Learning in Neural Networks: Detailed Models of Synaptic Plasticity

EPFL

Lecture 11

Synaptic plasticity

Wulfram Gerstner

EPFL, Lausanne, Switzerland

Reading for plasticity:
NEURONAL DYNAMICS
- Ch. 19.1-19.3

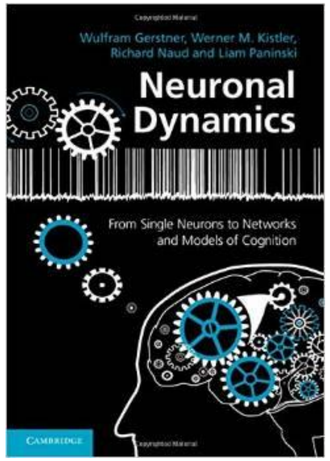

Cambridge Univ. Press

1. **Synaptic plasticity**
   motivation and aims
2. **Classification of plasticity**
   short-term vs. long-term
   unsupervised vs. reward modulated
3. **Model of short-term plasticity**
4. **Models of long-term plasticity**
   - Hebbian learning rules
   - Bienenstock-Cooper-Munro rule
5. **Spiking Models of plasticity**
   - STDP as Hebbian learning
   - Model of STDP: synaptic traces
6. **From STDP to rate models**
7. **Triplet STDP model**
8. **Clopath model**

1. Behavioral Learning – and Memory

Learning actions:

→ riding a bicycle

Remembering facts

→ previous president of the US

→ name of your mother

Remembering episodes

→ first day at EPFL

Building useful representations

→ beyond PCA/ICA/Clustering

→ RL with 3-factor rule needs ‘good representation!’

Review: Behavioral Learning – and synaptic plasticity

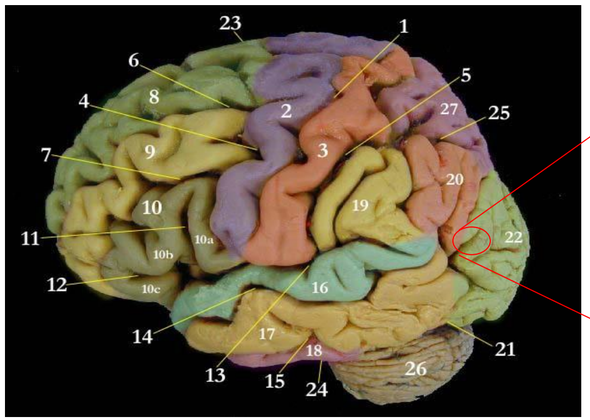

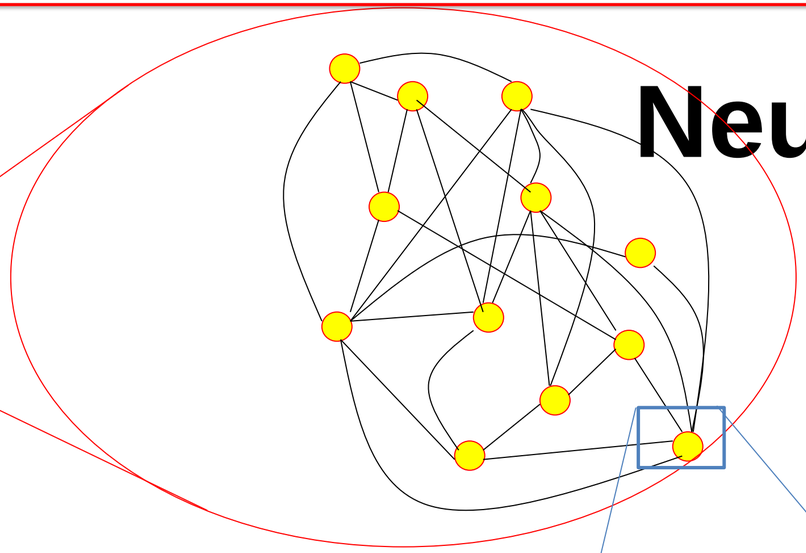

Neurons

Amplitude of Postsynaptic Potential (PSP)

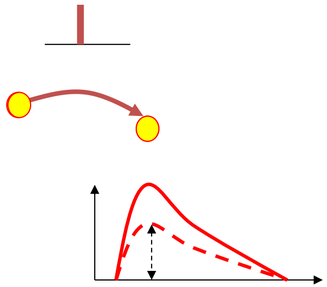

Synapse

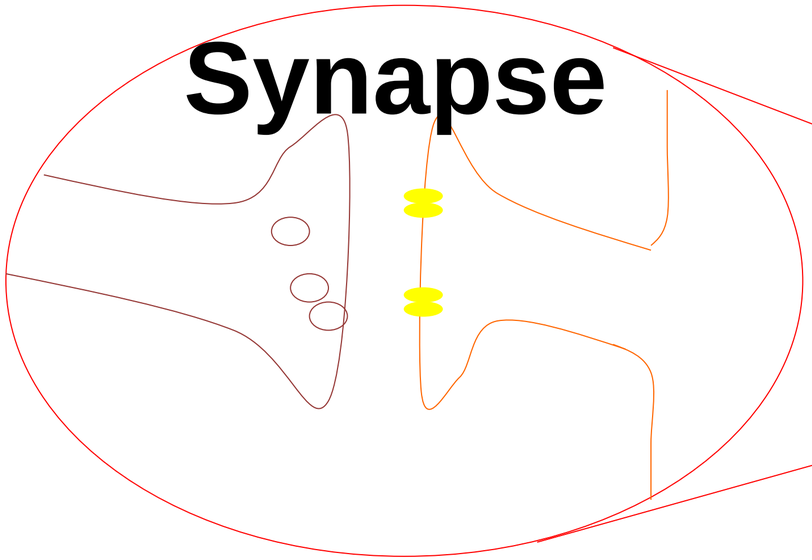

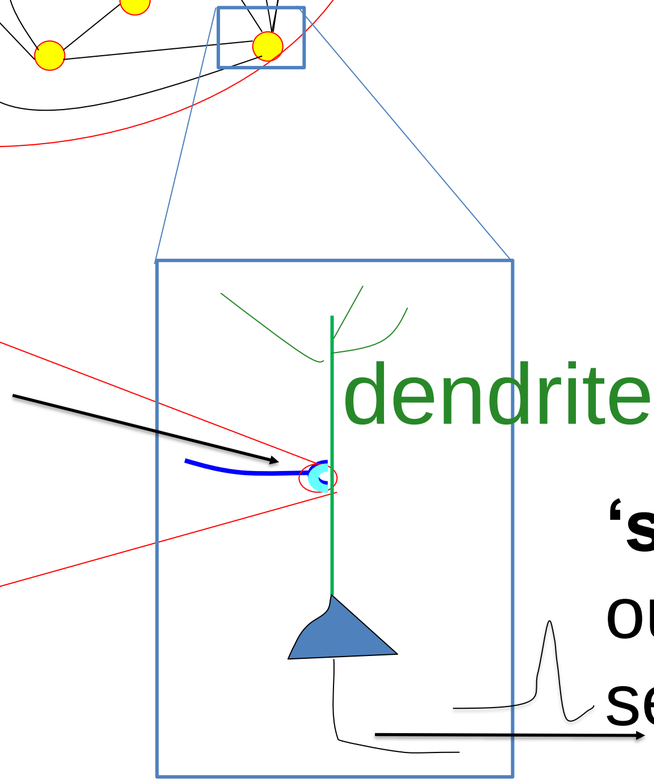

‘spike’:
output signal (pulse)
sent to other neurons

Synaptic Plasticity = Change in Connection Strength

Previous slide.

When we observe learning on the level of behavior (we get better at tennis), then this implies that something has changed in our brain:
The contact points between neurons (called synapses) have changed. Synaptic changes manifest themselves as a change in connections strength.

Synaptic plasticity describes the phenomena and rules of synaptic changes.

The connection strength can be measured by the
- amplitude of the postsynaptic potential (PSP)
- by physical size of the synapse (in particular the spine, see next slide)

Important:
Neurons communicate with each other by short electrical pulses, often called ‘spikes’.

Review: Synaptic plasticity – structural changes

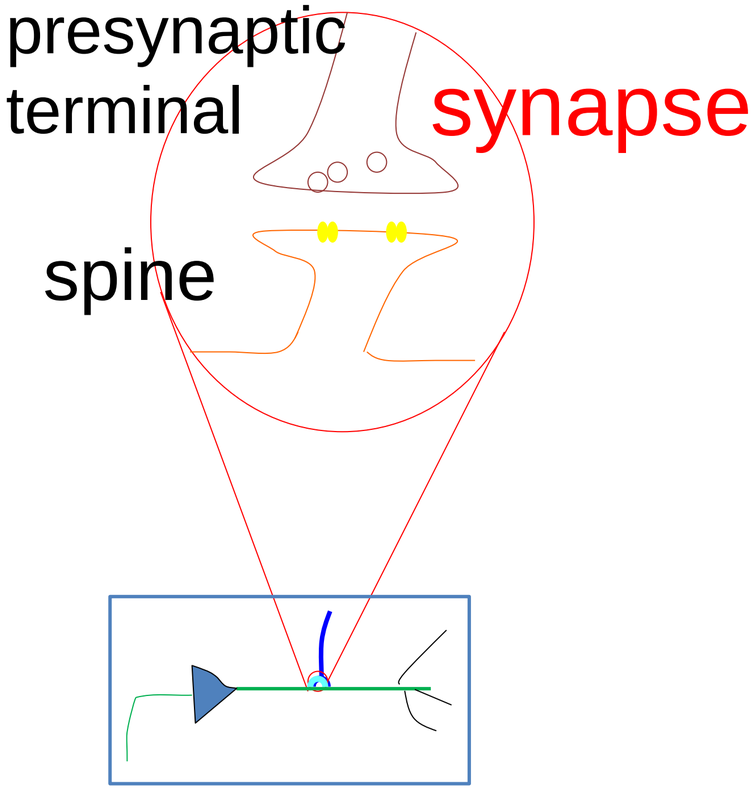

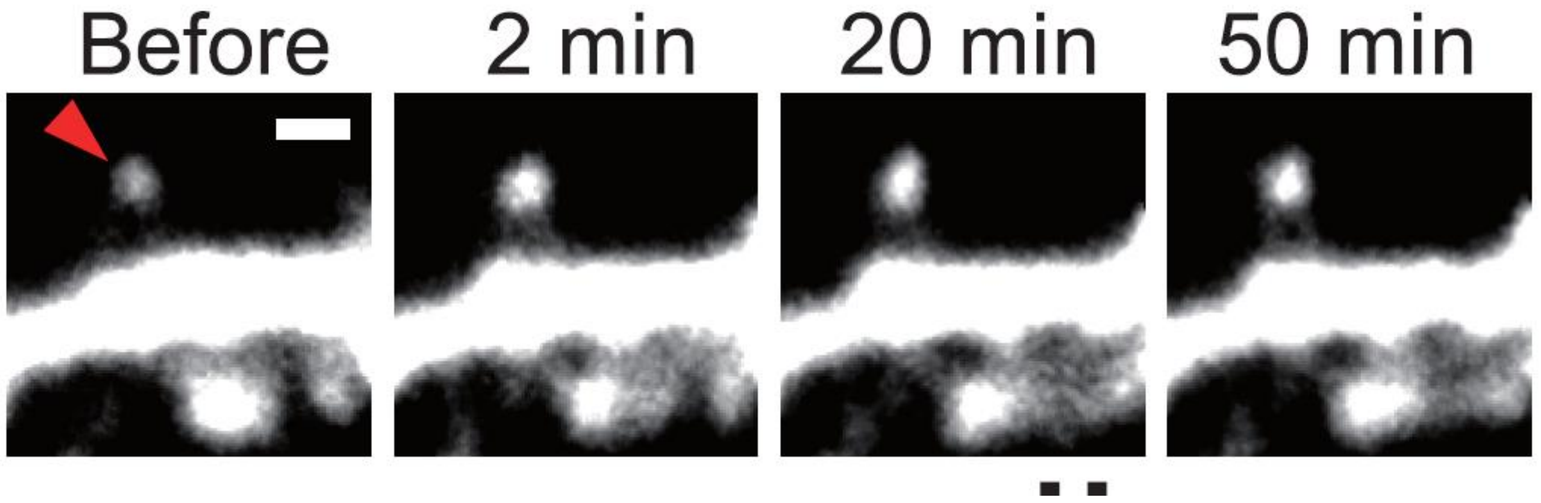

Yagishita et al.
Science, 2014

larger synapse/larger spine head
→ stronger synapse
→ changes induced by appropriate stimulation protocol

Previous slide.

The synaptic connection consists of two parts. The end of an axonal branch coming from the sending neuron; and the counterpart, a protrusion on the dendrite of the receiving neuron, called spine.

We refer to the sending neuron as presynaptic and to the receiving one as postsynaptic.

A change in the connection strength is observable with imaging methods as an increase in the size of the spine. The bigger spine remains big for a long time (here observed for nearly one hour).

1. synaptic plasticity – molecular changes

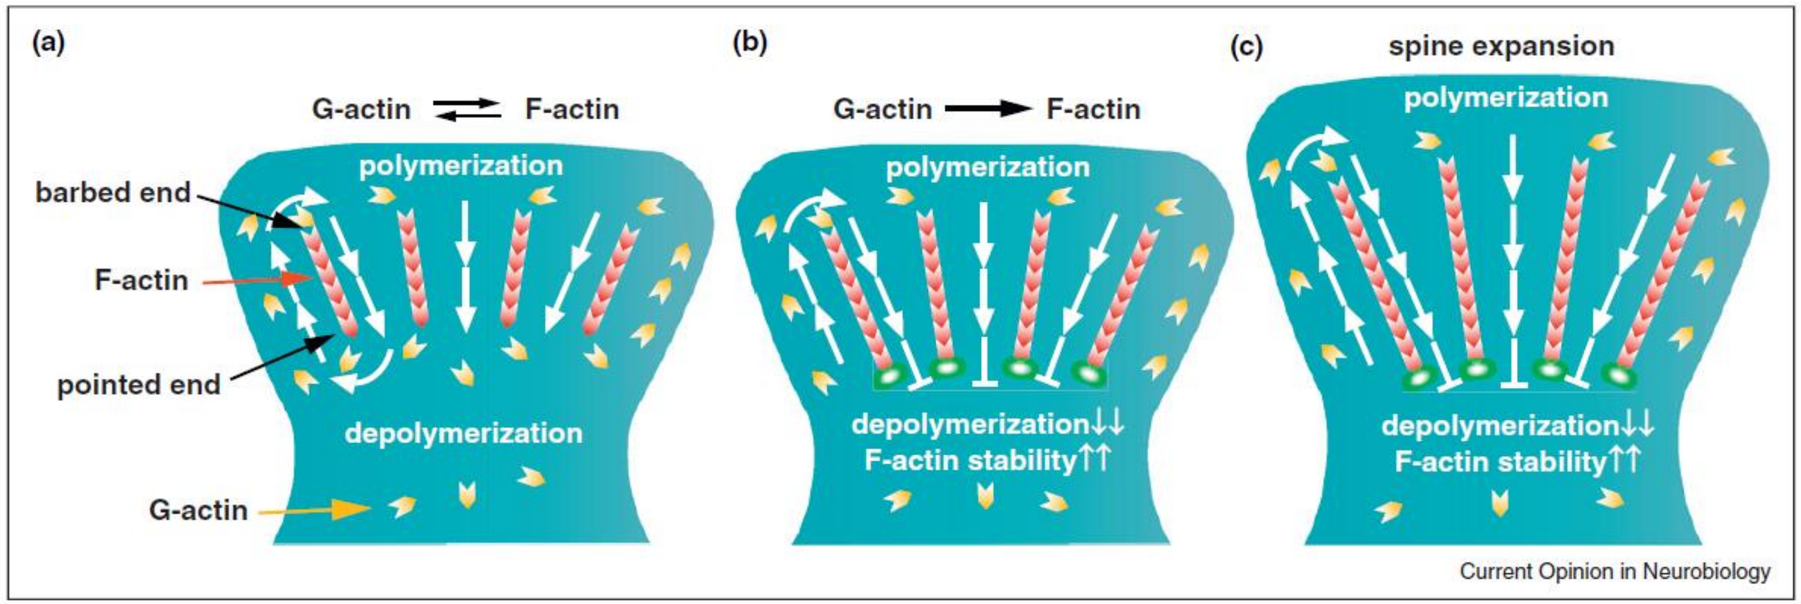

Bosch et al. 2012, Curr. Opinion Neurobiol.

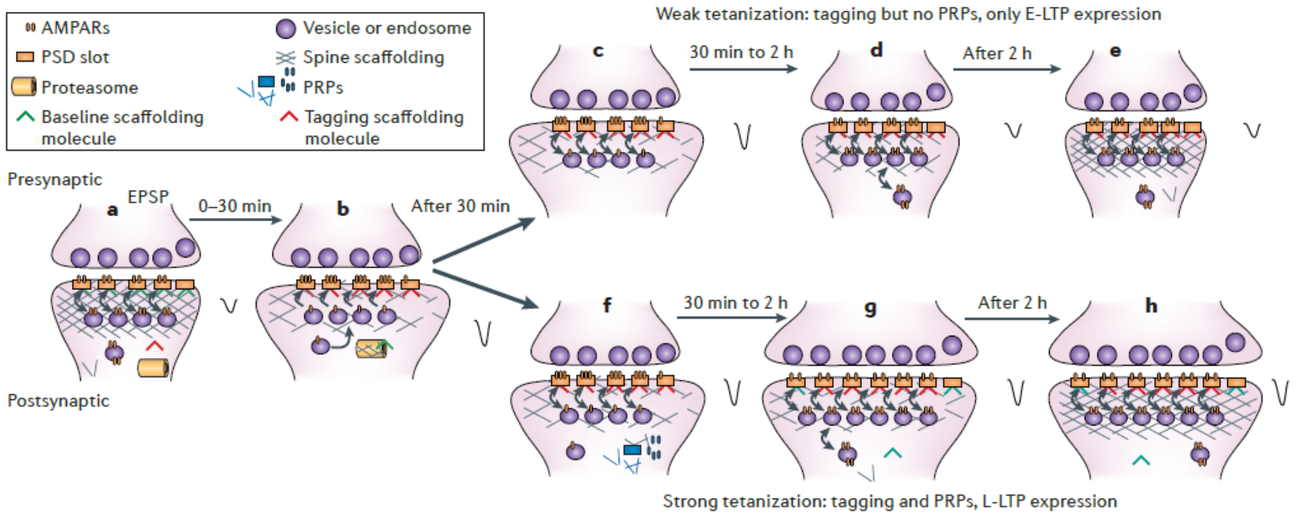

Redondo and Morris 2011, Nature Rev. Neurosci.

Synaptic changes are implemented by molecular pathways.

Previous slide.
The molecular processes involved in building a synapses are highly complex.
Every few years a review paper in a biological journal attempts to summarize the present understanding of molecular interactions.
The molecular processes are not part of this class.

Lower figure:
The form of the synapse is maintained by structural molecules (‘scaffolding molecules’) and making a synapses bigger also means restabilizing the new structure, indicated by the grid of diagonal stripes.

1. synaptic plasticity – molecular changes

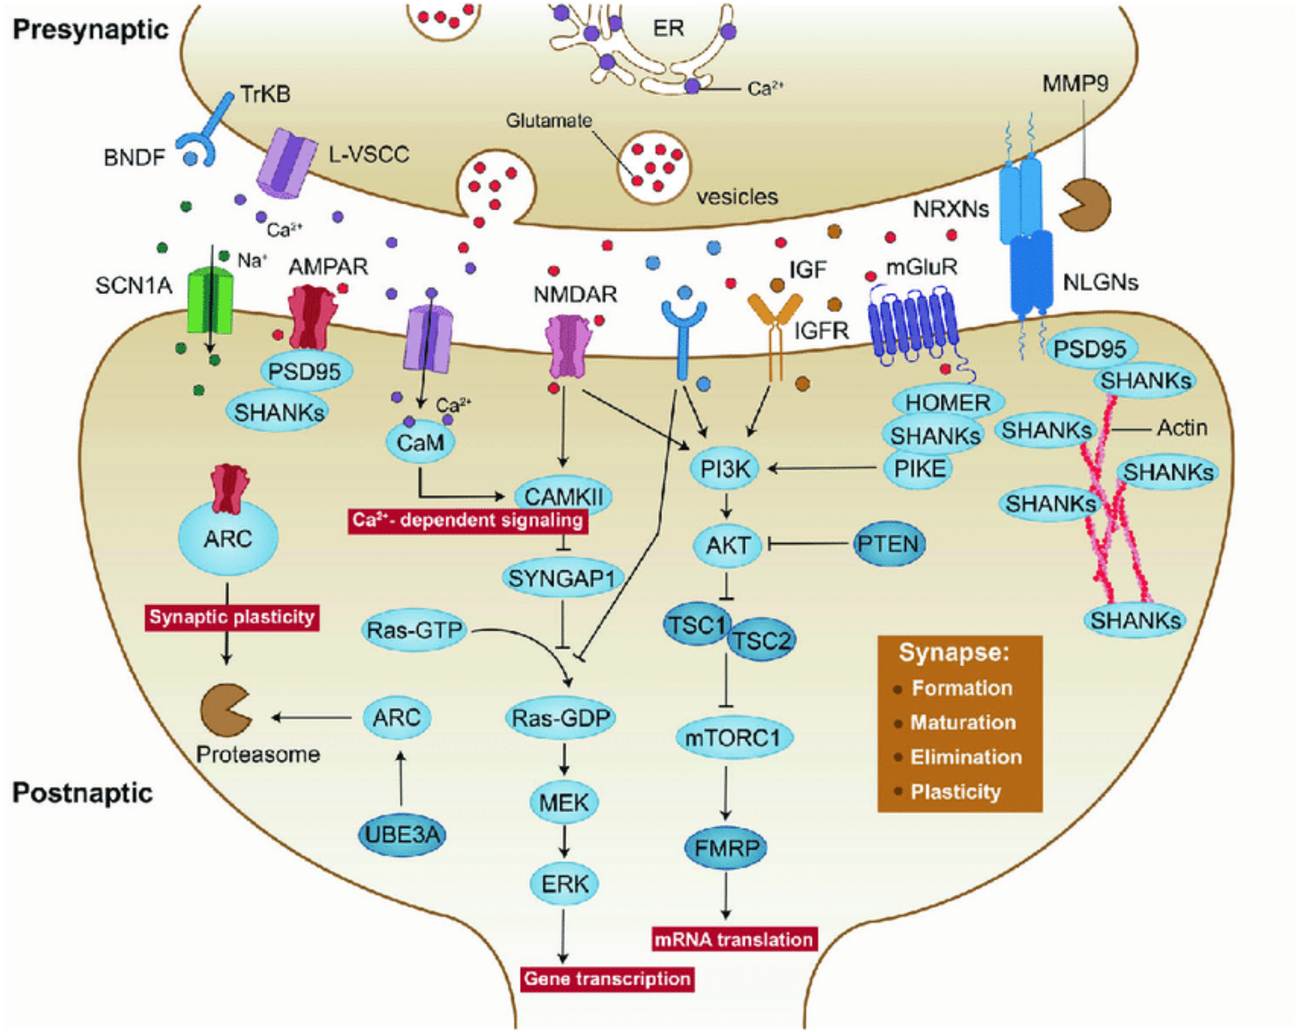

https://www.researchgate.net/publication/361914216/figure/fig3/AS:1176818883264513@1657586894637/Molecular-pathways-implicated-in-synaptic-function-for-ASD-At-the-excitatory-synapse.png

Previous slide.
The molecular processes involved in building a synapses are highly complex.

On the presynaptic side of an excitatory synapse, vesicles filled with the neurotransmitter glutamate are waiting. When a spike arrives at a presynaptic terminal, the increased voltage causes the vesicle to merge with the membrane and neurotransmitter is spilled into the synaptic cleft.

On the postsynaptic side of the synapse, there are two main types of channels awaiting the glutamate: AMPA and NMDA (these names have something to do with the type of chemical that reacts with the channel in addition to glutamate). If an AMPA channel opens, sodium enters. If a NMDA channel opens, sodium and calcium enter into the cell.

The names of a few important molecules that are involved in synaptic changes are:
- CaMKII,
- PI3K

PKMzeta and Shank molecules play a role during maintenance (long-term stability over many hours) of the synapse

Synaptic transmission: NMDA receptors as coincidence detectors

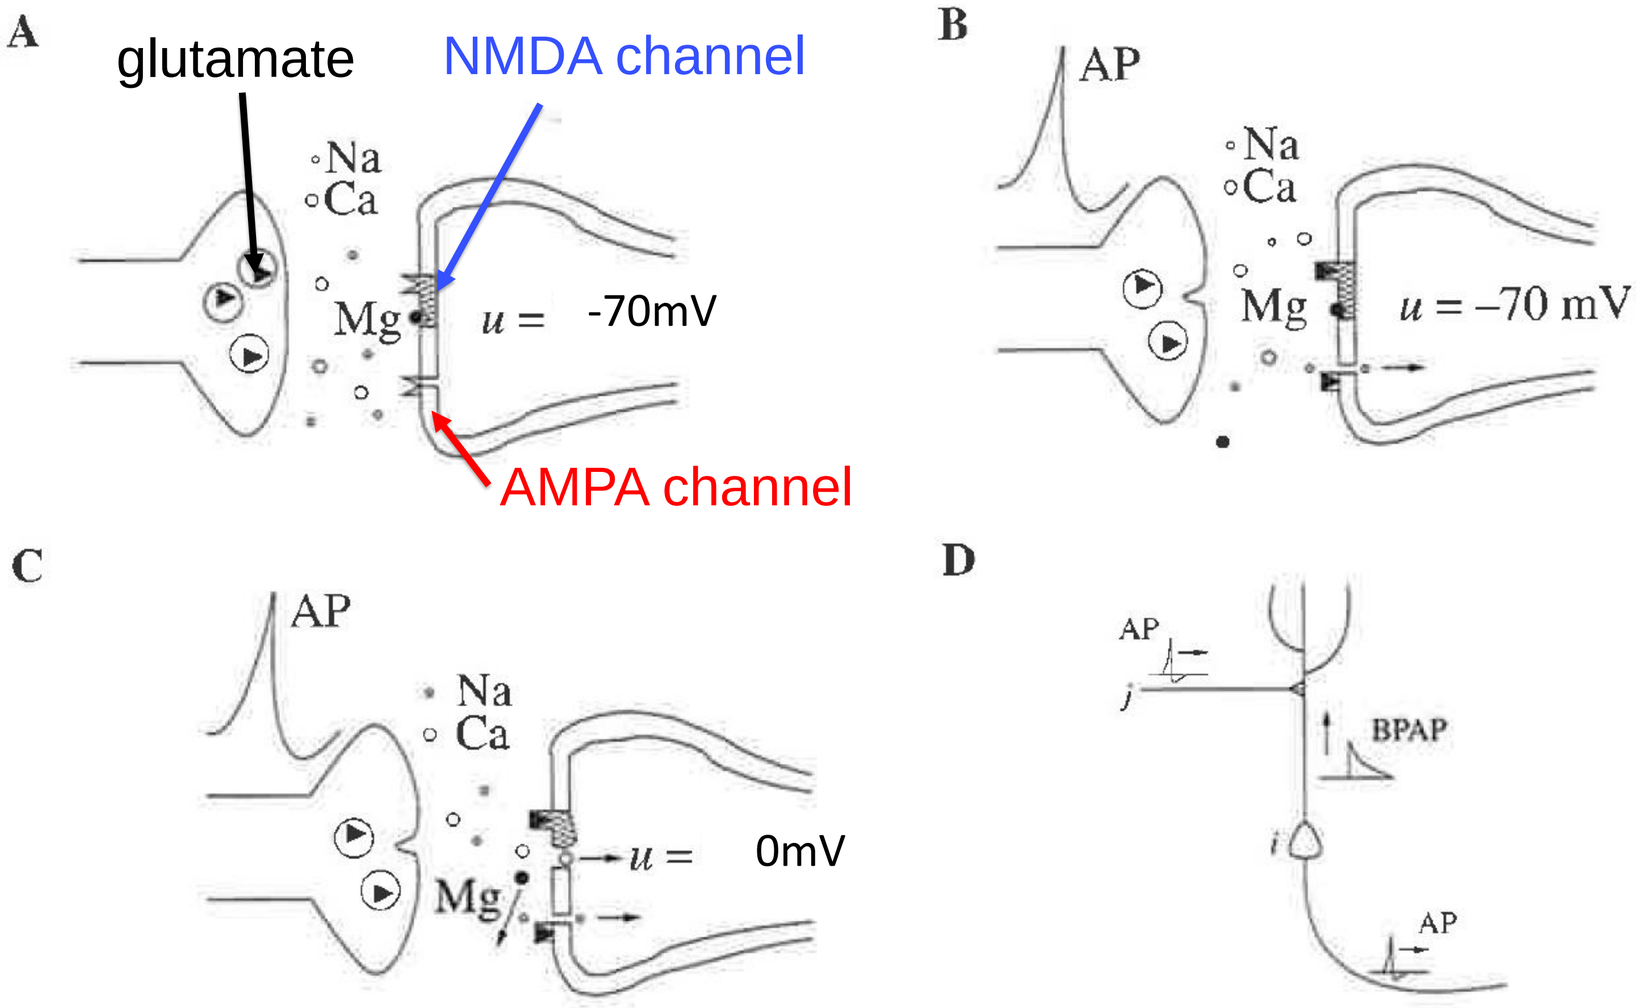

Previous slide.
AMPA channel open if glutamate docks onto the channel receptor.
NMDA channel open if
(i) glutamate docks onto the channel receptor.
(ii) High voltage removes the magnesium block.

The high voltage could be caused by **previous** activity of the postsynaptic neuron.

Coincidence detection of pre and post occurs in NMDA channels because,
Glutamate is only released into the synaptic cleft after **presynaptic spike arrival**.
The magnesium block is only removed of high **voltage of the postsynaptic neuron.**

Figure 10.10: NMDA-synapse. **A**. Vesicles in the presynaptic terminal contain glutamate as a neurotransmitter (filled triangles). At resting potential, the NMDA receptor mediated channel (hatched) is blocked by magnesium (filled circle). **B**. If an action potential (AP) arrives at the presynaptic terminal the vesicle merges with the cell membrane, glutamate diffuses into the synaptic cleft, and binds to NMDA and non-NMDA receptors on the postsynaptic membrane. At resting potential, the NMDA receptor mediated channel remains blocked by magnesium whereas the non-NMDA channel opens (bottom). **C**. If the membrane of the postsynaptic neuron is depolarized, the magnesium block is removed and calcium ions can enter into the cell. D. The depolarization of the postsynaptic membrane can be supported by the increased voltage of an earlier backpropagating action potential (BPAP)

2-factor versus 3-factor rules

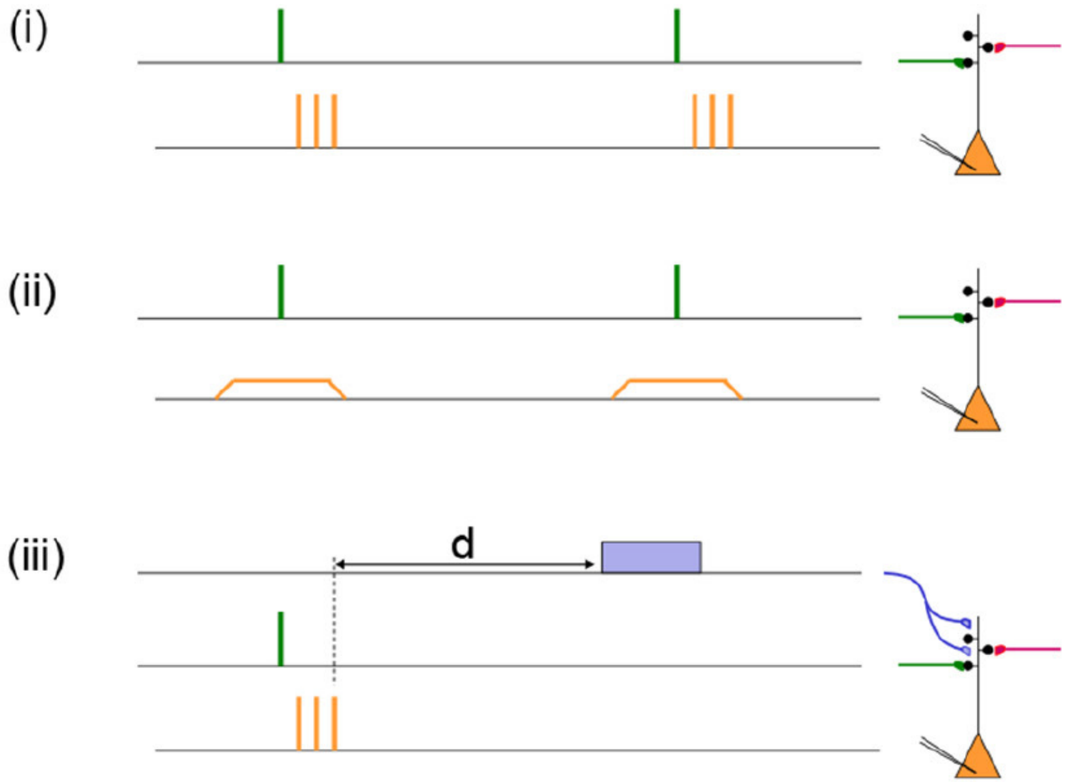

Hebbian: pre=spike
'post' = spikes

Hebbian:pre=spike
'post' = voltage

3-factor: pre+post+
neuromodulator (success/reward)

3-factor = Hebbian combined with
(potentiall delayed) Neuromodulator:
→ Reward based learning/reinforcement learning

Image: Gerstner et al. (2018, review paper in Frontiers)

Previous slide.

