Learning

Basic principles

Suggested reading:

- Chapter 8 in Dayan, P. & Abbott, L., Theoretical Neuroscience, MIT Press, 2001.

Learning

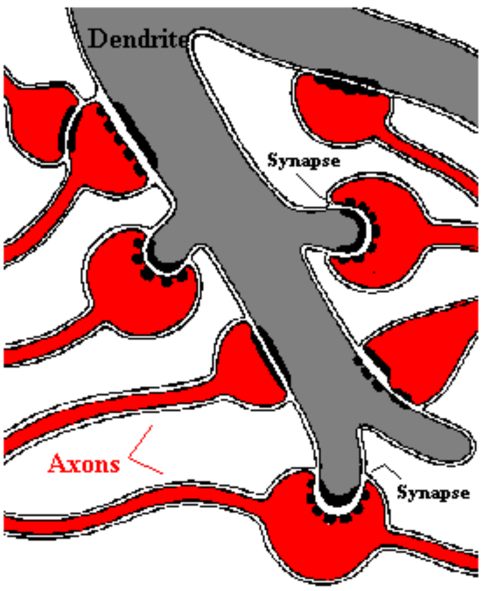

Contents:

- Types of learning
- Learning strategies
- Hebbian learning
- Hebbian learning and LTP
- Reinforcement learning
- Physiology of LTP
- Molecular mechanisms underlying LTP
- Basic Hebb rule
- Covariance rule
- Oja learning rule
- BCM rule
- Generalized Hebbian learning
- Anti-Hebbian learning
- Trace learning
- Spike-Timing Dependent Plasticity

Types of learning

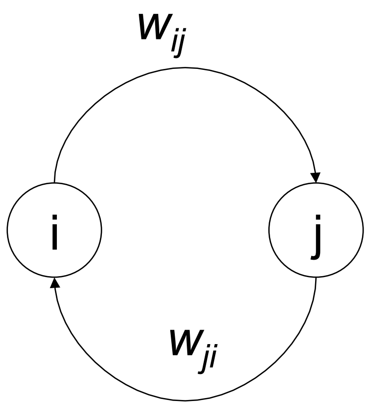

The memory in a neural network is stored in the connections (weights) among the neurons.

Learning means that the weights (w) are modified.

supervised: input $\mathbf{u}$, output $\mathbf{v}$, model $p(\mathbf{v}\,|\mathbf{u})$
adjustment to a target output

reinforcement: input $\mathbf{u}$ and scalar reward r ;
often associated with a temporal credit assignment problem

unsupervised: model $p(\mathbf{u})$
uses statistical properties in the input

Learning strategies

Bottom-up: rules of neural plasticity, e.g. Hebbian

Top-down: use of objective function $E(\mathbf{w})$ to determine the weight changes by the gradient: grad($E(\mathbf{w})$)

Hebbian learning

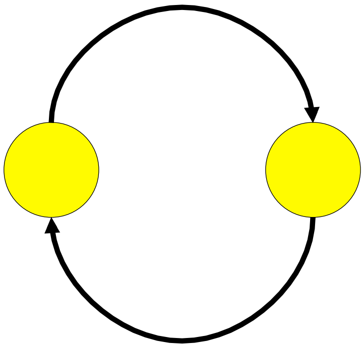

Neurons fire together
wire together.

Donald Olding Hebb (* July 22nd 1904 Kanada † August 20th 1985) is considered as the father of cognitive Psychobiology. In his well known book „The Organization of Behavior“, 1949, he focuses on the principles of learning in neural networks. His theory became known as Hebbian theory and the models which follow this theory are said to exhibit Hebbian learning.

Hebbian learning

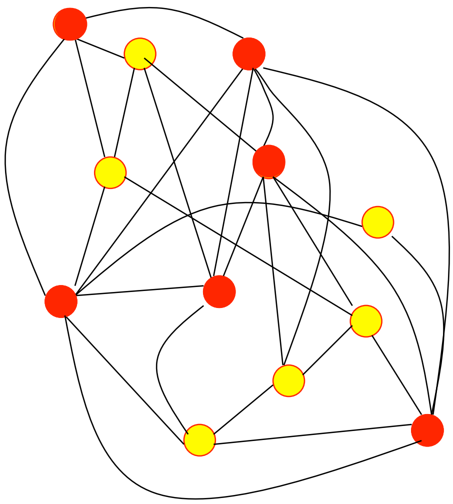

Hebbian learning

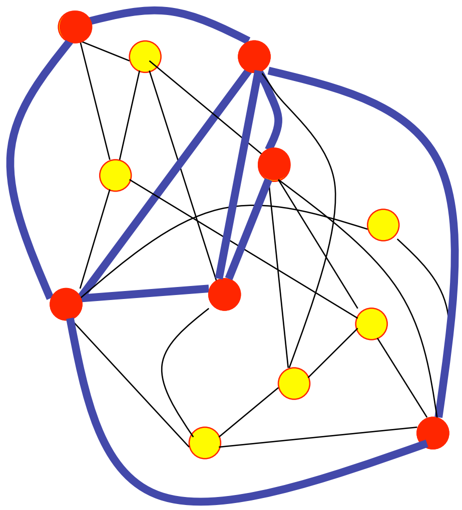

Item memorized

Hebbian learning

Recall:
Partial info

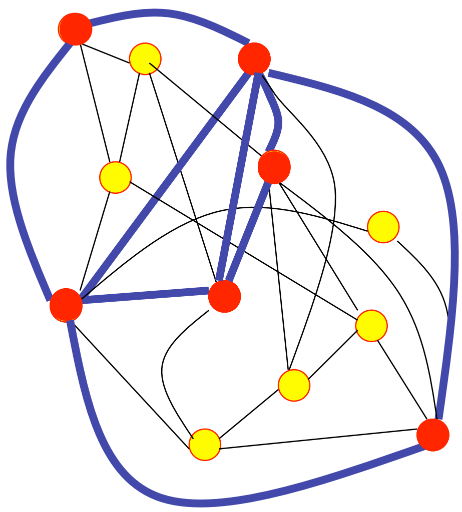

item recalled

Hebbian learning

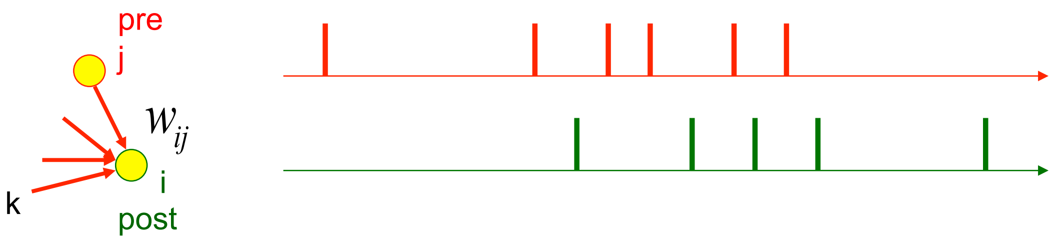

When an axon of cell *j* repeatedly or persistently takes part in firing cell *i*, then *j*'s efficiency as one of the cells firing *i* is increased

