

HOW MANY MODIFIABLE MECHANISMS DO MODIFIABLE SYNAPSES NEED?

Gardner-Medwin, AR, Department of Physiology, University College London, UK

Modifiable synapses (for example, those subject to LTP or LTD) can store a small amount of information about the history of local events. This information is probably expressed largely through changes in synaptic efficacy or weight, and one might therefore think that one modifiable parameter (determining this weight) would suffice for the synapse to compute any aspect of its history that it is able to express. Most theory-driven synaptic learning rules indeed assume just a single variable parameter, albeit often subject to changes that are quite complex functions of the local conditions (including states of pre- and post-synaptic terminals, neighbouring synapses and neuro-modulators, and the precise sequence and timing of events affecting these). This rather minimalist theoretical convention is actually a profound and unnecessary constraint on the computational power of a synapse. Recognition of the need for multiple modifiable mechanisms may help to throw some light on the diversity of mechanisms for long-term changes, both pre- and post-synaptic, observed in real synapses (1).

A single modifiable parameter easily provides a running tally (over what may be very long periods) of a frequency or probability - for example, frequency of near-simultaneous depolarisations of pre- and post-synaptic cells, as in the many postulated variants of the Hebb synapse. Suppose, however, that a synaptic weight should not just reflect the frequency with which a cell A has participated in the firing of cell B, as proposed by Hebb, but a true statistical association between pre- and post-synaptic activation. A large weight would then indicate that concurrent activity has happened more often than expected by chance coincidence. Such an association implies that $P(A\&B) > P(A)P(B)$, involving the comparison of 3 parameters that are altered in different ways by events in the history of the synapse. Computation based on such a sequence of events requires the continuous holding and updating of 3 variables, corresponding to the fact that there are 3 degrees of freedom in the contingency table for the joint probabilities of 2 events. Three separately modifiable physiological parameters must therefore exist if a synapse to be able to adapt to, and quantify, the inferences about postsynaptic activity that are deducible, on the basis of learning, from the presence or absence of presynaptic activity.

A Bayesian framework for combining such inferences, from relatively independent data derived from different inputs to a cell, suggests that a useful synaptic computation would be the log-likelihood ratio, or weight of evidence (2) for activity in B afforded by activity in A: $w = \log(P(A|B) / P(A|\text{not-B}))$. This is the statistic that sums linearly for independent data, and therefore approximately matches the neural summation of postsynaptic currents. If a synapse can do without information about the absolute frequencies of A and B, then this statistic can actually be estimated with 2 continuously modifiable parameters, but the additional discarded information (requiring a third modifiable parameter) would be necessary if, in a changing environment, the weight should reflect the history of events over a defined period of time.

Note that this need for multiple modifiable parameters arises even with a *single* simple statistic data expressed by the synaptic weight. In addition, synapses may with advantage store and express statistics accumulated over different timescales (for example, transient and consolidated memory using binary and graded mechanisms in series (3)), and there is scope for expression of more than one statistic through variable synaptic interactions and dynamics.

1. Malinow R, Maine ZF, Hayashi Y. LTP mechanisms: from silence to four-lane traffic. *Current Opinion in neurobiology*, 10:352-357, 2000
2. Good, IJ *Probability and the weighing of evidence*. London: Griffin, 1950.
3. Gardner-Medwin AR. Doubly modifiable synapses: a model of short and long-term auto-associative memory. *Proc. Roy. Soc. Lond. B* 238: 137-154, 1989.

HOW MANY MODIFIABLE MECHANISMS DO MODIFIABLE SYNAPSES NEED?

A. R. Gardner-Medwin

Three distinct concepts:

1.

How many (and which) local parameters are involved in causing long-lasting synaptic modifications?

Many - e.g. pre- and post-synaptic electrical and chemical conditions, precise timings of activity, and neuromodulators.

2.

How many (and which) independently modifiable mechanisms operate within a modifiable synapse?

Several - often with different timecourses and conditions for modification, and with complex interaction.

3.

How many parameters are needed to characterise changes of expressed function in a modified synapse?