Another important distinction is that between two-factor and three-factor rules, where the latter include the action of a neuromodulator.

Hebb rules are essentially two-factor rules.

A first model: Clopath model (1) – static voltage dependence

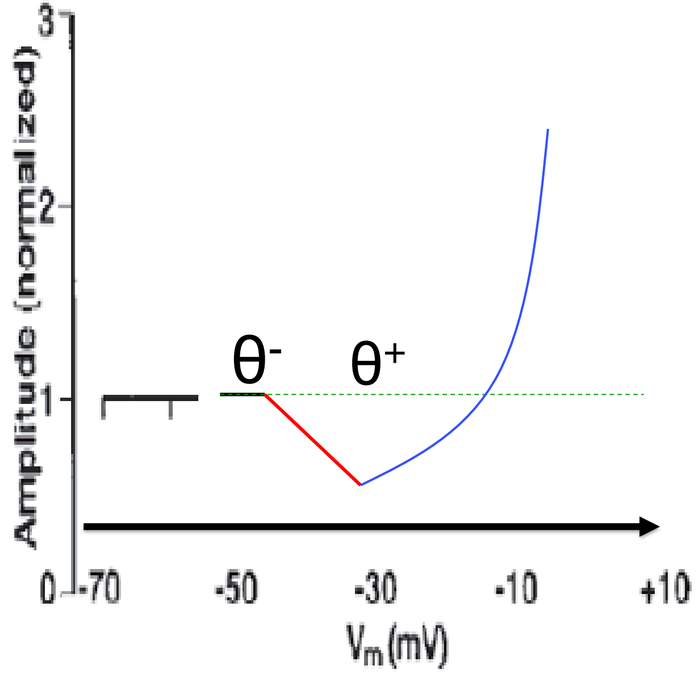

Depression term $-A^- \sigma^-(u)pre$

$-\sigma^-(u)$

θ^-

ReLU

Potentiation term $+A^+ [\sigma^+(u)]^2 pre$

$[\sigma^+(u)]^2$

θ^+

Ngezahayo et al. *J. of Neurosci.*, 2000

Artola, Bröcher, Singer. *Nature* 1990

Clopath et.al., Nature Neuroscience, 2010;

Previous slide.

There are influential experiments in two papers of Artola et al (1990) and Ngezehayo et al (2000) who caused presynaptic spike arrival (green) while the postsynaptic neuron was kept at a constant potential.

Suppose the resting potential of the orange cell is about -60 mV. Then a small constant increase in voltage combined with presynaptic spike arrival leads to depression (LTD) while a higher voltage leads to potentiation (LTP) of the synapse.

For constant voltage, the voltage dependence is described in the Clopath model by two ReLU function of the voltage u with different thresholds.

A first model: Clopath model (2) – dynamics of depression

$$\frac{d}{dt} w_j^- = -A_2^- X_j(t) \sigma^-(\bar{u}_-)$$

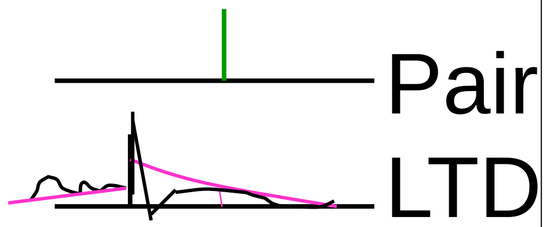

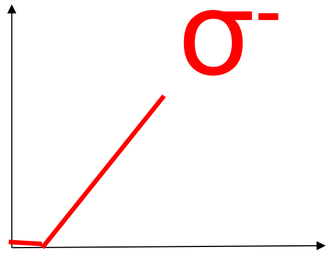

resting potential

pre

Some trace,
calcium?
second messenger?

where

$X_j(t) = \sum_i \delta(t - t_j^i)$ presynaptic spike train

$\tau_{u-} \frac{d}{dt} \bar{u}_- = -\bar{u}_- + u$ postsynaptic u low-pass (τ_{u-})

$\sigma^-(\bar{u}_-) = [\bar{u}_- - \theta^-]^+$ piecewise linear function (ReLU)

Previous slide.

For constant voltage, the voltage dependence is described in the Clopath model by two ReLU function of the voltage u with different thresholds.

If the voltage is NOT constant, the relevant quantity is a low-pass filter of the voltage u (called $\bar{u}_-$) which enters in the ReLU function $\sigma^-(\bar{u}_-)$.
In the Clopath model, each presynaptic spike depresses the synapse proportional to $\sigma^-(\bar{u}_-)$.

The differential equation for LTD (top line) can alternatively be written as

$\Delta w_j^- = -A_2^- \sigma^-(\bar{u}_-)$ if spike arrival at synapse
$\Delta w_j^- = 0$ else

A first model: Clopath model (3) – dynamics of potentiaion

$$\frac{d}{dt}w_j = -A_2^- X_j(t)\sigma^-(\bar{u}_-) + A_3^+ \bar{x}_j^{pre}\sigma^+(u)\sigma^-(\bar{u}_+)$$

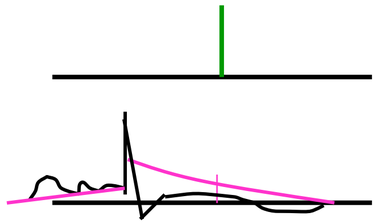

Pre: spike

Post: voltage-av

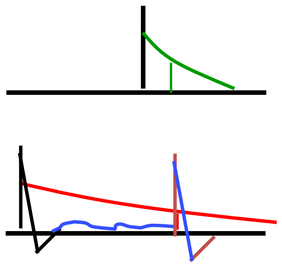

Pre: spike-trace

Post: voltage-now

voltage-av

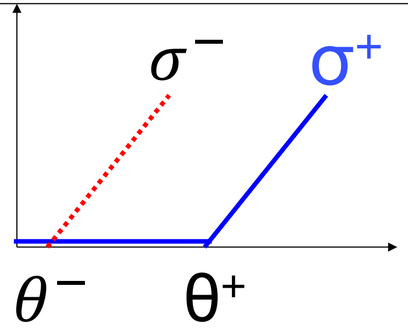

resting potential

firing threshold

Clopath et.al. Nat.NS 2010;

Previous slide.

Depression of synapses can be overwritten by potentiation of synapses.

For potentiation to occur, three conditions need to be met:

(i) A presynaptic spike has arrived a few milliseconds before. This condition is implemented by the trace $\overline{x}_j^{pre}$. The trace steps up by one unit, if a spike arrives and decays exponentially thereafter with a time scale of about 10ms.

(ii) The momentary voltage is above the firing threshold. This condition is implemented by the term $\sigma^+(u)$.

(iii) The low-pass filtered voltage is about the resting potential. This condition is implemented by the term $\sigma^-(\bar{u}_+)$.

A first model: Clopath model summary

Clopath et.al. Nat.NS 2010;

$$\frac{d}{dt}w_j = -A_2^- X_j(t)\sigma^-(\bar{u}_-) + A_3^+ \bar{x}_j^{pre}\sigma^+(u)\sigma^-(\bar{u}_+)$$

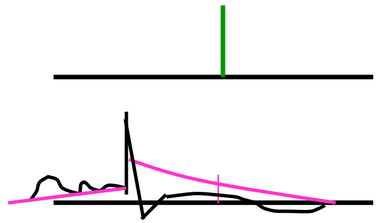

Pre: spike

Post: voltage-av

LTD

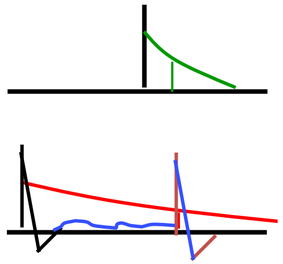

Pre: spike-trace

Post: voltage-now
voltage-av

LTP

LTP (positive change) needs: - presynaptic spike a few ms before
- voltage above threshold now
- voltage above rest previously

LTD (negative change) needs: - voltage > rest a few ms before
- presynaptic spike now

Previous slide. Summary of the Clopath model.

Together the conditions for LTP remove the magnesium block of the NMDA channel (average voltage in the recent past is above slightly above rest), open the channel (because of presynaptic spike arrival that causes glutamate release), and require activity of the postsynaptic neuron (voltage above firing threshold).

Without firing of the postsynaptic neuron only LTD is possible.

1. Synaptic plasticity

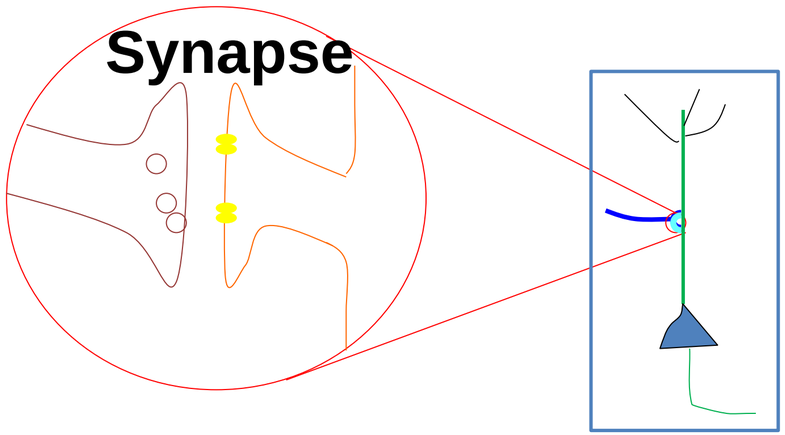

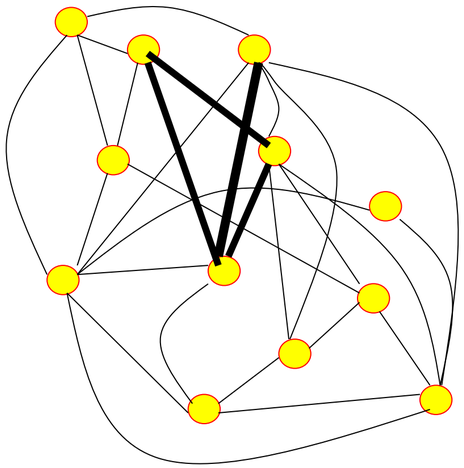

Should enable Learning

- adapt to the statistics of tasks and environments (receptive fields, allocate space, build representations etc)
- memorize facts and episodes
- learn motor tasks

‘control loop’

Should avoid:

- blow-up of activity
- unnecessary use of energy

Aim: models that capture the essence

Previous slide.
Hence synaptic plasticity is a multiscale process.

Moreover it should be useful for different tasks and, despite dynamic changes, does not lead to a catastrophic explosion.

1. Synaptic plasticity: summary

- Synaptic plasticity (= changes of synaptic contact points) are the basis of learning.

- Learning is necessary for a variety of different tasks.

- Learning leads to measurable changes in performance (you get better at a task) and to measurable changes in the brain.

- Calcium entry through the NMDA receptor is thought to be important for many forms of long-term potentiation

- The Clopath model is a first model

Previous slide.
In this lecture we study mathematical models of synaptic plasticity, but before looking at the models, we need to understand the main experimental phenomena.

Learning in Neural Networks: Detailed Models of Synaptic Plasticity

EPFL

Lecture 11

Synaptic plasticity

Wulfram Gerstner

EPFL, Lausanne, Switzerland

Reading for plasticity:
NEURONAL DYNAMICS
- Ch. 19.1-19.3

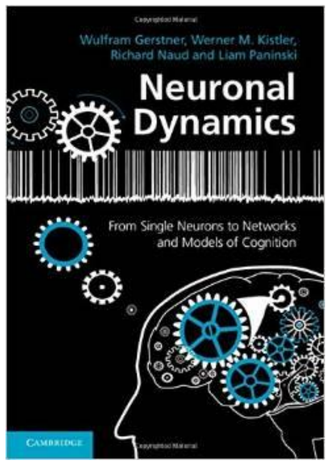

Cambridge Univ. Press

1. **Synaptic plasticity**
   motivation and aims
   Detour: Spiking Neurons
2. **Classification of plasticity**
   short-term vs. long-term
   unsupervised vs. reward modulated
3. **Model of short-term plasticity**
4. **Models of long-term plasticity**
   - Hebbian learning rules
   - Bienenstock-Cooper-Munro rule
5. **Spiking Models of plasticity**
   - STDP as Hebbian learning
   - Model of STDP: synaptic traces
6. **From STDP to rate models**
7. **Triplet STDP model**
8. **Clopath model**

Previous slide.
We need to talk a bit about signal transmission in the brain and introduce spiking neurons.

The brain: signal transmission by spikes

Signal:
action potential (short pulse)

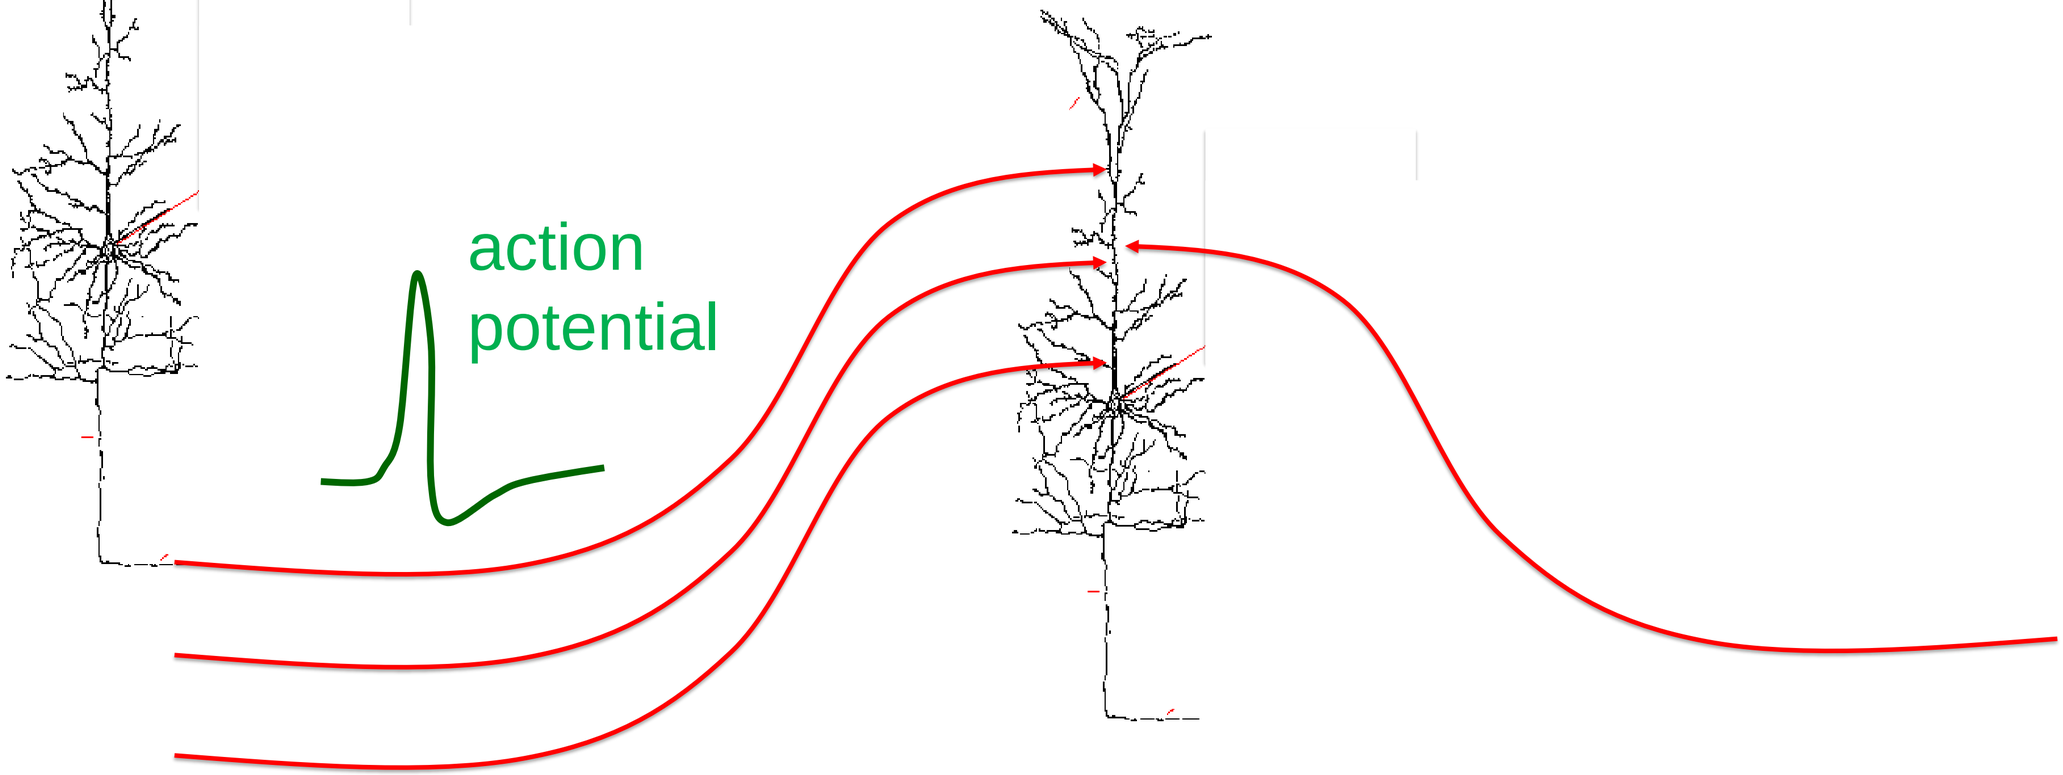

More than 1000 inputs

Previous slide.
Signals are transmitted along the wires (axons). These wires branch out to make contacts with many other neurons.

Each neuron in cortex receives several thousands of wires from other neurons that end in ‘synapses’ (contact points) on the dendritic tree.

Modeling: spiking neurons

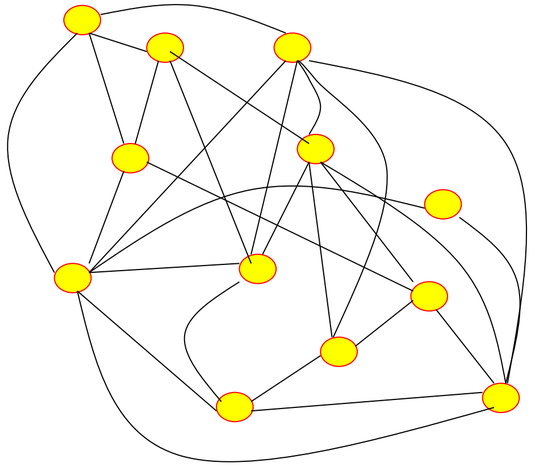

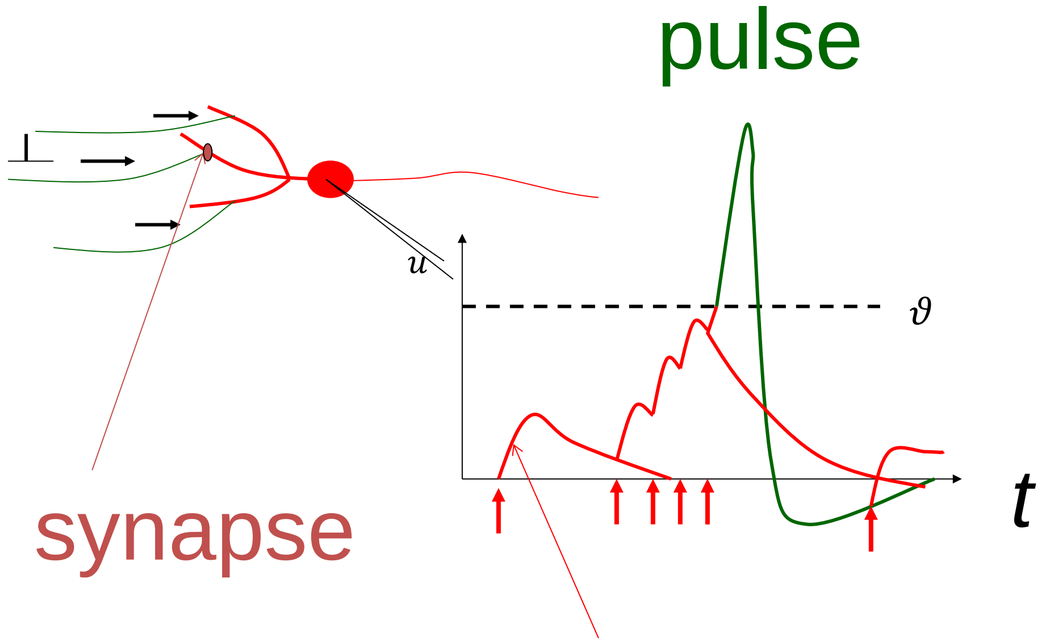

-responses are added
-pulses created at threshold
-transmitted to other

→ Mathematical description

Previous slide.
In the previous part we have seen that response are added and compared with a threshold.

This is the essential ideal that we keep for the abstract mathematical model in the following.

We drop the notion of pulses or spikes and just talk of neurons as active or inactive.

Review: modeling of spiking neurons

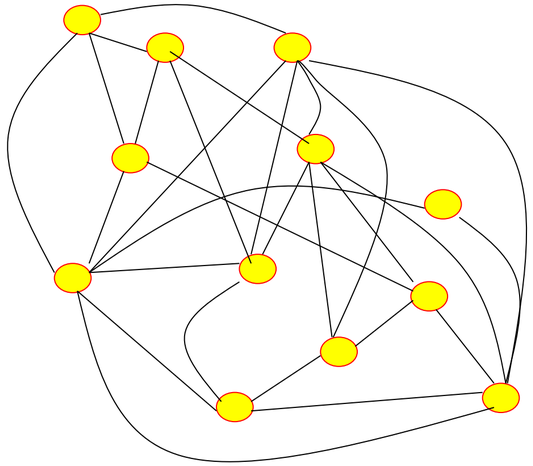

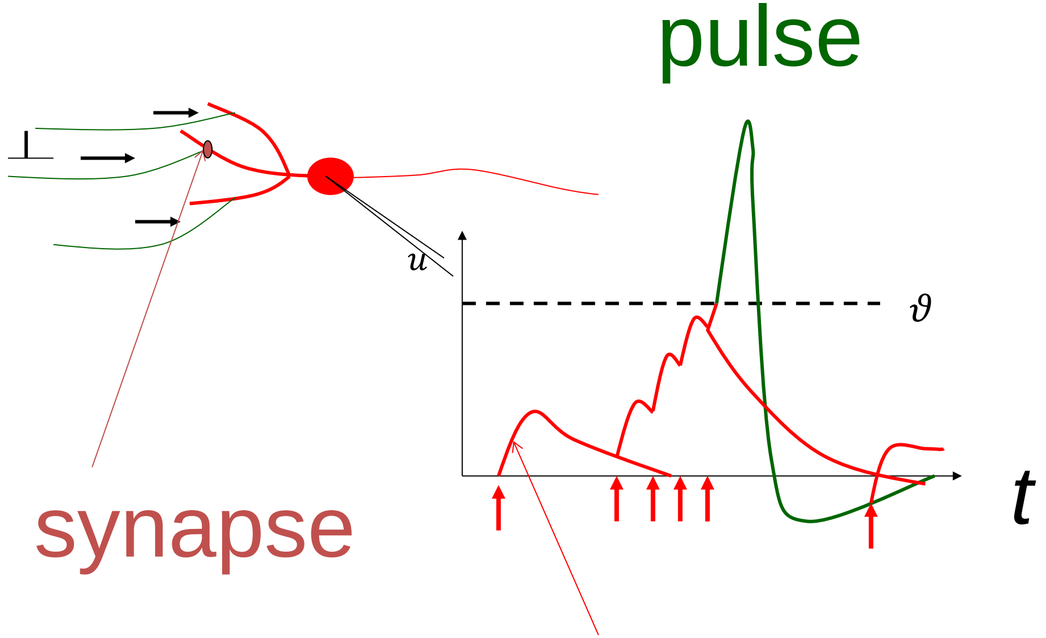

-responses are added
-pulses created at threshold
-transmitted to other
-pulses are ‘unitary events’: shape of spike irrelevant

→ Mathematical description:
Integrate-and-fire neuron

Previous slide.

The fact that responses are added and then compared with a threshold is an aspect that is shared between real neurons, integrate-and-fire neurons, and artificial neurons in ANNs.

This is the essential ideal that we keep for the abstract mathematical model in the following.

Note that spikes are formal events – their duration can be reduced to zero. What matters is the fact whether a pulse is transmitted, yes or no.

Spiking Neural Network – Leaky Integrate-and-Fire Model

(continuous time formulation)

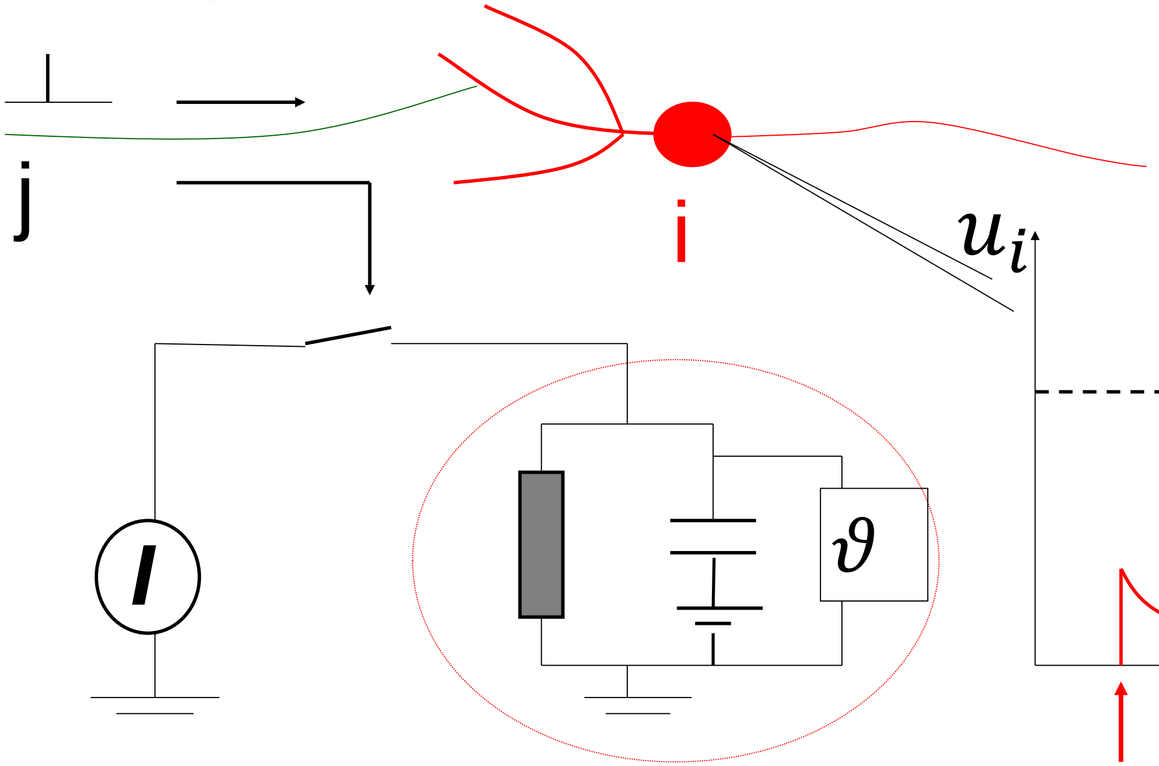

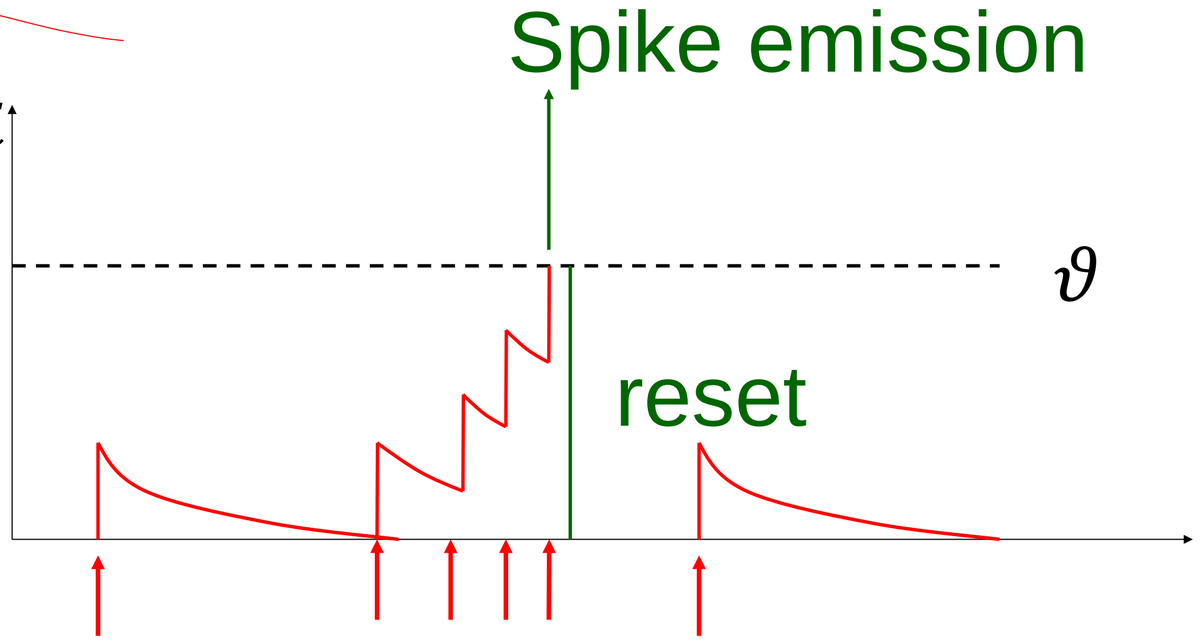

$$\mathrm{I_i(t)} = short\ pulses = \sum_j w_{ij}\ \delta(t - t_j^{pre})$$

$$\tau_m \frac{d}{dt} u_i = -u_i(t) + I_i(t)$$

linear

if $u_i(t) = \vartheta$ note spike+reset to $u(t) = 0$ threshold

Previous slide:

The Leaky integrate-and-fire model written in **continuous time** involves a LINEAR differential equation that can be interpreted as an electrical RC circuit charged by a current $I(t)$. This current $I(t)$ consists of short electrical pulses that present spike arrivals. The $\delta(t - t_j^{pre})$ denotes the Dirac delta function for each presynaptic spike arrival at times t_j^{pre} and w_{ij} are the weights. We can interpret w_{ij} as the charge delivered by the current pulse at time t_j^{pre}.