Hebb, 1949

→ strong element of *causality*

Hebbian Lerning and LTP

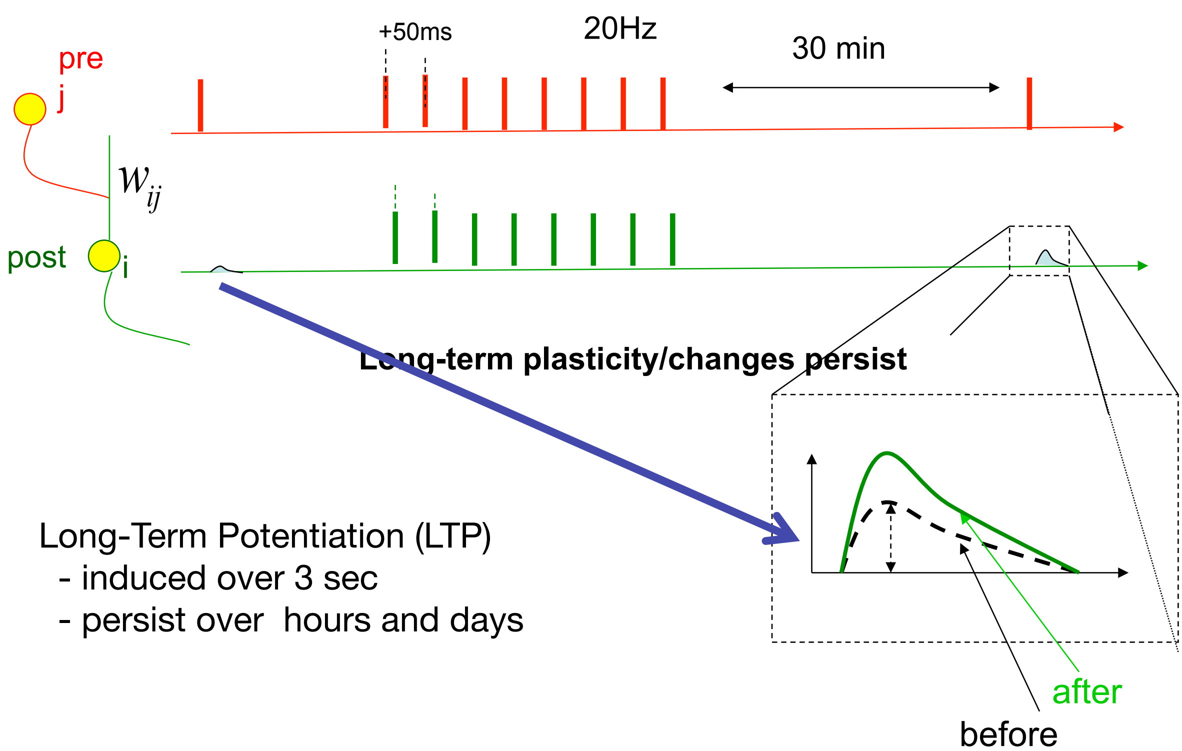

Long-Term Potentiation (LTP)
- induced over 3 sec
- persist over hours and days

Hebbian Learning and LTP

Hebbian learning is unsupervised learning.

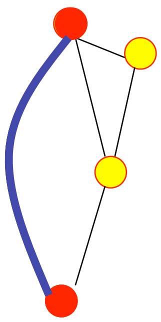

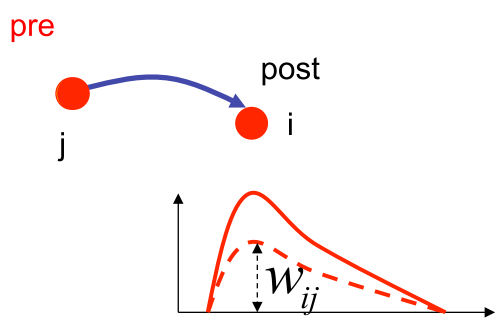

$$\Delta w_{ij} \propto F(pre, post)$$

Reinforcement Learning

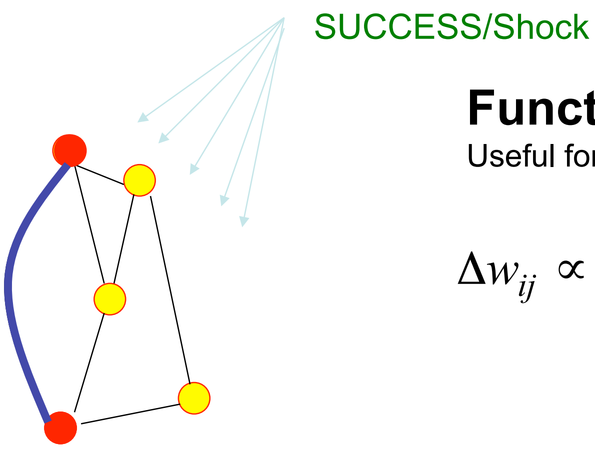

Functional Postulate

Useful for learning the important stuff

$$\Delta w_{ij} \propto F(pre, post, SUCCESS)$$

local global

Reinforcement Learning

Learning rule:

$$\Delta w = u \cdot v \cdot Dp$$

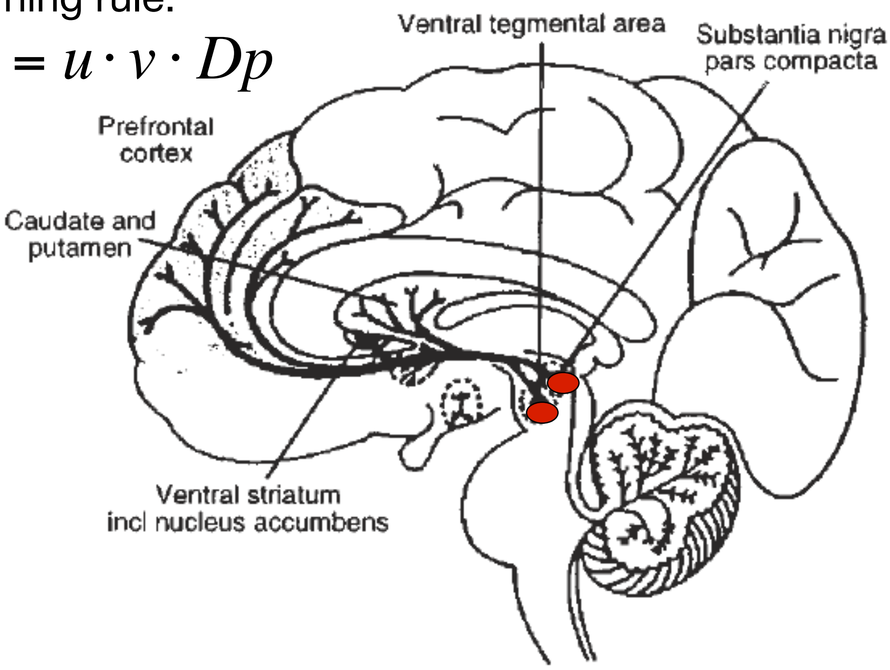

Dp=Dopamine = actual reward – expected reward

Physiology of LTP

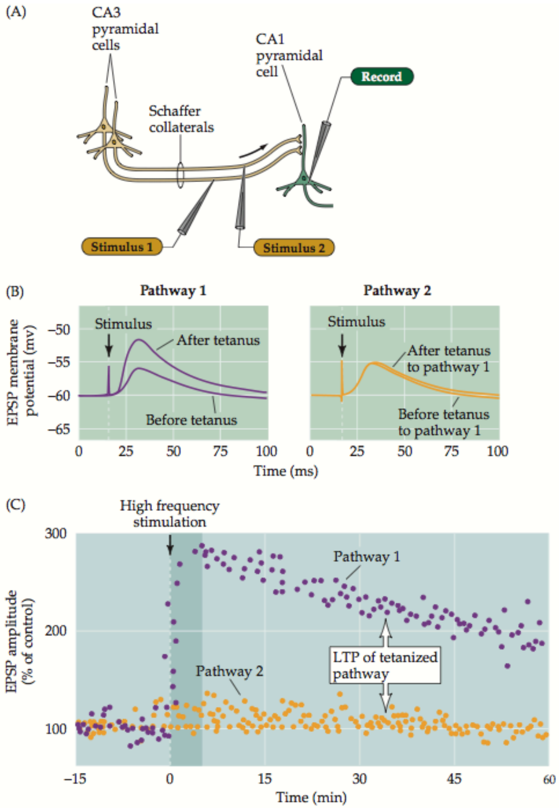

Long-term potentiation of Schaffer collateral-CA1 synapses in the Hippocampus.

(A) Two stimulating electrodes (1 and 2) each activate separate populations of Schaffer collaterals, thus providing test and control synaptic pathways.

(B) Left: Synaptic responses recorded in a CA1 neuron in response to single stimuli of synaptic pathway 1, minutes before and one hour after a high-frequency train of stimuli. Right: Responses produced by stimulating synaptic pathway 2, which did not receive high-frequency stimulation, is unchanged.

(C) The time course of changes in the amplitude of EPSPs evoked by stimulation of pathways 1 and 2. High-frequency stimulation of pathway 1 causes a prolonged enhancement of the EPSPs in this pathway (purple).

Molecular Mechanisms Underlying LTP

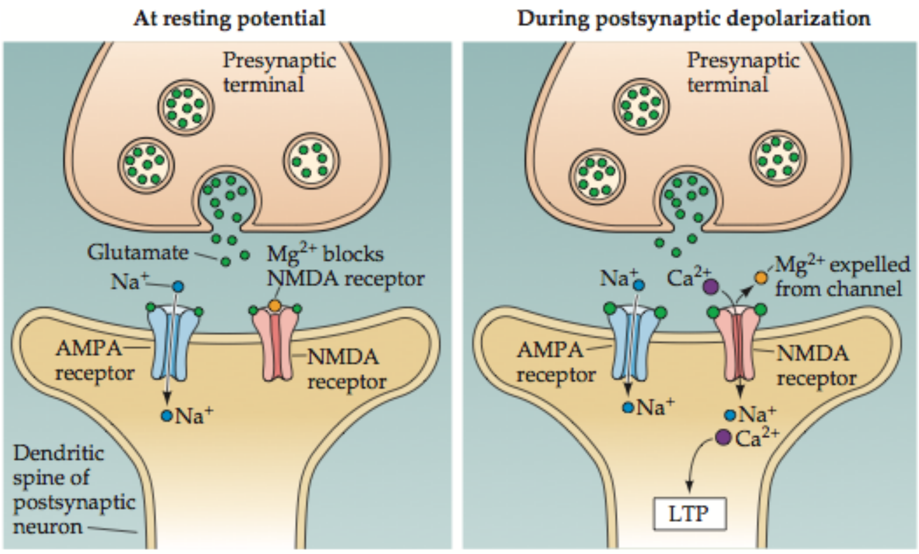

The NMDA receptor channel can open only during depolarization of the postsynaptic neuron from its normal resting level. Depolarization expels Mg^{2+} from the NMDA channel, allowing current to flow into the postsynaptic cell. This leads to Ca^{2+} entry, which in turn triggers LTP.

Basic Hebb rule

$$\tau_w \frac{d\mathbf{w}}{dt} = \mathbf{u} \cdot v$$

The right side of this equation can be interpreted as a measure of the probability that the pre- and postsynaptic neurons both fire during a small time interval.

τ_w is a time constant that controls the rate at which the weights change. This time constant should be much larger than the time constant of the neural dynamics.

Basic Hebb rule

Simple illustration of the Hebb rule:

Before:

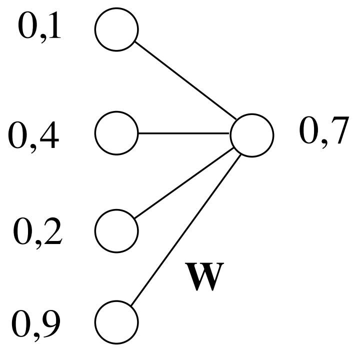

After:

$$\Delta w = \frac{1}{\tau_w} \cdot u \cdot v$$

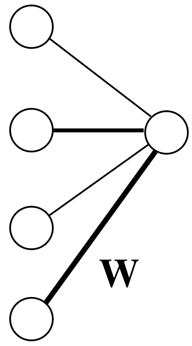

Basic Hebb rule

output v

weights $\mathbf{w}$

input $\mathbf{u}$

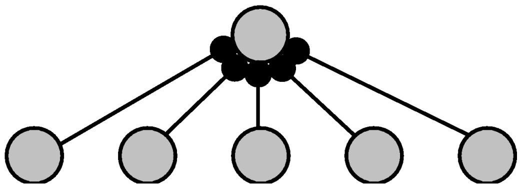

$$\tau_w \frac{d\mathbf{w}}{dt} = \mathbf{u} \cdot v$$

when averaged over the inputs used during training:

$$\tau_w \frac{d\mathbf{w}}{dt} = \langle \mathbf{u} \cdot v \rangle$$

$\langle \ \rangle$ denotes averages over the ensemble of input patterns presented during training

With v=**wu**:

$$\tau_w \frac{d\mathbf{w}}{dt} = \mathbf{Q}\mathbf{w} \quad \text{with} \quad \mathbf{Q} = \langle \mathbf{u} \cdot \mathbf{u} \rangle$$

The basic Hebb rule is called a correlation based learning rule

Covariance rule

The basic Hebb rule only explains long term potentation (LTP). The active decay of the weights, long term depression (LTD), can be modeled by a covariance rule:

$$\tau_w \frac{d\mathbf{w}}{dt} = (v - \theta_v)\mathbf{u} \qquad \tau_w \frac{d\mathbf{w}}{dt} = (\mathbf{u} - \boldsymbol{\theta}_u)v$$

θ_v and θ_u denote thresholds that can be determined according to a temporal or population mean with respect to *u* or *v*.