Possibly (and certainly in many models) as few as one - a synaptic 'efficacy' or 'weight' - though varying synaptic dynamics (e.g. facilitation and fatigue, etc.) may usefully express more than one.

Local synaptic computations

Modifiable synaptic parameters can only depend on the history of **local** conditions - not, for example, on patterns of activity across many cells. This presents an interesting and profound constraint on neural network computations, but from a theoretical standpoint there is no constraint on how many local parameters may be involved and how complex the functions may be (1, above).

A puzzle and a challenge:

Where modification is expressed by variation of a single synaptic weight (3, above) then it is tempting to think that a single modifiable storage mechanism (2, above) would suffice. Is this correct?

This poster aims to show that the answer is NO! Multiple independent storage mechanisms are sometimes necessary within a synapse to compute potentially important functions, even when these are expressed through only a single parameter.

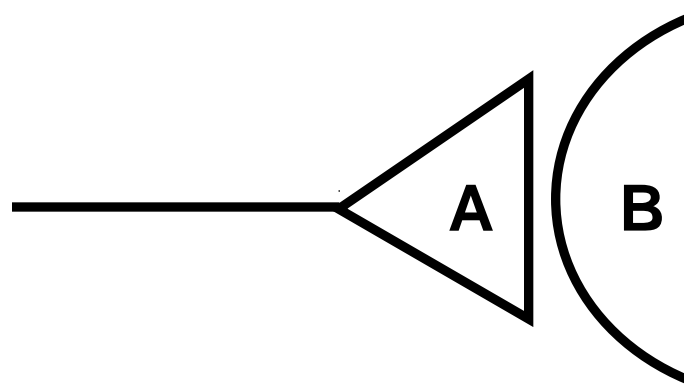
This argument is not the only reason why synapses might require multiple modifiable mechanisms - models have suggested useful roles for independent mechanisms that have different timecourses of memory retention and for ways in which the dynamics, as well as the strength, of a synapse may be varied. But the issue addressed here is particularly interesting because it may seem counter-intuitive.

What does a Hebb synapse compute?

The Hebb synapse (and its many variants) strengthen :




“when the presynaptic terminal contributes to firing the postsynaptic cell”

Such potentiation is often said to depend on pre- and post- **association**. But strictly, it depends not on a statistical association of pre- and post-synaptic firing, but on temporal coincidence (within some time-frame) of such firing, which may be due to chance. The distinction can be crucial when learning is to be used for inference.



Modulation of pre- & post-synaptic firing coincidences without statistical association

Blue lines show the probabilities of independent pre- and post-synaptic firing and the conjoint firing ($P_{A\&B} = P_A \cdot P_B$). Black lines show running synaptic frequency estimates based on single weight parameters that undergo fixed increments when the events occur, and exponential relaxation at other times (time constant 20 units). The graph based on coincidences (A&B) is analogous to a Hebb synapse, with substantial coincidence- dependent potentiation despite the absence of any pre & post- synaptic association.



What is an appropriate synaptic measure of statistical association?

Neurons make a '**decision**' about when to fire on the basis of '**evidence**' in the activity of their afferent axons. In many learning situations a **Bayesian** approach to this decision seems appropriate, where the evidence is used to establish the conditional probability that, with such evidence in the past, the postsynaptic cell has actually fired. When there is no association (i.e. pre- and post- synaptic firing have been statistically independent) then the pre-synaptic firing provides no evidence about whether firing should currently be elicited.

Since simple dendrites tend often to sum synaptic currents approximately linearly, the appropriate synaptic strength should on this basis be an evidence function (ϵ) that sums linearly for different (sufficiently independent) pieces of evidence, to compute a conditional probability. This is the log likelihood ratio:-

$$\epsilon = \text{Evidence for firing of B, given firing of A} = \log (P(A | B) / P(A | \text{not-B})) \quad [1]$$

where $P(A|B)$ means the conditional probability of A, given B. The summed synaptic influence, given such a measure of association, is the increment (above an *a priori* level without any evidence) for $\log(P/(1-P))$ for the firing of cell B, known as a belief function β or log-odds :-


$$\beta = \log (P(B) / (1 - P(B))) = \sum (\epsilon \text{ from afferent fibres}) + \beta_0 \quad [2]$$

