Computational Neuroscience: Neuronal Dynamics of Cognition

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Synaptic Plasticity and Learning

Wulfram Gerstner

EPFL, Lausanne, Switzerland

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

Cambridge Univ. Press

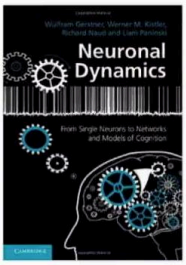

1. Synaptic plasticity
   motivation and aims
2. Classification of plasticity
   - short-term vs. long-term
   - unsupervised vs. reward modulated
3. Model of short-term plasticity
4. Models of long-term plasticity
   - Hebbian learning rules
   - Bienenstock-Cooper-Munro rule
5. Spike-Timing Models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

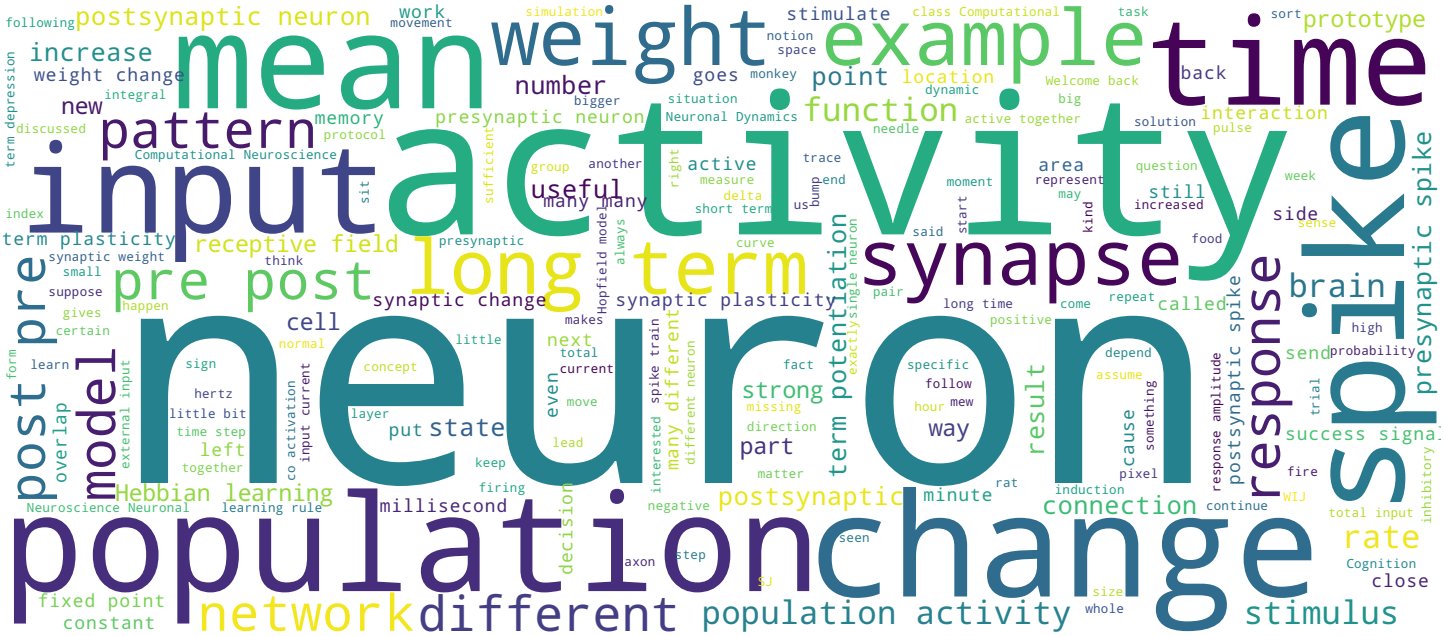

Search MOOC

Video

EPFL

2. Classification of synaptic changes: Short-term plasticity

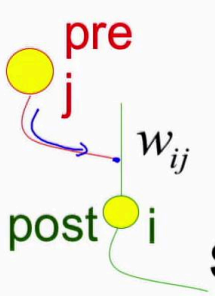

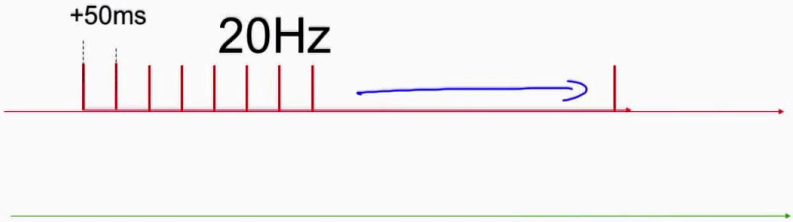

Short-term plasticity/fast synaptic dynamics

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

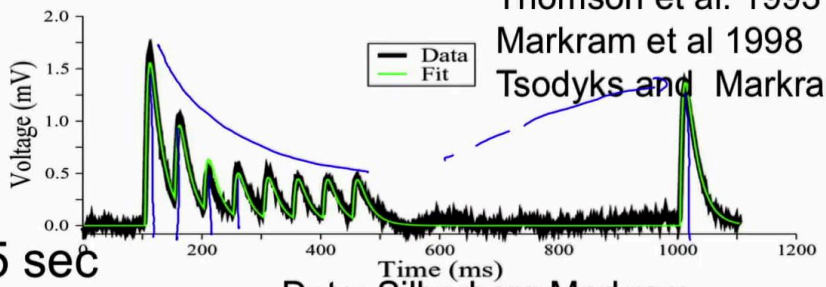

Changes
- induced over 0.5 sec
- recover over 1 sec

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

Welcome back to the class, Computational Neuroscience, Neuronal Dynamics of Cognition. We talked about synaptic plasticity and learning and I would like to discuss now different forms of plasticity that happen in the brain, so one way to cause synaptic plasticity is to send out a spike along the axon from a presynaptic neuron J to a receiving neuron or postsynaptic neuron. The neuron that sits after the synapse, the postsynaptic neuron called I and now a first spike will evoke a response and it has a certain amplitude. And the second spike 50 milliseconds later will again cause a response and it has a reduced amplitude so you see that the response amplitude which for us is a measure of the synaptic weight. This response amplitude goes down from one spike to the next, now if I weight something like 500 milliseconds and then give another spike then the response amplitude is basically back to normal. This phenomenon is called short-term plasticity, they are fast synaptic changes the amplitude goes down rapidly but it also recovers rapidly back to normal. The specific case here is called synaptic depression but there's also synaptic facilitation so that over a few spikes the response would get stronger, and again the recovery is rapid.