The linear equation is combined with a NONLINEAR FIRE-and-RESET process. If the variable u ('membrane potential of the neuron') reaches the threshold theta, then u is reset to zero.

Side Note: An electrical RC circuit consists of a capacitance C and a resistor R and has a time constant $\tau = RC$. Therefore after each short current pulse, the voltage (membrane potential) decays exponentially back to zero with time constant $\tau = RC$.

Spiking Neural Network – Leaky Integrate-and-Fire Model

(discrete time formulation)

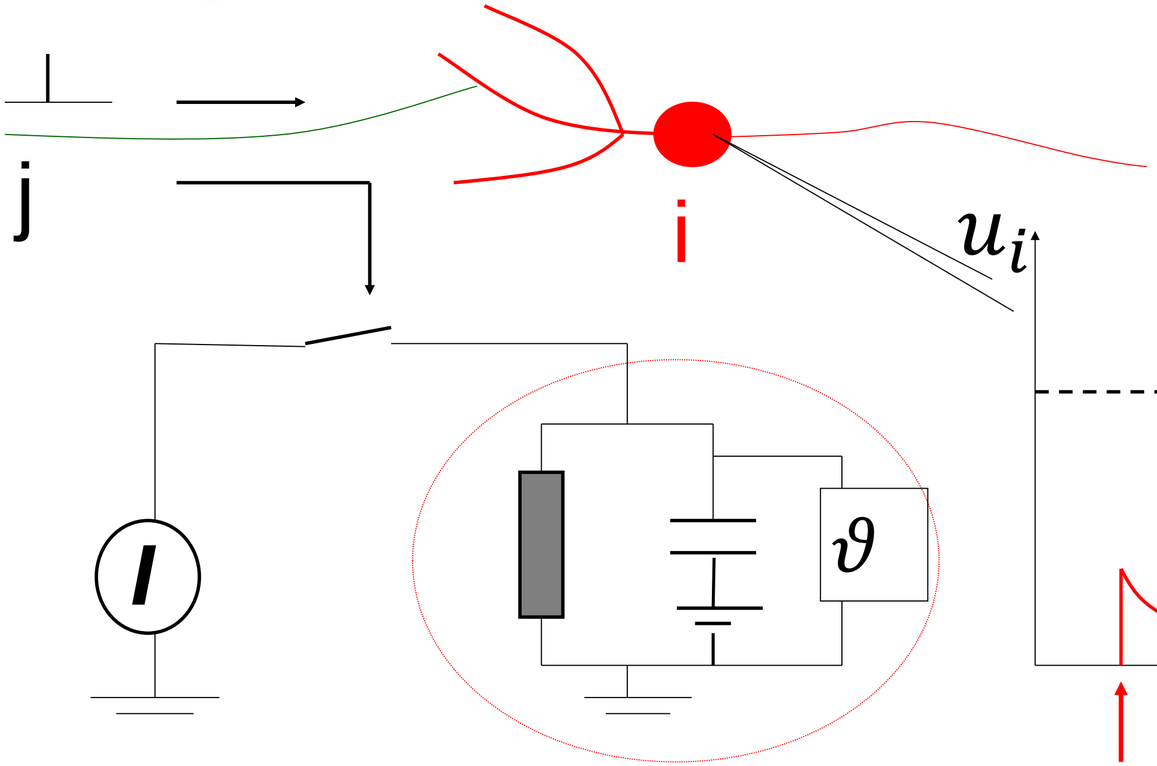

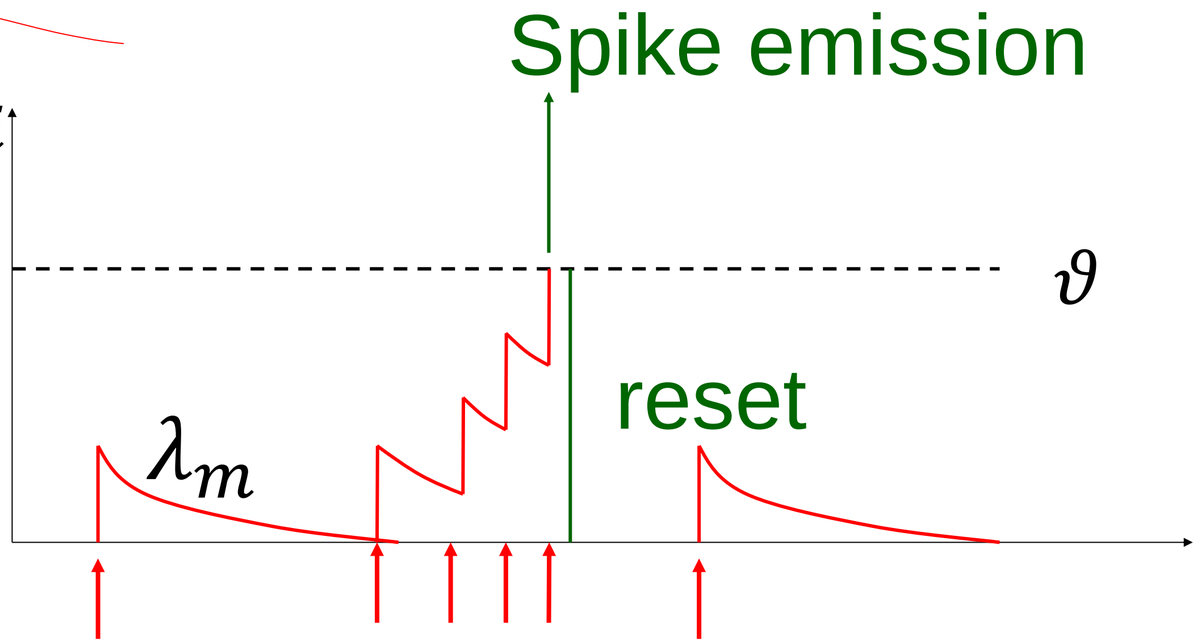

discrete time steps

$$\Delta u_i = w_{ij} \quad if \quad t = t_j^{pre}$$

$$u_i \leftarrow \boldsymbol{\lambda_m} u_i$$

$$if \quad u_i = \vartheta \qquad u_i \leftarrow 0$$

linear, voltage jump

linear, decay with parameter $\boldsymbol{\lambda_m}$

threshold → fire+reset

Previous slide:

The Leaky integrate-and-fire model written in discrete time (say time step Δt = 1ms)
has two linear update steps:

- each presynaptic spike causes a jump of the voltage (membrane potential) by the synaptic weight w_{ij}.
- In each time step the membrane potential decays with a factor λ_m<1. (Asice: If we compare with the previous equation in continuous time, we find that the factor is $\lambda_m = 1 - \left(\frac{\Delta t}{\tau_m}\right)$ where Δt is the time step.)

These linear update steps are combined with a NONLINEAR FIRE-and-RESET process. If the variable u ('membrane potential of the neuron') reaches the threshold theta, then u is reset to zero.

Leaky Integrate-and-Fire Model

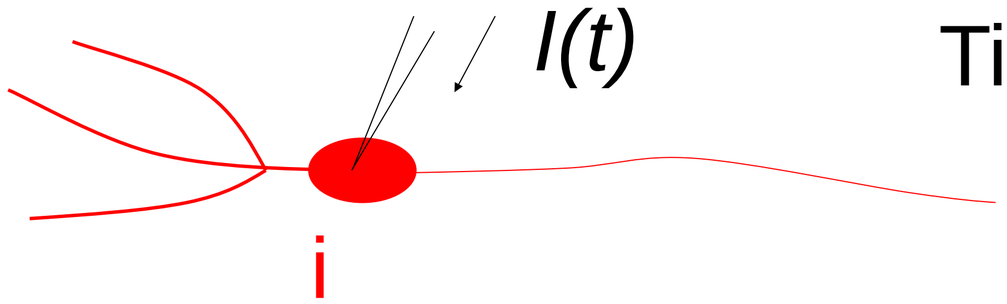

Time-dependent input

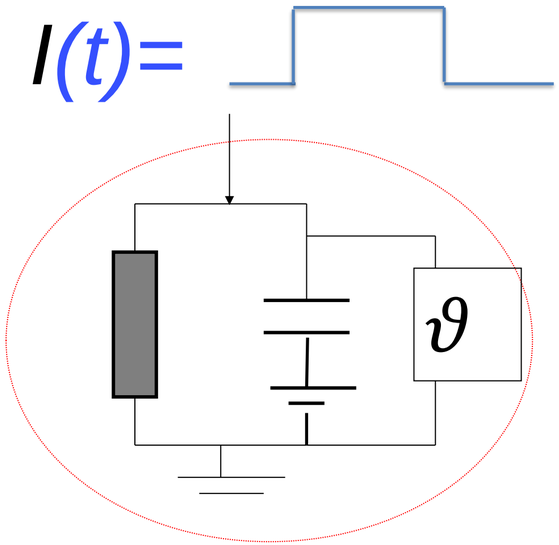

repetitive
spikes

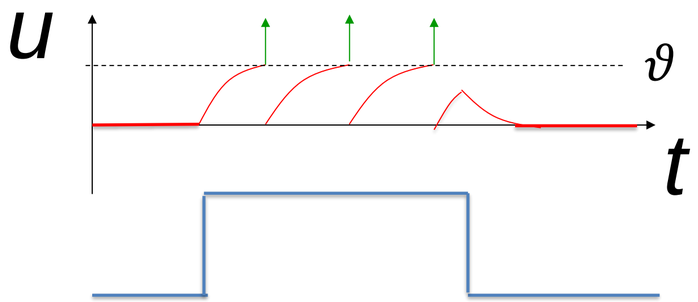

-spikes are events
-triggered at threshold
-spike/reset/refractoriness
-repetitive firing for sustained input

Previous slide.
If an experimentalist injects a current that steps from zero to some large value and then after one second back to zero, the neuron has time to fire several spikes.

Computational Neuroscience: Neuronal Dynamics

EPFL

Lecture 14

Synaptic plasticity and Learning

Wulfram Gerstner

EPFL, Lausanne, Switzerland

Reading for plasticity:
NEURONAL DYNAMICS
- Ch. 19.1-19.3;

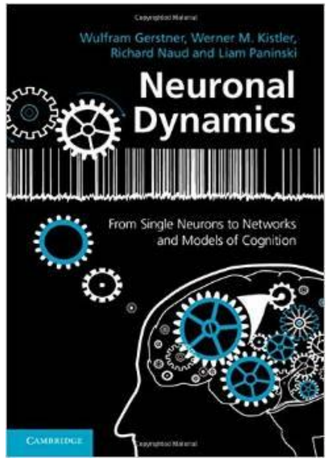

Cambridge Univ. Press

1. **Synaptic plasticity**
   motivation and aims
2. **Classification of plasticity**
   short-term vs. long-term
   unsupervised vs. reward modulated
3. **Model of short-term plasticity**
4. **Models of long-term plasticity**
   - Hebbian learning rules
   - Bienenstock-Cooper-Munro rule
5. **Spiking Models of plasticity**

Previous slide.
Plasticity phenomena have different temporal scales.
Moreover, they have different functions.

2. Classification of synaptic changes: Short-term plasticity

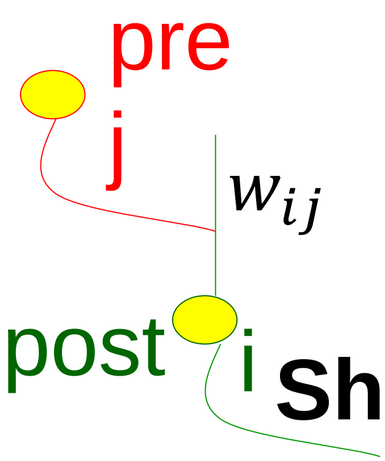

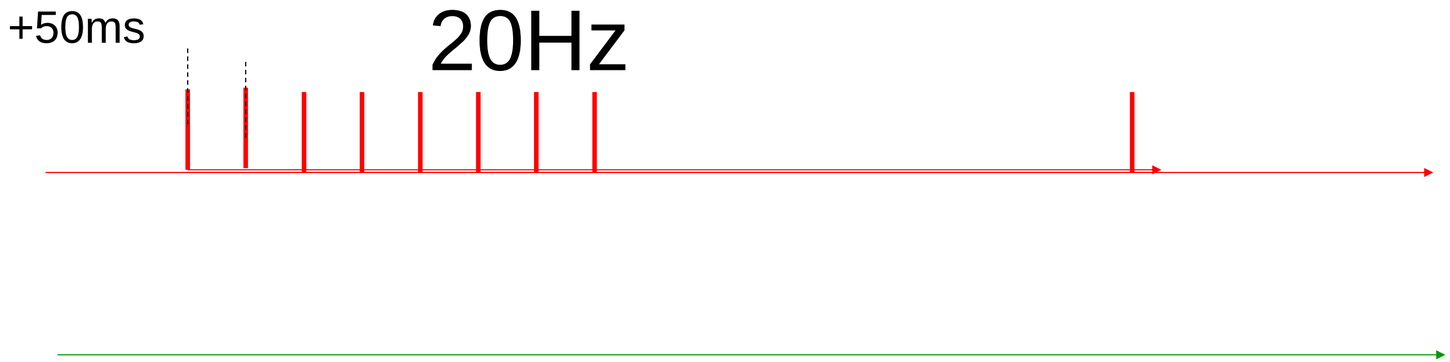

Short-term plasticity/fast synaptic dynamics

Thomson et al. 1993
Markram et al 1998
Tsodyks and Markram 1997

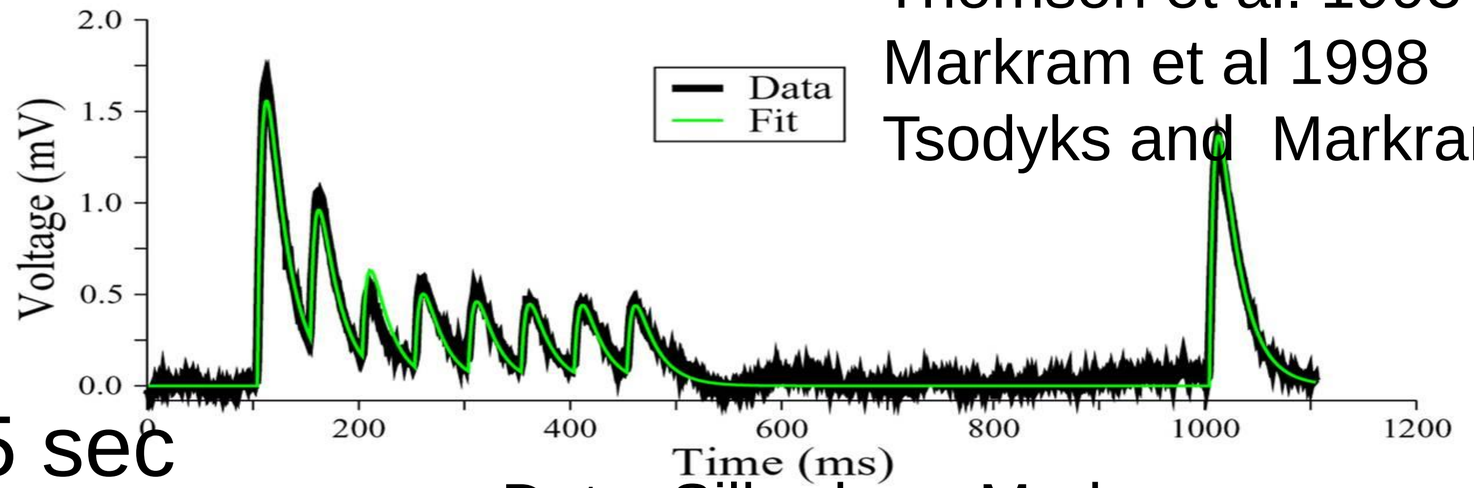

Changes
- induced over 0.5 sec
- recover over 1 sec

Data: Silberberg,Markram
Fit: Richardson (Tsodyks-Markram model)

Previous slide.
Plasticity phenomena have different temporal scales.
Short-term plasticity is induced over 0.5 seconds. Importantly, the induced changes decay back to the baseline within roughly one second.
In other words, they do not cause long-term memory.

Short-Term Plasticity (STP) can induce either a depression (as in this image) or a facilitation of synapses (not shown), or even a combination thereof.

2. Classification of synaptic changes: Long-term plasticity

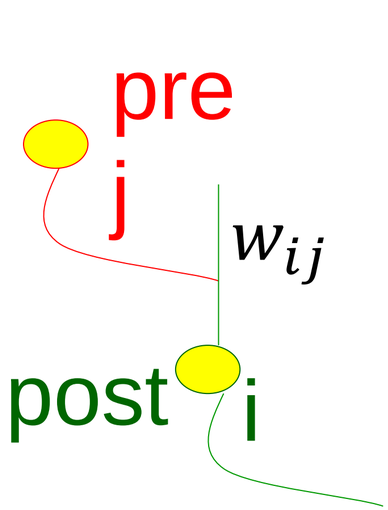

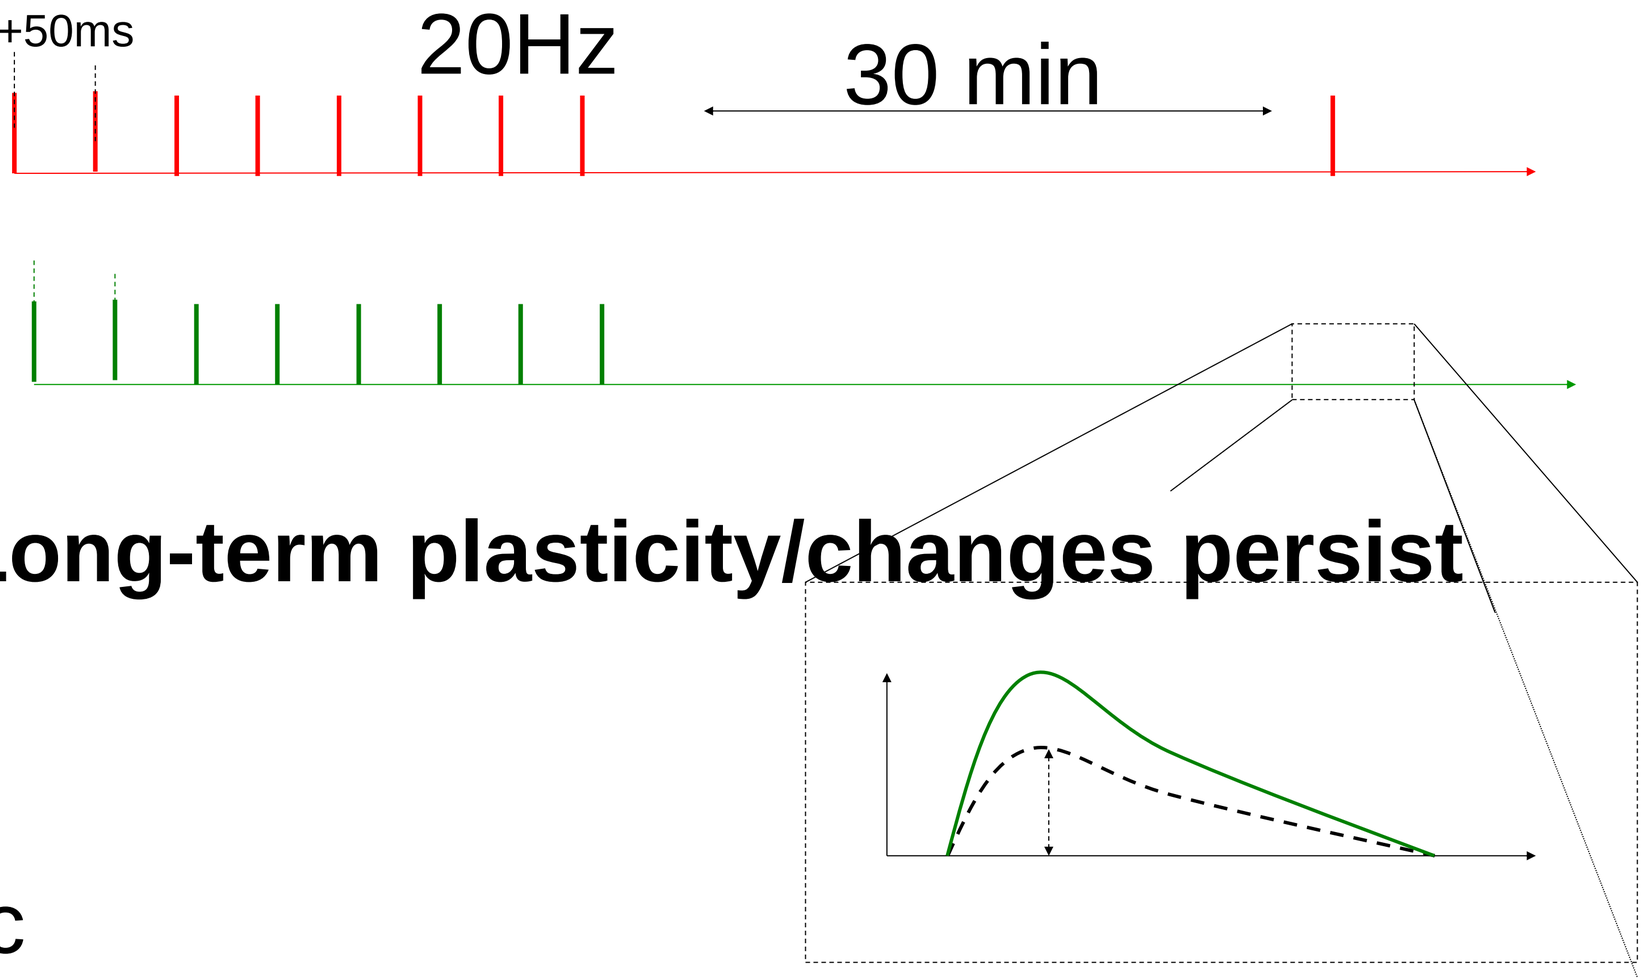

Long-term plasticity/changes persist

Changes
- induced over 3 sec
- persist over 1 – 10 hours (or longer?)

Previous slide.
Plasticity phenomena have different temporal scales.
Long-term plasticity can also be rapidly induced, for example within 1 to 3 seconds. Importantly, the induced changes persist for several hours.
Therefore, they are a good candidate to be the basis of long-term memory.

Long-term plasticity comes in two flavors: Long-Term potential (LTP) and Long-Term Depression (LTD).

2. Classification of synaptic plasticity: STP vs LTP

Short-Term (STP) vs/ Long-Term (LTP)

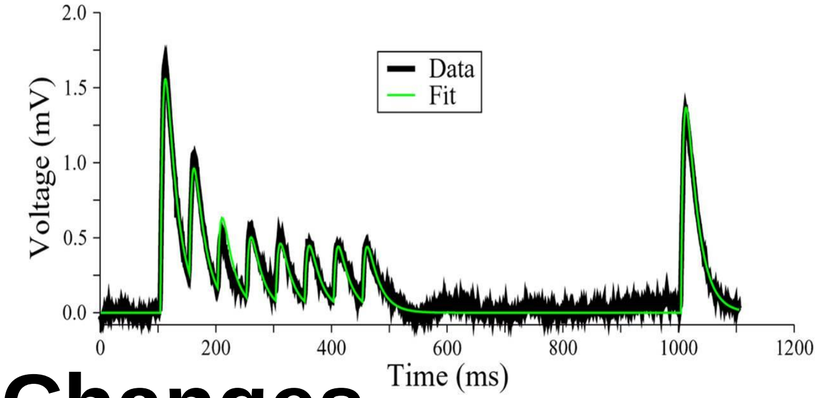

Changes

- induced over 0.1-0.5 sec
- recover over 1 sec

Protocol

- presynaptic spikes

Model

- well established

Tsodyks, Pawelzik, Markram, Neur. Comput. (1998) 10: 821–835
Dayan and Abbott, Theoretical Neuroscience (2001), MIT Press

LTP/LTD/Hebb

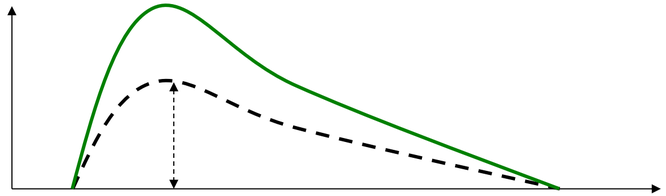

Changes

- induced over 0.5-5sec
- remains over hours

Protocol

-presynaptic spikes + …

Model

- we will see

Previous slide.
Plasticity phenomena have not only different temporal scales, but also different induction protocols.

While Short-Term Plasticity (STP) does not depend on the state of the postsynaptic neuron, but Long-Term Plasticity does.

2. Classification of synaptic changes

Induction of changes

- fast (if stimulated appropriately)
- slow (homeostasis)

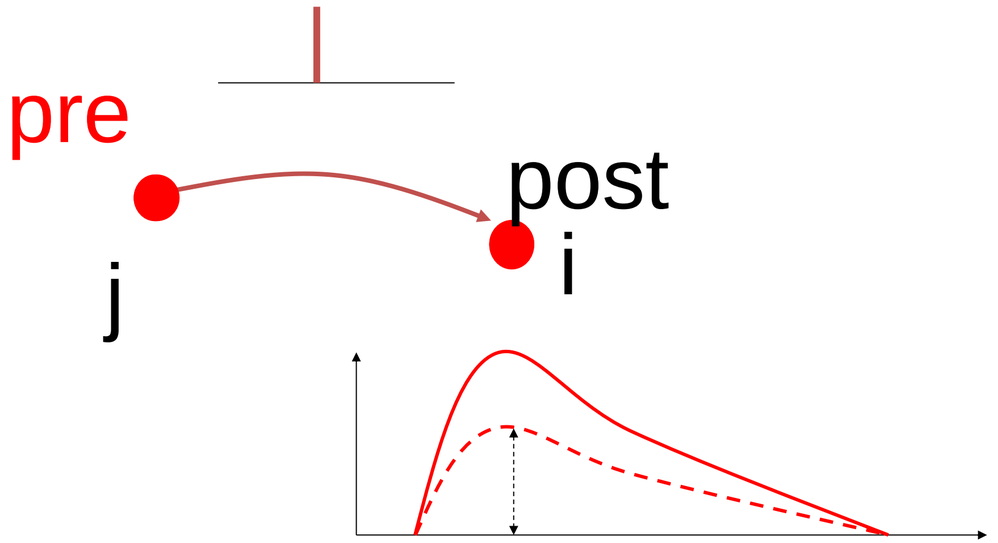

Persistence of changes

- long (LTP/LTD): induced changes stay at least for 1-2 hours.
- short (short-term plasticity, STP) induced changes decay back within 1s.

Previous slide.

Coming back to the time scale:

There are also slower processes (sometimes called homeostatic processes) that are induced on a time scale of hours. These often depend only on some long-term firing history of the postsynaptic neuron.

2-factor versus 3-factor rules

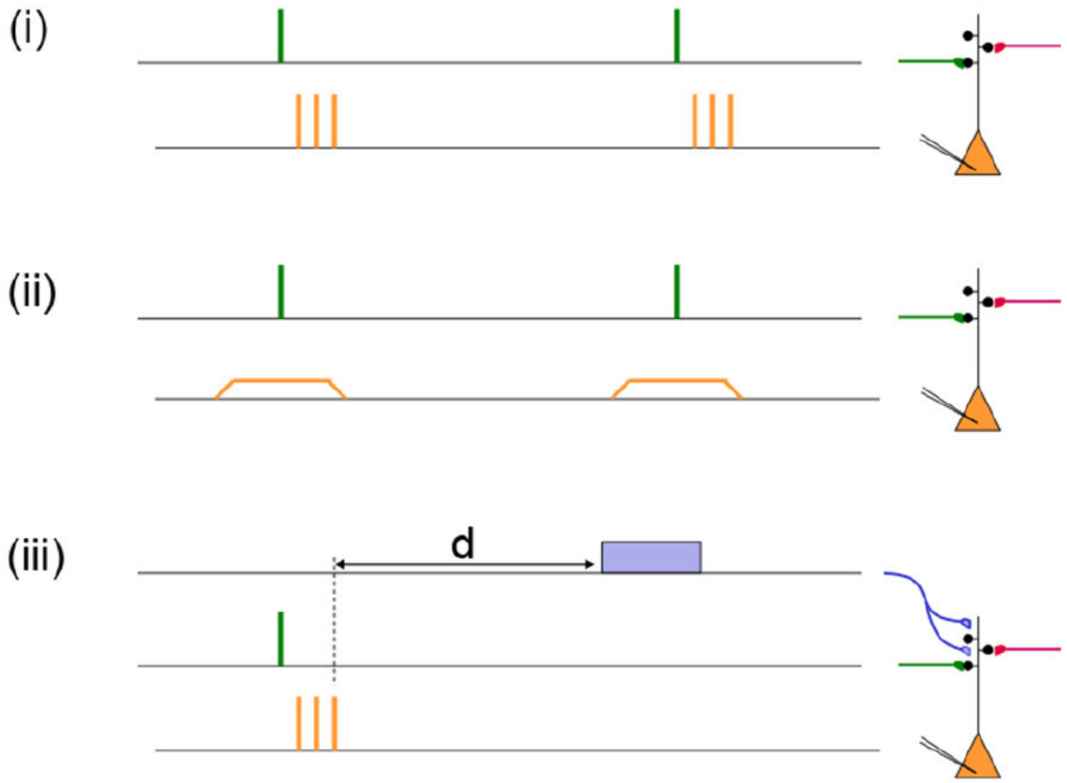

Hebbian: pre=spike
'post' = spikes

Hebbian:pre=spike
'post' = voltage

3-factor: pre+post+
neuromodulator (success/reward)

3-factor = Hebbian combined with
(potentiall delayed) Neuromodulator:
→ Reward based learning/reinforcement learning

Image: Gerstner et al. (2018, review paper in Frontiers)

Previous slide.

Another important distinction is that between two-factor and three-factor rules, where the latter include the action of a neuromodulator.

Hebb rules are essentially two-factor rules.

2. Review: Hebb rule

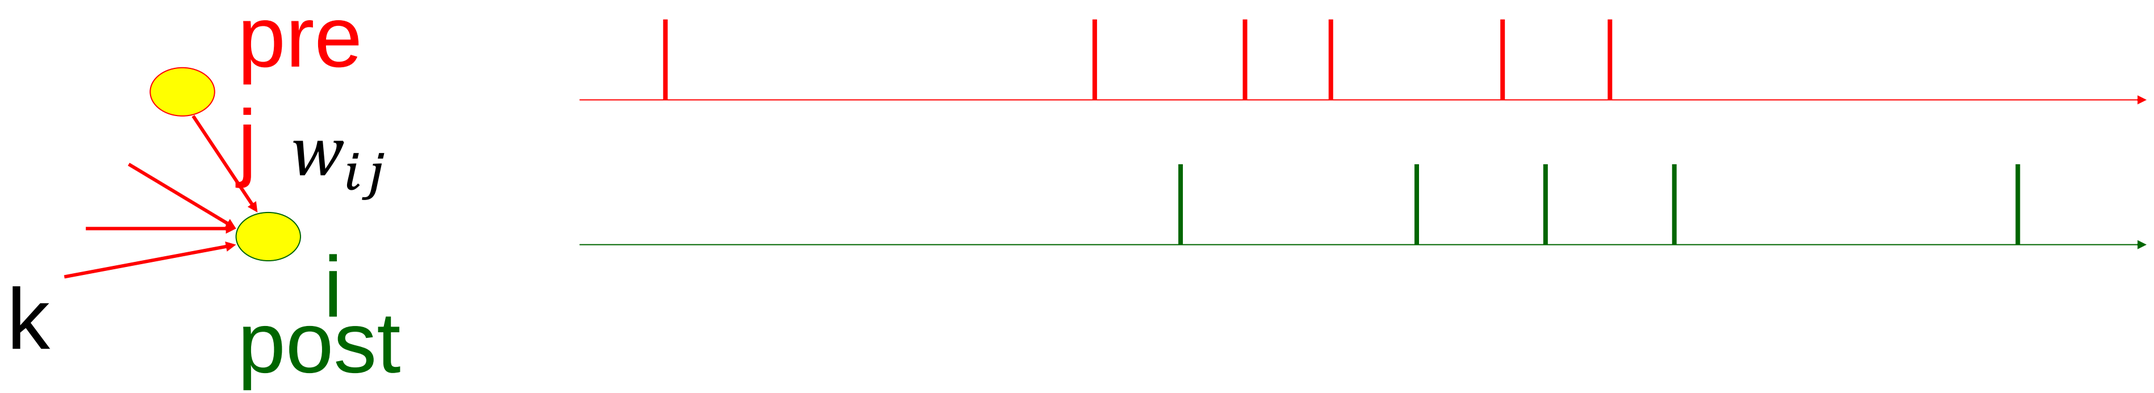

When an axon of cell **j** repeatedly or persistently takes part in firing cell **i**, then j's efficiency as one of the cells firing i is increased

Hebb, 1949

- local rule
- simultaneously active (correlations)

Previous slide.