It is also possible to combine the presynaptic and postsynaptic threshold rules by subtracting thresholds from both the **u** and v terms, but this has the undesirable feature of predicting LTP when pre- and postsynaptic activity levels are both low.

Thus, determine LTD dependent on the presynaptic activity only.

$$\tau_w \frac{d\boldsymbol{w}}{dt} = (v - \theta_v)^+ (\boldsymbol{u} - \boldsymbol{\theta}_u)$$

Covariance rule

With v=**wu** we can write the covariance rule as:

$$\tau_w \frac{d\mathbf{w}}{dt} = \mathbf{C}\mathbf{w} \quad \text{with } \mathbf{C} = \left\langle (\mathbf{u} - \langle \mathbf{u} \rangle) \cdot (\mathbf{u} - \langle \mathbf{u} \rangle) \right\rangle = \langle \mathbf{u}\mathbf{u} \rangle - \langle \mathbf{u} \rangle^2$$

where **C** is the input covariance matrix.

Note, this does not mean that both covariance learning rules are identical.

Problems of basic Hebb and covariance rules

Basic Hebb rule:

$$\tau_w \frac{d\mathbf{w}}{dt} = \mathbf{u} \cdot v$$

The weights can increase without bound!

Solution: Additional constraint for the weights.

$$\sum_j w_{ij} = C \qquad \sqrt{\sum_j w_{ij}^2} = C$$

- induces also competition among the weights (which is not the case when the weight is just limited by an upper bound)
- Putative biological mechanisms unclear

Hebb with normalization (Oja learning rule)

$$\tau_w \frac{d\mathbf{w}}{dt} = \mathbf{u} \cdot v - \alpha v^2 \mathbf{w}$$

- This rule implicitly provides a constraint that the sum of the squares is constant.
- The normalization it imposes is called multiplicative because the amount of modification induced by the second term is proportional to ***w***.
- This rule extracts the largest principal component (after the mean) as shown by Oja (1982).

Stability analysis (dot product with 2**w**)

$$\tau_w \frac{d|\mathbf{w}|^2}{dt} = 2v^2(1 - \alpha|\mathbf{w}|^2)$$

- $|\mathbf{w}|^2$ will relax over time to the value $1/\alpha$.
- The Oja learning rule also induces competition between the different weights because, when one weight increases, the maintenance of a constant length for the weight vector forces other weights to decrease.

BCM rule

The Bienenstock, Cooper, Munro (BCM) learning rule introduces a dynamic threshold θ_w to stabilize learning.

output v

weights $\mathbf{w}$

input $\mathbf{u}$

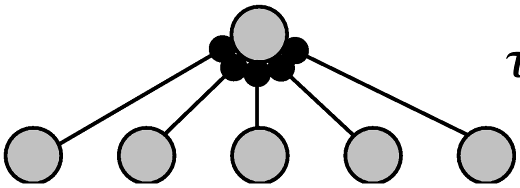

$$\tau_w \frac{d\mathbf{w}}{dt} = \phi(v,\theta_w)\mathbf{u}$$

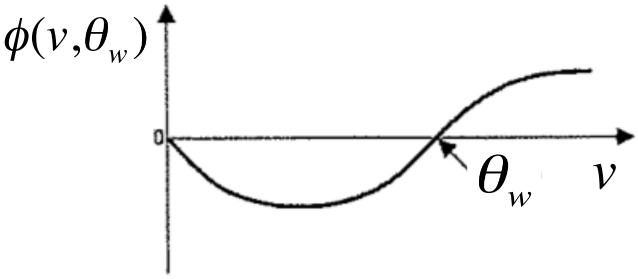

The function ϕ should be negative for small arguments and positive for larger ones (see left).

For example:

$$\phi(v,\theta_w) = v(v-\theta_w)$$

BCM rule

The originally proposed adaptation of the threshold was the square of the temporal average over the past history of the cell (Bienenstock, et al., 1982):

$$\theta_w = (\bar{v})^2 \qquad \text{with} \qquad \bar{v} = \frac{1}{\tau_\theta}\int_{-\infty}^{t} v(t')e^{-\frac{t-t'}{\tau_\theta}}dt'$$

A more stable solution has been later given by Intrator and Cooper (1992):

$$\theta_w = \overline{v^2} = \frac{1}{\tau_\theta}\int_{-\infty}^{t} v^2(t')e^{-\frac{t-t'}{\tau_\theta}}dt' \qquad \text{or} \qquad \tau_\theta \frac{d\theta_v}{dt} = v^2 - \theta_v$$

$$\tau_\theta < \tau_w$$

BCM rule - LTD

The BCM learning rule predicts monosynaptic Long-Term Depression (LTD), i.e. the weakening of synapses not simply as a result of an increase elsewhere due to normalization (heterosynaptic LTD).

