assume that the soma receives time-varying conductance inputs g_E and g_I, and that the weights of dendritic inputs change in order for the dendritic voltage to reproduce the resulting somatic voltage. Note that, due to the "teaching" inputs being conductances, the teaching current vanishes (and can thus be removed) as soon the soma follows a matching potential U_M set by the inputs:

$$U_M(t) = \frac{g_E(t)E_E + g_I(t)E_I}{g_E(t) + g_I(t)}$$

A learning rule assuring this dynamic of dendritic input "explaining away" somatic inputs can be achieved by following a gradient on a suitable error function, resulting in a weight update rule

$$\dot{w} \propto (S(t) - \phi(V^*))PSP(t), \qquad (1)$$

where $S(t)$ is the somatic spike train (generated by an inh. Poisson process with rate $\phi(U)$, V^* is the attenuated dendritic potential and $PSP(t)$ is the postsynaptic potential evoked by a presynaptic spike. In order to make direct analogies to experiments, we enhance the biophysical plausibility of the published model with the following modifications: (i) bidirectional current flow between somatic and dendritic compartments, (ii) conductance based inputs in the dendrites, resulting in low-pass filtered version PSP' of the postsynaptic potential entering the equation above, and (iii) adding explicit spike currents (Na^+ and K^+), resulting in a stereotypic spike shape (see Fig. 1a inset). We will show in the next sections that this extended model is able to reproduce plasticity experiments while retaining useful functional capabilities.

Results I: Plasticity experiments

As can be seen from Eq. 1, a synapse is only potentiated if its EPSP is shortly followed by a postsynaptic spike. This characteristic lies at the heart of the pervasive spike-timing dependence of plasticity (STDP) [1]. Indeed, an STDP-like plasticity curve readily emerges as we simulate common STDP plasticity protocols (Fig 1a). The LTD part

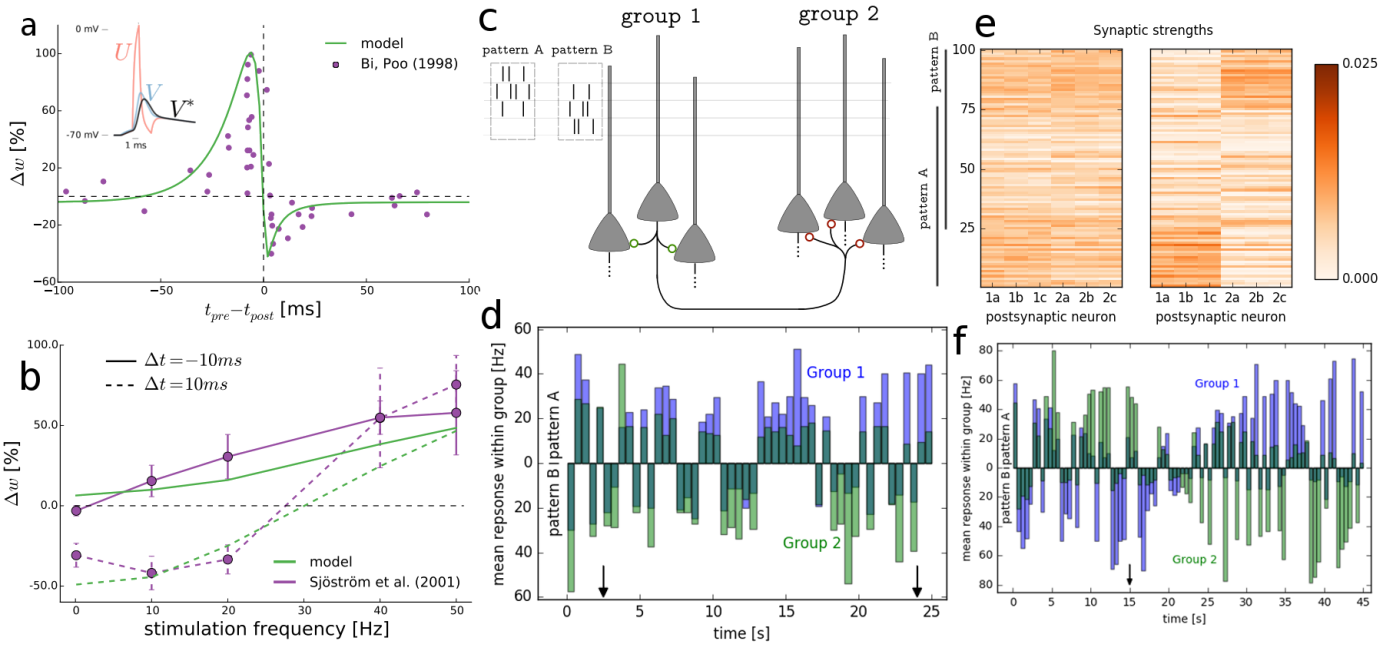

of the curve is driven by the correlation of the back-propagating action potential (Fig. 1a inset) and a subsequent EPSP. We stress that the sharp transition from LTP to LTD is largely independent of parameter choices. STDP is well-known to be a complex and malleable process. Due to space constraints we concentrate on stimulation frequency dependence[1] in both experiments [2] and under the model (Fig 1. b): As we increase stimulation frequency, post-pre-post triplets start to emerge and turn an LTD post-pre protocol into and LTP one.

Results II: Computational Task

At its heart, the learning rule is a supervised one where somatic voltage acts as the teacher for dendritic voltage. Given sufficiently rich inputs onto the dendrite (for all results here, we use 100 inputs with frozen Poisson spike trains), the learning rule will force the dendritic voltage to follow a trajectory that subsequently reproduces U_M in the soma. We refer to the previous publication ([3] Fig. 1) for a simple demonstration and focus on more interesting settings here.

A well-known technique for unsupervised learning is to use the activity itself as a teaching signal. We thus construct a small network where the outputs are fed back as a somatic teaching signal. For neurons within the same group, these connections have a high reversal potential, whereas the connections across groups are inhibitory (see Fig. 1c). The two groups thus learn to respond in a structured winner-take-all manner to two partially overlapping input patterns (see Fig. 1d): The learning rule picks out connections that are exclusively part of one of the two patterns (Fig. 1e, weights are depicted for time points marked by arrows in Fig 1. 1d).

Extending this to reinforcement learning, we can enforce one of the two mappings from patterns to groups by integrating Eq. 1 with a long time constant to get eligibility traces e_{ij}. After each trial, weights are changed by the eligibility trace modulated by reward R: $\Delta w \propto Re_{ij}$, where $R = 1$ if the correct group was more active and $R = -1$ otherwise. Fig. 1f shows the course of learning where the reward mapping was flipped at $t = 15s$ (indicated by arrow) .

References

[1] Guo-qiang Bi and Mu-ming Poo. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *The Journal of neuroscience*, 18(24):10464–10472, 1998.

[2] Per Jesper Sjöström, Gina G Turrigiano, and Sacha B Nelson. Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron*, 32(6):1149–1164, 2001.

[3] Robert Urbanczik and Walter Senn. Learning by the dendritic prediction of somatic spiking. *Neuron*, 81(3):521–528, 2014.

[1] other reproduced characteristics include voltage-dependence and the influence of synapse location in the dendritic tree