Notes

Summary

0m 16s

2. Classification of synaptic changes: Long-term plasticity

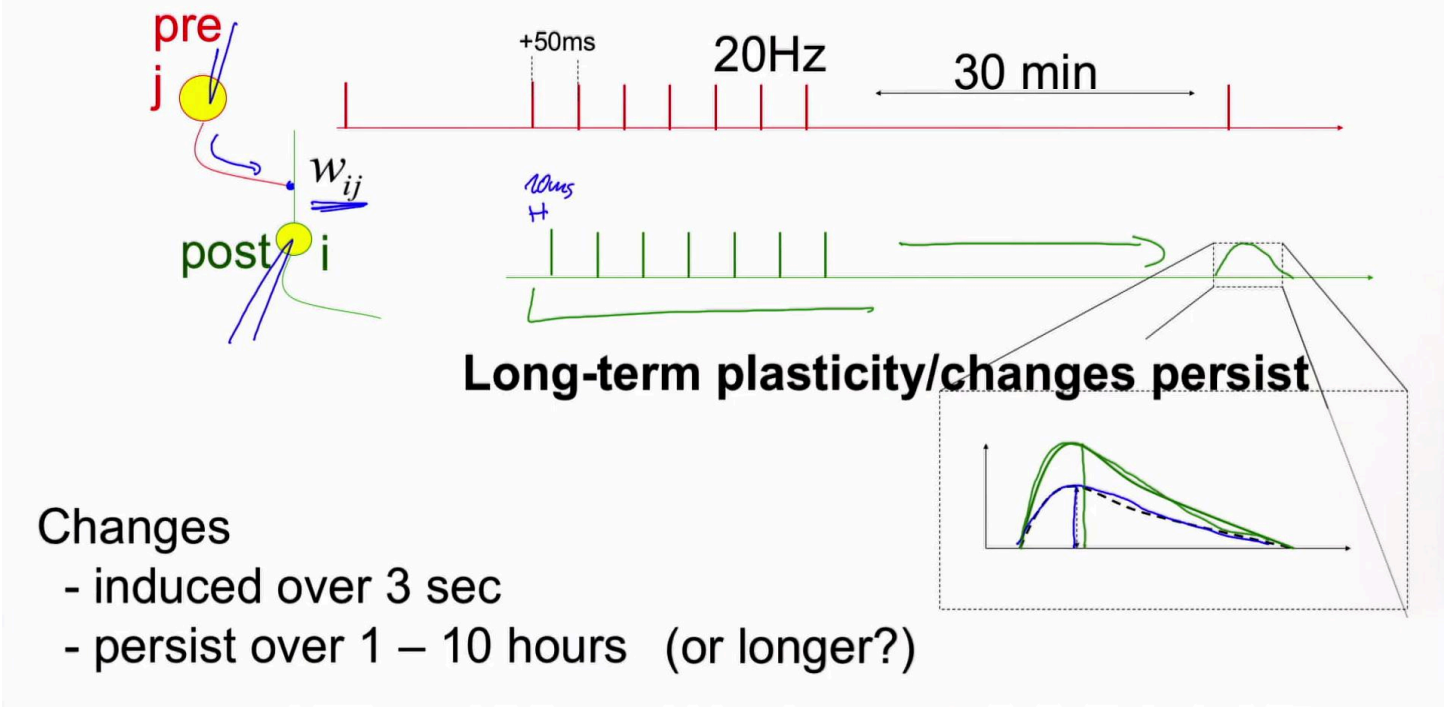

Now here's a different scenario of stimulating, previously I just stimulated the presynaptic neuron, here will stimulate two neurons so the way the experiment proceeds very schematically. I sent over a spike over the axon it goes through the synapse to the postsynaptic neuron and causes a response. And this response has a certain amplitude which is my measure of the synaptic weight W IJ from J to I, and then I stimulate both neurons.I stimulate this neuron here I gives an input pulse, and then I stimulate this neuron here. Ten milliseconds later I stimulate the next neuron and then they repeat the whole thing again, pre then post, pre-post, pre-post, pre-post. Maybe 60 or 80 repetitions and the synopsis will change, now I will measure this change for example 10 minutes afterward or 30 minutes afterwards. So I give a pulse here, and it will again cause a response but now this response is much bigger, that's this curve here. This response is much bigger, the amplitude has increased, so long-term plasticity means that the changes that are induced by this protocol pre-post, pre-post. A protocol there two neurons active together, these changes will last for a long time, for example, 30 minutes or 60 minutes or two hours or even 24 hours or longer.

Notes

Summary

1m 32s

2. Classification of synaptic changes: Long-term plasticity

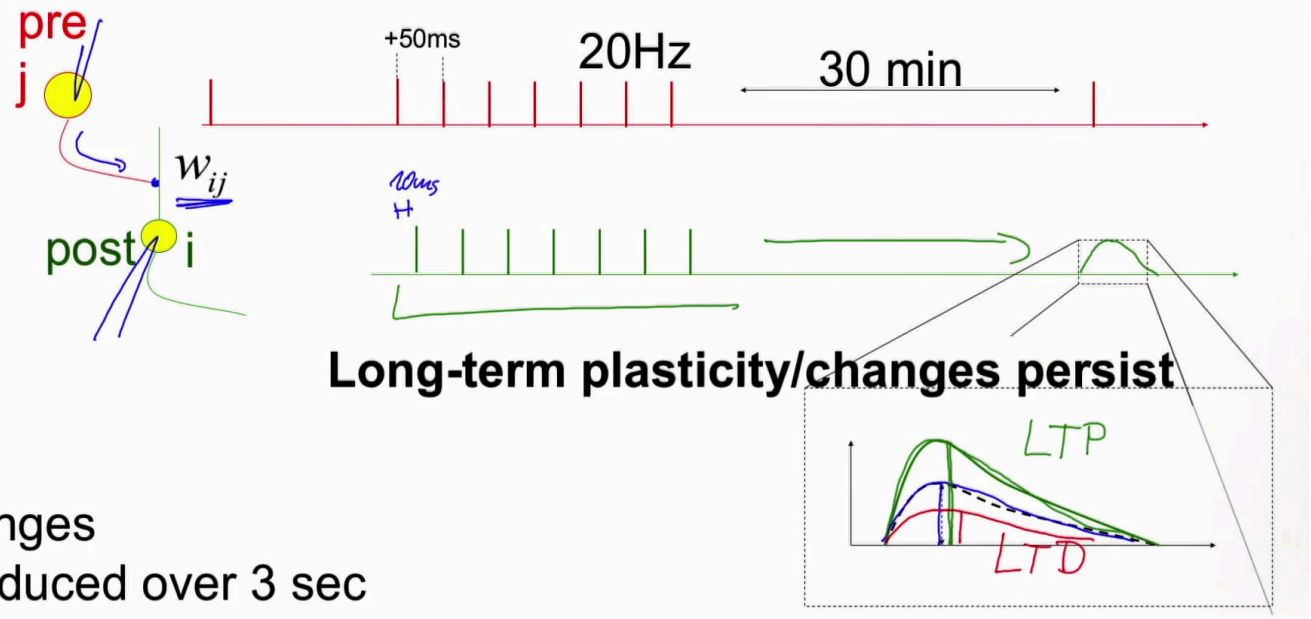

Changes
- induced over 3 sec
- persist over 1 – 10 hours (or longer?)