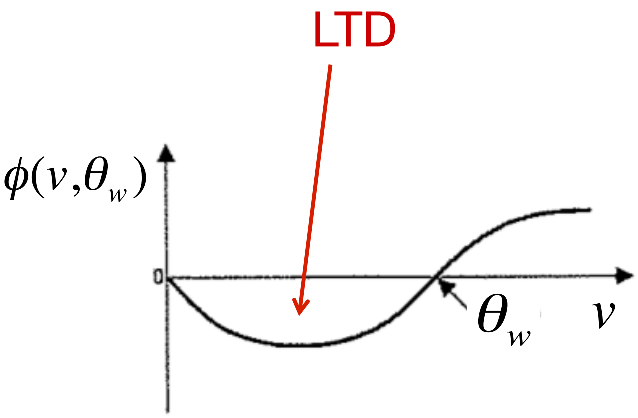

$$\tau_w \frac{d\mathbf{w}}{dt} = \phi(v,\theta_w)\mathbf{u}$$

BCM rule - LTD

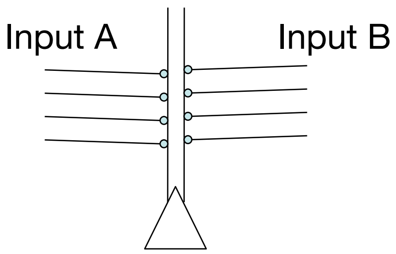

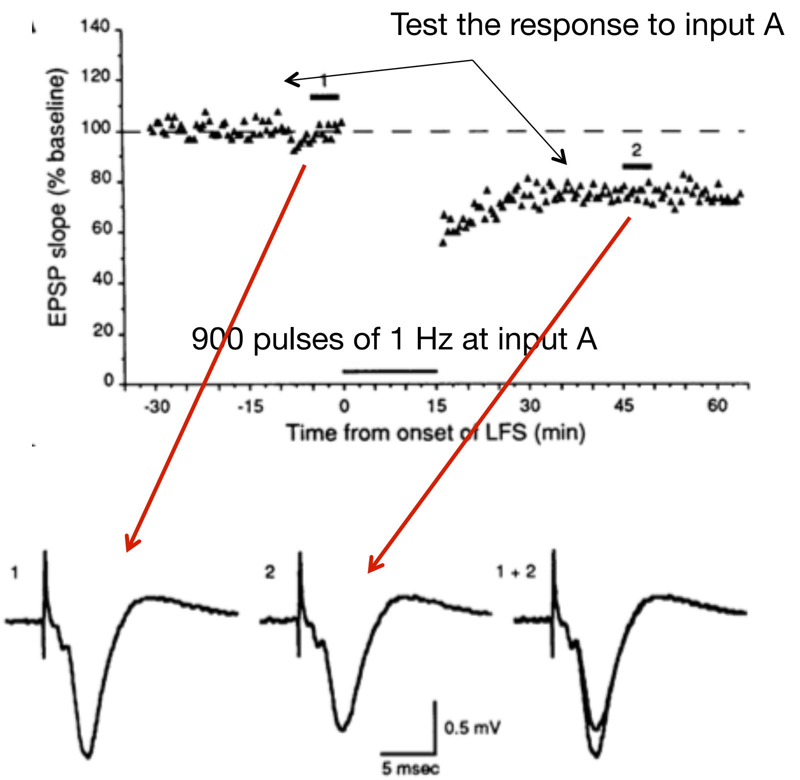

Dudek & Bear, PNAS, 1992

BCM rule - LTD

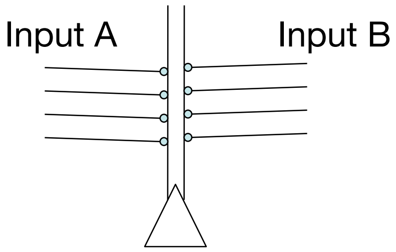

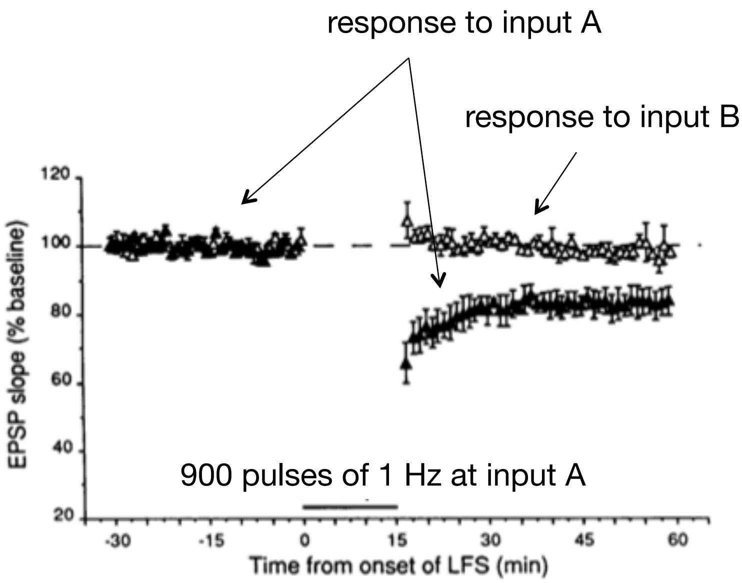

Test for homosynaptic LTD

Dudek & Bear, PNAS, 1992

BCM rule - LTD

Model:

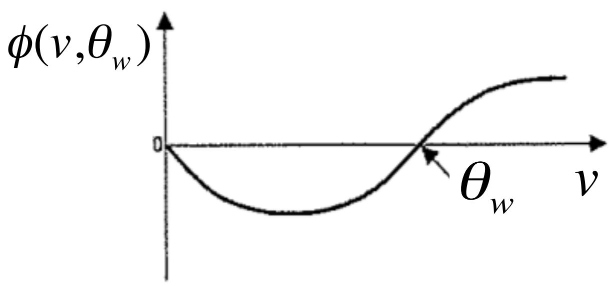

Data:

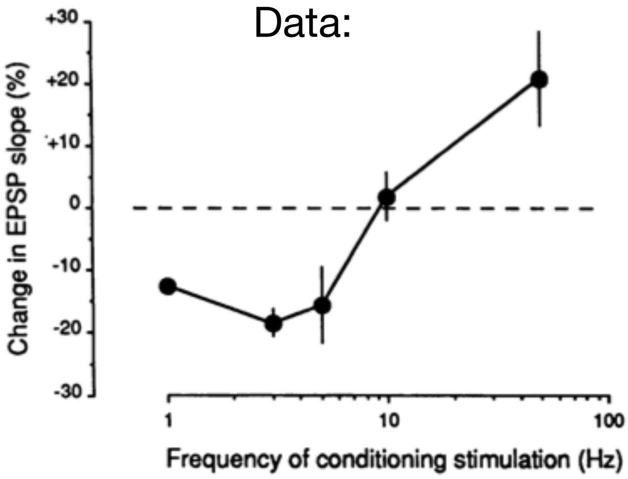

Comparison of model and experiment. Data: Mean (and SEM) effect of 900 pulses delivered at various frequencies on the response measured 30 min after the pulses.

Dudek & Bear, PNAS, 1992

BCM rule – modifiable threshold

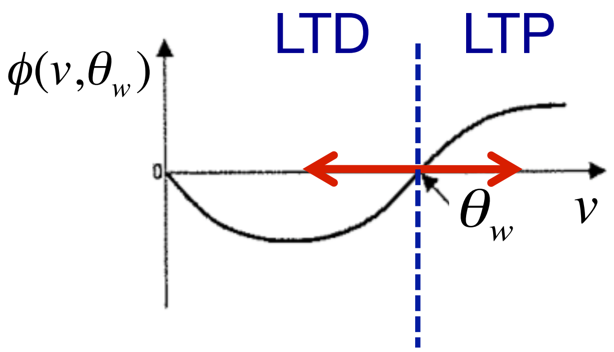

The model predicts that the modification threshold depends on the previous history of the neuron.