Computation of an evidence function

Evidence [1, above] is fairly easily computed, but depends on the full 3 degrees of freedom of the contingency table for the combined probabilities of two random variables (pre- and post- synaptic firing). It requires either 3 or (with loss of information about the rate at which data has been collected - ok if associations are assumed to be unvarying) at least 2 modifiable synaptic mechanisms for storage of independent variables. Simply storing the current evidence function ϵ itself is not sufficient, because the way it changes in response to a particular contingency, like the joint firing of A and B, depends not just on the current value of ϵ , but on the separate values of other parameters, such as the conditional probabilities $P(A|B)$ and $P(A|\text{not-B})$.



Dept. of Physiology, University College London, London WC1E 6BT, UK



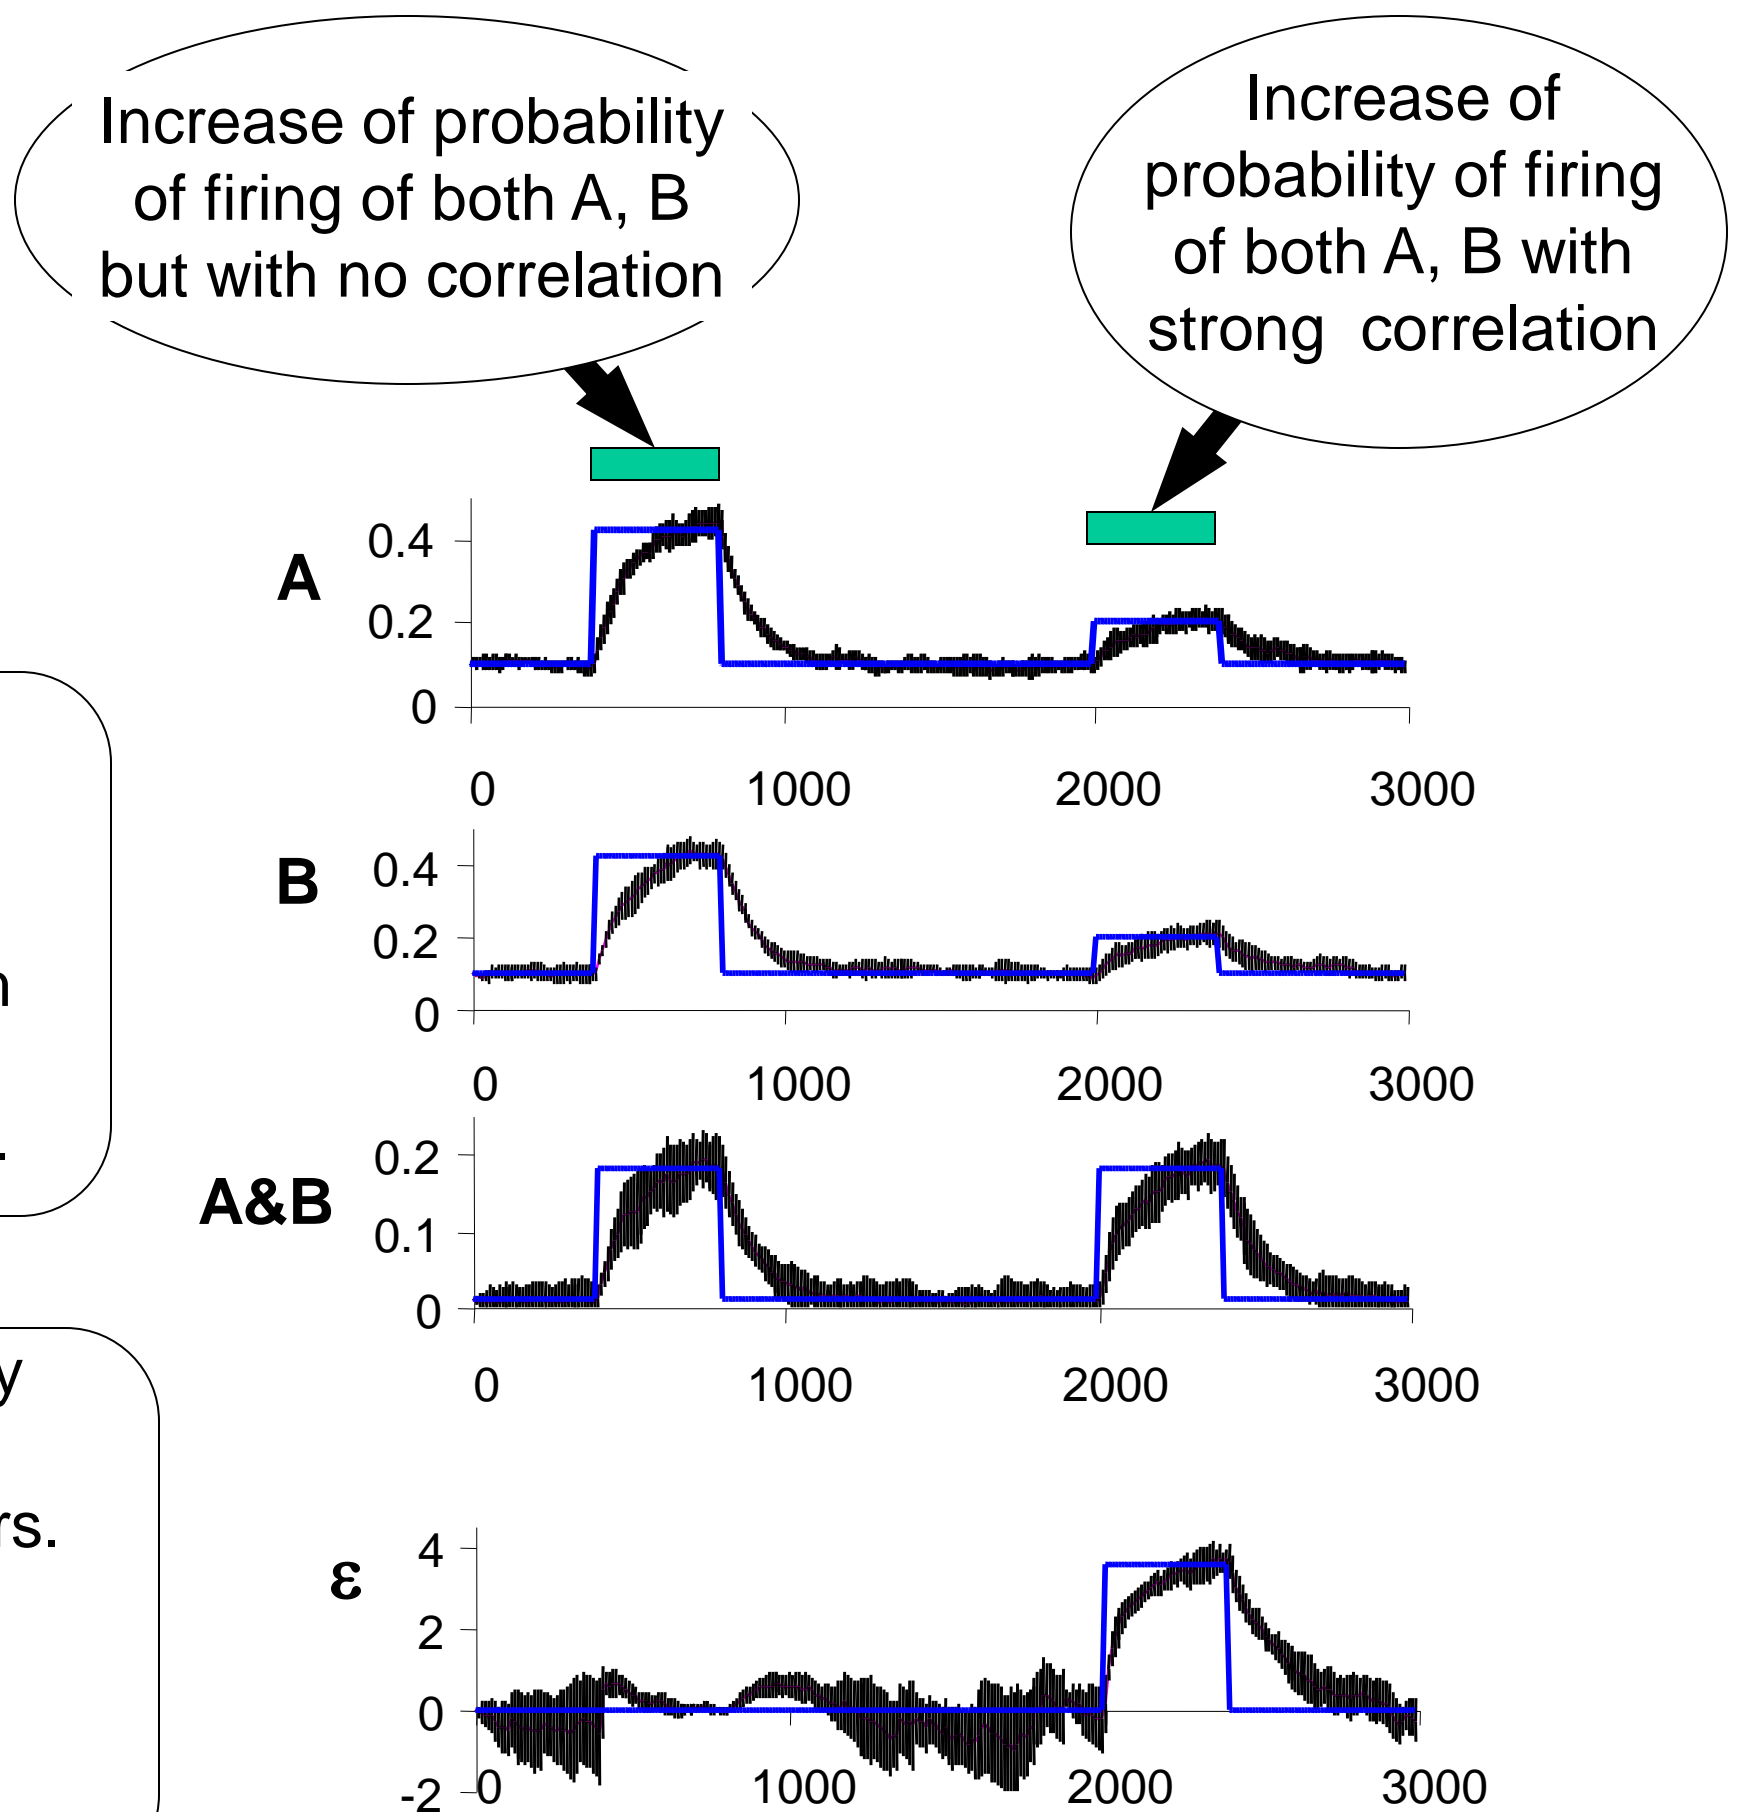
Simulation results (mean \pm s.d. from 10 simulations)