Maybe 10 years, so long term plasticity comes in two flavors one is an increase, it's called long term potentiation or LTP. And with a different manipulation here, it could also be a decrease of the response with reduced amplitude which would be called long term depression or LTD. The induction of long term plasticity can be quite rapid just a few seconds of stimulation, stimulation of both neurons together but importantly these changes persist for a long time.

Notes

Summary

3m 05s

2. Classification of synaptic changes

Short-Term vs/ Long-Term

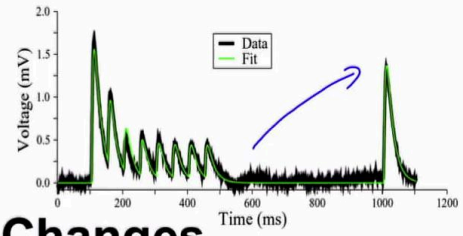

Changes

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

Protocol

- presynaptic spikes

Model

- well established

(Tsodyks, Pawelzik, Markram Abbott-Davan)

LTP/LTD/Hebb

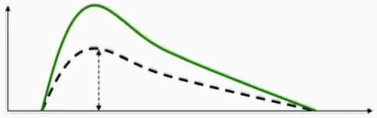

Changes

- induced over 0.5-10sec
- remains over hours

Protocol

-presynaptic spikes + …

Model

- we will see

So we arrive at the first classification of synaptic changes, there's short-term plasticity, changes are induced rapidly but they recover rapidly. Here changes are also induced rapidly but they remain over hours, to induce these changes it's sufficient to just work with presynaptic stimulation. On this side you need to combine the presynaptic stimulation with postsynaptic activity, postsynaptic deprivation.

Notes

Summary

3m 48s

2. Classification of synaptic changes

Induction of changes

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

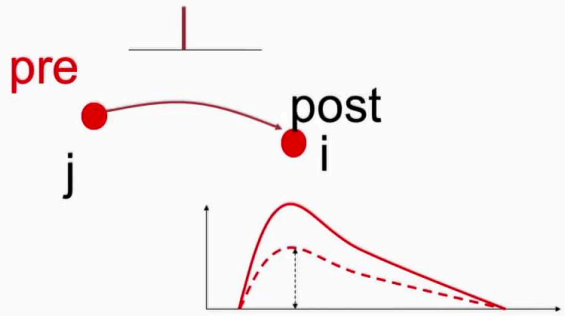

Persistence of changes

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

And we'll see how this works, so the main difference between short-term plasticity is the persistence of the changes. The induction of the changes can be fast in both cases if the stimulus is appropriate, now there are other changes related to homeostasis. This slow control mechanism and in this case the induction of changes is also typically slow.

Notes

Summary

4m 21s

2. Review: Hebb rule

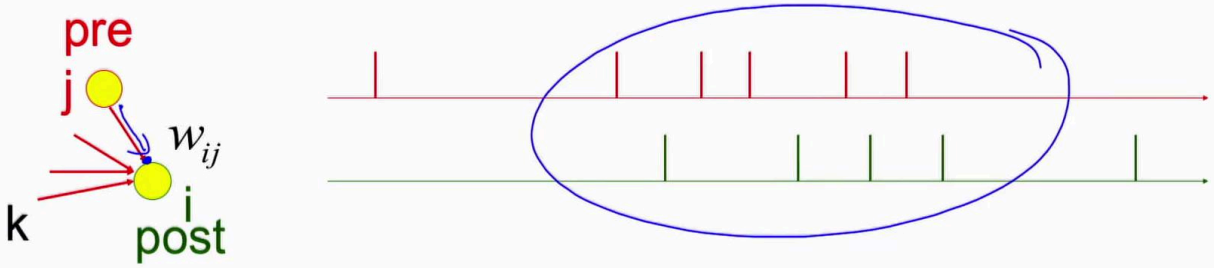

When an axon of cell **j** repeatedly or persistently takes part in firing cell **i**, then j's efficiency as one of the cells firing i is increased

Hebb, 1949

- local rule
- simultaneously active (correlations)

Long term potentiation is closely linked to the Hebb Rule; Hebb was a theoretical Psychologist who wrote up his thoughts about synaptic changes. And what he said is that if there is an axon coming from a presynaptic cell J, and if this axon carries spikes so that the presynaptic neuron takes part in firing the post in a big one. Then some change happens so that the efficiency of J as one of the cells firing I is increased, now if I asked what this means? It's a local rule for this synapse here, what matters is the activity of neurons J and the activity of neuron I but not the activity of other neurons K. And somehow this notion of takes part in firing the postsynaptic cell means that the two neurons have to be active together, very schematically a Hebbian learning experiment runs as this.