Abraham et al., PNAS, 2001

BCM rule – modifiable threshold

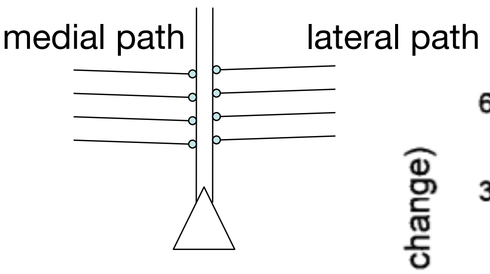

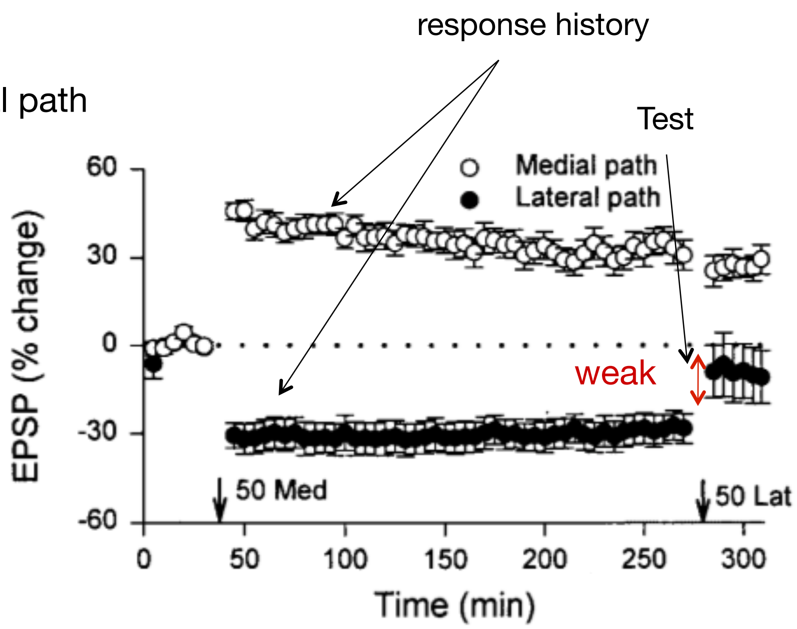

The initial activation of the medial path did not allow strong learning due to an activation of the lateral path.

Abraham et al., PNAS, 2001

BCM rule – modifiable threshold

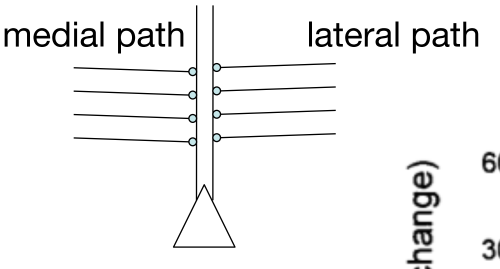

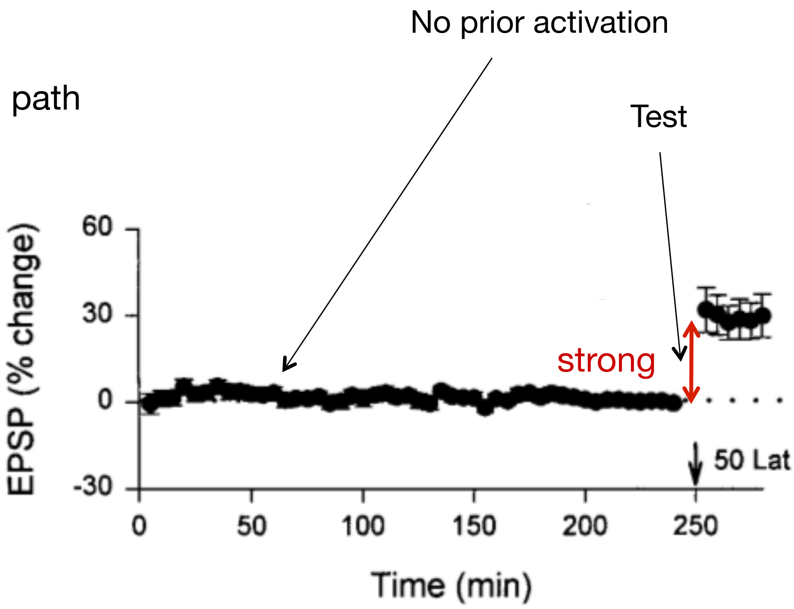

The control experiment shows a stronger increase in weight.

Abraham et al., PNAS, 2001

BCM rule – modifiable threshold

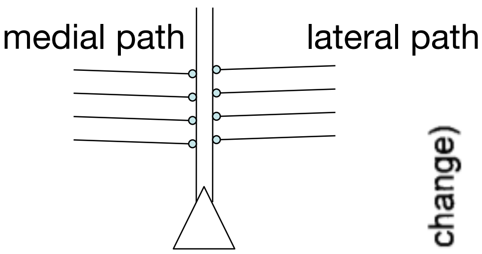

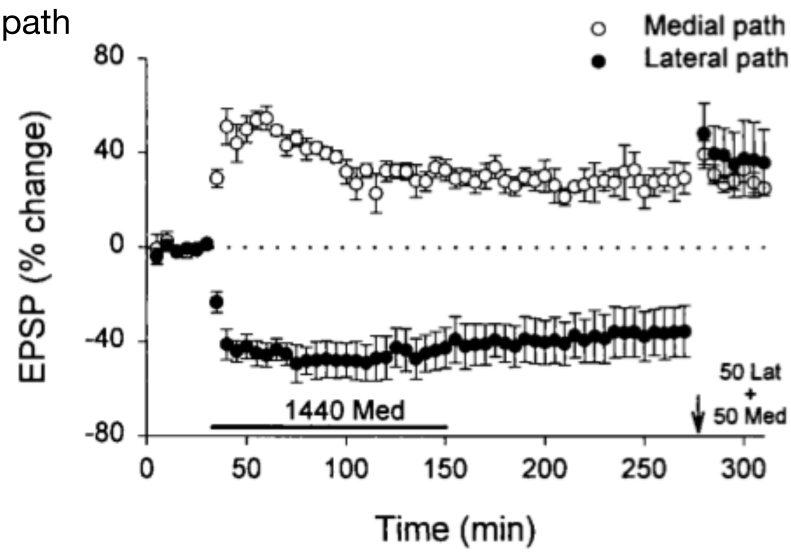

Abraham et al., PNAS, 2001

Pulses at both sites allow a strong increase similar to the level without prior activation.

Extension to multiple output neurons

The previous discussion has been extremely simplified, since in most cases we are interested in multiple output neurons.

Take the Oja rule for example. Simply adding another output neuron does not give you the second largest principal component, but again the first one.

Thus, additional mechanisms must ensure that different output neurons learn different features from the input data. For example:

- Larger components can be removed from the input of the neurons that learn the smaller components (Sanger, 1989).
- The postsynaptic cells that show a correlated response develop lateral inhibitory weights.

Generalized Hebbian Learning

Sanger (1989) proposed a parallel iterative learning procedure according to which the larger components are removed from the input of the neurons that learn the smaller components.

$$\tau_w \frac{dw_{ij}}{dt} = \left(u_i - \sum_{k=1}^{j} w_{ik} v_k \right) v_j$$

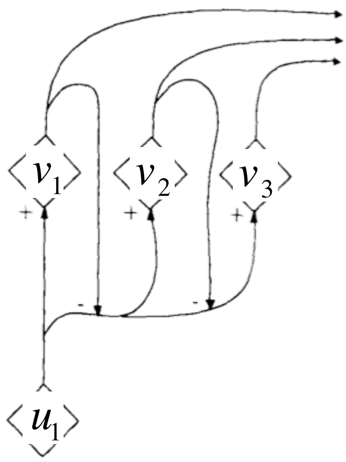

The generalized Hebbian Learning algorithm allows to learn the principal components (Sanger, 1989).

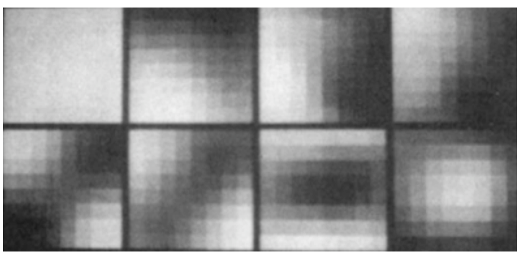

16 components learned from 8x8 image patches (from Sanger, 1989).

