# Computational Neuroscience: Neuronal Dynamics of Cognition

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

## Attractor Networks and Generalizations of the Hopfield model

Wulfram Gerstner

EPFL, Lausanne, Switzerland

*Reading*:
NEURONAL DYNAMICS
- Ch. 17.2.5 – 17.4

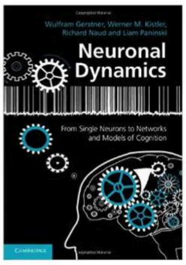


Cambridge Univ. Press

1. Attractor networks
2. Stochastic Hopfield model
3. Energy landscape
4. Towards biology (1)
   - low-activity patterns
5. Towards biology (2)
   - spiking neurons

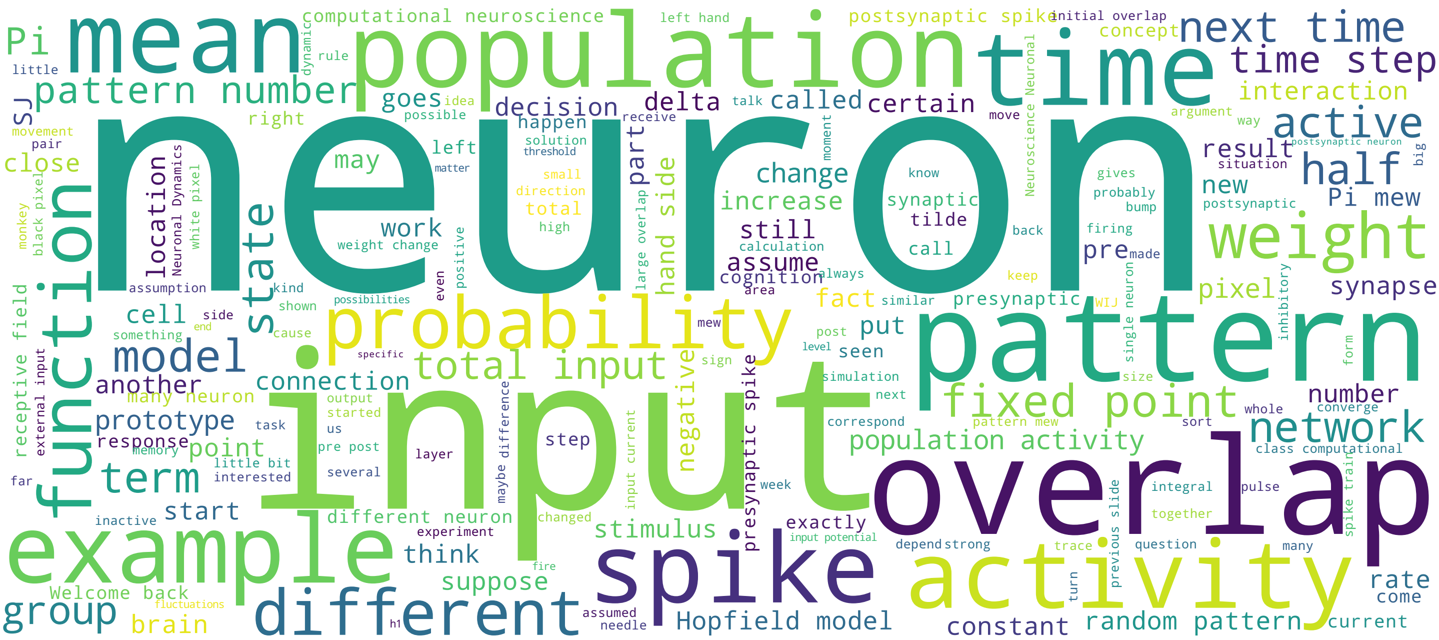


EPFL

Search MOOC

Video

# 2. Stochastic Hopfield model

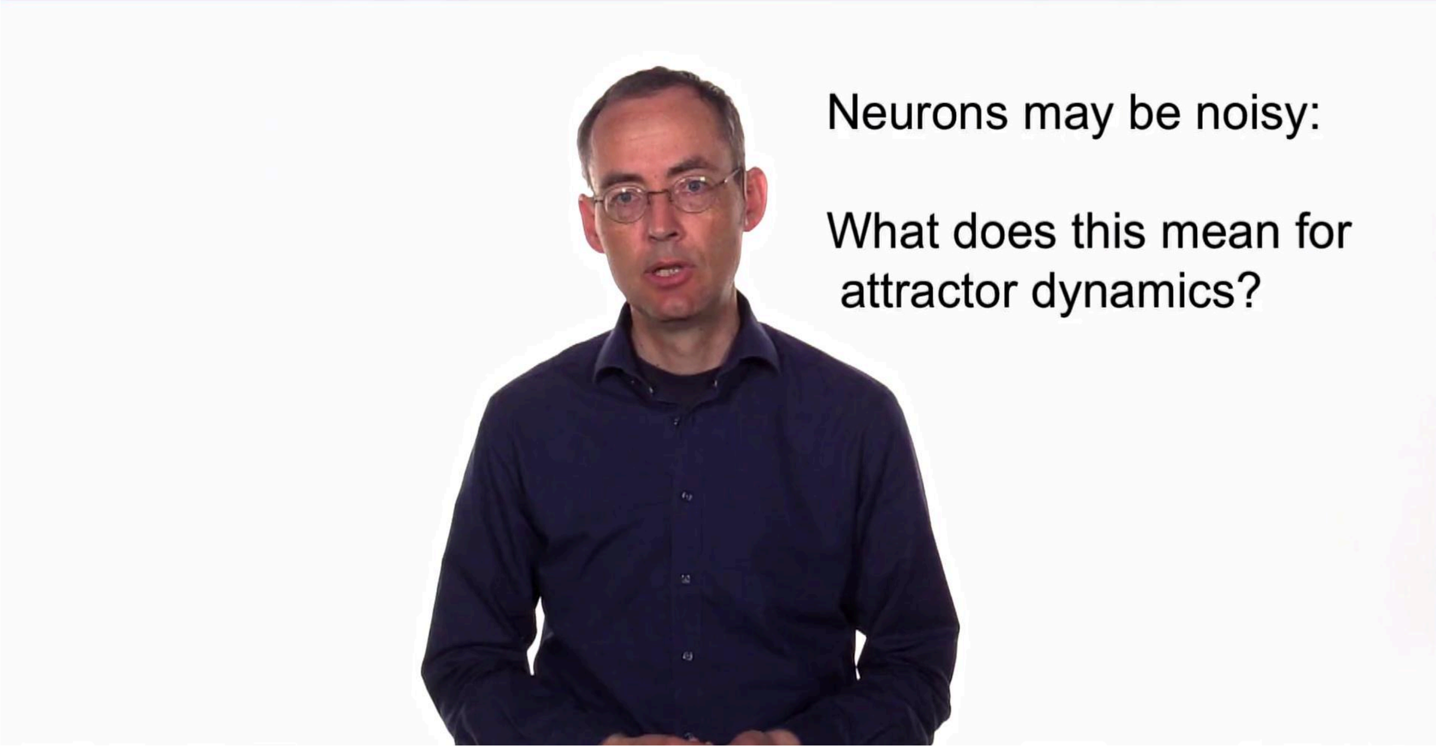


Welcome back to the class computational neuroscience, neural dynamics of cognition. We have looked at the tractor nerks. And so far we have been considering only the deterministic Hopfield model deterministic Hopfield model means that the output of a neuron is just the sign function of its input. But neurons may be noisy. So what would that mean? What would that imply for the attractor dynamics?

Notes

Summary

0m 04s

# 2. Stochastic Hopfield model

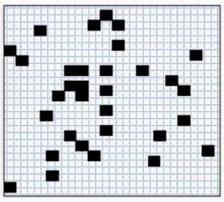

Prototype $\vec{p}^1$

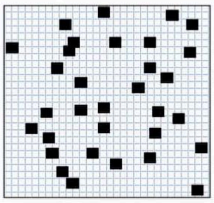

Prototype $\vec{p}^2$

Random patterns

Interactions (1) $w_{ij} = \frac{1}{N}\sum_{\mu=1}^{M} p_i^\mu p_j^\mu$

Dynamics (2)

$$\Pr\{S_i(t+1) = +1 \mid h_i\} = g[h_i] = g\Big[\underbrace{\sum_{j=1}^{N} w_{ij} \overset{\pm 1}{S_j}(t)}_{h_i}\Big]$$

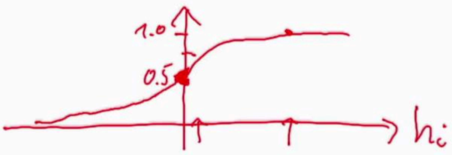


So as before, we have our interactions, you have stored random patterns prototypes, P 1 P 2 P three, maybe a total of M prototypes. And these prototypes are represented in the interaction matrix Wij. And as before, I sum of all the inputs, I sum of all neurons j equal 1 over and I take these prototypes in the form of the interaction matrix I take the sum with a momentary state as j Sj. There Sj is plus-minus 1 as before, and this gives the total input. But now the difference is, it's no longer just a deterministic signed function. Now we have a probability that the next stage is plus 1, which depends on this total input. And the total input is this. So what could this function look like? Well, for example, we could take a function where I bought here the input, and then say, if the input is zero, I'm at O point five. And if the input is big, I will somehow approach 1, there's no need that this is very symmetric. So it could also be asymmetric like this a little wiggle, this is just a function g of h. And so the probability, you have a strong input that the next output is plus 1 will be very close to 1. If I have a big input, if I a slightly positive input, it may, for example, only open six. So the probability increases as we increase the total input. And it decreases if you go to negative inputs.