Notes

Summary

4m 48s

2. Synaptic plasticity: Long-Term Potentiation (LTP)

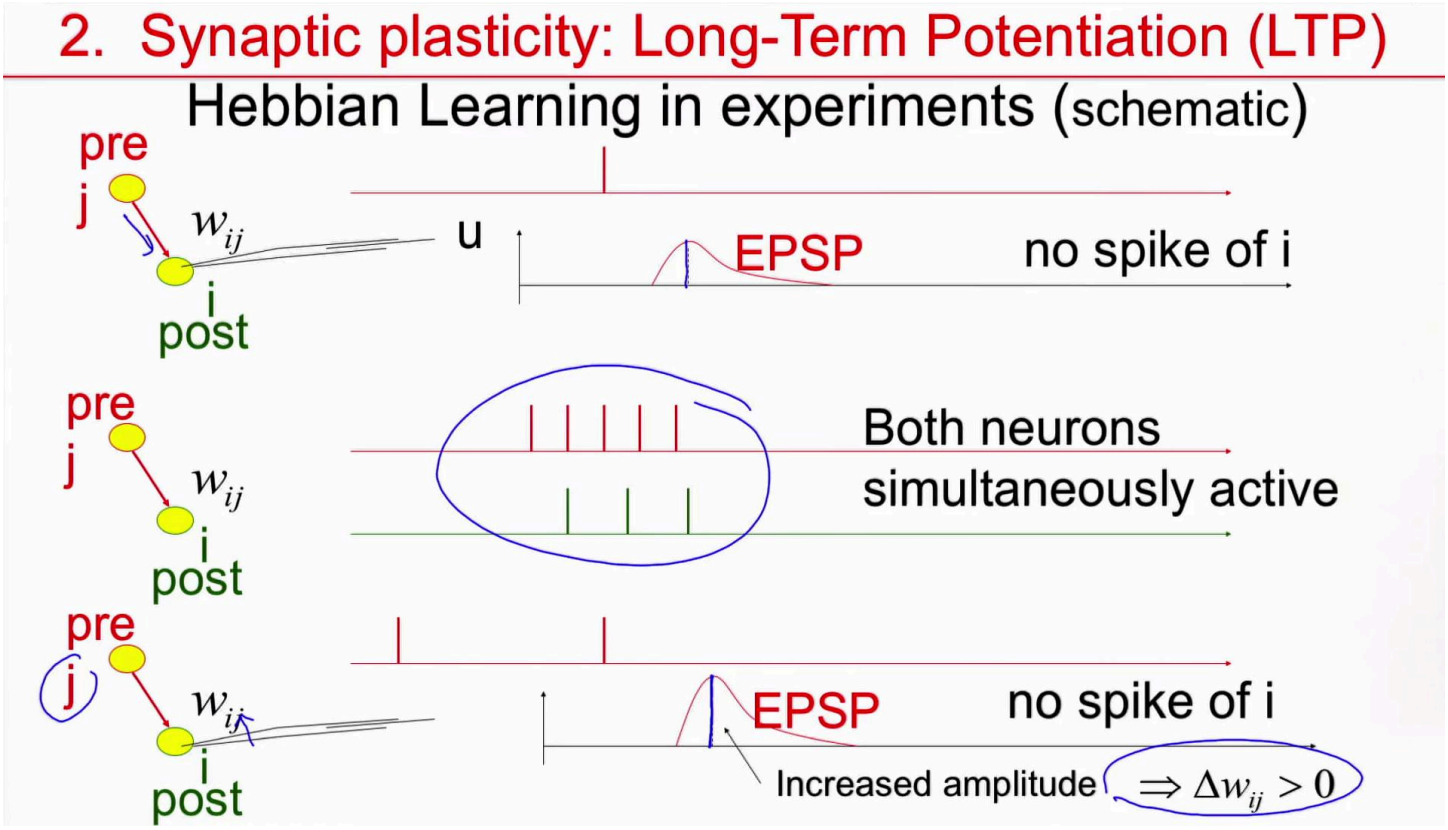

You first send out a test spike to see how strong the weight currently is, then you do the Hebbian protocol where both neurons are active together. And then you send out another test pulse and you may see that the weight has increased, so this increase implies in our notation a change of the synaptic weight the weight from neuron J. This is this index here to neuron I, WIJ.