Review of Hebb’s formulation.

Hebb rules are essentially two-factor rules.

Synaptic plasticity: Long-Term Potentiation (LTP)

Hebbian Learning in experiments (schematic)

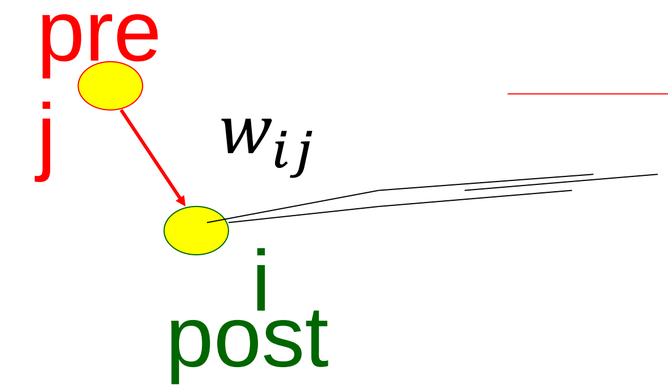

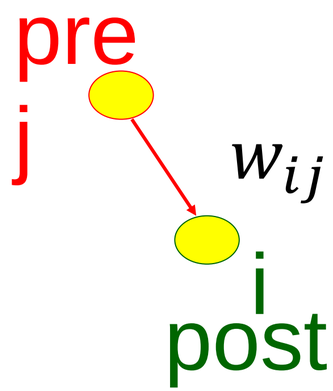

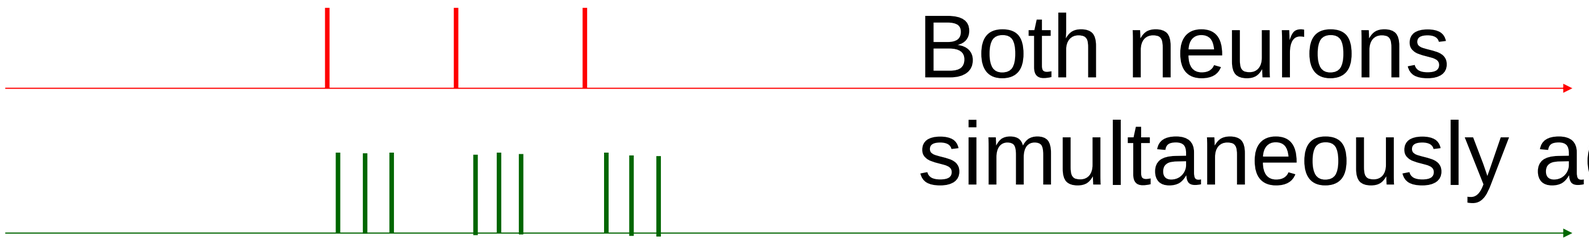

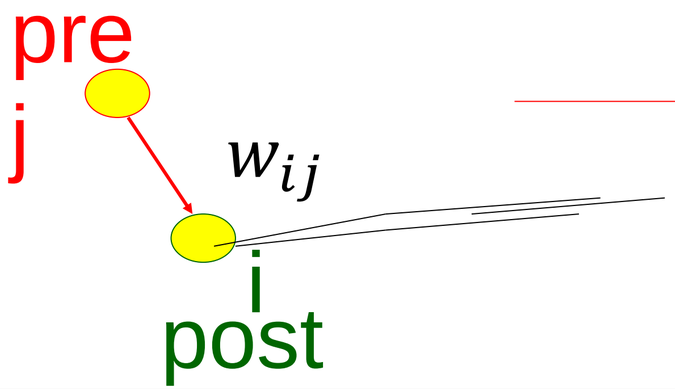

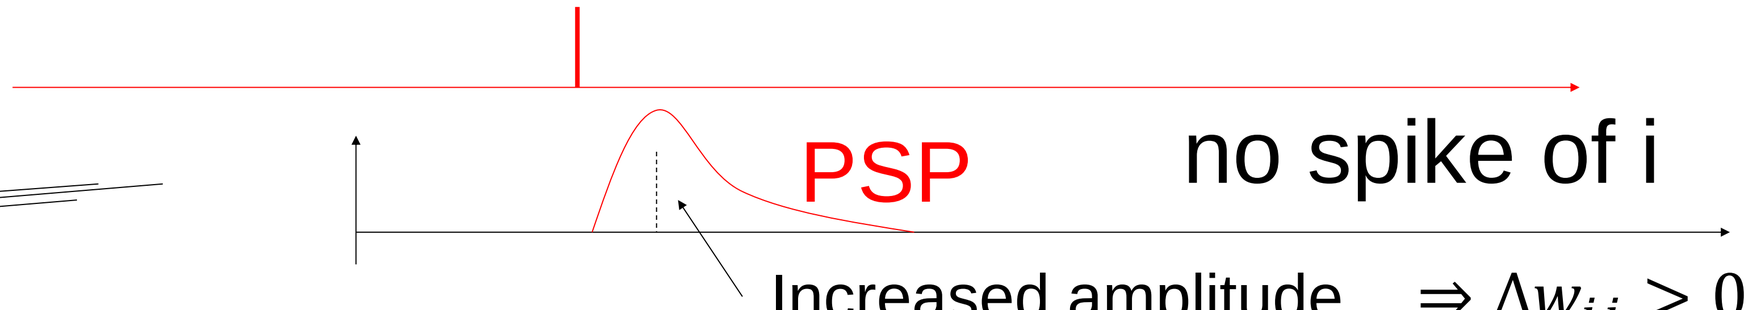

Previous slide.

In a schematic experiment (shown already in the first week)

1) You first test the size of the synapse by sending a pulse from the presynaptic neurons across the synapses. The amplitude of the excitatory postsynaptic potential (EPSP) is a convenient measure of the synaptic strength. It has been shown that it is correlated with the size of the spine.

2) Then you do the Hebbian protocol: you make both neurons fire together

3) Finally you test again the size of the synapse. If the amplitude is bigger you conclude that the synaptic weight has increased.

In the next slides we consider different experimental induction protocols in more detail.

Classical paradigm of LTP induction: 1. Strong extracellular stimulation

extracellular electrode

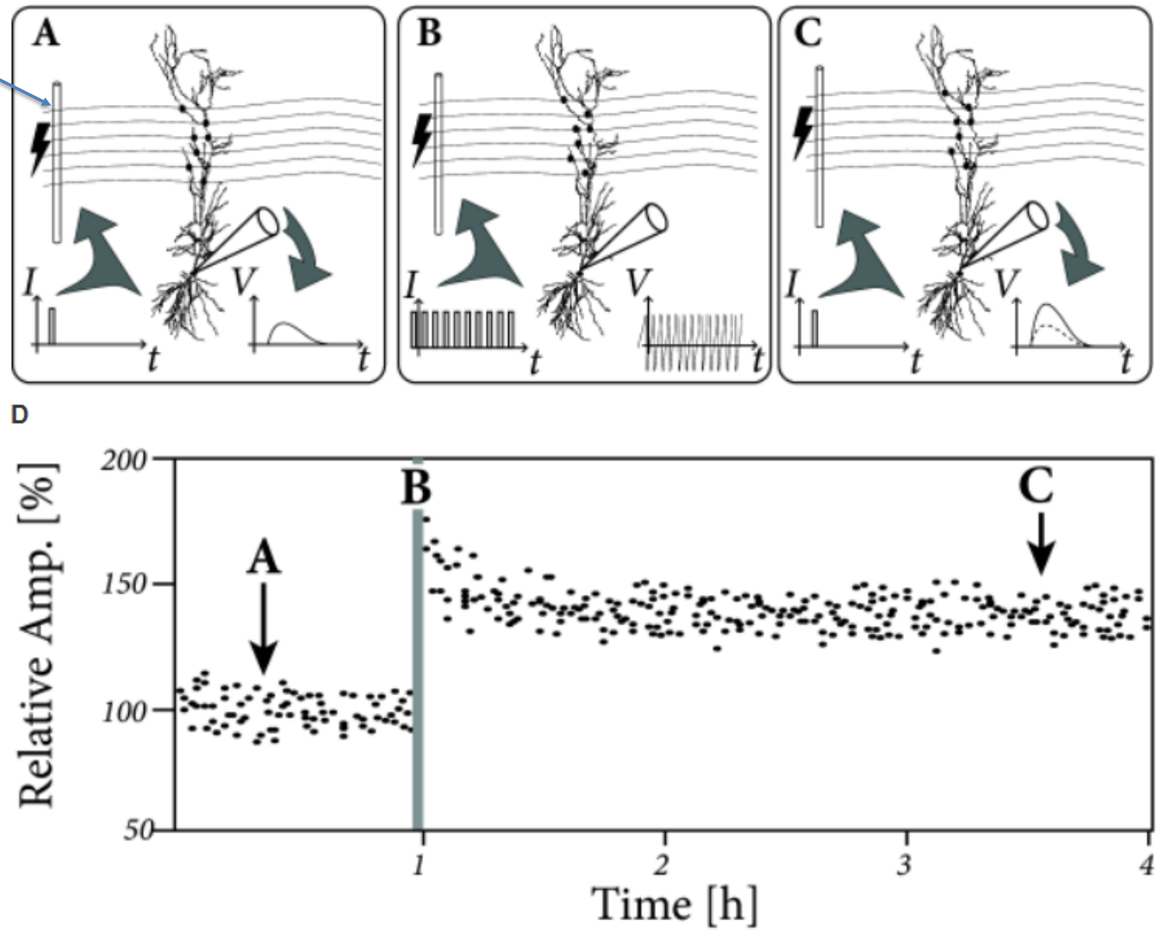

C: Late LTP (consolidation)

Image (schematic): Neuronal Dynamics, Gerstner et al., MIT Press (2014)

Previous slide.

Here the stimulation is done with an extracellular electrode which unspecifically stimulates a bundle of fibers (horizontal wiggly lines).
A recording electrode measures the response. In a schematic experiment (shown already in the first week)

A) You first test the size of the synapse by sending a pulse from the presynaptic neurons across the synapses. The amplitude of the excitatory postsynaptic potential (EPSP) is a convenient measure of the synaptic strength.

B) Then you repeatedly stimulate with the extracellular electrode strong enough that the postsynaptic neuron emits spikes. Essentially this represents a Hebbian protocol. It is sometimes called a tetanus stimulation (for example stimuli are repeated at 100Hz)

C) Finally you test again the size of the synapse. If the amplitude is bigger you conclude that the synaptic weight has increased.

Fig. 19.2: Schematic drawing of a paradigm of LTP induction. **A**. A weak test pulse (left) evokes the postsynaptic response sketched on the right-hand side of the figure. **B**. A strong stimulation sequence (left) triggers postsynaptic firing (right, the peak of the action potential is out of bounds). **C**. A test pulse applied some time later evokes a larger postsynaptic response (right; solid line) than the initial response. The dashed line is a copy of the initial response in **A**. **D**. The relative amplitude as measured with the test pulses illustrated in **A** and **C** is increased after the strong stimulation at $t = 1$ h. (Schematic figure.)

Classical paradigm of LTP induction: 2. Pairing tetanus and depolarization

Resting state u_{eq}:

- 75mV< u_{eq} <-60mV

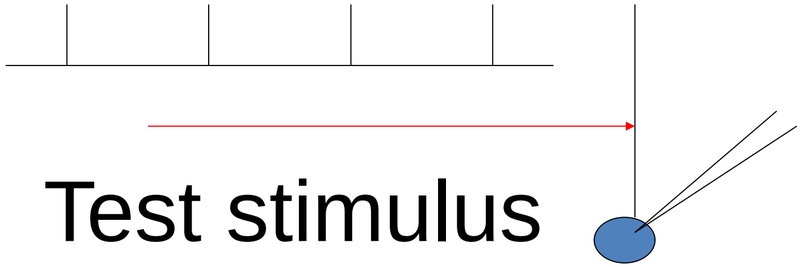

Test stimulus

At 0.1 Hz neuron at -70mV

LTP induction:

tetanus at 100Hz

neuron depolarized to -40mV

Standard LTP

PAIRING experiment

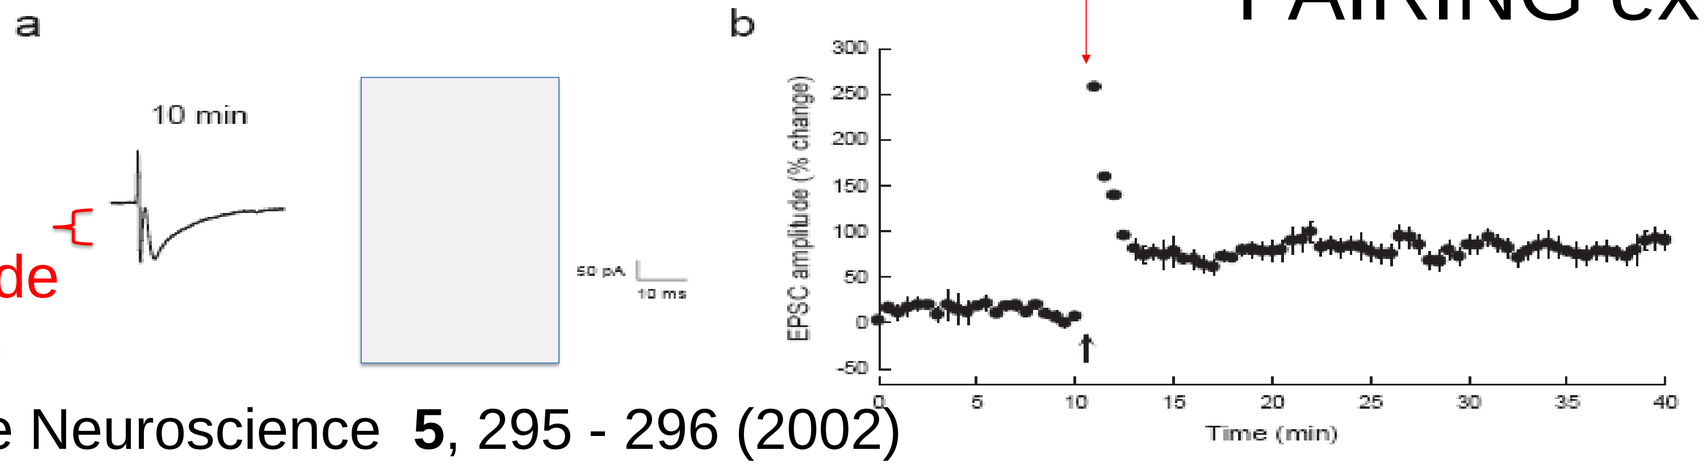

EPSC amplitude (sign-inverted)

Fig. from Nature Neuroscience **5**, 295 - 296 (2002)

D. S.F. Ling, … & Todd C. Sacktor

See also: Bliss and Lomo (1973), Artola, Brocher, Singer (1990), Bliss and Collingridge (1993)

Previous slide.

A weak extracellular stimulation at 0.1 Hz serves as a test stimulus.
An intracellular (patch) electrode was used to measure the size of the EPSC (excitatory postsynaptic current).

TEST: To avoid any opening of NMDA receptors, the neuron was kept close to the resting potential at -75mV.

INDUCTION (called pairing): Three bursts of extracellular pulses (each burst 100 pulses at 100Hz) where applied while the postsynaptic cell was held at -40mV (which is above the firing threshold of the cell).

TEST: The test pulses of the same size as before now evoked a stronger EPSC which lasted for 30min, hence LTP was induced.

Classical paradigm of LTP induction: 3. Voltage-dependence

Low-frequency paired
With depolarization
100 pulses at 2Hz
many synapses

Voltage dependence

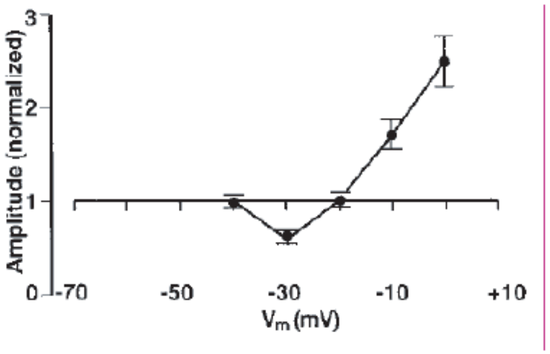

Ngezehajo et al. 2000,
Artola et al.

Extracellular electrode (stimulus)

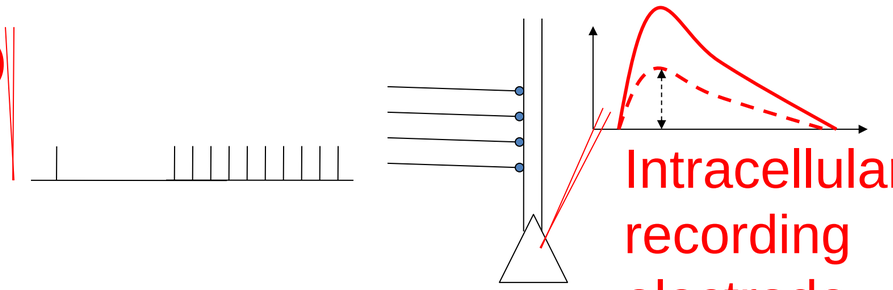

Intracellular recording electrode,
Current injection/ voltage clamp

LTP

EPSP amplitude

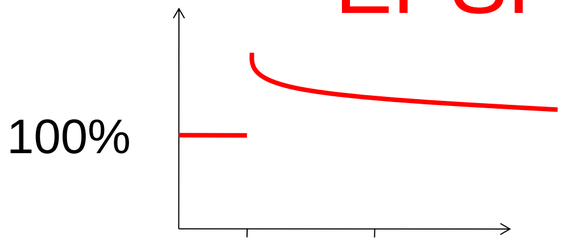

0 30 min

stim

Previous slide.

LTP is induced with 100 pulses at 2Hz if the postsynaptic cell is depolarized to -10mV (which is above the firing threshold)

LTD is induced with 100 pulses at 2Hz if the postsynaptic cell is depolarized to -30mV.

Note that voltage measurements are never reliable because of various artifacts. However, the sequence ‘no change/LTD/LTP’ for increasing voltage is probably reliable.

Classical paradigm of LTP induction: 4. STDP

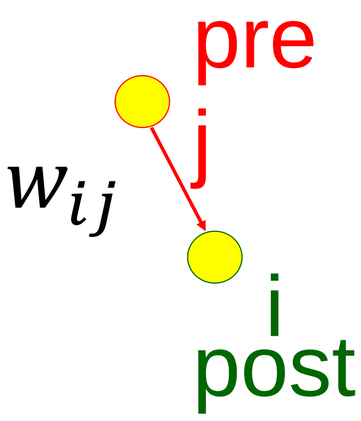

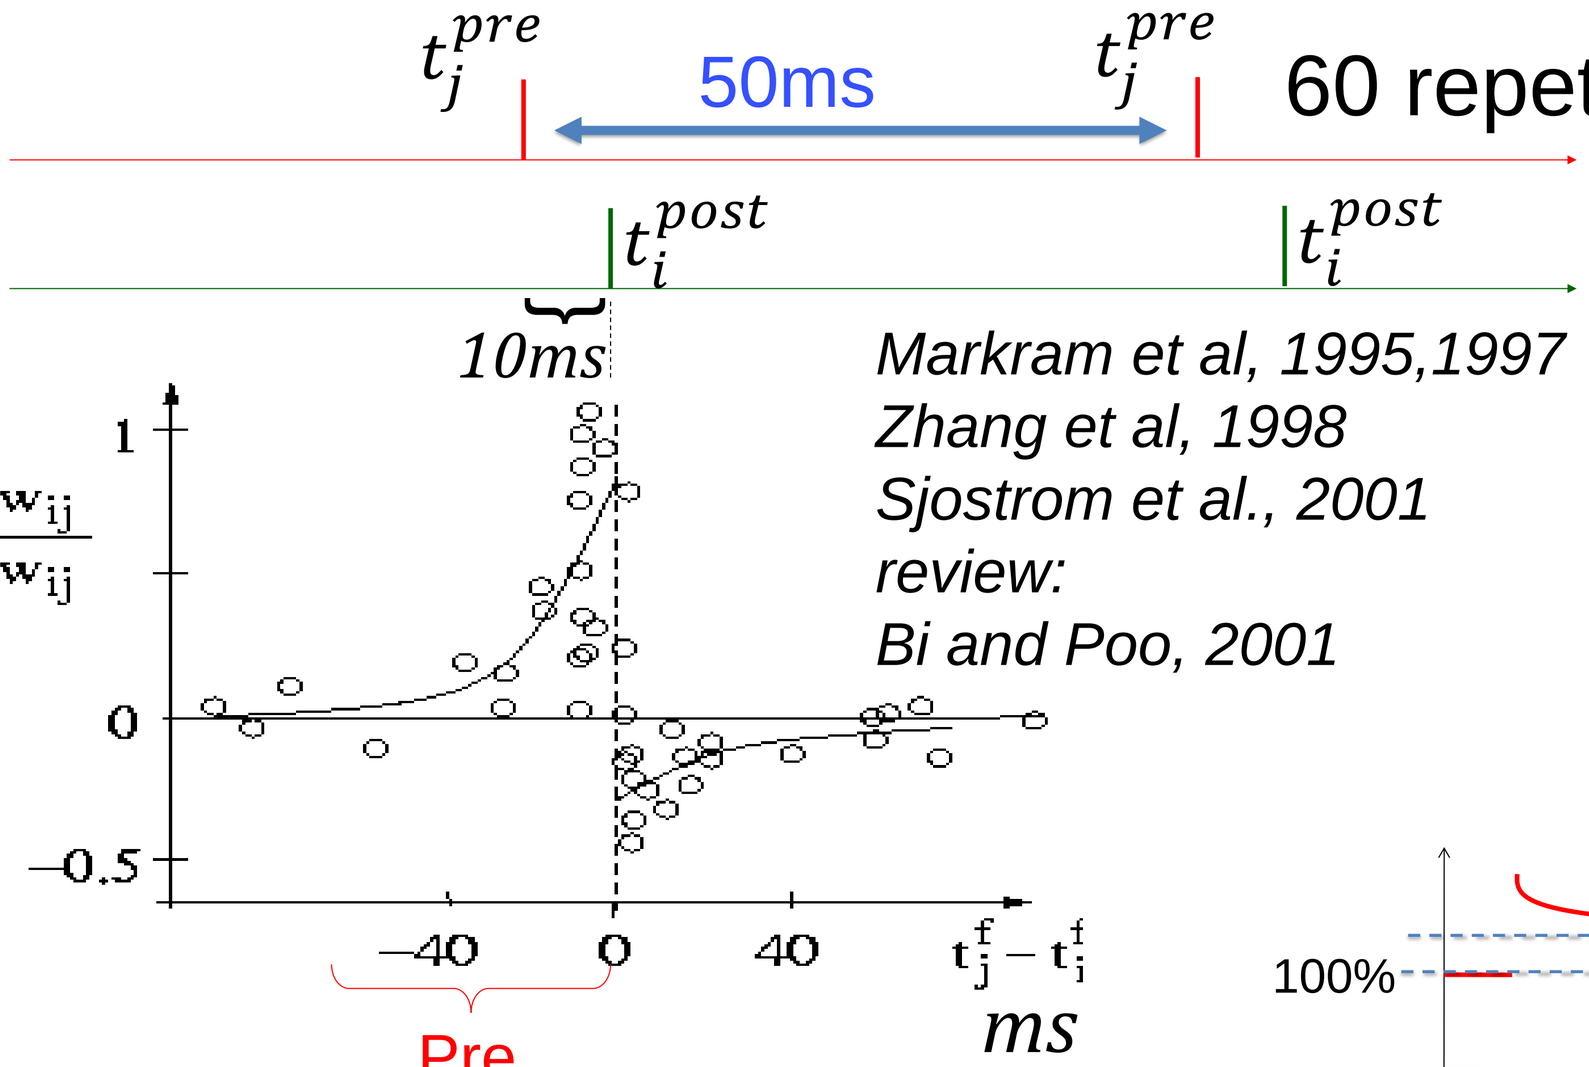

Markram et al, 1995,1997
Zhang et al, 1998
Sjostrom et al., 2001
review:
Bi and Poo, 2001

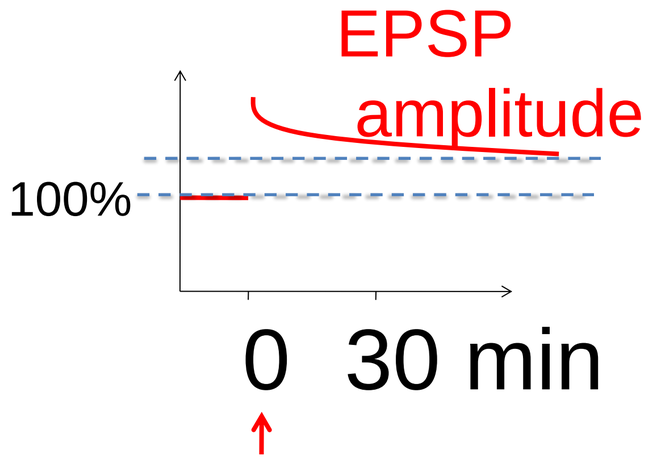

Previous slide.

Classical paradigm of LTP induction: 4. Interaction of voltage and timing

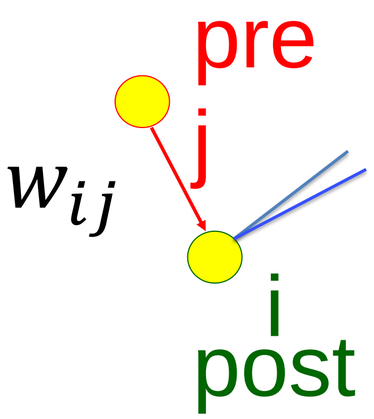

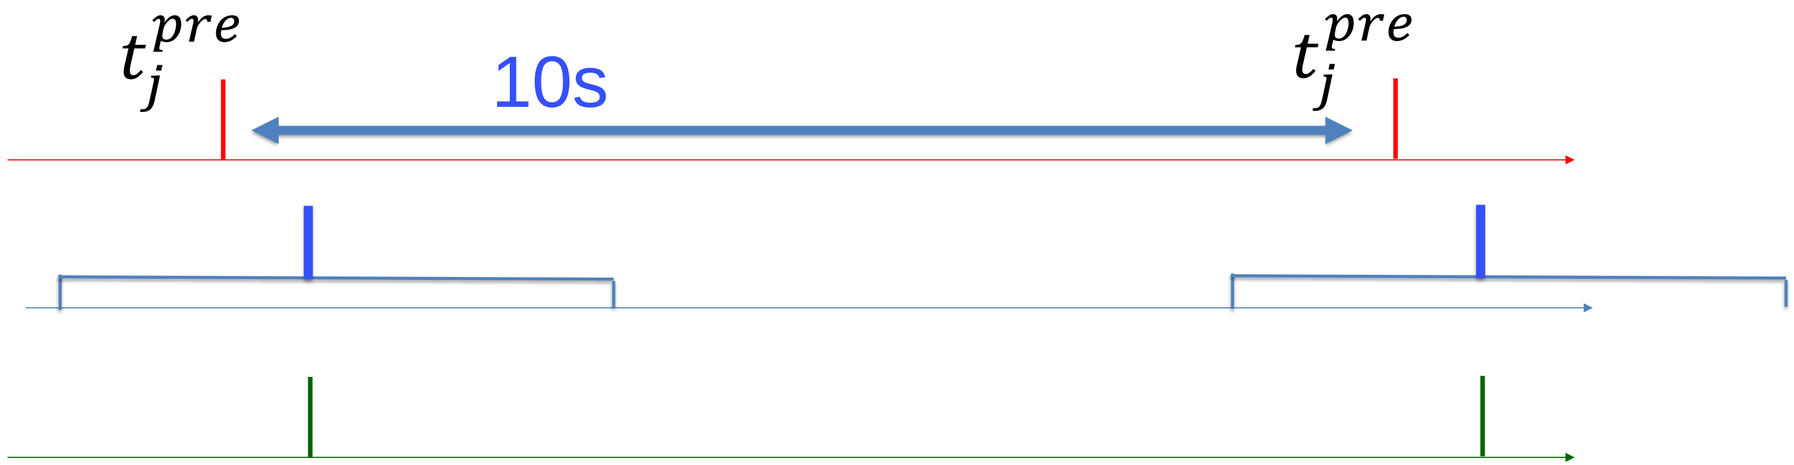

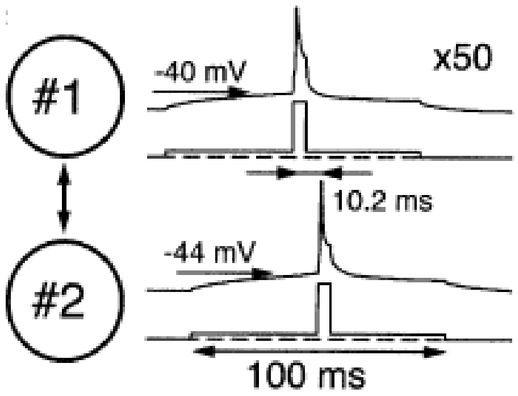

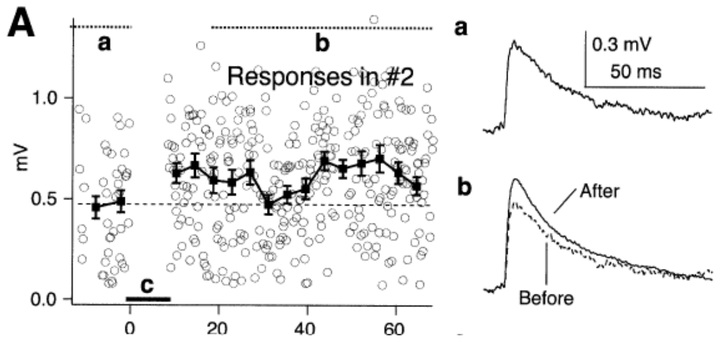

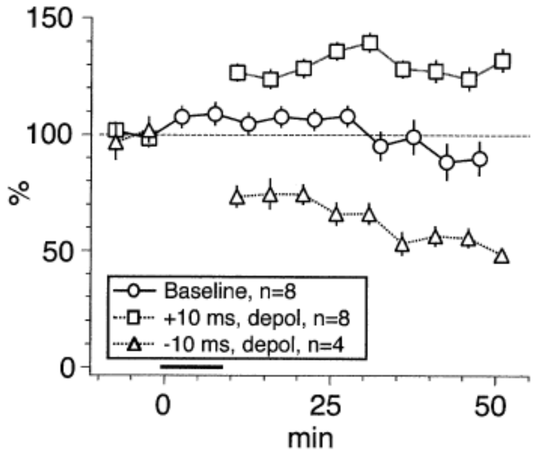

Sjostrom et al., Neuron, 2001

Previous slide.
Normally, induction of LTP with a standard STDP protocol (+10ms) does not work at 0.1 Hz, i.e., if the interval between pre-post repetitions is 10s. But if combined with depolarization, LTP is possible!

However, LTD can be induced with an STDP protocol (-10ms), even if repetitions occur at very low frequency. Depolarization does not change this.