1. Estimates of pre-, post- and paired firing probability per time unit, with a relaxation time constant of 100 units. True probabilities shown in blue. Each estimate would require one modifiable mechanism and one stored parameter.

2. Evidence for firing of B conveyed by firing of the pre-synaptic axon A, calculated from the above 3 parameters.

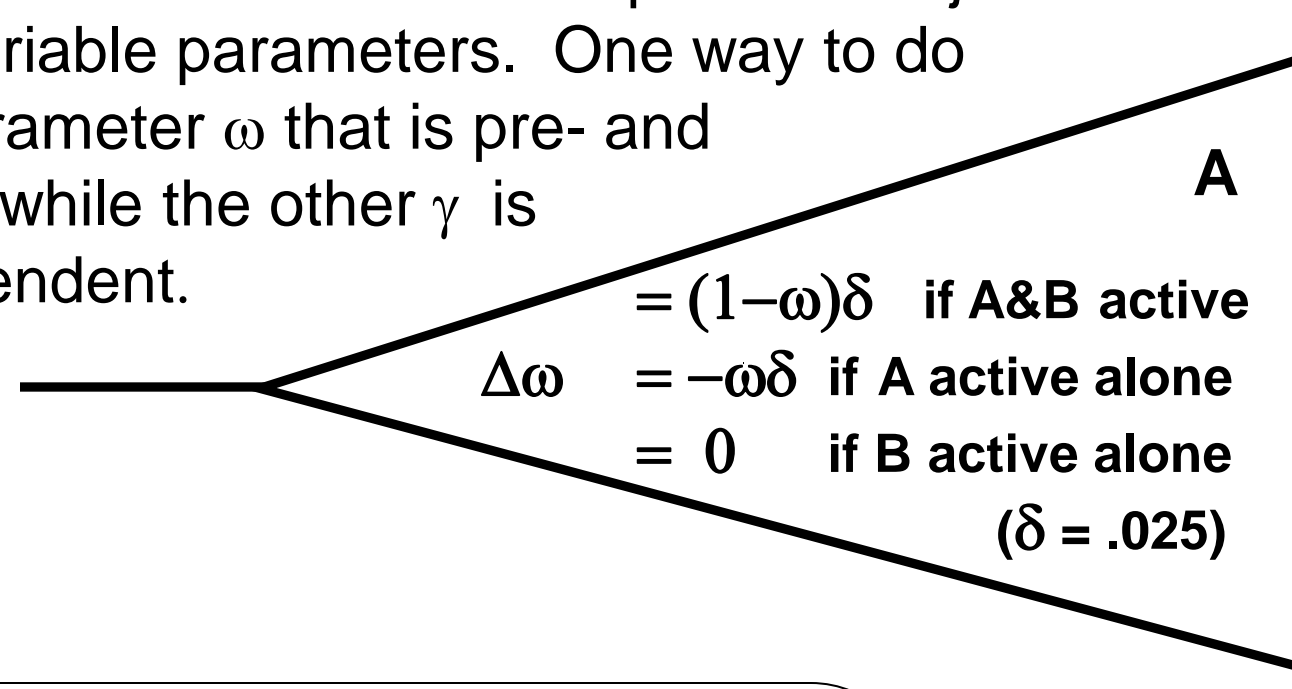
$$\varepsilon = \ln \left(\frac{P(A\&B) (1-P(B))}{(P(A)-P(A\&B)) P(B)} \right)$$

Note reduced s.d. during periods with more information.



Evidence estimated with just 2 modifiable parameters


In a steady state, evidence can be computed from just 2 independently variable parameters. One way to do this uses one parameter ω that is pre- and post- dependent while the other γ is purely post- dependent.



$\Delta\gamma = (1+\gamma)\delta'$ if B active ($\delta'=0.01$)
 $\Delta\gamma = -\gamma(1+\gamma)\delta'$ if B inactive
 evidence(B|A): $\varepsilon = \ln(\omega/\gamma)$
 $[\cong 2(\omega-\gamma)/(\omega+\gamma)]$
 summed over active pre- axons

The simulation uses the odds ratio for firing of B given A [$\omega = P(A\&B)/(P(A)-P(A\&B))$] and the odds ratio for firing of B itself [$\gamma = P(B)/(1-P(B))$]. Computation equations are above. Note that $1/\gamma$ rather than γ is graphed, to be analogous to a component of synaptic efficacy, though γ itself could be modelled by spine conductance.

Evidence to be summed across active synapses is computed as $\ln(\omega/\gamma)$, or approximated by simpler functions. With only 2 stored parameters, changes of probabilities, even with no statistical association, can lead to marked transient errors of evidence estimation, as at *.



Summary

- ◆ Every expressed synaptic parameter that is modifiable during learning may (depending on an aspect of the complexity of its computation*) require two or more separately variable storage mechanisms within the synapse for its computation and correct updating.
- ◆ A Bayesian approach to synaptic computation, in which the manipulated parameters are probabilities, can give insight into the possible nature and complexity of elementary synaptic learning processes.
- ◆ Appropriate manipulation of probability estimates depends on the statistical model of underlying causes (especially in a changing environment) and may require modulation of elementary synaptic computation for its optimisation.
- ◆ Constraints of realistic physiology (for example the fact that synapses probably do not switch between excitation and inhibition - analogous to evidence *for* and *against* activation) provide interesting challenges for efficient design.
- ◆ There is seldom talk of ways that modifiable synapses might adaptively change the effect they have on dendrites when they are *not* active. This might be:-
 - (i) a trick that evolution missed (failing to convey useful evidence based on when an axon is *silent*),
 - or (ii) quantitatively unimportant (because axons are silent most of the time),
 - or (iii) something simply experimentally less tractable than changes of the response to stimulation.