Notes

Summary

5m 55s

2. Spike-timing dependent plasticity (STDP)

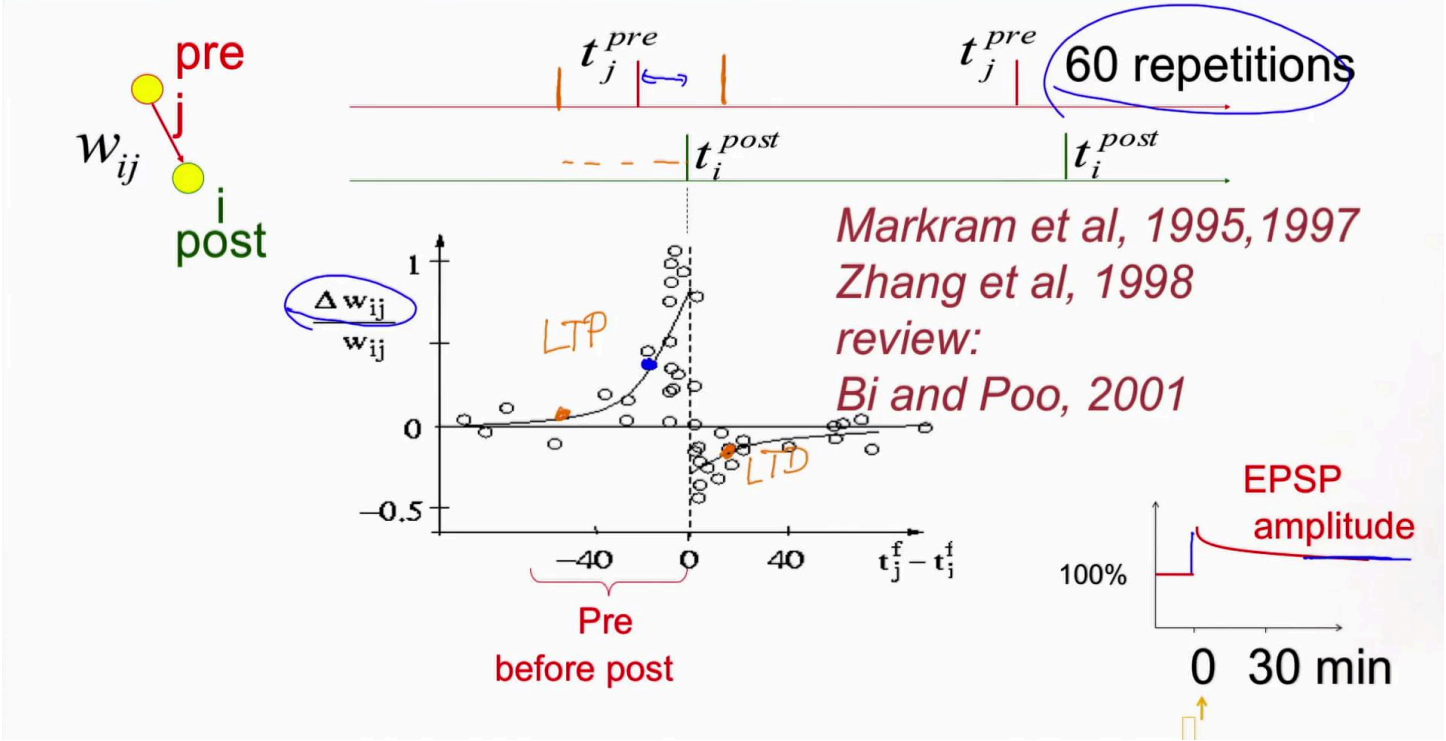

And here's the result of a specific protocol, that's the protocol I mentioned before. I have the two neurons firing together but each time the, neuron the presynaptic neuron fires a little bit before the postsynaptic one. And then we have 60 repetitions, pre-post, pre-post, pre-post, pre-post. Now, if you look at the EPSP amplitude in a situation where you do this pre-post with the time difference of few milliseconds here for example about 20 milliseconds. And you find a change of the weight, of this weight that's the weight change and you can follow this weight change. That's the weight change you'll see that it persists for a long time, thirty minutes or one hour, and so this final value is the value that's plotted here. Now if the presynaptic spike comes earlier with respect to the postsynaptic spike then I would have a point maybe here. If the presynaptic spike comes a few milliseconds after the postsynaptic one and again I do several repetitions, so this will be post-pre, post-pre, post-pre. Then I get a result here, so on this side, we have long term potentiation and on this side, we have long term depression or LTD.

Notes

Summary

6m 26s

2. Spike-timing dependent plasticity (STDP)

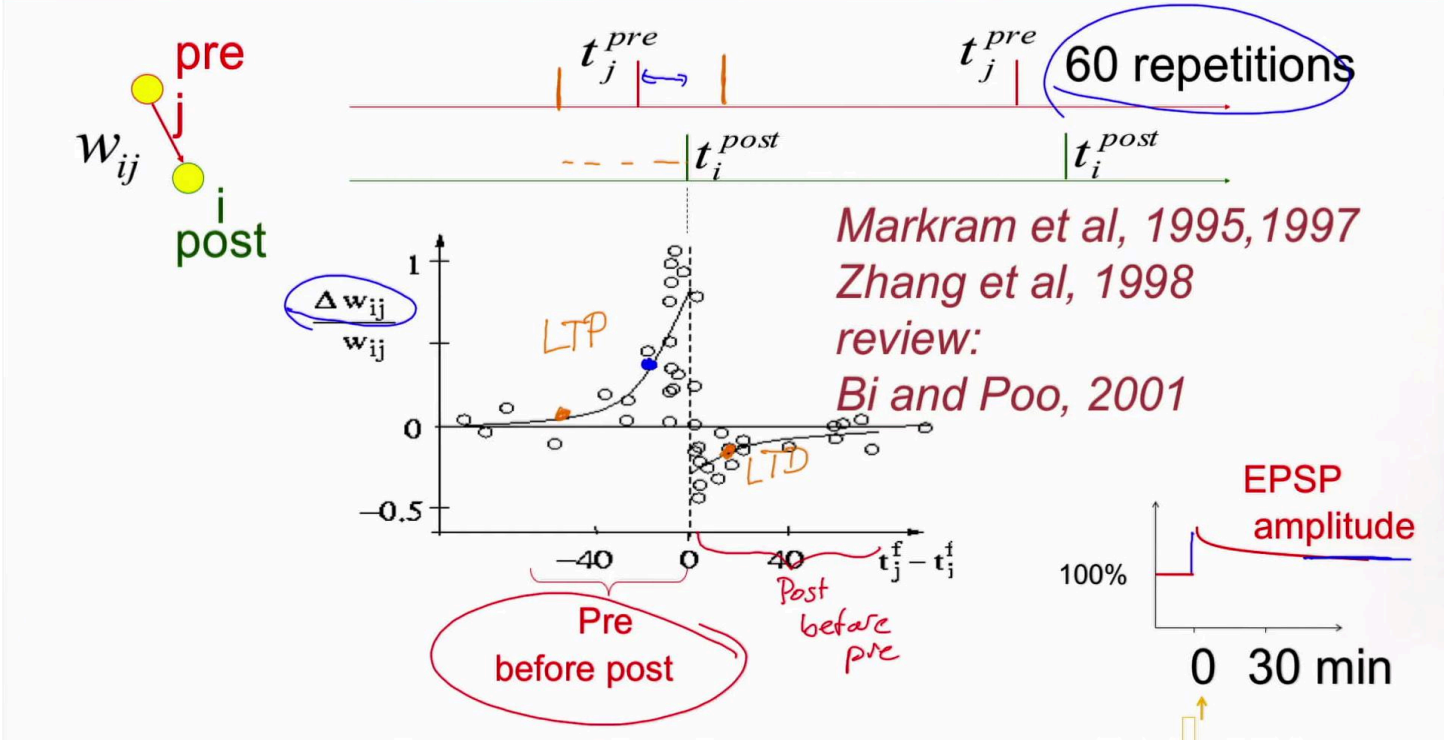

So long term potentiation occurs for the situation that pre before post repeated several times, long term depression is post before pre. Because the result depends on the precise relative timing of the presynaptic spike with respect to the postsynaptic one. This is called spike-timing-dependent plasticity, this is just one manifestation of long term potentiation and long-term depression, there are other means.