Notes

Summary

0m 28s

# 2. Stochastic Hopfield model: firing probability

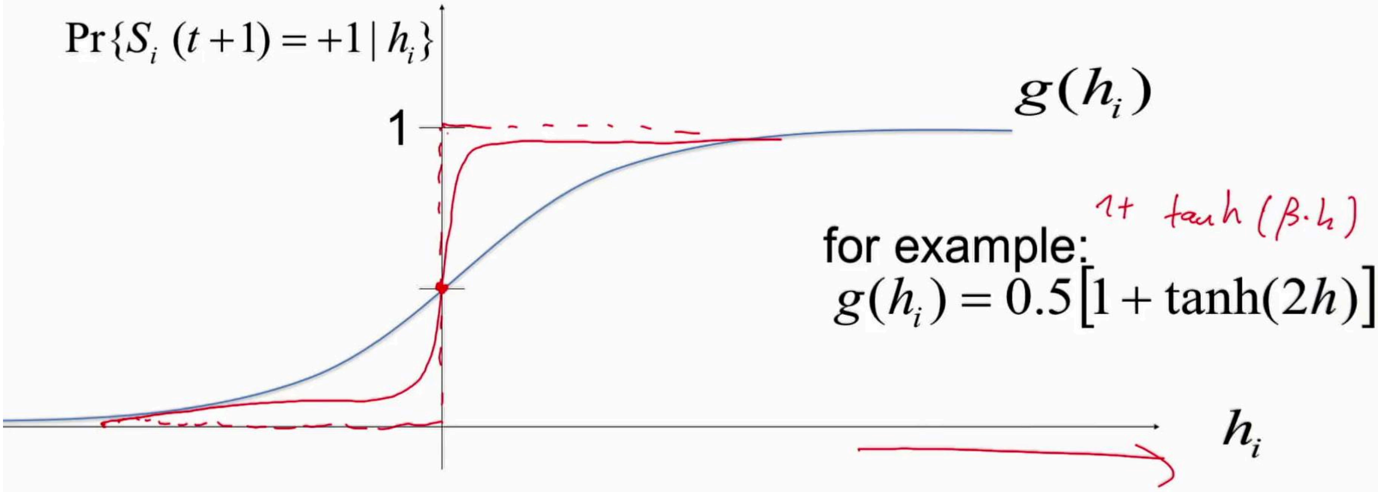


$$\Pr\{S_i(t+1)=+1 \mid h_i\} = g[h_i] = g\left[\underline{\textstyle\sum_j w_{ij} S_j(t)}\right] = g\left[\underbrace{\textstyle\sum_\mu p_i^\mu m^\mu(t)}\right]$$

So this is 1 of the examples. For example, you could take a hyperbolic tangent function with a slope factor to the slope is factor 2 is the slope at this location. If you generalize and say, I hear hyperbolic tangent of beta age, and I take beta equal 10, or 20, then it will become steeper and steeper. And if I take beta to infinity, then it really goes back to a step function. And this is the situation that would correspond to the deterministic model. So the total input is this age, this is this axis here, that's my total input, the total input is the sum of all in neurons. But since the weights are the same as before since the state of the neuron is still plus-minus 1, I can still define my overlaps as before, and I still get the same equation as before. So the only thing that has changed is the output.