Anti-Hebbian Learning

Goodall (1960) proposed to decorrelate the different output units by keeping the forward weights stable and adjusting the lateral weights c_{ij}.

$$\tau_c \frac{d\mathbf{c}}{dt} = -\mathbf{c} + \mathbf{I} - \mathbf{w}\mathbf{u} \cdot \mathbf{v}$$

The identity matrix pushes the diagonal elements to one and ensures that the responses will be decorrelated and whithened. Since this rule assumes a linear model, only second order redundancies are removed.

Anti-Hebbian Learning

Földiák (1990) combined Hebbian and Anti-Hebbian learning mechanisms to learn decorreleted responses.

$$\tau_c \frac{dc_{ij}}{dt} = -v_i v_j - p^2$$

Where p is a constant that was adjusted to match the probability that a component appears in the input pattern.

Here, p=1/8

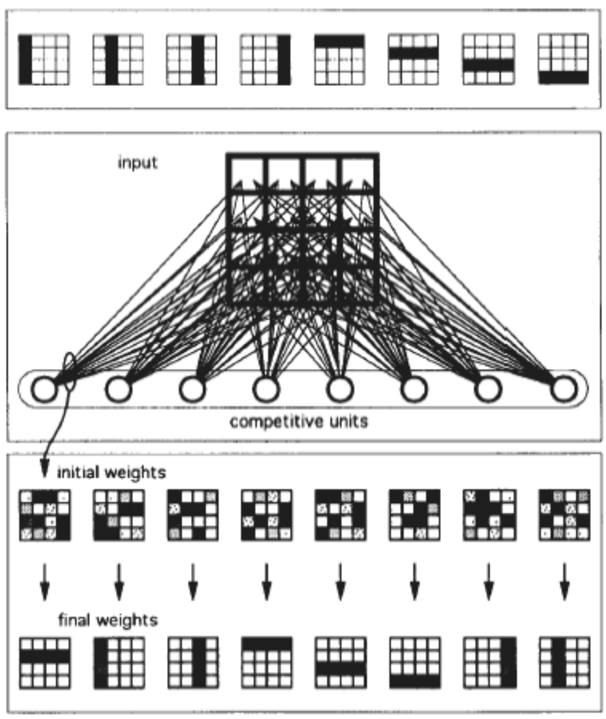

From: P. Foldiak, Learning constancies for object perception. In Perceptual Constancy: Why things look as they do, eds. V Walsh & J J Kulikowski, Cambridge, U.K.: Cambridge Univ. Press , pp. 144-172, 1998.

Anti-Hebbian Learning

Anti-Hebbian learning is a very flexible mechanism that can also be used in non-linear neurons. In this case, higher order dependencies are also considered which will lead to largely independent responses.

Learning invariant representations - the trace learning rule

From: Logothetis NK, Pauls J, Poggio T (1995) Curr Biol., 5:552-63.

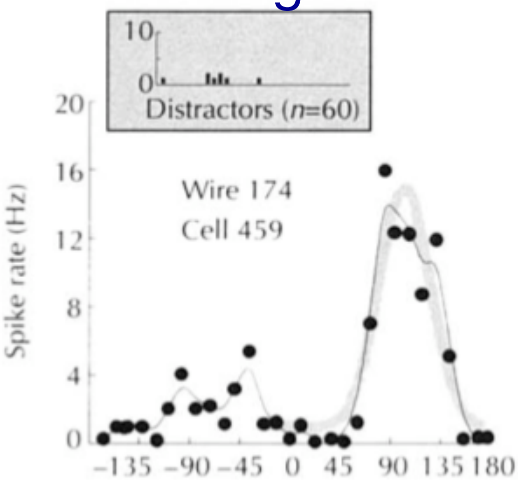

Cells have been found which are selective for a range of views of objects (left: a cell in area IT). This requires that neurons become selective for a more broad range of input patterns from a particular object.

This can be achieved by linking different views of the same object through time. For static neurons this amounts to linking the new input with the activity of the previous input (Földiák, 1991; Stringer & Rolls, 2002):