Notes

Summary

7m 43s

2. Classical paradigm of LTP induction – pairing

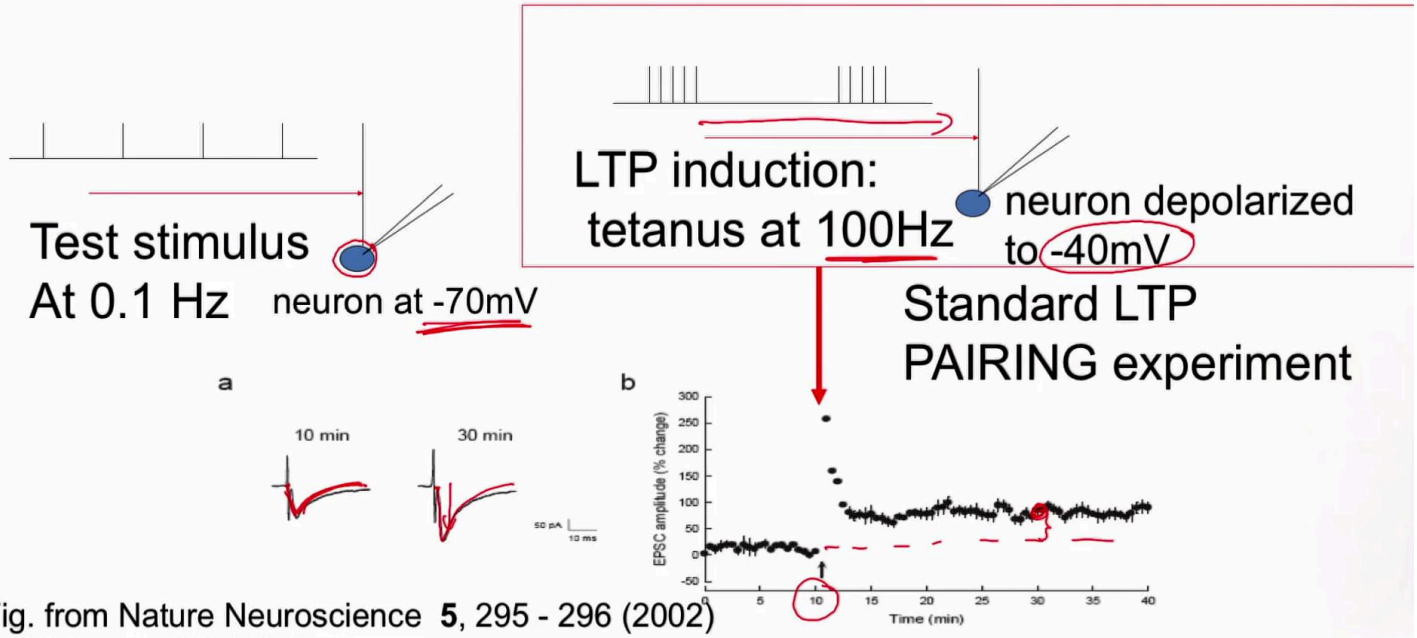

Fig. from Nature Neuroscience **5**, 295 - 296 (2002)
D. S.F. Ling, … & Todd C. Sacktor
See also: *Bliss and Lomo (1973), Artola, Brocher, Singer (1990), Bliss and Collingridge (1993)*

There are other means to induce synaptic plasticity so in this case, the postsynaptic neuron does not fire at all, the neuron is first held at its resting potential close to minus-70 millivolt. And we give test pulses and every 10 seconds of pulse, and each pulse causes a postsynaptic current which is brought down here but it's an excitatory current. Then we give a protocol and during the protocol, the neuron is kept at minus 40 mil, so it's not spiking but this is the value that's very close to the spiking threshold or even above the spiking threshold. So the neuron is sort of active but without firing it's kept at a high voltage, and violets kept at a high voltage the presynaptic synapses are stimulated with a high-frequency burst. A couple of bursts at 100 Hertz so it's really boom-boom, it's a strong stimulus such a stimulus is sometimes called a Tetanus or Titanic Stimulation. And then we can follow, so this stimulus was given at the time, 10 minutes after 10 minutes of this testing and then we can follow. And we can, for example, look at 30 minutes and at 30 minutes the EPSC is much bigger. And so this increase in the size of the EPSP, that's the signature of long-term potentiation, an increase that persists for many, many minutes.

Notes

Summary

8m 18s

2. Classification of synaptic changes

Induction of changes

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

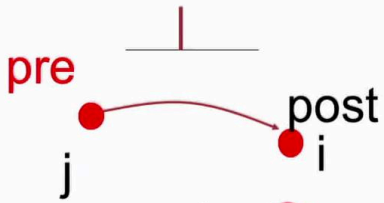

Persistence of changes

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

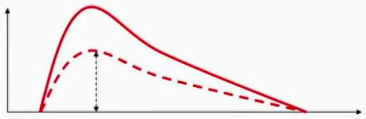

Functionality

- 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 coding

So long-term potentiation can be induced by many different protocols and it means that the changes persist for a long time, now the question which comes up is, what's the function of this? Is this really a form of synaptic plasticity that is useful for learning a new behavior? Is it useful for development firing receptive fields? Is it useful for activity control for homeostasis? Is it useful for coding for energy conservation?