Repetition : NMDA receptors needs glutamate and postsynaptic voltage

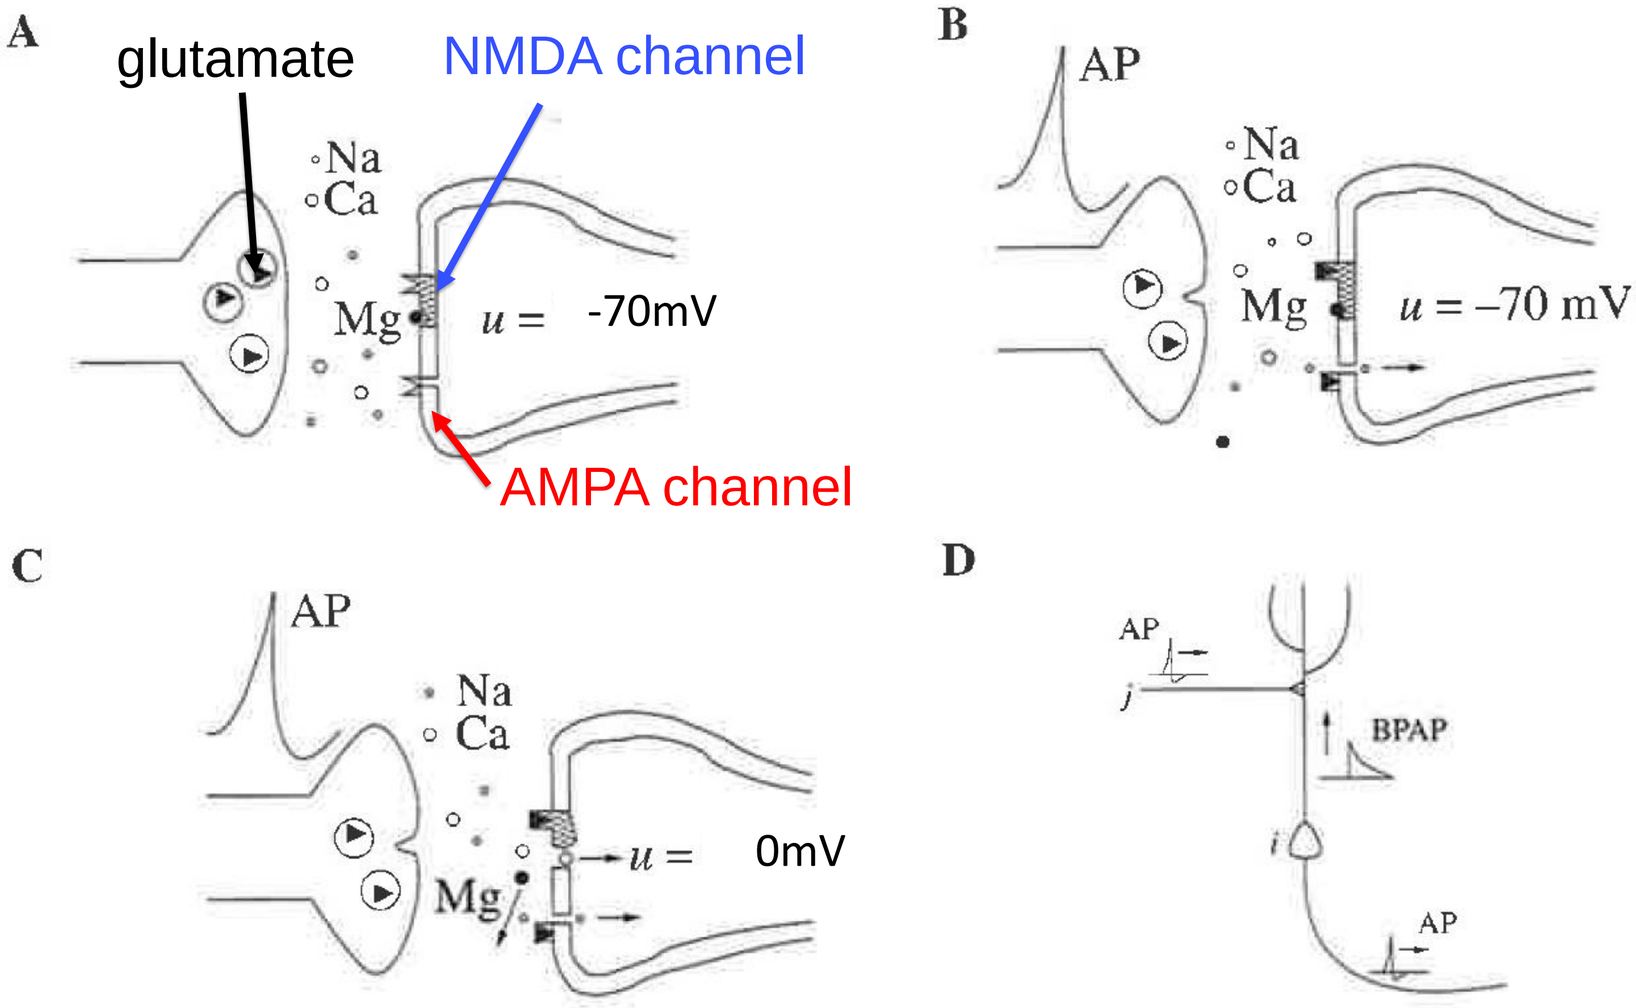

Previous slide.

The picture and function of the NMDA channel (discussed earlier) is compatible with the interaction of voltage and spike timing observed in experiments.

Three-factor rules in striatum: eligibility trace and delayed DA

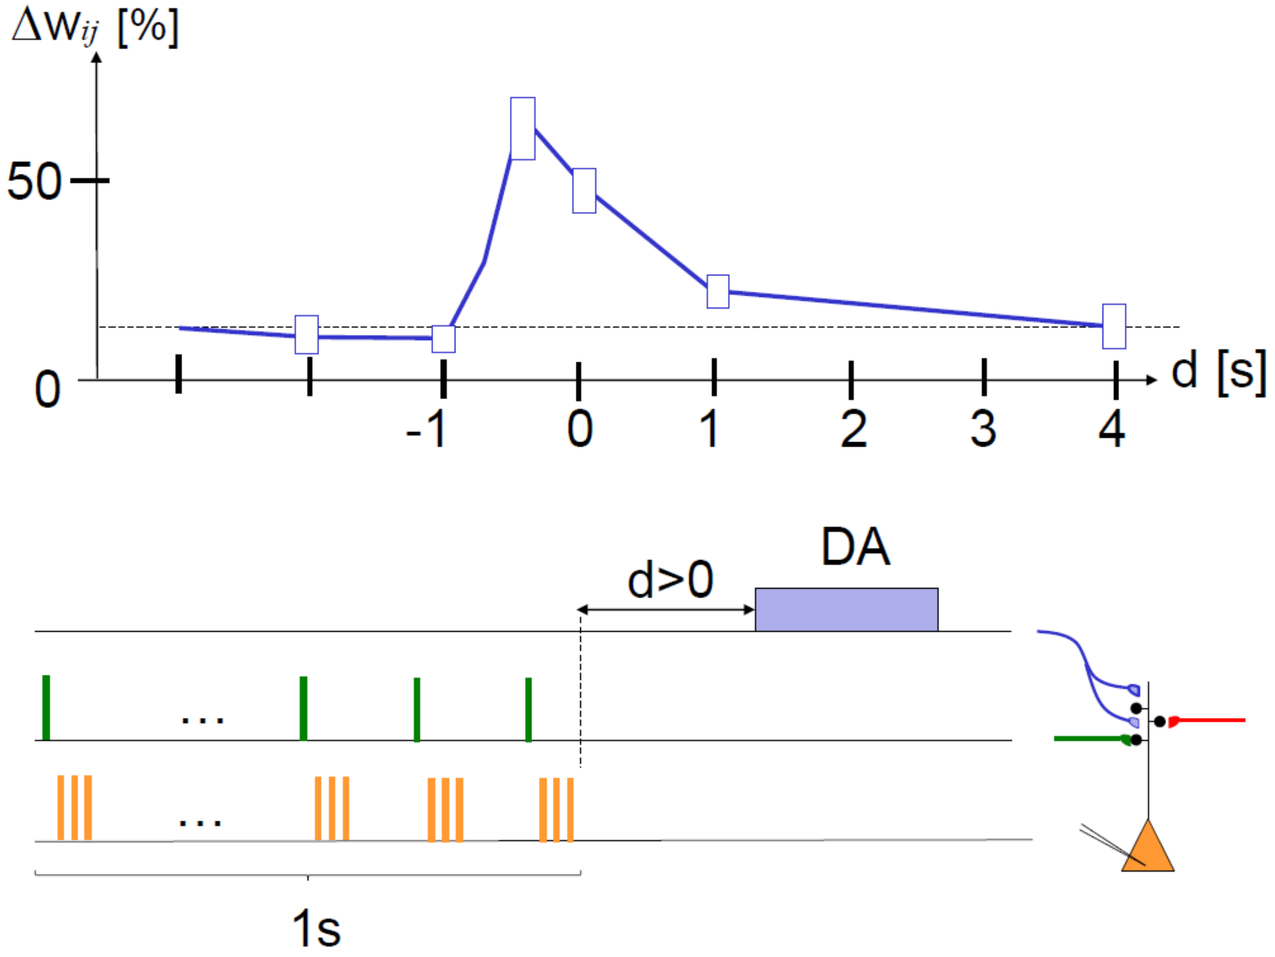

Yagishita et al. 2014, SCIENCE
Kasai lab

Striatum involved in action selection (later today)

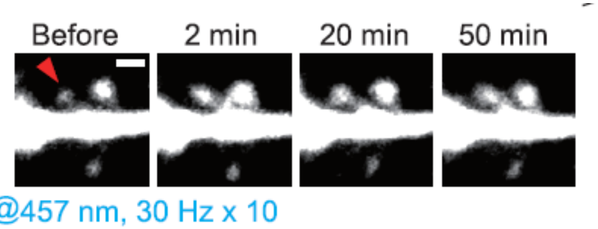

-Dopamine (DA) can come with a delay of 1s
-Long-Term stability over at least 50 min.

3. Three-factor rules in striatum: eligibility trace and delayed Da

Yagishita et al. 2014

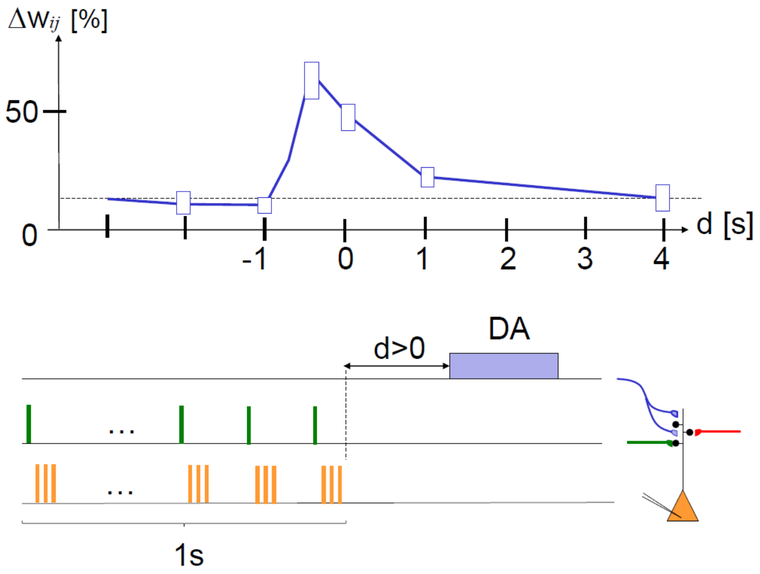

In striatum medial spiny cells, stimulation of presynaptic glutamatergic fibers (green) followed by three postsynaptic action potentials (STDP with pre-post-post-post at +10ms) repeated 10 times at 10Hz yields LTP if dopamine (DA) fibers are stimulated during the presentation (d < 0) or shortly afterward (d = 0s or d = 1s) but not if dopamine is given with a delay d = 4s; redrawn after Fig. 1 of (Yagishita et al., 2014), with delay d defined as time since end of STDP protocol.

Lower left: the image from the beginning of this lecture comes from this experiment of Yagishita. This image demonstrates the Long-Term Stability over at least 50 min

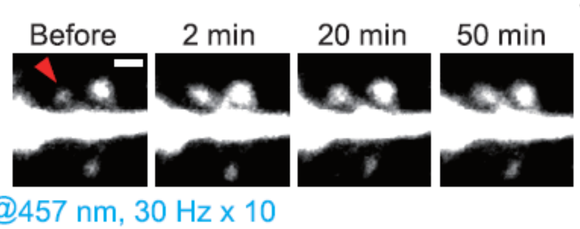

Three-factor rules in cortex: eligibility trace and delayed NE

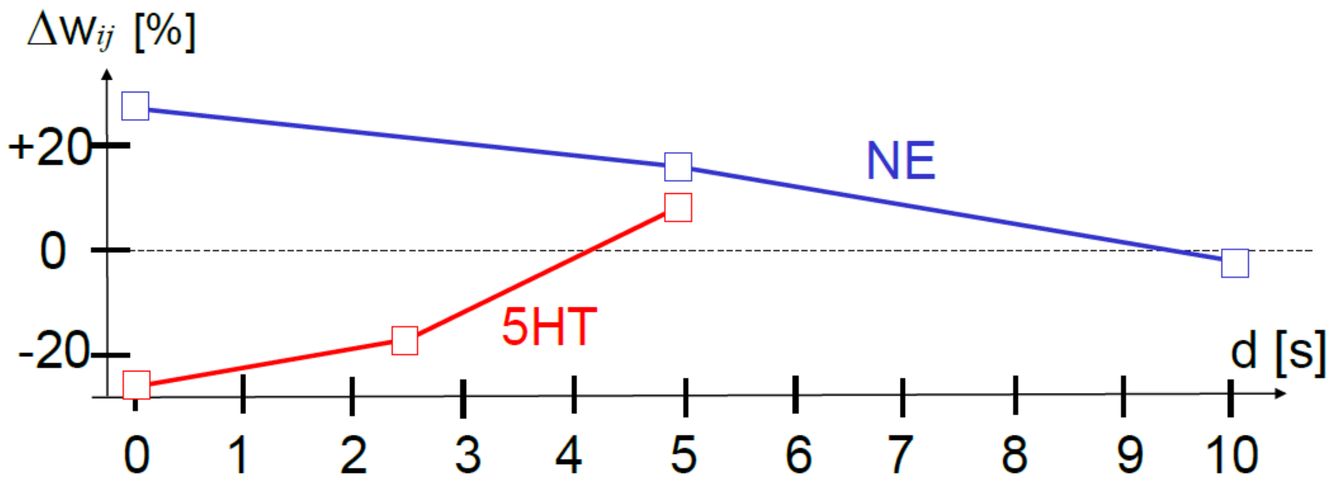

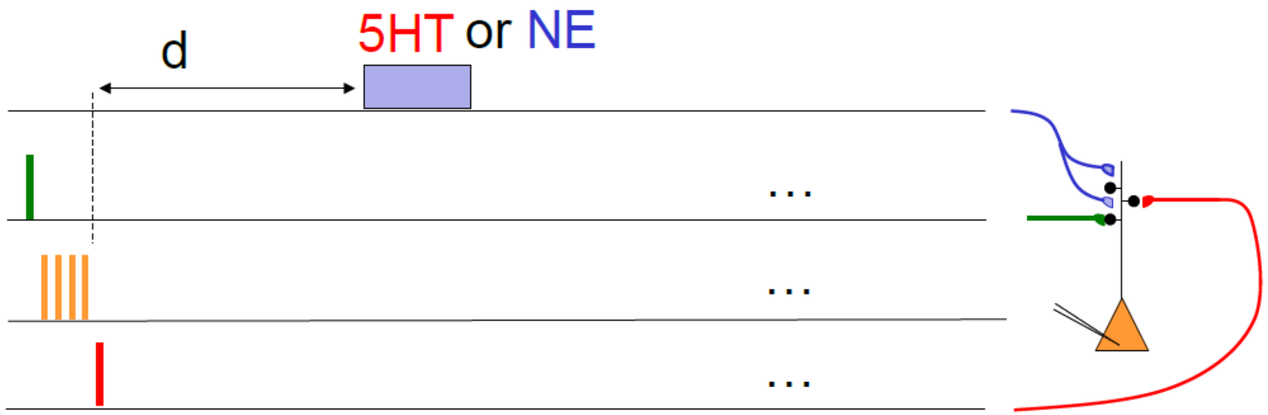

He et al., 2015, NEURON
Kirkwood lab.

NE = norepinephrine

5HT=serotonin

second example

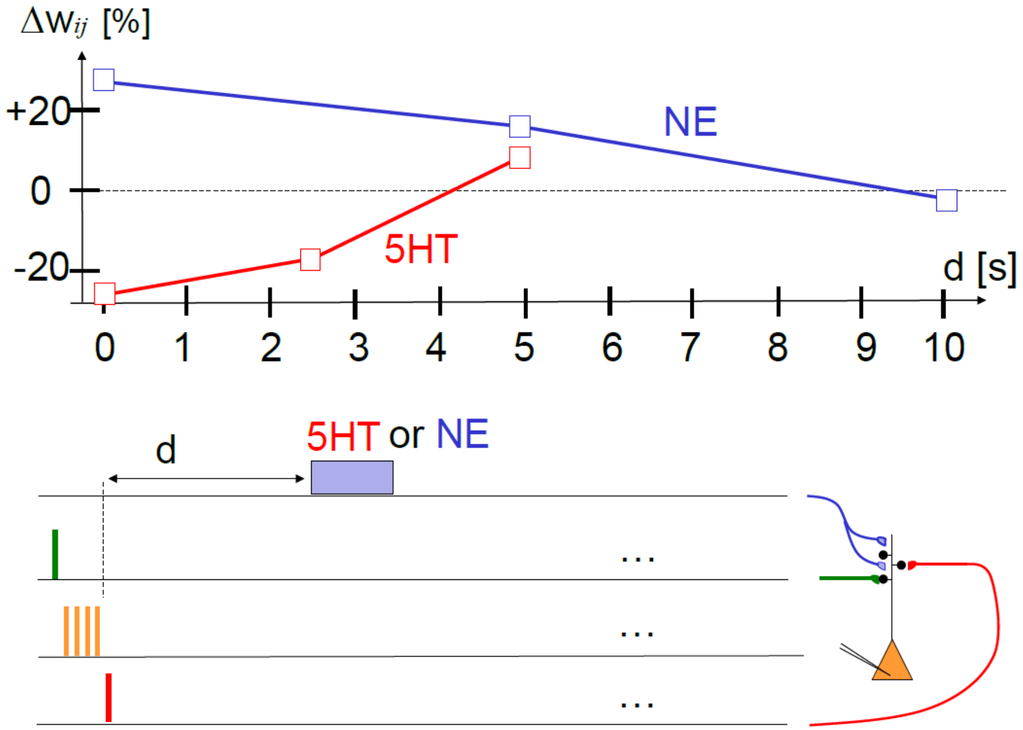

In cortical pyramidal cells, stimulation of two independent presynaptic pathways (green and red) from layer 4 to layer 2/3 by a single pulse is paired with a burst of four postsynaptic spikes (orange).

If the pre-before-post stimulation was combined with a pulse of norepinephrine (NE) receptor agonist isoproterenol with a delay of 0 or 5s, the protocol gave LTP (blue trace).

If the post-before-pre stimulation was combined with a pulse of serotonin (5-HT) of a delay of 0 or 2.5s, the protocol gave LTD (red trace).

(He et al., 2015).

Three-factor rules in hippocampus: Behavioral time scale plasticity

Complex spike can come with a delay of 1-5 seconds

One-shot learning: no repetition necessary

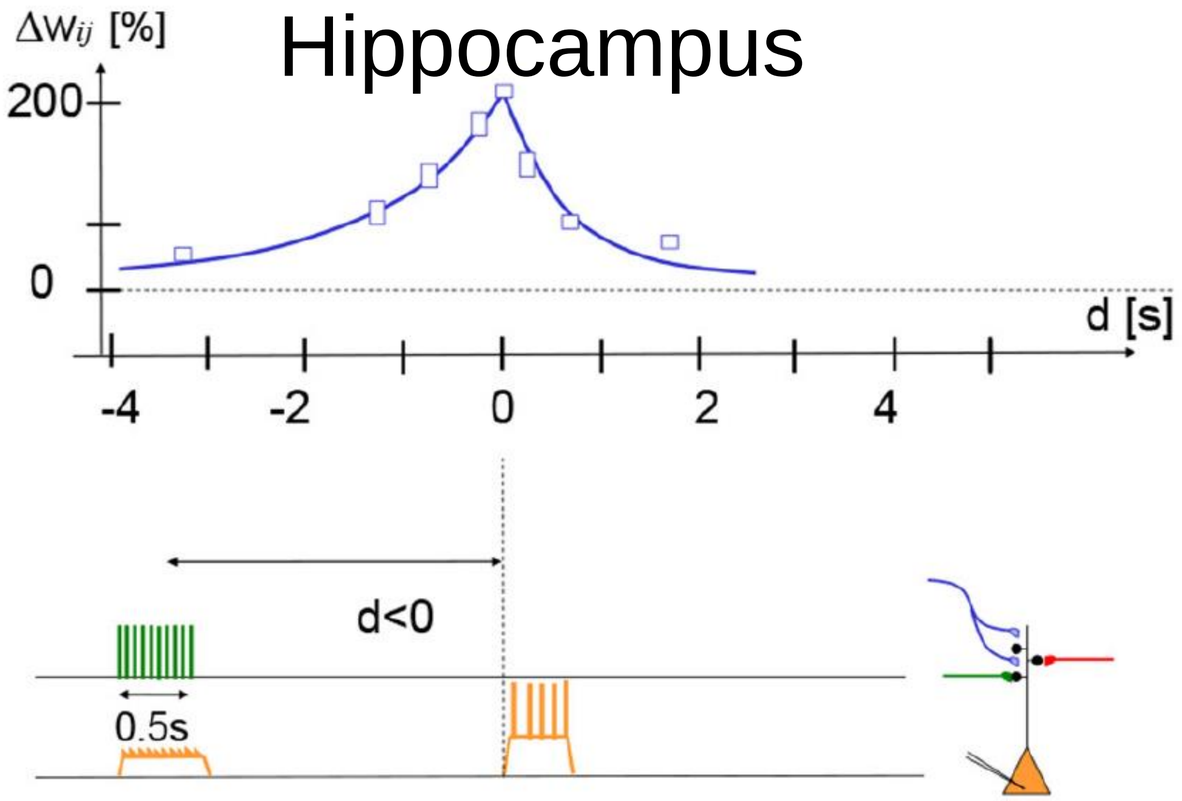

Bittner *et al., 2017* (Magee lab)

Complex spike
- indicates surprise(?)
- or novelty (?)
- triggered by ?

Image: Gerstner et al. (2018, review paper in Frontiers)

Previous slide.

In 2017, Bittner et al from the lab of Jeff Magee discovered a novel form of plasticity which was observed both in vivo and in vivo:

(i) In vivo: neuronal activity of many neurons in hippocampus was recorded while a mouse was running forward on a moving closed cyclic tape (which had a length of several meters). The authors observed that a new place cell appeared, for example in round 7, and then stayed on for the remaining rounds. It is single-shot learning since the place cells had immediately full strength, and nearly its final shape. The appearance of the place cell coincided with a complex spike in the neuron

(ii) In vitro (figure). A pyramidal cell in area CA1 (orange) receives input from many cells in CA3. Stimulation of fibers in the pathway from CA3 to CA1 (green) caused depolarization of the postsynaptic cell, but no spiking. Then the authors stimulated the postsynaptic cell such that it emitted a complex spike (visible as a burst of spikes). The combination of presynaptic stimulation and postsynaptic complex spike leads to a strengthening of the synapses even if the two events are separated by 2 seconds.

The authors interpreted this as a novel two-factor rule, with long synaptic traces. I prefer to see this as a three-factor rule where the pre-spike together with postsynaptic depolarization sets the trace, and the complex spike is a manifestation of the action of the third factor. Jeff Magee mentioned that in vivo they had the impression that a place cell emerged when the mouse was puzzled.

2024: **Learning rule**: Feedforward synapse on basal dendrite depend on lateral/feedback input (on apical dendrite)

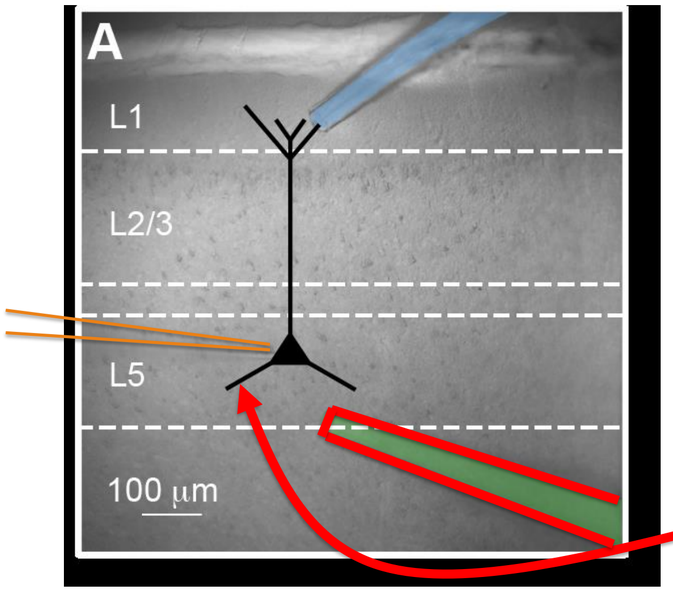

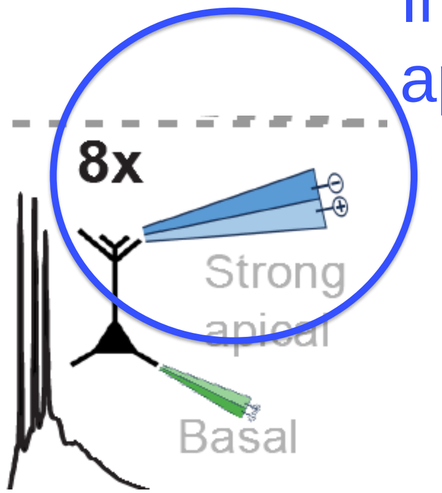

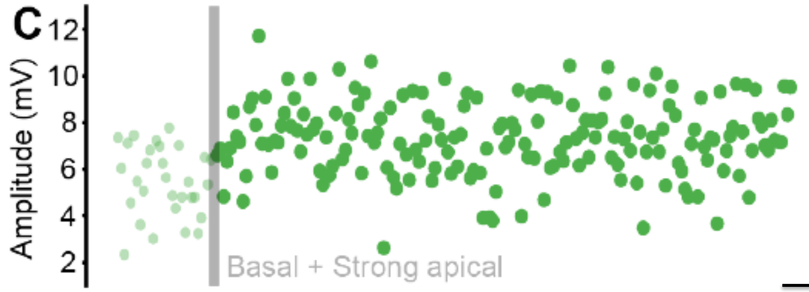

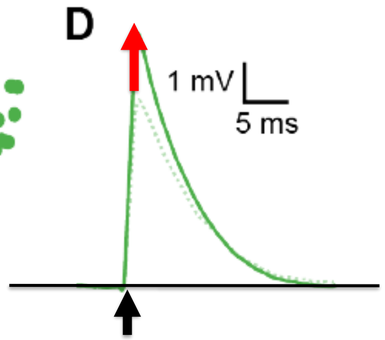

synapse on basal dendrite → LTP

Aceituno, …, Grewe, bioRxiv (2024)
https://doi.org/10.1101/2024.04.10.588837;

Recent experiments in L5: Grewe group (2024)
Experiments in L2/3: Williams and Holtmaat (2019)

Consistent with voltage-dependent plasticity (*Sjostrom et al 2001*)
and Clopath model (*Clopath et al. 2010)*

Such a rule useful to learn ‘good’ representation!

Experiments in Mouse Frontal Cortex, L5 cells, slice, from the Grewe lab (INI, Zurich). Two electrodes are used for extracellular stimulation at the basal dendrite (red-green) and apical dendrite (blue). Voltage is recorded with the brown electrode (A).
Initially, EPSPs are evoked by small-amplitude pulse stimuli (strength s1) with the red-green electrode yielding an EPSP of a few mV. Then the stimulation amplitude is increased (strength s2) so that the firing threshold is reached, and the postsynaptic neuron fires an isolated spike. After 8 repetitions (at 0.1Hz) no change in the EPSP amplitude is found. Thereafter the stimulation of basal synapses (with strength s2) is paired with stimulation of the apical dendrite, causing a short burst of spikes and a prolonged voltage response. After 8 repetitions (at 0.1Hz) the EPSP amplitude in response to stimulus s1 is increased (C and D).

These findings are consistent with experiments of J. Sjostrom (2001) and the voltage-dependent plasticity model of C. Clopath (2010): synaptic changes require either multiple postsynaptic spikes are a prolonged depolarization of the postsynaptic neuron, or a combination of both.

2. Classification of synaptic changes

Induction of changes

- fast (if stimulated appropriately)
- slow (homeostasis, over hours)

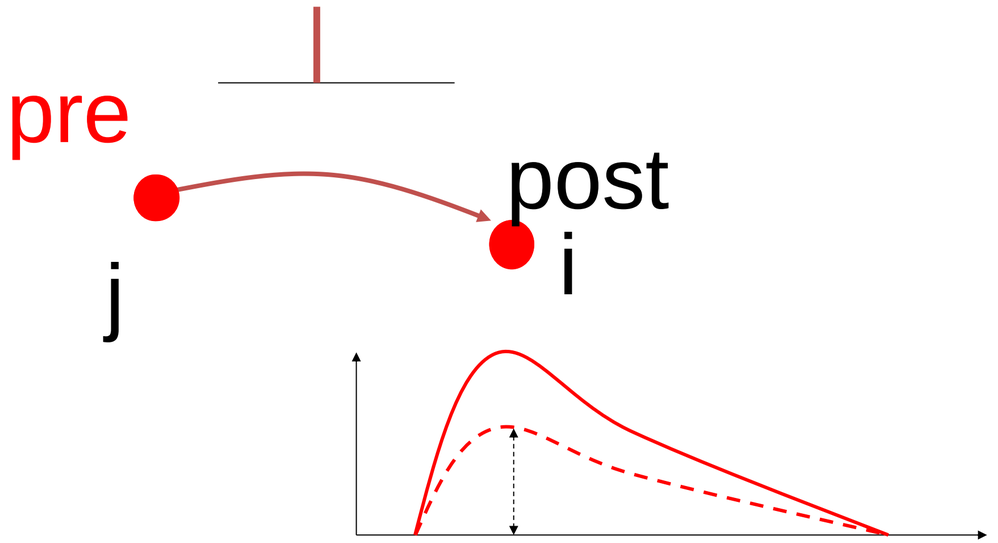

Persistence of changes

- long (LTP/LTD)
- short (short-term plasticity, STP)

Functionality

- useful for learning a new behavior/forming new memories
- useful for development (wiring for receptive field development)
- useful for activity control in network: **homeostasis**
- useful for representation learning

2. Summary: Classification of synaptic changes

Several categories can be used to classify synaptic changes:

1) Do changes **last for a long time** (hours: Long-Term Potentiation) or do they decay rapidly back to baseline (around a second: Short-Term Potentiation)?

2) Do changes depend mainly on presynaptic and postsynaptic activity (**Hebbian learning/2-factor rule**), or also on the presence of a neuromodulator **(three-factor rule)?**

3) Learning paradigm: is the learning scenario just exploiting input statistics (**unsupervised learning**/no teacher, no reward); or does it also involve notions of ‘reward’ or ‘success (**reinforcement learning/neuromodulator**)

4) Is synaptic plasticity induced by an **artificial protocol**, or by real-world learning experience? What is the protocol (STDP, bust-STDP, pairing pre-voltage, …)

Quiz 1. Synaptic Plasticity and Learning Rules

Long-term potentiation

[ ] has an acronym LTP
[ ] takes more than 10 minutes to induce
[ ] lasts more than 30 minutes
[ ] depends on presynaptic activity, and on spikes of the postsynaptic neuron
[ ] depends on presynaptic activity, and on the voltage of the postsynaptic neuron

Short-term potentiation

[ ] has an acronym STP
[ ] takes more than 10 minutes to induce
[ ] lasts more than 30 minutes
[ ] depends on presynaptic activity, but not on state of postsynaptic neuron

Synaptic changes

[ ] Hebbian learning depends on presynaptic activity and on state of postsynaptic neuron
[ ] Reinforcement learning depends on neuromodulators such as dopamine indicating reward

Computational Neuroscience: Neuronal Dynamics

EPFL

Lecture 14

Synaptic plasticity and Learning

Wulfram Gerstner

EPFL, Lausanne, Switzerland

Note done in class

1. **Synaptic plasticity**
   motivation and aims
2. **Classification of plasticity**
   short-term vs. long-term
   unsupervised vs. reward modulated
3. **Model of short-term plasticity**
4. **Models of long-term plasticity**
   - Hebbian learning rules
   - Bienenstock-Cooper-Munro rule
5. **Spiking Models of plasticity**

3. Model of short-term plasticity

See Week X on MOODLE or See week 3 on:

http://lcn.epfl.ch/~gerstner/NeuronalDynamics-MOOC1.html

Synapses, dendrites and the cable equation

Part 1 - Synapses (15 min)

Part 2 - Synaptic short term plasticity (9 min)

https://www.youtube.com/watch?v=iEz_SUsJMJ8

Reading for STP:
NEURONAL DYNAMICS
- Ch 3.1.3.