$$\Delta \mathbf{w}(t) = \frac{1}{\tau_w} \mathbf{u}(t) \cdot v(t-1)$$

Spike-based Hebbian Learning

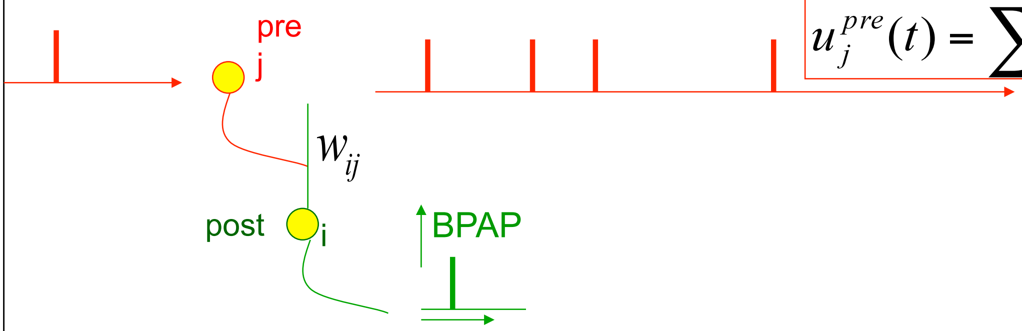

$$u_j^{pre}(t) = \sum_f \delta(t - t_j^f)$$

Spike-based Hebbian Learning

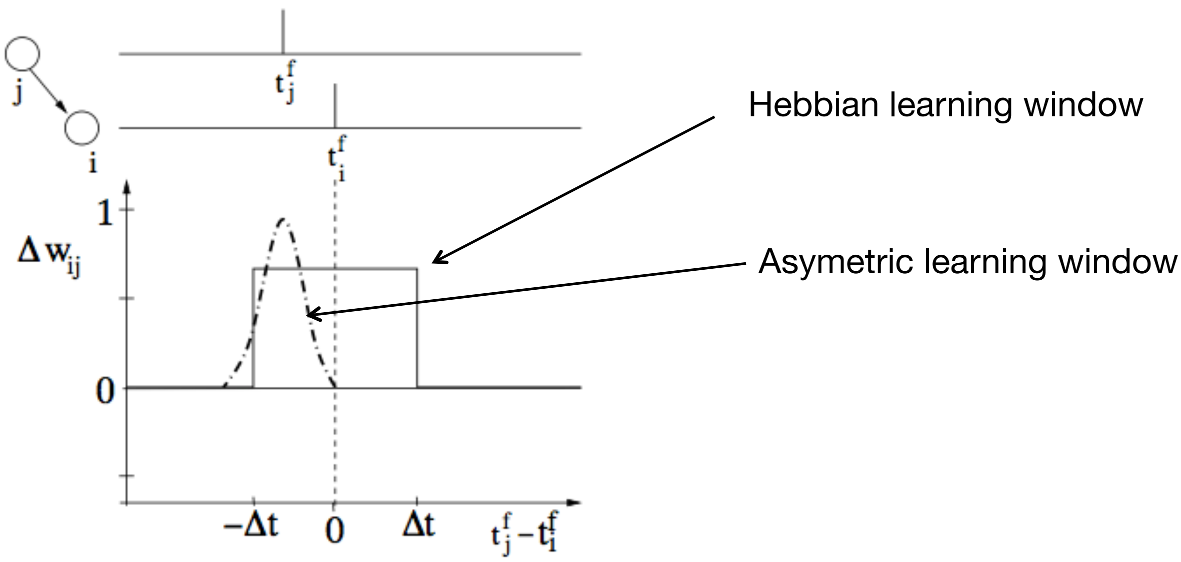

In learning with an asymetric learning window the synapse is strengthened only if the presynaptic spike arrives slightly before the postsynaptic one and is therefore partially 'causal' in firing it.

Spike-Timing Dependent Plasticity (STDP)

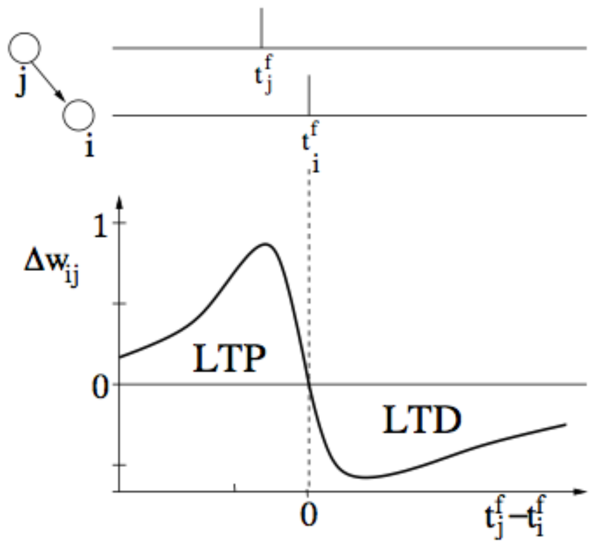

In an asymmetric bi-phasic learning window a synapse is strengthened (long-term potentiation, LTP), if the presynaptic spike arrives slightly before the postsynaptic one, but is decreased (long-term depression LTD), if the timing is reversed.

Spike-Timing Dependent Plasticity (STDP)

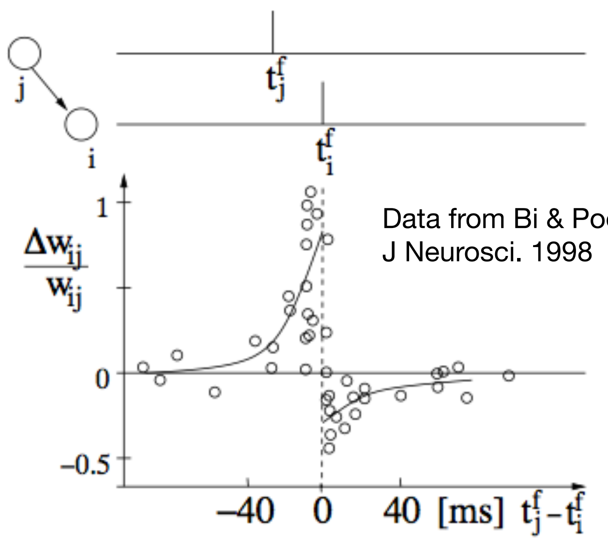

Data from Bi & Poo, J Neurosci. 1998

In an asymmetric bi-phasic learning window a synapse is strengthened (long-term potentiation, LTP), if the presynaptic spike arrives slightly before the postsynaptic one, but is decreased (long-term depression LTD), if the timing is reversed.

Spike-Time Dependent Plasticity

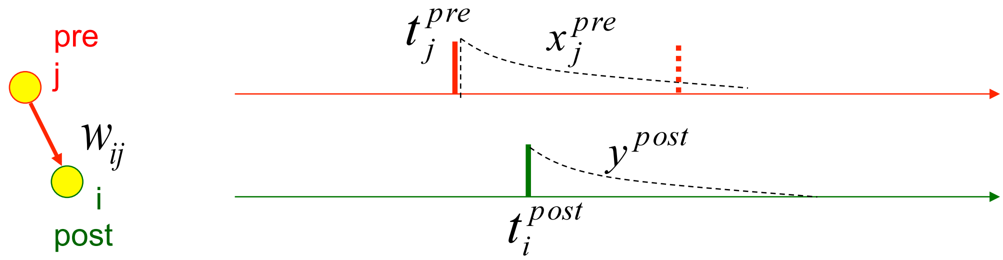

$$\tau_+ \frac{d}{dt} x_j^{pre} = -x_j^{pre} + \delta(t - t_j^{pre})$$ jump at presyn. spike

$$\tau_- \frac{d}{dt} y^{post} = -y^{post} + \delta(t - t_i^{post})$$ jump at postsyn. spike

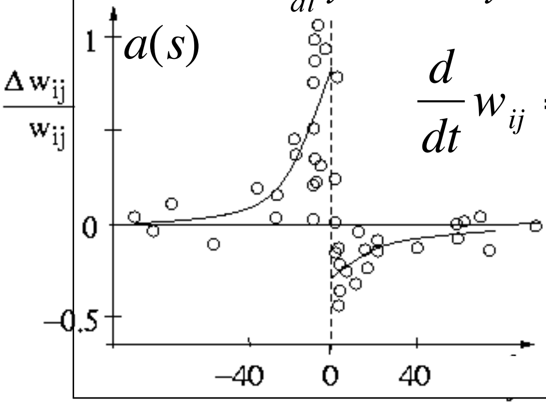

$$\frac{d}{dt} w_{ij} = a(s) x_j^{pre} \delta(t - t_i^{post}) + a(s) y_j^{post} \delta(t - t_j^{pre})$$

pre-before-post post-before-pre

$$s = t_j^{pre} - t_i^{post}$$

References:

- Oja, E (1982) A simplified neuron model as a principal component analyzer. J. Mathematical Biology 15:267-273.
- Bienenstock, EL, Cooper, LN, Munro, PW (1982) Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J Neurosci.* 2:32-48.
- Cooper, L. N., Intrator, N., Blais, B. S., and Shouval, H. Z. (2004). Theory of cortical plasticity. World Scientific, New Jersey.
- Földiák P (1990) Forming sparse representations by local anti-Hebbian learning. Biol. Cybern. 64:165-70.
- Földiák P (1991) Learning invariance from transformation sequences. Neural Computation, 3:194-200.
- Sanger, TD (1989) Optimal unsupervised learning in a single-layer linear feedforward neural network. Neural Networks 2:459-473.
- Stringer SM, Rolls ET (2002) Invariant object recognition in the visual system with novel views of 3D objects. Neural Comput., 14:2585-96.