Notes

Summary

9m 59s

2. Classification of synaptic changes: unsupervised learning

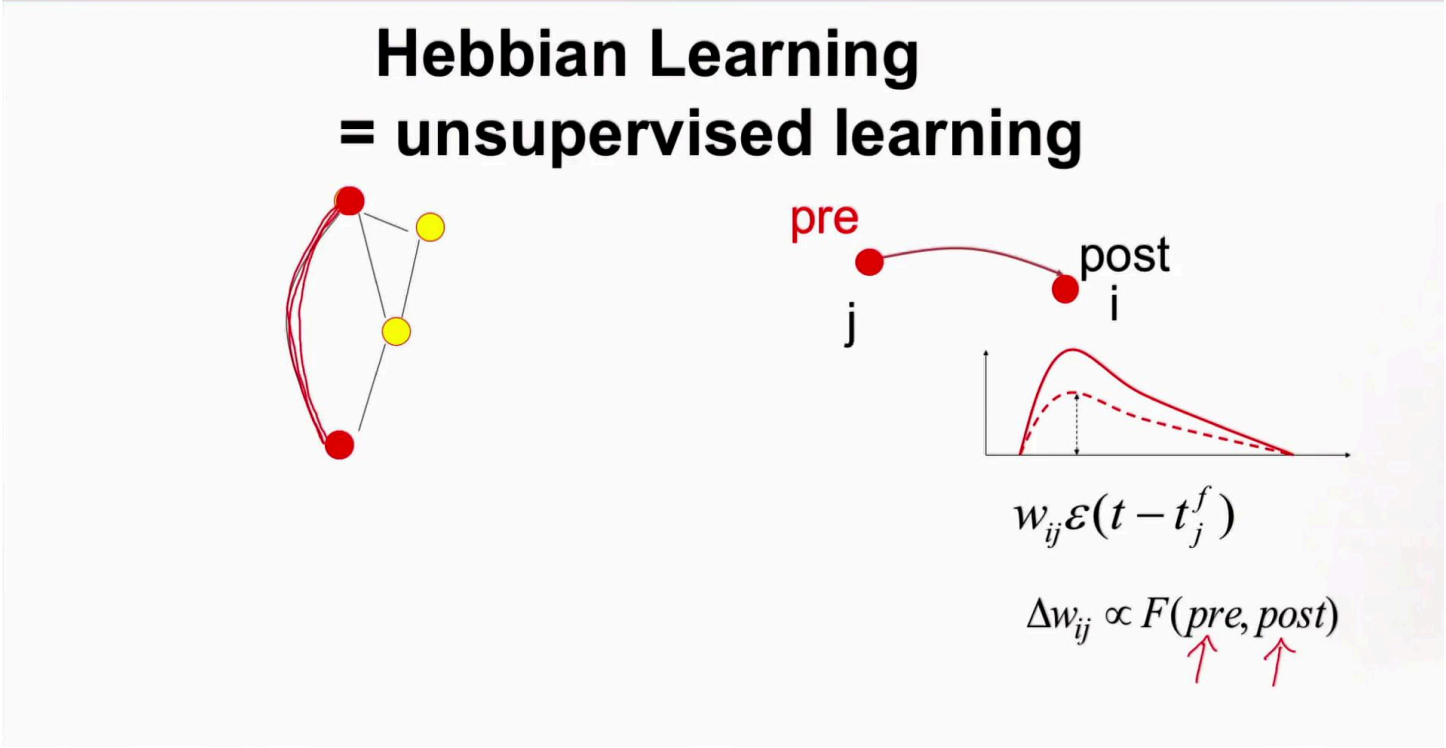

So Hebbian learning as I described it is essentially unsupervised learning, it's a local learning rule, If I have two neurons that are active together then this connection will be strengthened. There is no notion whether it's good or bad to strengthen this connection. I just need the activity of a presynaptic neuron some signature of an activated state of the postsynaptic neuron and this leads to a change of the synapse.

Notes

Summary

10m 33s

2.Limits of unsupervised learning

Is Hebbian Learning sufficient?
No!

Image: Gerstner et al. NEURONAL DYNAMICS,

Eligibility trace:
Synapse keeps memory of pre-post Hebbian events

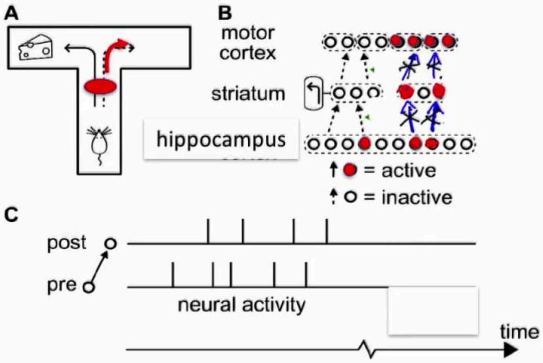

Dopamine:
Reward/success

Schultz et al. 1997; Waelti et al., 2001;

→ Reinforcement learning: success = reward – (expected reward)

TD-learning, SARSA, Policy gradient (book: Sutton and Barto, 1997

Now Hebbian learning per se as described is not sufficient for the task we could be interested in, for example, a rat sits in a T maze. The food is to the left but actually in the first trial it makes a right turn, now first trial at this point here in the hippocampus that represents the space. Remember the London taxi drivers, a certain subset of neurons is active maybe these neurons. Now the right term is encoded in some other area of the brain and it means coding for the program turn right means there again some active neurons. Now Hebbian learning means that this co-activation would lead to a strengthening of these synapses and the result is that after the first trial it's even more likely if the rat is back here that it would again turn to the right. Which is the wrong action if the rat would like to find the food, so what is missing? What's missing is this notion of success, this notion of reward. So the way we can think about this is that the co-activation of pre and postsynaptic neurons, neurons here only marks the synapse as eligible for a change. But if no success signal comes then these synapses will decay back to their normal strength afterward, Hebbian co-activation leaves an eligibility trace but per se it does not yet change the weights.

Notes

Summary

11m 04s

2.Limits of unsupervised learning

Is Hebbian Learning sufficient?
No!

Image: Gerstner et al. NEURONAL DYNAMICS,

Eligibility trace:
Synapse keeps memory of pre-post Hebbian events

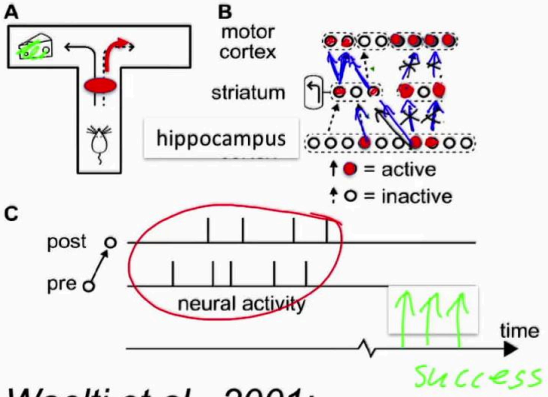

Dopamine: Reward/success

Schultz et al. 1997; Waelti et al., 2001;

→ Reinforcement learning: success = reward – (expected reward)

TD-learning, SARSA, Policy gradient (book: Sutton and Barto, 1997

However if in some other trial at the same stage here, the rat activates the other program, the one that leads to a left turn. Then I have again co-activation and in this case, it would find the food which means some time later there's a success signal. And in this case the synaptic changes will actually be strong, it will be implemented as strong changes. So there's a whole field of reinforcement learning also called reward-based learning which comes in different flavors, policy gradient,[13:29 inaudible] TD learning. And this whole field uses learning rules that basically work with the principles I just discussed, the synapse keeps the memory of the joint activity of pre and posting of the neurons. But the actual change needs a success signal and this success signal would be reward minus expected reward transmitted for example by the new model, Dopamine.

Notes

Summary

12m 48s

2. Classification of synaptic changes: Reinforcement Learning

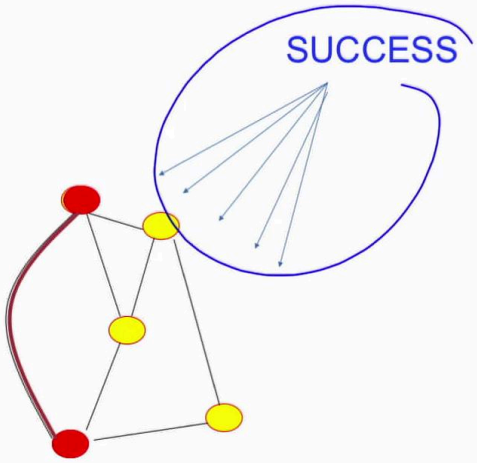

Reinforcement Learning
= reward + Hebb

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

local global

broadly transmitted signal:
neuromodulator

Now to put this back in our framework as before I've two neurons if they are activated together it encodes a Hebbian signal, this signal is stored in the synapse for a very short time but we need this additional success signal to actually make the synapse change to actually finalize the synaptic change and learn.