Cambridge Univ. Press

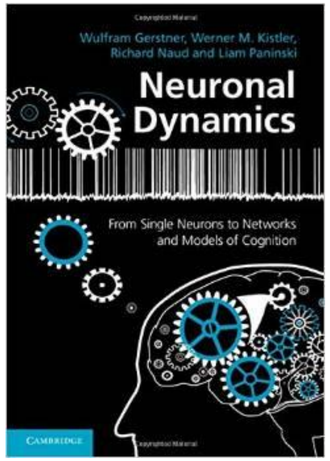

Computational Neuroscience: Neuronal Dynamics

EPFL

Lecture 14

Synaptic plasticity and Learning

Wulfram Gerstner

EPFL, Lausanne, Switzerland

Reading for plasticity:
NEURONAL DYNAMICS
- Ch. 19.1-19.3

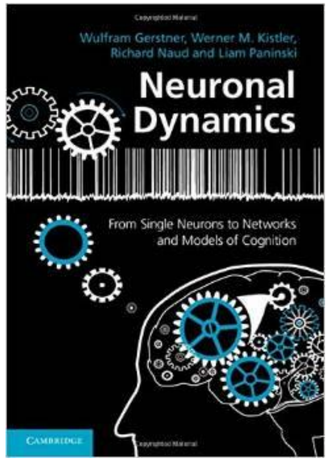

Cambridge Univ. Press

1. **Synaptic plasticity**
   motivation and aims
2. **Classification of plasticity**
   short-term vs. long-term
   unsupervised vs. reward modulated
3. **Model of short-term plasticity**
4. **Models of long-term plasticity**
   - Hebbian learning (rate model)
   - Bienenstock-Cooper-Munro rule
5. **Spiking Models of plasticity**
   - STDP as Hebbian learning
   - Model of STDP: synaptic traces
6. **From STDP to rate models**
7. **Triplet STDP model**
8. **Online learning of memories**

4. Repetition: Hebbian Learning (rate models)

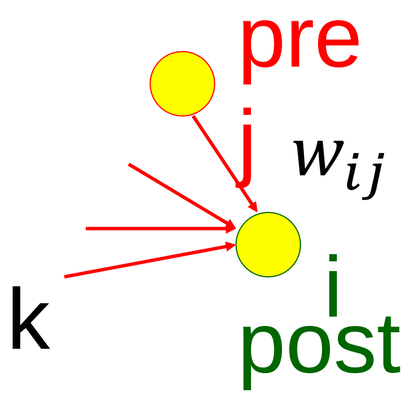

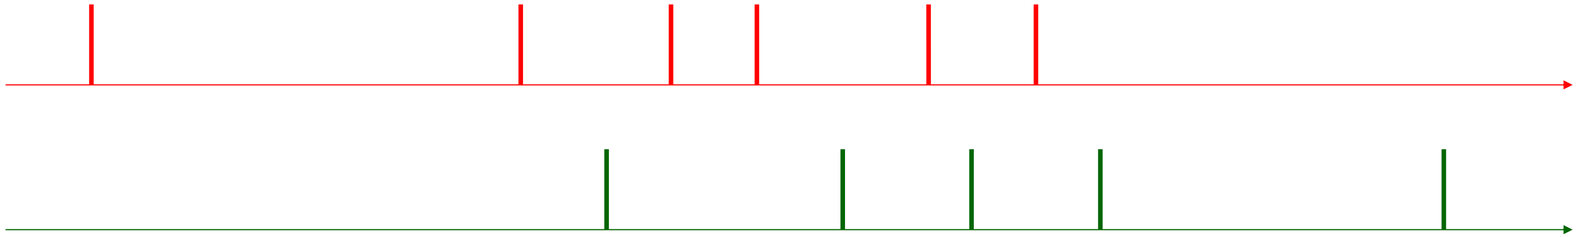

When an axon of cell **j** repeatedly or persistently takes part in firing cell **i**, then j's efficiency as one of the cells firing i is increased

Hebb, 1949

- local rule
- simultaneously active (correlations)

Rate model:

active = high rate = many spikes per second

4. Repetition: Rate-based Hebbian Learning

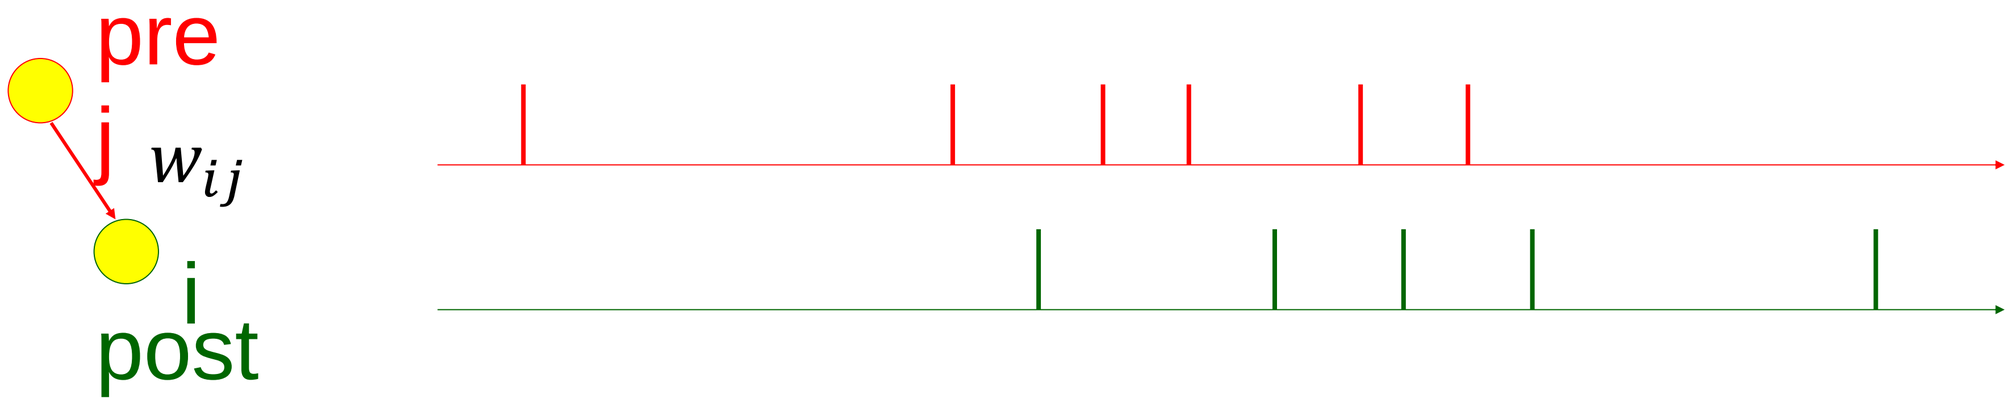

Local rule:

$$\frac{d}{dt}w_{ij} = F(w_{ij}, MOD; \nu_j^{pre}, \nu_i^{post})$$

Modulator MOD=const

Taylor expansion:

$$\frac{d}{dt}w_{ij} = a_0 + a_1^{pre}\nu_j^{pre} + a_1^{post}\nu_i^{post} + a_2^{corr}\nu_j^{pre}\nu_i^{post} + \ldots$$

$a = a(w_{ij})$

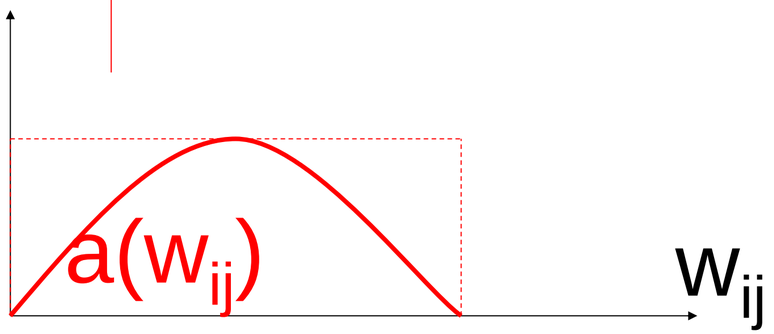

4. Repetition: Bienenstock-Cooper-Munro rule

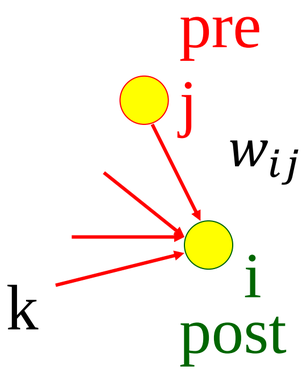

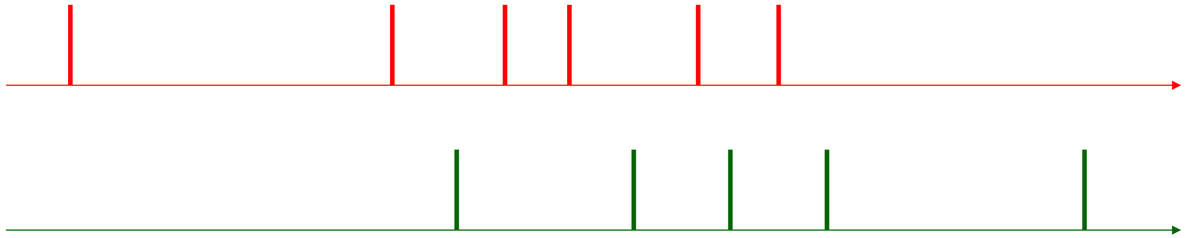

Bienenstock, Cooper Munro, 1982

presynaptically gated

$$\frac{d}{dt}w_{ij} = a_2^{corr}(v_i^{post} - \vartheta)v_j^{pre}$$

BCM: 3rd order ('triplet')

$$\frac{d}{dt}w_{ij} = bv_i^{post}(v_i^{post} - \vartheta)v_j^{pre}$$

$$= b(v_i^{post})^2 v_i^{pre} - b\vartheta v_i^{post} v_j^{pre}$$

triplet pair

$$\frac{d}{dt}w_{ij} = bv_i^{post}(v_i^{post} - \vartheta)v_j^{pre}$$

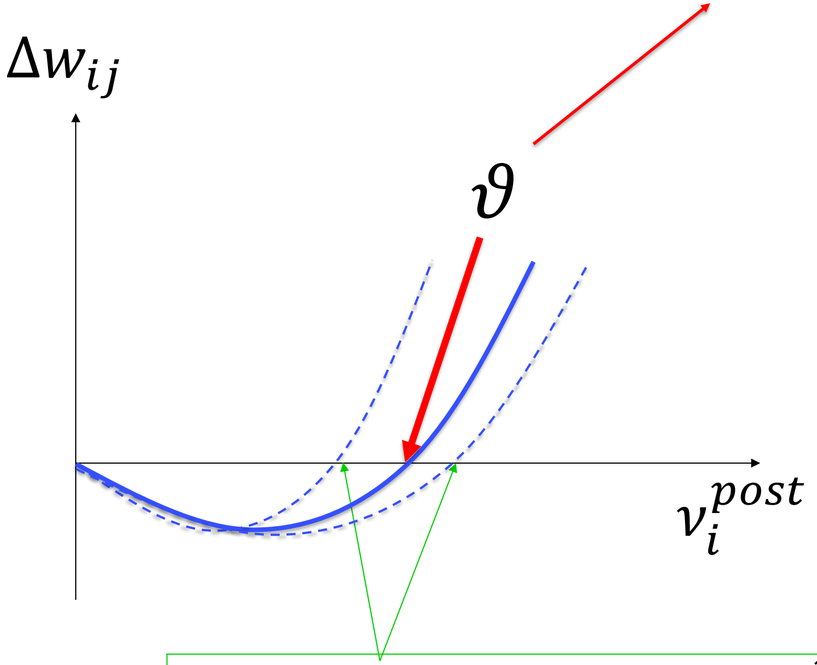

assume $v_j^{pre} > 0$

Homeostatis $\vartheta = f(\bar{v}_i^{post})$

Previous slide.
We saw in an earlier lecture that a nonlinear model of synaptic plasticity such as the BCM model can perform ICA. In the context of neuroscience it has been used to model the development of receptive fields. So far, all our plasticity models were rate-based models: the state of a neuron is described by a continuous variables $\nu_i^{post},\ and\ \nu_j^{pre}$

Later in this lecture we will see a spiking model of plasticity that is closely related to the BCM model. The terminology ‘pair’ and ‘triplet’ will become clear in the context of the spiking model.

4. Repetition: Development of Cortex

Initial:
random
connections

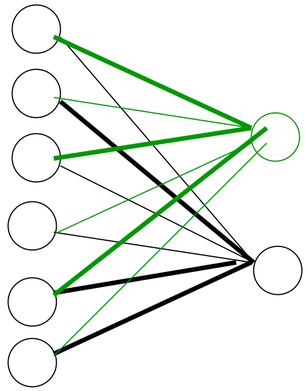

unselective neurons

BCM leads to specialized Neurons (developmental learning);
Bienenstock et al. 1982

Development and learning rules:
Willshaw&Malsburg, 1976
Linsker, 1986
K.D. Miller et al., 1989

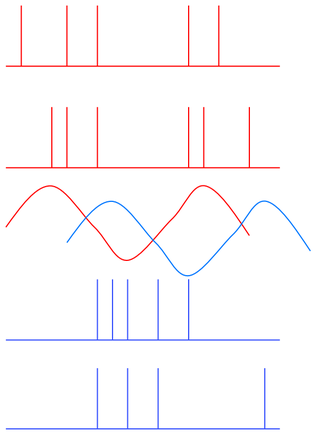

output neurons

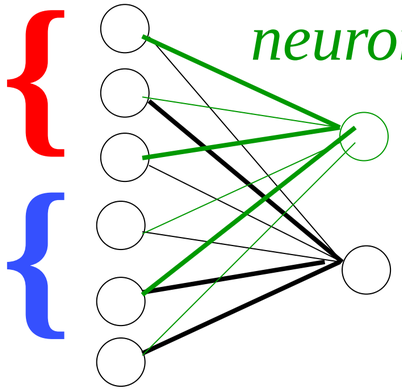

Correlated input

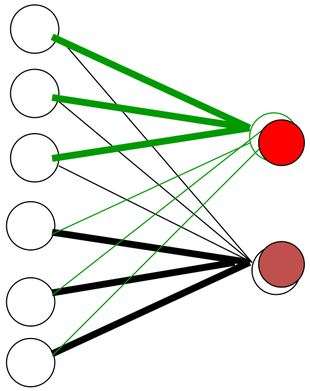

output neurons specialize:
Receptive fields

4. Repetition: Models for Hebbian Long-Term-Plasticity

- Many ‘Hebbian’ rules
- LTP and LTD
- Can describe RF development and ICA
- BCM is a well-known example
- Competition: some synapses grow at the expense of others

4. Summary: Models for Hebbian Learning

- Hebbian learning refers to a **family of learning rules**, rather than one specific rule.

- Rules can be classified by mapping them to a **Taylor explansion**.

- Terms with a negative coefficient induce **Long-Term Depression (LTD).**

- A clever combination of LTP and LTD can explain the **development of receptive fields (RF).**

- A clever combination of LTP and LTD leads to synaptic **competition: some synapses grow at the expense of others.** A well-known example of a Hebbian rule is the Bienenstock-Cooper-Munro (BCM) rule (*Exercise 1*).

Computational Neuroscience: Neuronal Dynamics

EPFL

Lecture 14

Synaptic plasticity and Learning

Wulfram Gerstner

EPFL, Lausanne, Switzerland

Reading for plasticity:
NEURONAL DYNAMICS
- Ch. 19.1-19.3

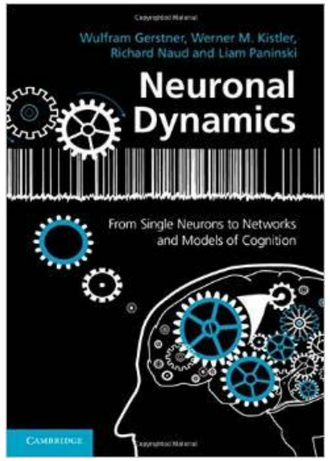

Cambridge Univ. Press

1. **Synaptic plasticity**
   motivation and aims
2. **Classification of plasticity**
   short-term vs. long-term
   unsupervised vs. reward modulated
3. **Model of short-term plasticity**
4. **Models of long-term plasticity**
   - Hebbian learning rules
   - Bienenstock-Cooper-Munro rule
5. **Spiking Models of plasticity**
   - STDP as Hebbian learning
   - Model of STDP: synaptic traces
6. **From STDP to rate models**
7. **Triplet STDP model**
8. **Clopath Model**

5. Review: STDP experiments

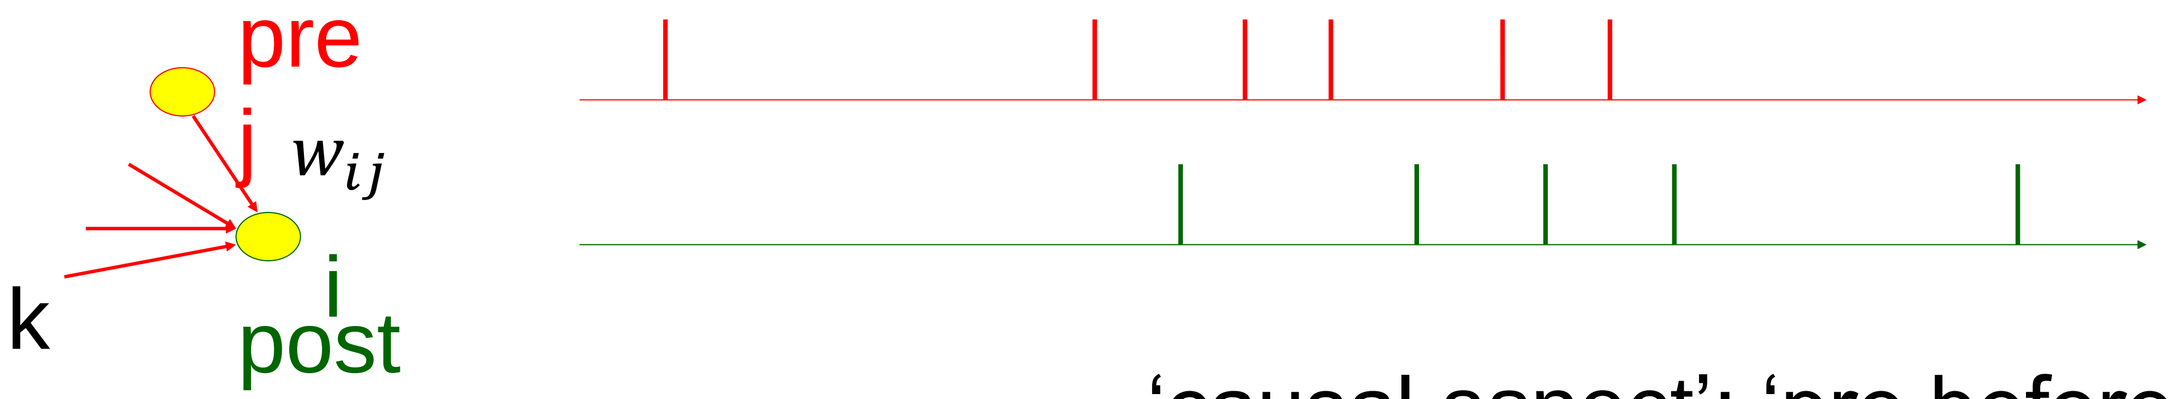

‘causal aspect’: ‘pre before post’

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**

Experiments: Levy and Stewart, 1983, …
Markram et al. 1997, Bi and Poo, 1998, Sjostrom et al. 2001…

Reviews: Sjostrom et al. 2008...
Markram et al. 2011, ...

Previous slide.

We have seen the citation of the text by Hebb several times. It is interesting to focus on the formulation ‘takes part in firing’ the postsynaptic neuron. This formulation suggests a causal aspect.

But causality requires that the presynaptic spike occurs before the postsynaptic one. Hence the exact timing of pre- and postsynaptic spikes should matter.

5. STDP as Hebbian Learning

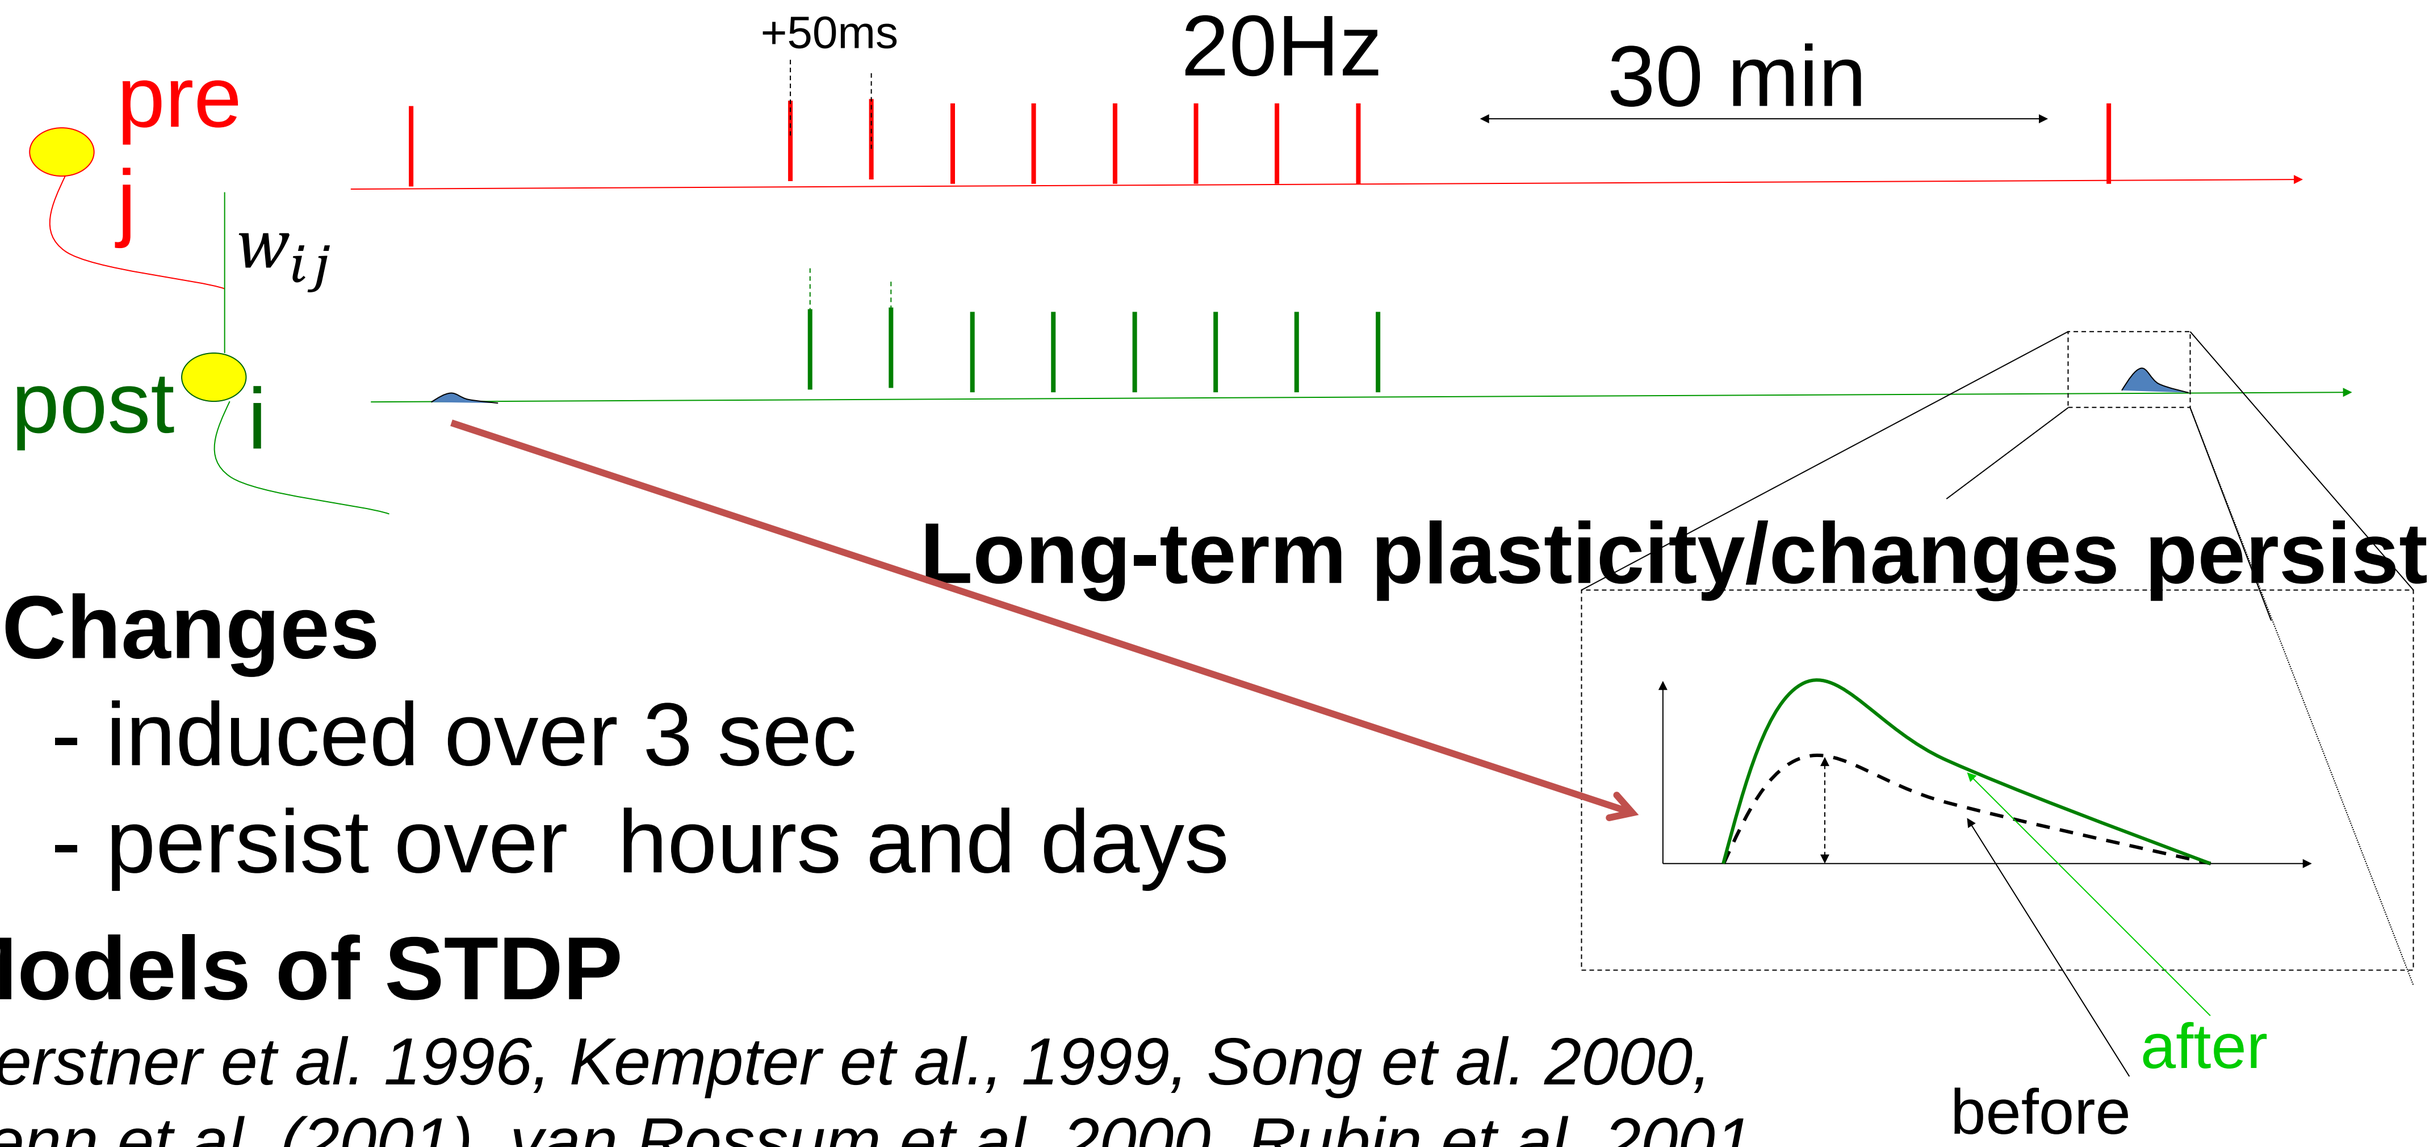

Changes

- induced over 3 sec
- persist over hours and days

Models of STDP

Gerstner et al. 1996, Kempter et al., 1999, Song et al. 2000, Senn et al. (2001), van Rossum et al. 2000, Rubin et al. 2001 Shouval et al. (2002), Clopath et al. (2010)

Previous slide.

In this example of an STDP experiment, both neurons are active, but the presynaptic neuron fires 10ms before the postsynaptic one. The pairs pre-post are repeated several times.

The question arises how the sequence of timing can be detected. The answer are synaptic traces that we see on the next slide.