* Can anyone put this in more precise mathematical terminology ?

HOW MANY MODIFIABLE MECHANISMS DO MODIFIABLE SYNAPSES NEED?

Gardner-Medwin, AR, Dept of Physiology, University College London, London WC1E 6BT, UK

Modifiable synapses (for example, those subject to LTP or LTD) can store a small amount of information about the history of local events. The expression of this information is often assumed to be through long-term changes of a single variable (the synaptic efficacy or weight), interacting with the short-term dynamic properties of synapses and neural codes (4). Given this assumption, one might think that long-term storage of only one variable parameter would be required, since only one is expressed. Most theory-driven synaptic learning rules indeed assume just one long-term variable, albeit subject to changes that may be complex functions of the local conditions (including states of pre- and post-synaptic terminals, neighbouring synapses and neuro-modulators, as well as precise relative timing of their changes). This restriction is actually a profound constraint on the computational power of a synapse, even in models where the expression of stored information is limited to a simple weight. The value of multiple modifiable mechanisms in this context may help to throw light on why there is a diversity of physiological mechanisms for long-term changes, both pre- and post-synaptic, in real synapses (3).

A single modifiable parameter easily provides a running tally (over what may be very long periods) of a frequency or probability - for example, the frequency of near-simultaneous depolarisations of pre- and post-synaptic cells, as in the many postulated variants of the Hebb synapse. Suppose, however, that a synaptic weight should not just reflect the frequency with which a cell A has participated in the firing of cell B (as proposed by Hebb), but a true statistical association between pre- and post-synaptic activation. A large weight should then indicate that concurrent activity has been more frequent than expected by chance coincidence. Such an association implies that $P(A\&B) > P(A)P(B)$, involving comparison of 3 parameters that are altered in different ways by events in the history of the synapse. There are 3 degrees of freedom in the joint probabilities for 2 events, and correct updating requires the continuous holding of 3 variables. There must therefore be 3 separately modifiable physiological parameters for a synapse to be able to adapt to, and quantify correctly, the inferences about postsynaptic activity that are deducible, on the basis of learning, from the presence or absence of presynaptic activity.

A Bayesian framework for combining such inferences, from relatively independent data arriving at different inputs to a cell, suggests that a useful synaptic computation would be the log-likelihood ratio, or weight of evidence (2) for activity in B afforded by activity in A: $w = \log(P(A|B) / P(A|not-B))$. This is the statistic that sums linearly for independent data, and therefore approximately matches the neural summation of postsynaptic currents. If a synapse can do without information about the absolute frequencies of A and B, then this statistic can be estimated with just 2 continuously modifiable parameters, but the additional discarded information (requiring a third modifiable parameter) would be necessary if, in a changing environment, the weight is to reflect the history of events over a defined period of time.

This need for multiple modifiable parameters arises even with a *single* statistic expressed as the synaptic weight. In addition, synapses may usefully store statistics accumulated over different timescales (e.g. transient and consolidated memory expressed through binary and graded mechanisms in series (1)), while variation of the parameters of synaptic dynamics (4) offers scope to express several statistics, requiring additional modifiable mechanisms.

1. Gardner-Medwin AR. Doubly modifiable synapses: a model of short and long-term auto-associative memory. Proc Roy Soc Lond B 238: 137-154, 1989.
2. Good, IJ Probability and the weighing of evidence. London: Griffin, 1950.
3. Malinow R, Maine ZF, Hayashi Y. LTP mechanisms: from silence to four-lane traffic. Current Opinion in Neurobiology, 10:352-357, 2000
4. Tsodyks M, Pawelzik K & Markram H. Neural networks with dynamic synapses. Neural Comp 10: 821-835, 1998



**takes two
(at least)
to tango**