Notes

Summary

14m 00s

2. Classification of synaptic changes

unsupervised vs reinforcement

LTP/LTD/Hebb
Theoretical concept

- passive changes
- exploit statistical correlations

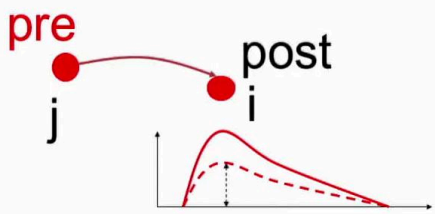

Functionality

-useful for development
(wiring for receptive field)

Reinforcement Learning
Theoretical concept

- conditioned changes
- maximise reward

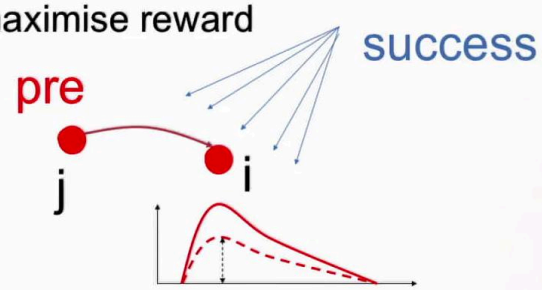

Functionality

- useful for learning
a new behavior

So, the following I will talk mainly about this class of unsupervised learning, Hebbian learning. This is not useful through learning a specific task, a specific behavior but it is useful as we will see for development, for buying receptive field, for allocating space in cortex.

Notes

Summary

14m 29s

2. Three-factor rule of Hebbian Learning

= Hebb-rule gated by a neuromodulator

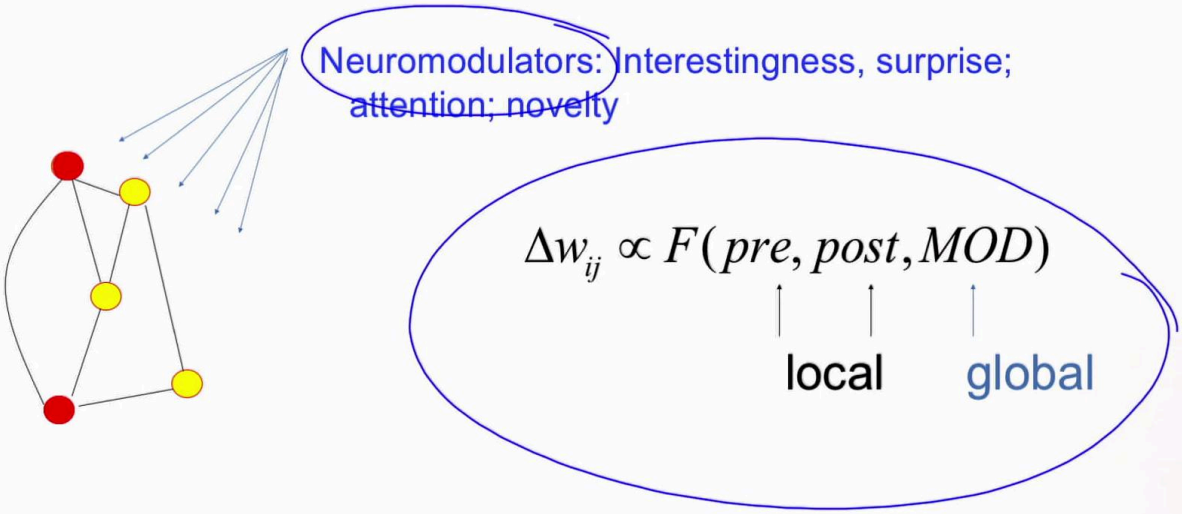

For the general case, we want to learn a specific task, we need to have neuromodulators that encode the success in the sense of reward man is expected to reward or just interestingness. Wow, this is surprising, this is novel, I would like to learn this.

Notes

Summary

14m 51s

Neuromodulator projections

- 4 or 5 neuromodulators
- near-global action

Dopamine/reward/TD:
Schultz et al., 1997,
Schultz, 2002

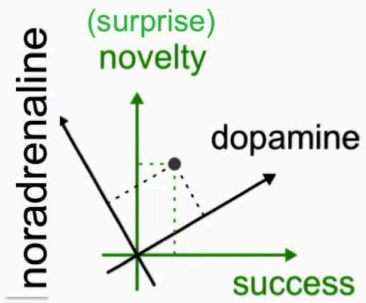

Image:
Fremaux and Gerstner, Frontiers (2016)

Image: Biological Psychology, Sinauer

Dopamine

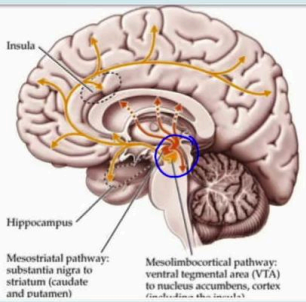

Noradrenaline

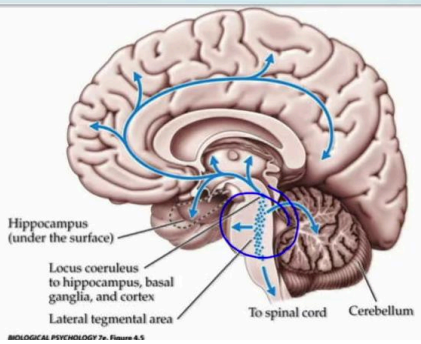

And in our formalism I would just write this down by a global signal, now what this really refers to are neuromodulatory signals. These signals are generated in nuclei that sits somewhere deep in the brain, and then they have axons that ramify and reach nearly every area of the cortex, maybe not primary sensory cortex for Dopamine. But it's close to a globally distributed signal; similar noradrenalin also starts in some deep nuclei and sends out signals via axonal ramifications to many, many different brain areas to many, many different neurons to many, many different synapses. The exact role of these neuromodulators is still unclear and under research, Dopamine is clearly involved in reward signaling nor eternally might be related to the multi and surprise but note that that does not have to be a one-to-one correspondence. For example, if I spin a space with two axes, one is success, the other novelty. Then it could very well be that Dopamine and Noradrenalin have a slightly shifted configuration so I discussed a lot of biological aspects and before we continue I would like to encourage you to look at the following quiz.

Notes

Summary

15m 18s

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, but not on state of 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

Learning rules

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

Notes

Summary

16m 27s