5. Spike-timing dependent plasticity: ‘traces’ for STDP

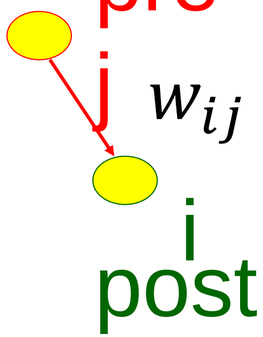

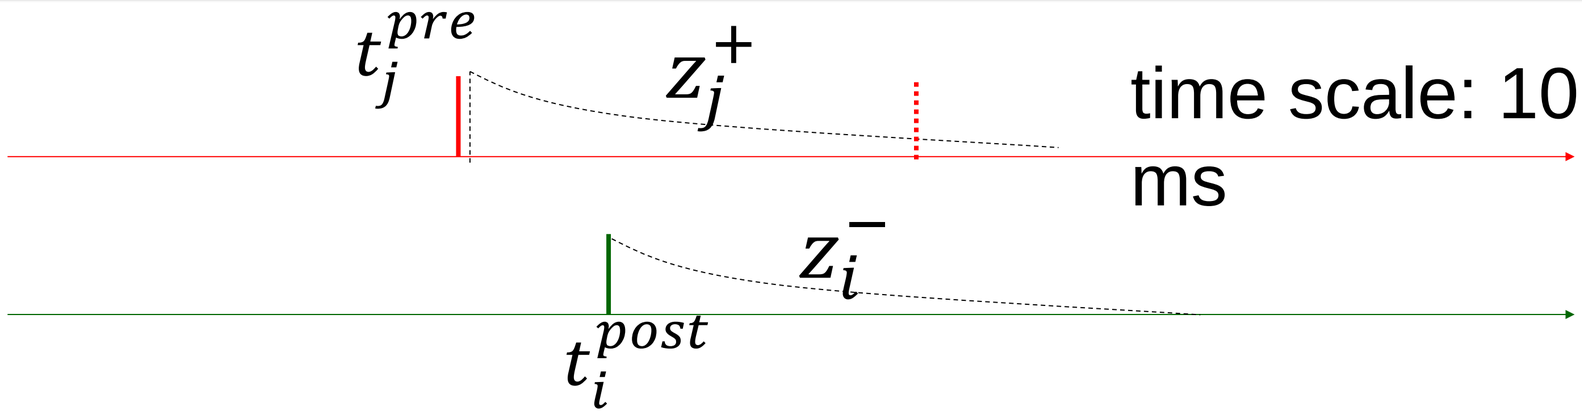

$$\tau_+ \frac{d}{dt} z_j^+ = -z_j^+ + \delta(t - t_j^{pre})$$

jump at presyn. spike

$$\tau_- \frac{d}{dt} z_i^- = -z_i^- + \delta(t - t_i^{post})$$

jump at postsyn. spike

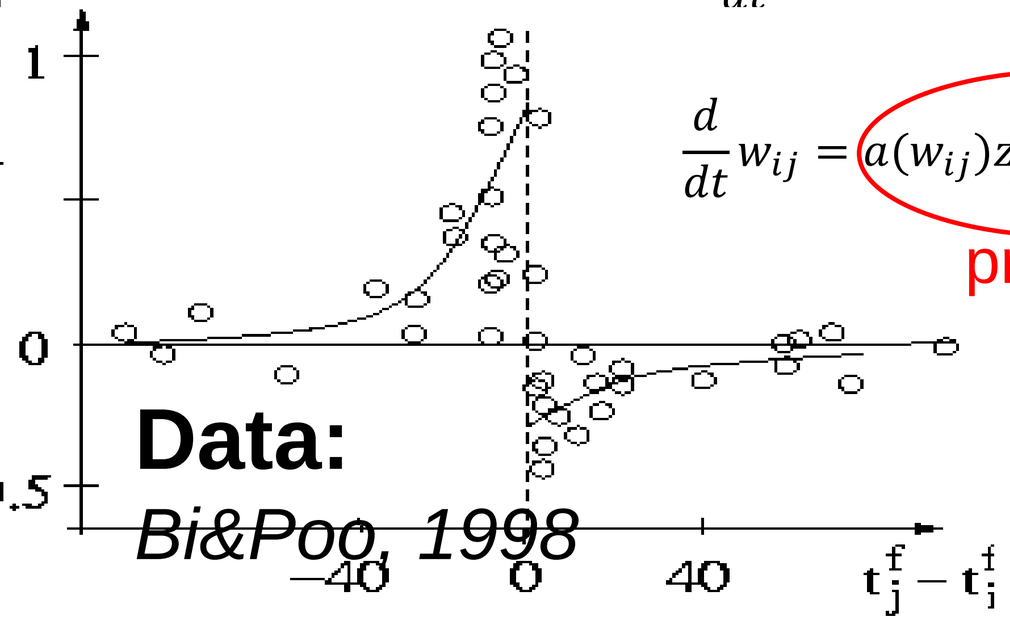

Data:
Bi&Poo, 1998

$$\frac{d}{dt} w_{ij} = a(w_{ij}) z_j^+ \delta(t - t_i^{post}) - b(w_{ij}) z_i^- \delta(t - t_j^{pre})$$

pair

pre-before-post

post-before-pre

→ **Simple STDP model**
(Gerstner et al. 1996,
Song-Miller-Abbott 2000, etc)

Previous slide.
A presynaptic spike at time t_j^{pre} leaves a trace z_j^+ at the synapse which decays over a few milliseconds. This trace could correspond to the fraction of glutamate bound to the postsynaptic channel.

A later postsynaptic spike also leaves a trace z_i^- . This trace could be the running average of the postsynaptic voltage or something else.

LTP happens if a postsynaptic spike occurs while the presynaptic trace is still non-zero; LTD happens if a presynaptic spike occurse while the postsynaptic trace is non-zero.

The relative timing of the pair of spikes yields the STDP window shown at the bottom left.

The effects of all pairs are simply added (even though variants exist).

5. Summary: Spike-timing dependent plasticity (STDP)

STDP is a form of Hebbian learning induced by spikes.
For a phenomenological model, we can take the view that each spike arriving at the presynaptic terminal leaves a trace at the synapse (e.g., amount of glutamate in the synaptic cleft, or bound to the postsynaptic receptor).
If a spike of the postsynaptic neuron coincides with the trace left by the presynaptic spike, a change happens (proportional to the momentary value of the trace).
The total weight change after many pairs of spikes is

$$\Delta w_{ij} = \sum_{f,f'} W(t_i^f - t_j^{f'})$$

Exercise

The goal of this exercise is to show that it is possible to account for the asymmetry in the STDP window using a simple microscopic model of synaptic plasticity.

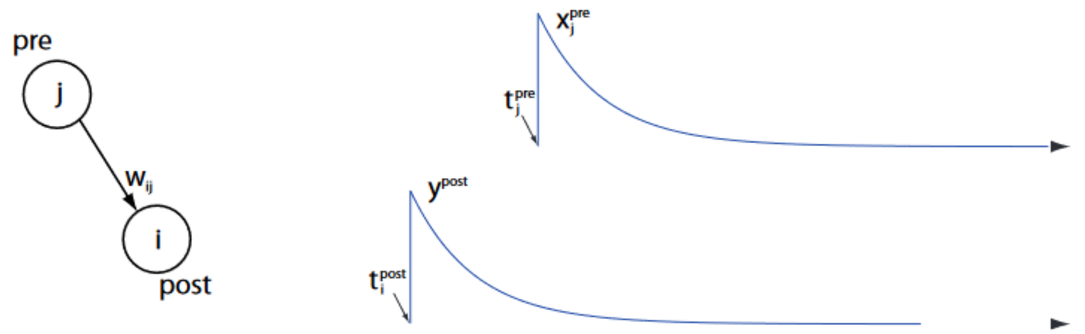

Figure 2: Memory traces of pre- and post-synaptic spike trains.

Suppose that the change in synaptic weight is controlled by the local concentration of two molecules $x^{\rm pre}$ and $y^{\rm post}$. The substance $x^{\rm pre}$ acts as a memory trace of presynaptic spikes in the sense that each presynaptic spike triggers an increase in the concentration of $x^{\rm pre}$:

$$\tau_+ \frac{d}{dt} x_j^{\rm pre} = -x_j^{\rm pre} + \delta(t - t_j^{\rm pre})\,. \tag{3}$$

Similarly, $y^{\rm post}$ is the trace left by the postsynaptic spike train,

$$\tau_- \frac{d}{dt} y^{\rm post} = -y^{\rm post} + \delta(t - t_i^{\rm post})\,. \tag{4}$$

Calculate the form of the learning window $\Delta w = f(\Delta t)$ – where $\Delta t = t_j^{\rm pre} - t_i^{\rm post}$ assuming that the synaptic weights are updated according to the rule

$$\frac{d}{dt} w_{ij} = a_+ x_j^{\rm pre} \delta(t - t_i^{\rm post}) - a_- y^{\rm post} \delta(t - t_j^{\rm pre})\,. \tag{5}$$

The constants a_+ and a_- are both positive.

Hint: Calculate the weight change for a pair of pre/post spikes. Consider the two cases $\Delta t > 0$ and $\Delta t < 0$.

Computational Neuroscience: Neuronal Dynamics

EPFL

Lecture 14

Synaptic plasticity and Learning

Wulfram Gerstner

EPFL, Lausanne, Switzerland

Reading for plasticity:
NEURONAL DYNAMICS
- Ch. 19.1-19.3

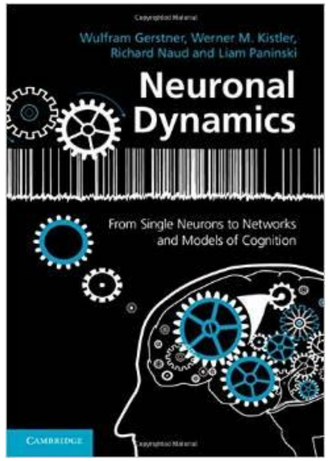

Cambridge Univ. Press

1. **Synaptic plasticity**
   motivation and aims
2. **Classification of plasticity**
   short-term vs. long-term
   unsupervised vs. reward modulated
3. **Model of short-term plasticity**
4. **Models of long-term plasticity**
   - Hebbian learning rules
   - Bienenstock-Cooper-Munro rule
5. **Spiking Models of plasticity**
   - STDP as Hebbian learning
   - Model of STDP: synaptic traces
6. **From STDP to rate models**
7. **Triplet STDP model**
8. **Online learning of memories**

6. from STDP to rate models

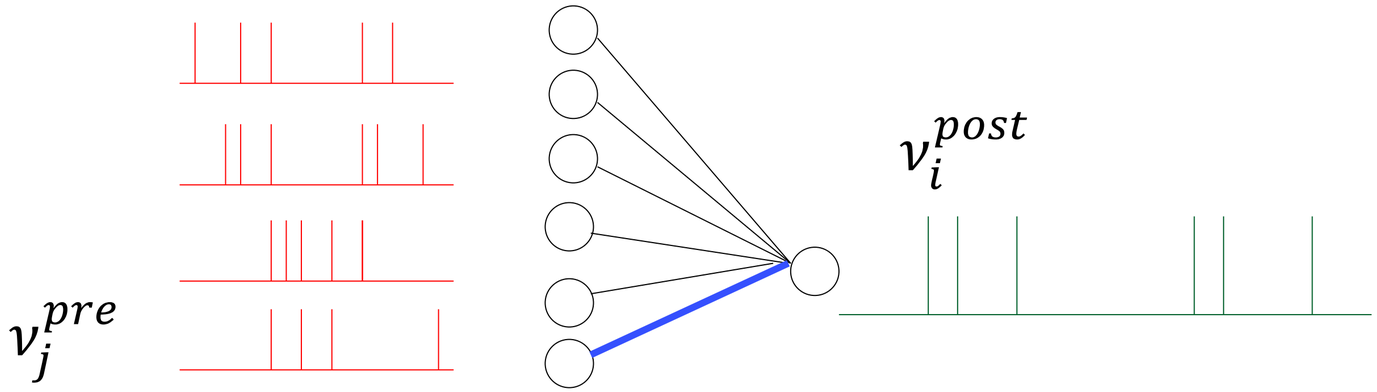

Blackboard

$$\Delta w_{ij} = \sum_{f,f'} W\left(t_i^f - t_j^{f'}\right) \quad (1)$$

$$\frac{1}{T}\Delta w_{ij} = \frac{1}{T}\int_0^T \int_{-\infty}^{\infty} W(s) S_i(t) S_j(t+s) ds \quad (2)$$

Preparation: STDP driven by Poisson spikes

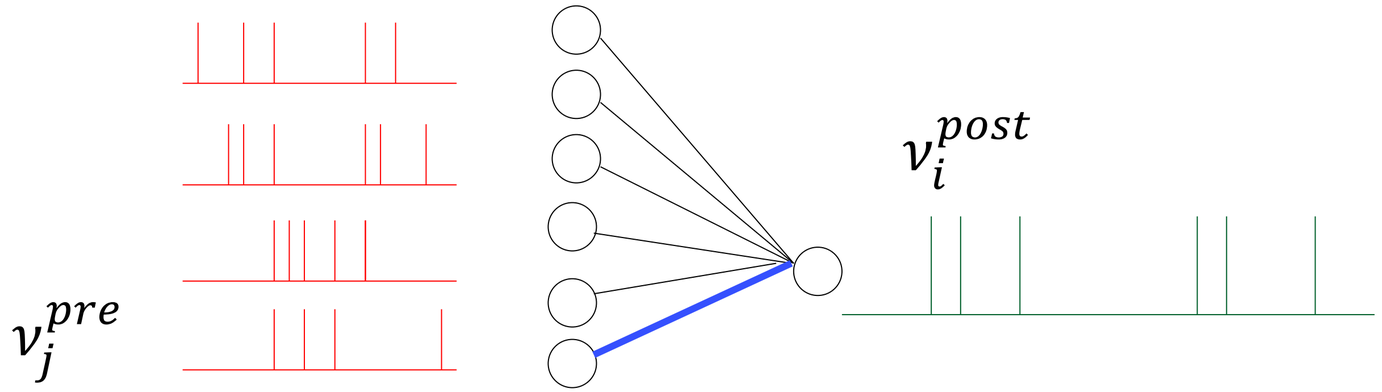

$$\Delta w_{ij} = \sum_{f,f'} W\left(t_i^f - t_j^{f'}\right) \quad (1)$$

Assume presynaptic spikes are generated by Poisson process with rate ν_j^{pre}

Assume postsynaptic spikes are generated by Poisson process with rate ν_i^{post} . Important spikes post are independent of pre.

What is the expected change of weights in a time T ?

$(T >> \tau^{LTP}, \tau^{LTD})$

6. from STDP to rate models

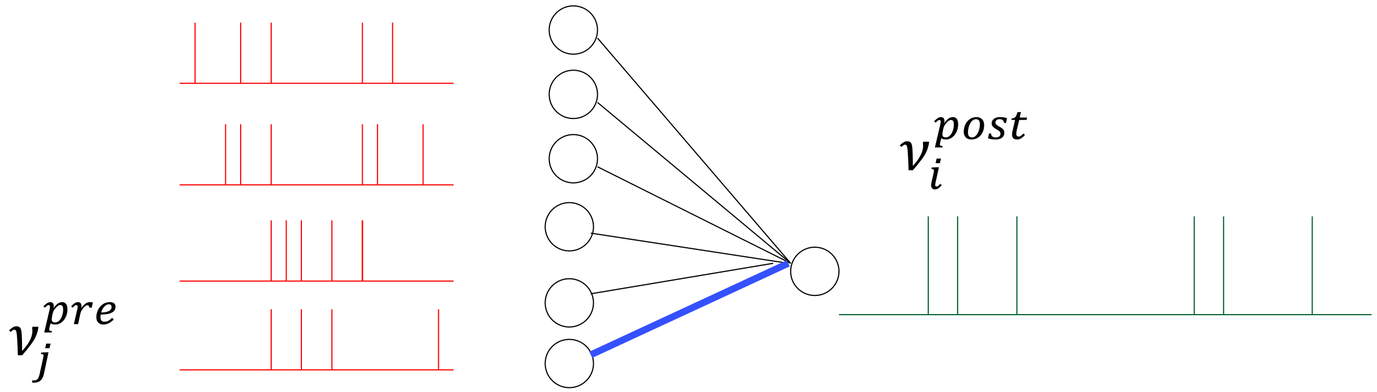

$$\left\langle \frac{1}{T}\Delta w_{ij} \right\rangle = \frac{1}{T}\int_0^T \int_{-\infty}^{\infty} W(s)\langle S_i(t)S_j(t+s)\rangle ds$$

$$= \frac{1}{T}\int_0^T \int_{-\infty}^{\infty} W(s)\langle S_i(t)\rangle\langle S_j(t+s)\rangle ds$$

$$= \nu_j^{pre}\,\nu_i^{post}\int_{-\infty}^{+\infty} W(s)ds \quad (4)$$

Previous slides.
Eq. (1) In the pair-based plasticity rule, the total weight change is the sum over all pair of spikes, inserted into the ‘STDP window’ $W\left(t_i^f - t_j^{f'}\right)$.

The same equation can be written as an integral, where S(t) denotes the spike train, that is the sequence of pulses. This gives Eq. (2).

We now assume that all pulses are generated by Poisson processes. Moreover, since the postsynaptic neuron receives spike trains from thousands of different input neurons, we assume that the spikes of the output spike trains are not correlated with those of the input spike trains (beyond correlations arising from the rates, i.e., if one neuron switches to firing at a higher rate, the output neuron will also fire at a slightly higher rate; yet we assume here that there are no spike-spike correlations)

Then, Eq. (4), the expected weight change (averaged over the statistics of the Poisson process), is just the expectations of the individual Poisson processes. This results in a plasticity rule for firing rates. If the firing rates are constant, then the integral over the STDP window yields the pair based of the Taylor expansion of the rate-based Hebbrule.

6. Summary: from STDP to rate models

In an STDP model, changes of synapses depend on the exact timing of pre- and postsynaptic spikes.
However, if we assume that both presynaptic and postsynaptic spike trains are generated by a homogeneous Poisson Process (with stationary firing rates ν_i and ν_i), we can translate the effect induced by STDP after many spikes into an equivalent rate model by evaluating the expected change.

The standard STDP window gives then a rate model

$$c\ \nu_i\ \nu_j$$

where c is the integral over the STDP window W(s).

Expectations and Correlations of Poisson spike train:

see my other class, or

Watch video ‘Membrane Potential fluctuations’ on:
http://lcn.epfl.ch/~gerstner/NeuronalDynamics-MOOCall.html

Computational Neuroscience: Neuronal Dynamics

EPFL

Synaptic Plasticity and Learning

Wulfram Gerstner

EPFL, Lausanne, Switzerland

Reading for plasticity:
NEURONAL DYNAMICS
- Ch. 19.1-19.3

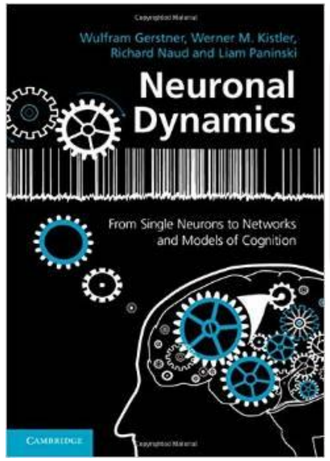

Cambridge Univ. Press

1. **Synaptic plasticity: aims**
2. **Classification of plasticity**
3. **Model of short-term plasticity**
4. **Models of long-term plasticity**
5. **Spike-Timing Models of plasticity**
   - STDP as Hebbian learning
   - Model of STDP: synaptic traces
6. **From spiking models to rate models**
7. **Triplet STDP model**
   - Relation to experiments
   - Relation to BCM rule
8. **Clopath model**

7. Why do we need a Triplet STDP model?

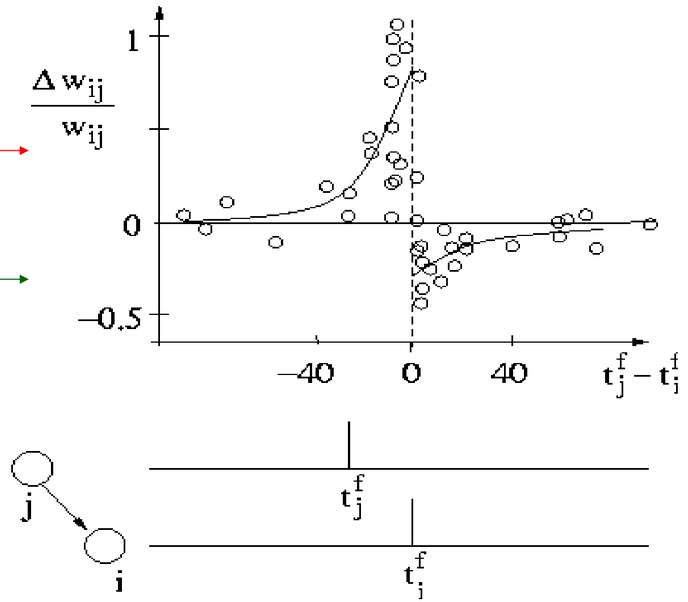

60 repetitions of spike-pairs ‘pre-post’:
Amount of LTP depends on repetition frequency:

→ 1) STDP window is only part of story

→ 2) Pair-based STDP model is not sufficient

Previous slide.
The STDP window (shown on the right-hand side) gives a biased view of the underlying plasticity process since it focuses only on the relative timing within one pair of spikes.
However, induction of LTP requires several repetition of pre-post-pairs and the repetition frequency matters!

7. frequency dependence of STDP

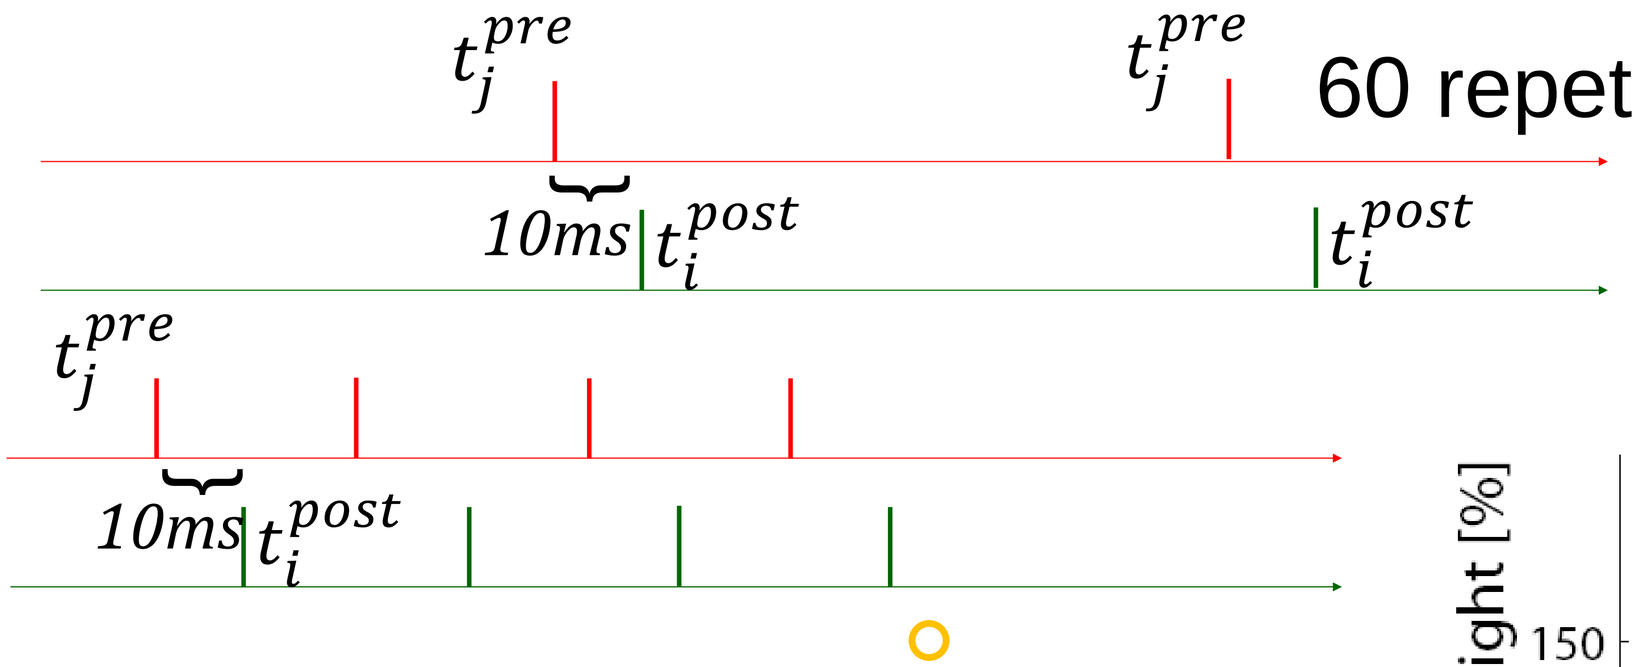

increase of repetition frequency ρ

Sjostrom et al. 2001
See also:
Markram et al. 1997,
Senn et al. 2001,

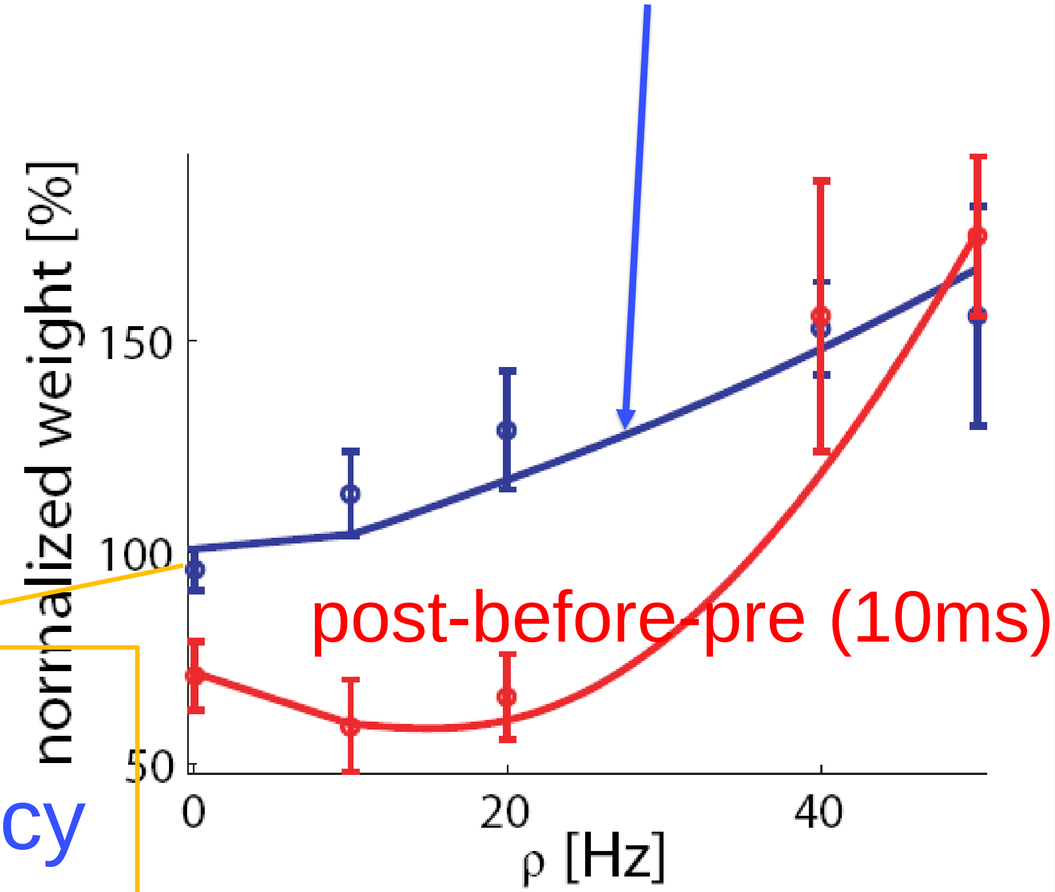

No plasticity
At low frequency

Previous slide.
The standard STDP window is evaluated at a repetition frequency ρ of 20 Hz.

At a repetition frequency of 50Hz, the pre-post-pre-post … sequence is symmetric (except for the very first spike). Therefore, it is expected that the LTD protocol and the LTP protocol give the same result. What is not expected is that the net effect if positive!

If the repetition frequency is reduced to 0.1 Hz, there is no LTP even though the same number of pre-post pairs (with distance 10ms) was used as at the other repetition frequencies. This shows that LTP is NOT induced by pairs of spikes!

The triplet STDP model discussed on the next slides is able to account for these observations.

7. Triplet STDP model

Triplet
LTP

$$\frac{d}{dt} w_j^+ = +A^+ \, z_j^+ \, z_i^{slow} \, \delta(t - t_i^{post})$$

pre

post

post

Triplet

fast exponential trace:

z_j^+

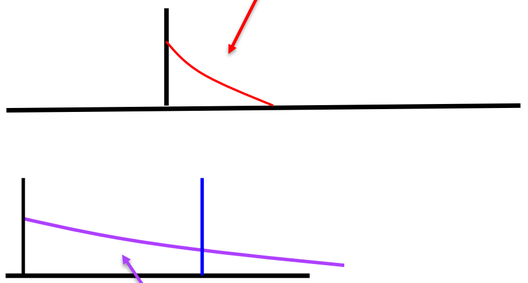

slow exponential trace:

z_i^{slow}

Previous slide.

LTP needs not a pair of spikes, but a triplet of spikes post-pre-post (or pre-post-post).

The first post-spike leaves slow trace. The pre-spike leaves a faster trace (e.g. glutamate bound to the postsynaptic terminal). The next postsynaptic spike causes the weight change proportional to
(i) The momentary value of the slow postsynaptic trace.
(ii) The momebtary value of the presynaptic trace

7. Triplet STDP model

Pfister and Gerstner, 2006

$$\frac{d}{dt} w_j =$$

$$-B z_i^- \delta(t - t_j^{pre})$$

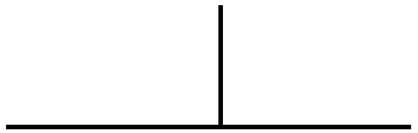

Pre: spike

fast exponential trace: z_i^-

Post: spike-trace

$$+A^+ z_j^+ z_i^{slow} \delta(t - t_i^{post})$$

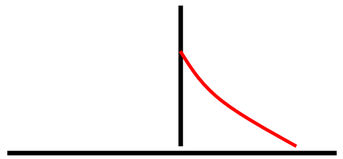

Pre: spike-trace

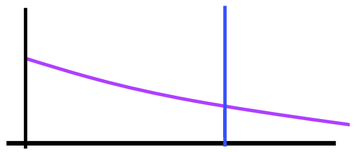

Post: spike-now

spike-trace

Previous slide.

The triplet LTP term is combined with a pair-based LTD term (similar to the standard STDP model).

7. Triplet STDP model

Pfister and Gerstner, 2006

Similar triplet mechanism in
Senn et al. 2001,
Rubin et al. 2005,
Clopath et al. 2010

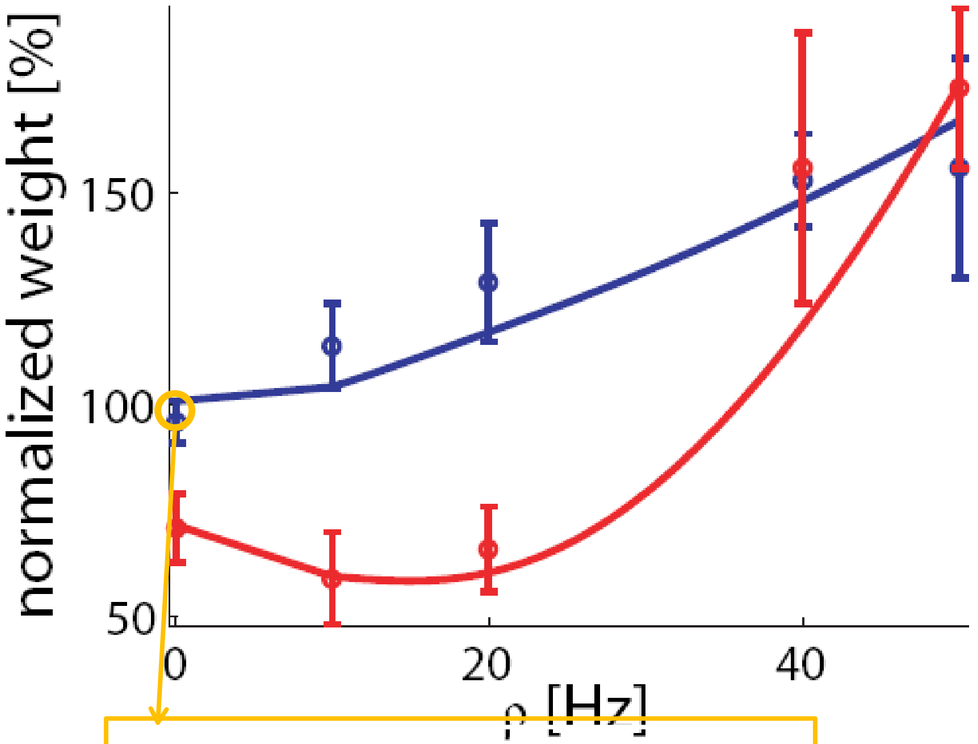

No plasticity
At low frequency

Previous slide.
The result is the triplet STDP model. This model is closely related to earlier models, in particular an elegant model of Walter Senn et al. (2001) and a calcium-based model of Jonathan Rubin et al (2005).

7. Triplet STDP model → BCM model

$$\frac{d}{dt}w_j = -Bz_i^-\delta(t - t_j^{pre})$$

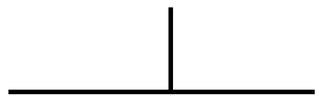

Pre: spike

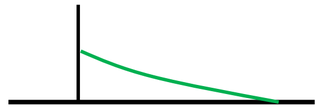

Post: spike-trace

$$+A^+ z_j^+ z_i^{slow}\delta(t - t_i^{post})$$

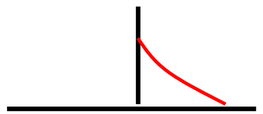

Pre: spike-trace

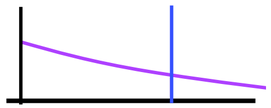

Post: spike-now AND
Post: slow spike-trace

Assume Poisson firing

Pfister and Gerstner, 2006

$$\frac{d}{dt}w_j = c^+A^+\nu_j^{pre}(\nu_i^{post})^2 - c^-B \quad \nu_j^{pre}\nu_i^{post}$$

Bienenstock, Cooper, Munro, 1982

Previous slide.
If we assume that all spikes are generated by Poisson processes, we can now use the same arguments as for the pair-based STDP model.

The resulting weight change is then exactly the one of the BCM model.

An important side result is that the parameters of the (ad hoc) BCM model can hence be extracted from STDP experiments!

7. Summary: Triplet STDP → BCM model

Triplet STDP model

- parameters can be extracted from experimental data

- for Poisson spikes closely related to rate-based BCM

- but captures additional spike-timing effects (pre-before-post)

- simple pair-based STDP model is not sufficient,
  because STDP depends also on repetition frequency
  (and not only on relative timing of pairs of spikes).

Computational Neuroscience: Neuronal Dynamics

EPFL

Synaptic Plasticity and Learning

Wulfram Gerstner

EPFL, Lausanne, Switzerland

Reading for plasticity:
NEURONAL DYNAMICS
- Ch. 19.1-19.3

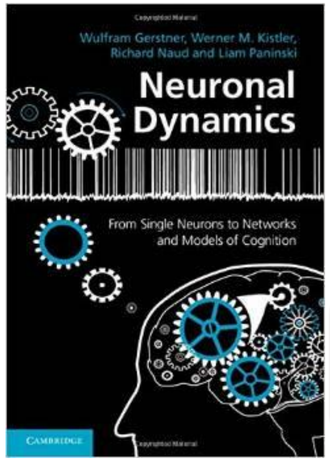

Cambridge Univ. Press

1. **Synaptic plasticity: aims**
2. **Classification of plasticity**
3. **Model of short-term plasticity**
4. **Models of long-term plasticity**
5. **Spike-Timing Models of plasticity**
   - STDP as Hebbian learning
   - Model of STDP: synaptic traces
6. **From spiking models to rate models**
7. **Triplet STDP model**
   - Relation to experiments
   - Relation to BCM rule
8. **Clopath model**

Experimental induction protocols (1)

STDP

60 repetitions at 20Hz

5-10 synapses

20Hz

Intracellular electrode (Pulse injection)

+50ms

+10ms

Intracellular recording electrode,

Pulse injection

Hebbian interpretation:

Pre-post = causal relation

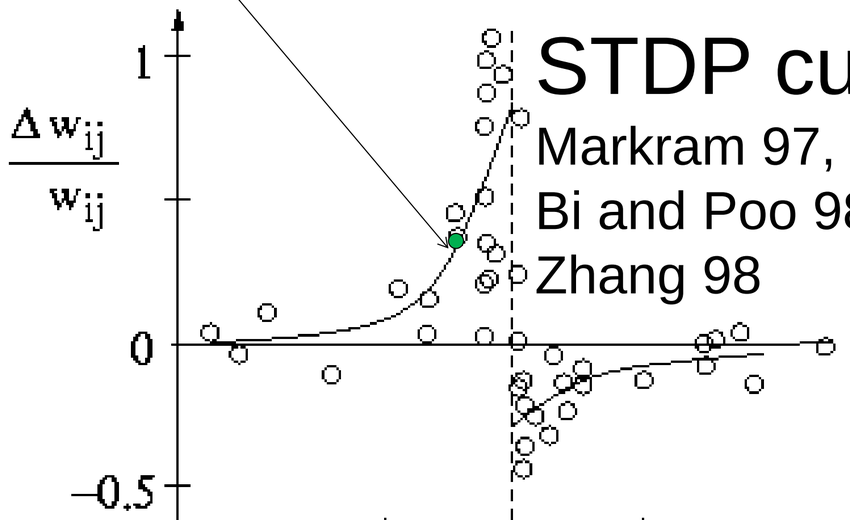

STDP curve,

Markram 97,

Bi and Poo 98,

Zhang 98

EPSP amplitude

100%

stim

0

30 min

Experimental induction protocols (2)

Extracellular electrode (stimulus)

Low-frequency paired
With depolarization
100 pulses at 2Hz
50 synapses

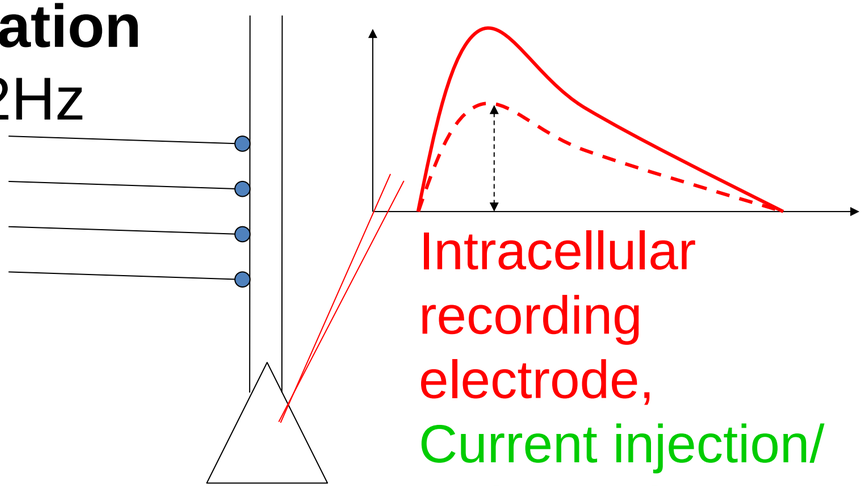

Intracellular recording electrode,
Current injection/ voltage clamp

Hebbian interpretation:

Depolarization similar to activity of postsyn. neuron, together with presyn. spike arrival

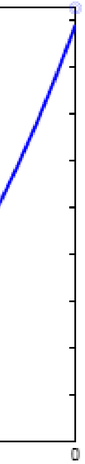

Voltage dependence

Ngezehajo et al. 2000,
Artola et al.

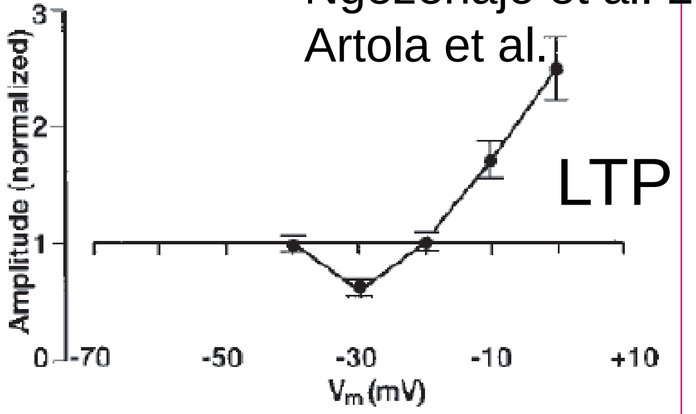

LTP

EPSP amplitude

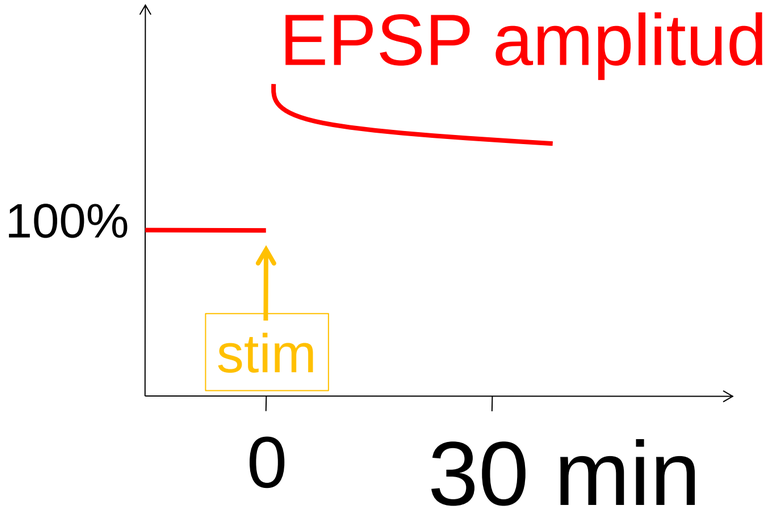

Reminder: Clopath model (1) – static voltage dependence

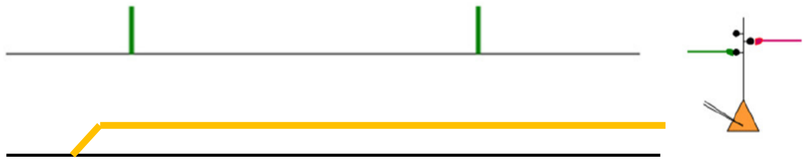

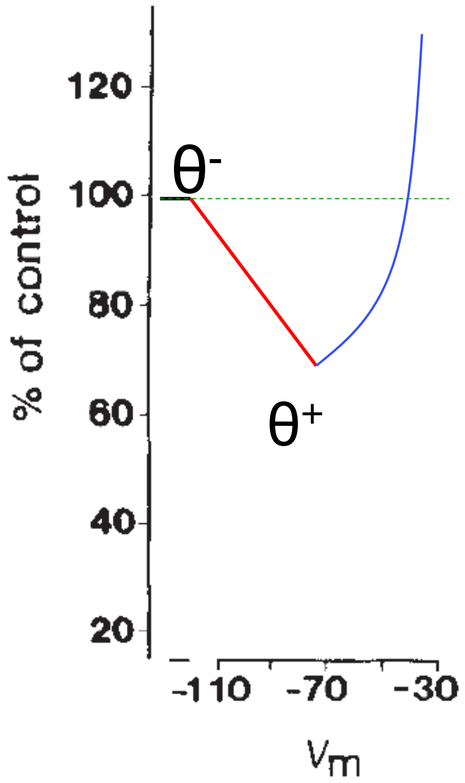

Depression term $-A^- \sigma^-(u) pre$

$-\sigma^-(u)$

θ-

ReLU

Potentiation term $+A^+ [\sigma^+(u)]^2 pre$

$[\sigma^+(u)]^2$

θ+

Ngezahayo et al. *J. of Neurosci.*, 2000

Artola, Bröcher, Singer. *Nature* 1990

Clopath et.al., Nature Neuroscience, 2010;

Previous slides.
The triplet model is an excellent model of STDP induction, but – by design - it cannot account for voltage dependence of plasticity.

The Clopath model can be seen as a voltage-based formulation of the triplet model.

As a reminder, spikes – as used in STDP experiments – are short voltage pulses.
Hence each spike leaves a trace in the Clopath model.
Moreover, by construction, the Clopath model has explicit voltage dependence as found in the experiments of Ngezehayou et al. (2000) or Artola et al (1990).

Reminder: Clopath model (2) – dynamics of depression

$$\frac{d}{dt} w_j^- = -A_2^- X_j(t) \sigma^-(\bar{u}_-)$$

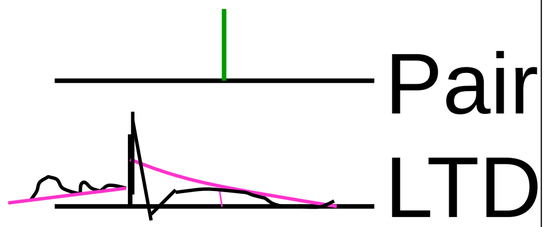

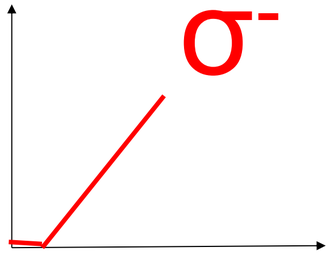

resting potential

pre

Some trace,
calcium?
second messenger?

where

$$X_j(t) = \sum_i \delta(t - t_j^i)$$ presynaptic spike train

$$\tau_{u-} \frac{d}{dt} \bar{u}_- = -\bar{u}_- + u$$ postsynaptic u low-pass (τ_{u-})

$$\sigma^-(\bar{u}_-) = [\bar{u}_- - \theta^-]^+$$ piecewise linear function (ReLU)

Reminder: Clopath model summary

Clopath et.al. Nat.NS 2010;

$$\frac{d}{dt}w_j = -A_2^- X_j(t)\sigma^-(\bar{u}_-) + A_3^+ \bar{x}_j^{pre}\sigma^+(u)\sigma^-(\bar{u}_+)$$

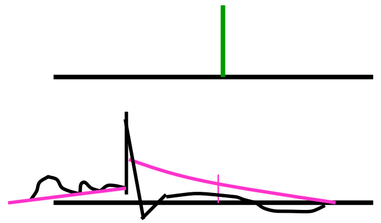

Pre: spike

Post: voltage-av

LTD

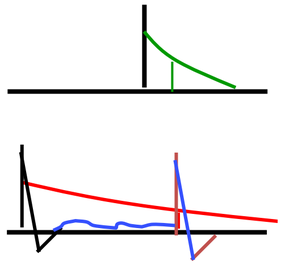

Pre: spike-trace

Post: voltage-now
voltage-av

LTP

LTP (positive change) needs: - presynaptic spike a few ms before
- voltage above threshold now
- voltage above rest previously

LTD (negative change) needs: - voltage > rest a few ms before
- presynaptic spike now

Clopath model applied to STDP experiments

Clopath et.al. Nat.NS 2010;

$$\frac{d}{dt}w_j^+ = +A_3^+ \bar{x}_j^{pre} \sigma^+(u)\sigma^-(\bar{u}_+)$$

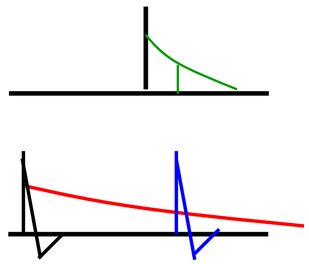

Triplet LTP

pre post post

Triplet

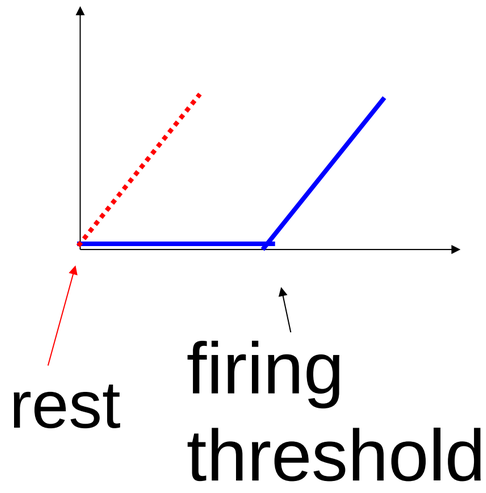

where

$$\tau_r \frac{d}{dt}\bar{x}_j^{pre} = -r_j^{pre} + X_j$$

trace of presyn. spike: glutamate?

$$\tau_{u+} \frac{d}{dt}\bar{u}_+ = -\bar{u}_+ + u$$

postsynaptic u low-pass: Ca concentration?

$$\sigma^{-/+}(x) = [x - \theta^{-/+}]^+$$

piecewise linear function

Previous slide.
We suppose that spikes are artificially induced by the experimentalist using injection of short current pulses. Then the voltage trajectory is essentially described by isolated spikes.

In this case, the traces left by the voltage are identical to the traces left by the spikes in the triplet model.

Clopath model applied to STDP experiments

Clopath et.al. Nat.NS 2010;

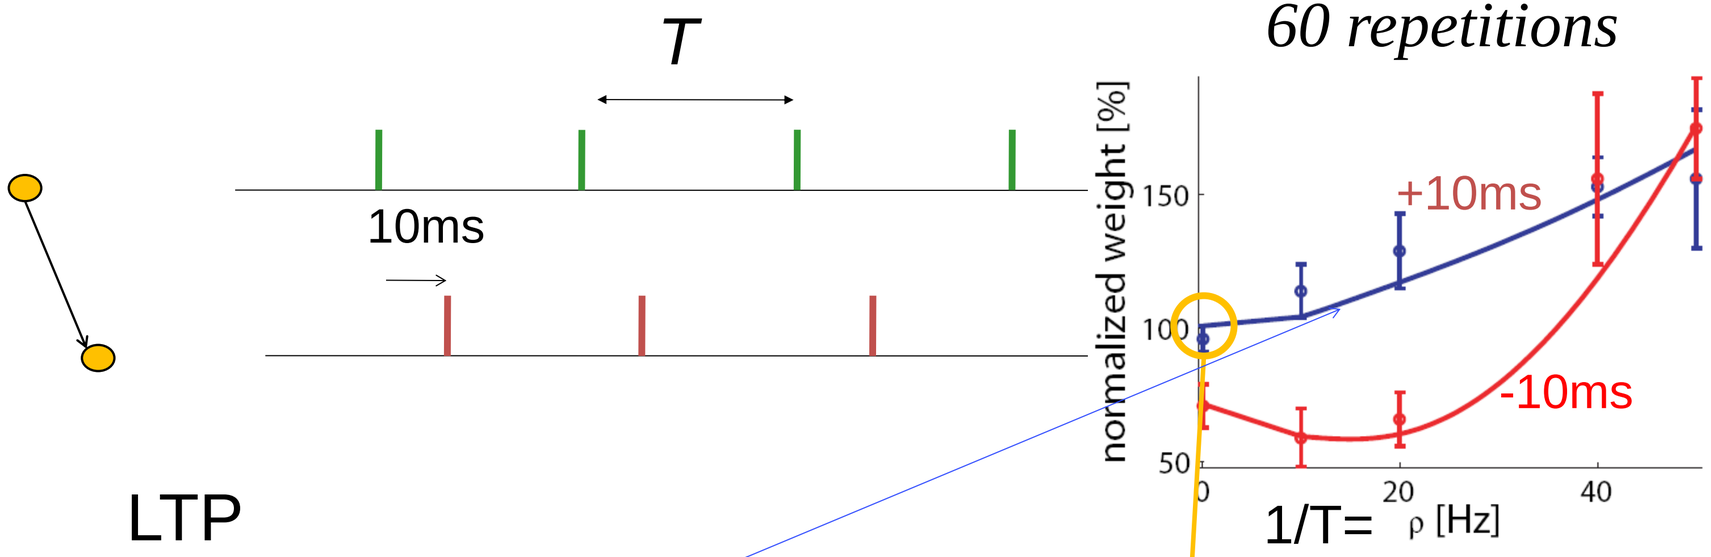

LTP
+10ms

Data: Sjostrom
See also:
Senn et al. 2001

Model

Clopath et.al., Nat. Neurosci. 2010
See also:
Pfister et al., J. Neurosci. 2006
Sjostrom et al., Neuron, 2001
Senn et al. 2001

No plasticity
at low frequency

Previous slide.

Hence it is expected, that the Clopath model explains the frequency dependence of STDP. The interval T between two repetitions yields the repetition frequency

$\rho = 1/T$

Clopath model applied to STDP experiments

Clopath et.al. Nat.NS 2010;

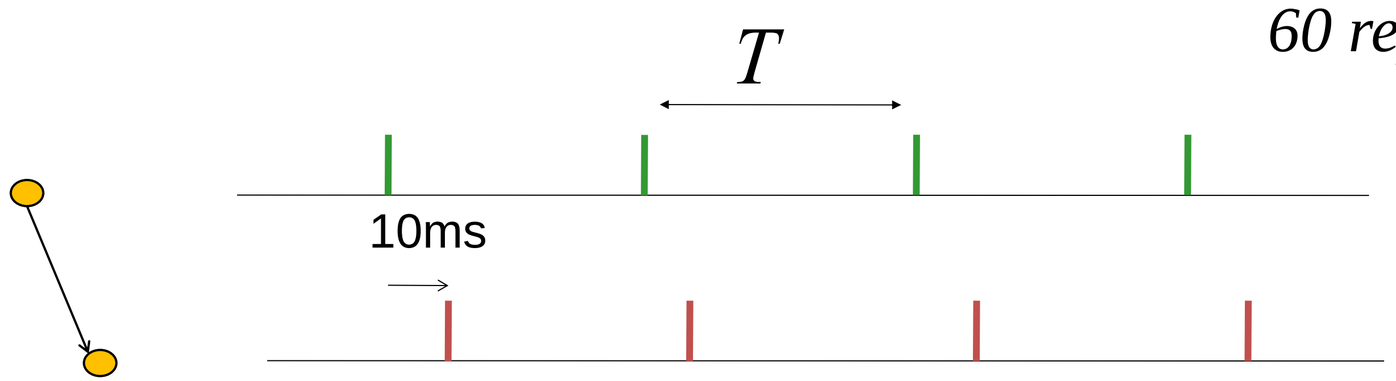

LTP
+10ms

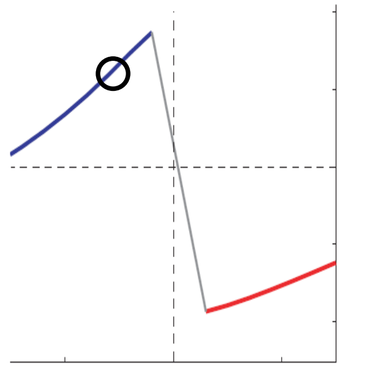

Plasticity is voltage dependent

Postsynaptic spike=voltage peak

Previous slide.
Moreover, the Clopath model explains the spike-timing dependence as summarized in an STDP window.

STDP interacts with voltage: Data and Clopath model

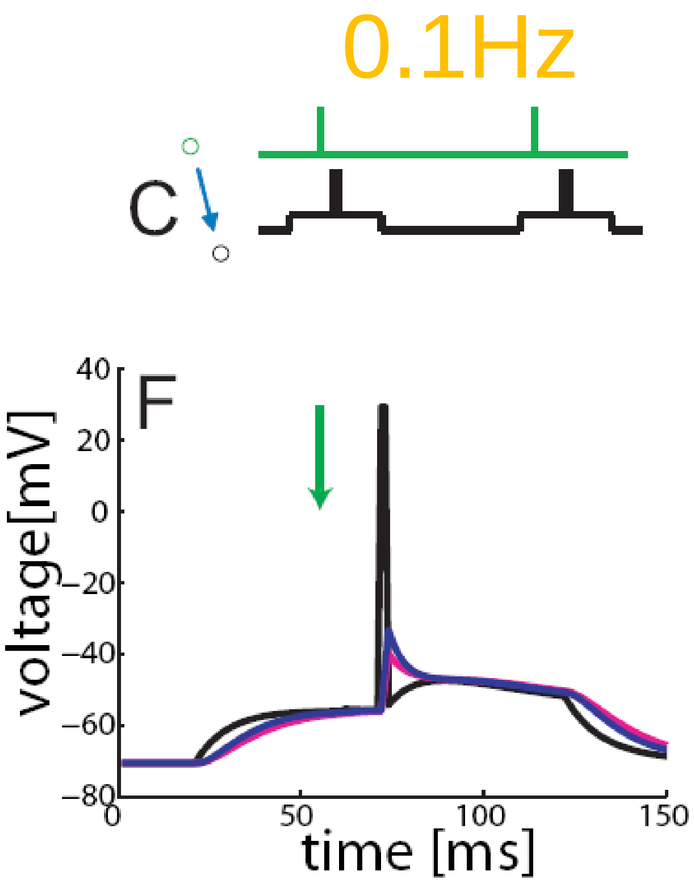

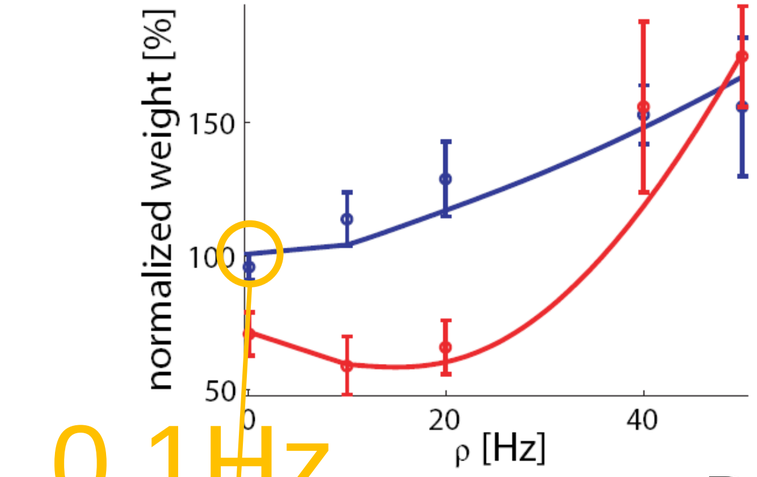

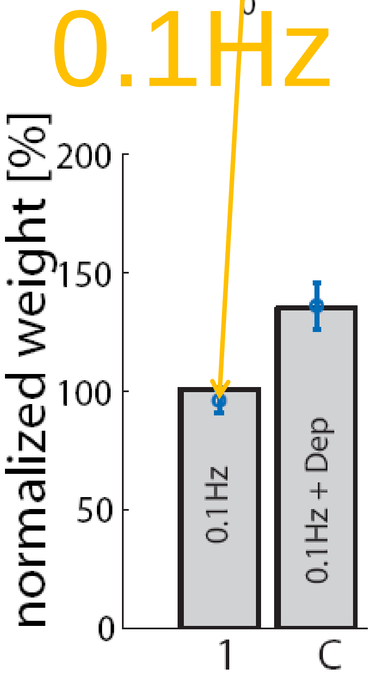

Data:

Sjoestroem et al. 2001

Model:

Clopath et.al. Nat.NS, 2010

Previous slide.

Finally, the Clopath model can also explain the interaction of spikes and voltage. One experiment of Sjostrom et al (2001) used a pre-post (10ms) sequence at 0.1 Hz repetition frequency that normally gives no LTP.

However, when the postsynaptic spikes are embedded in a weak subthreshold current that puts the membrane potential of the postsynaptic neuron above rest (but below the firing threshold), then the same pre-post (10ms) sequence at 0.1 Hz repetition frequency gives LTP.

Bottom right: Blue dots with error bars are data; gray histogram bars are results of the Clopath model.

Interaction of voltage-spike-LTP: consequences

- Sjostrom et al (2001): local voltage interaction inside postsynaptic neuron

- Aceituno et al. (2024): voltage interaction from apical to basal dendrite.

2024: **Learning rule**: Feedforward synapse on basal dendrite depend on lateral/feedback input (on apical dendrite)

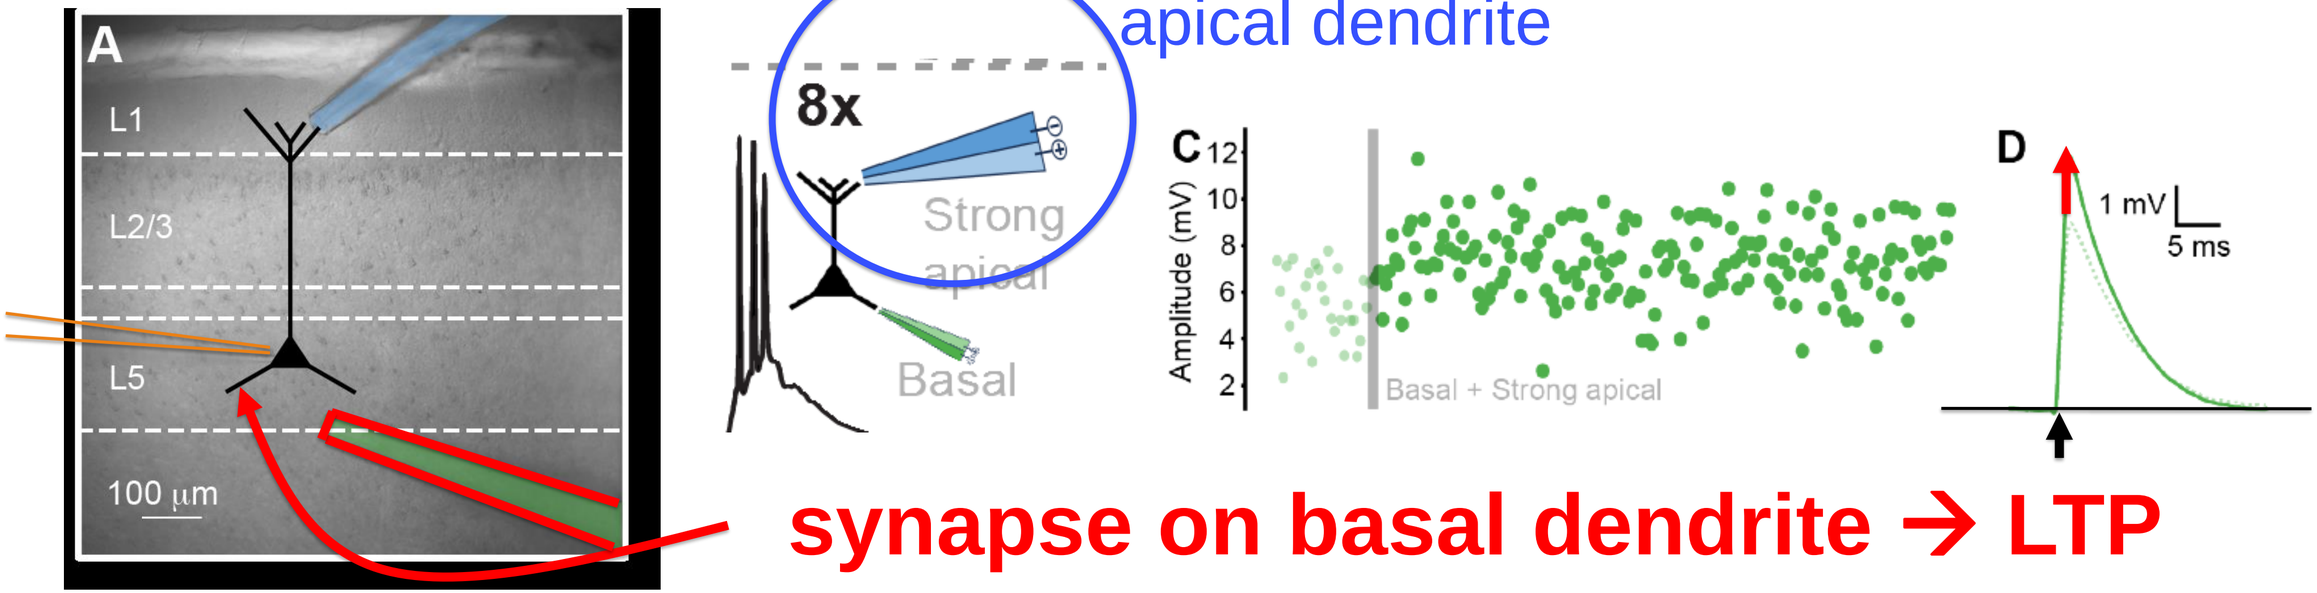

synapse on basal dendrite → LTP

Aceituno, …, Grewe, bioRxiv (2024)
https://doi.org/10.1101/2024.04.10.588837;

Recent experiments in L5: Grewe group (2024)
Experiments in L2/3: Williams and Holtmaat (2019)

Consistent with voltage-dependent plasticity (*Sjostrom et al 2001*)
and Clopath model (*Clopath et al. 2010)*

Such a rule useful to learn 'good' representation!

Experiments in Mouse Frontal Cortex, L5 cells, slice, from the Grewe lab (INI, Zurich). Two electrodes are used for extracellular stimulation at the basal dendrite (red-green) and apical dendrite (blue). Voltage is recorded with the brown electrode (A).
Initially, EPSPs are evoked by small-amplitude pulse stimuli (strength s1) with the red-green electrode yielding an EPSP of a few mV. Then the stimulation amplitude is increased (strength s2) so that the firing threshold is reached, and the postsynaptic neuron fires an isolated spike. After 8 repetitions (at 0.1Hz) no change in the EPSP amplitude is found. Thereafter the stimulation of basal synapses (with strength s2) is paired with stimulation of the apical dendrite, causing a short burst of spikes and a prolonged voltage response. After 8 repetitions (at 0.1Hz) the EPSP amplitude in response to stimulus s1 is increased (C and D).

These findings are consistent with experiments of J. Sjostrom (2001) and the voltage-dependent plasticity model of C. Clopath (2010): synaptic changes require either multiple postsynaptic spikes are a prolonged depolarization of the postsynaptic neuron, or a combination of both.

The end