Notes

Summary

2m 13s

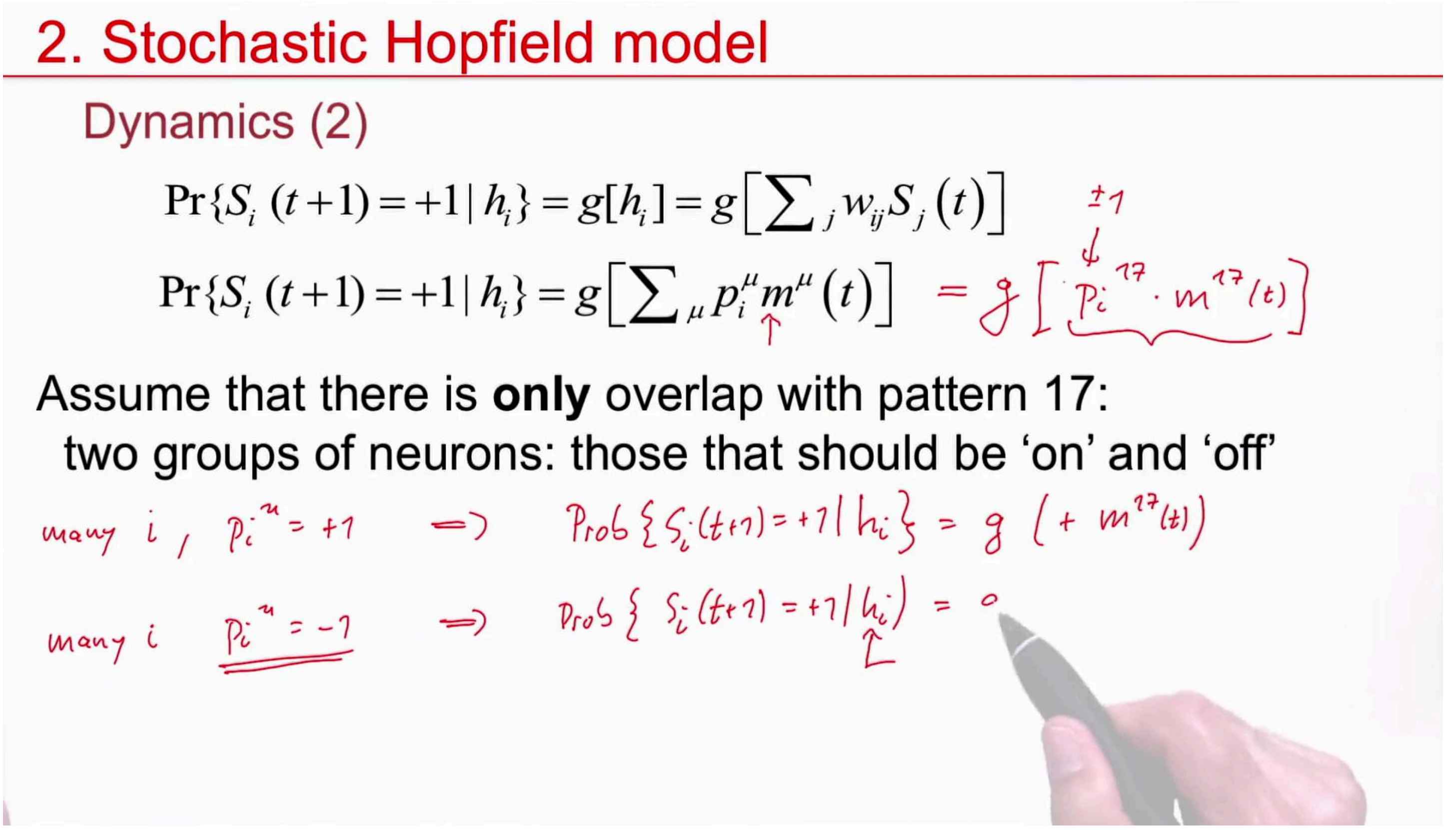


So now if you look at this, so this is the total input sum of all neurons, rewritten as a sum of overlaps. Now assume that there's only overlap with pattern number 17. What does that mean? Well, it means that the overlap this all other patterns is zero. And therefore, in this case, I would have g off, some of them you gust, so I only have pattern number 17. So it's Pi 17 M 17 of T. So now you see that there are only 2 possibilities Pi 17 What's this, this is the value of the pixel i in pattern number 17. So Pi 17 is plus-minus 1, it's black pixels or white pixels. Now think of all these different neurons I've many neurons I which have Pi mew equal plus 1 for this specific pattern, mew equals 17. Now that road Pi mew, but I really mean PI 17. So let's correct this, that say four mews equal 17 for this specific pattern. Now for all these neurons, with Pi 17 equal plus 1, the probability that Si in the next step is plus 1, given this input, is just G off. Well, this is plus 1, so plus M 17 of T. And then if other neurons, I've many neurons, i, their Pi mew is minus 1, for all these other neurons, the probability that the next state is plus 1, given their input, but it's a different input, the input of these neurons, they will have a bike, this condition here, they will have a PI 17 minus 1.

Notes

Summary

3m 12s

## 2. Stochastic Hopfield model

### Dynamics (2)

$$\Pr\{S_i(t+1)=+1\,|\,h_i\}=g[h_i]=g\left[\sum_j w_{ij}S_j(t)\right]$$

$$\Pr\{S_i(t+1)=+1\,|\,h_i\}=g\left[\sum_\mu p_i^\mu m^\mu(t)\right]$$

Assume that there is only overlap with pattern 17:
two groups of neurons: those that should be 'on' and 'off'

for $p_i^{17}=+1$

$$\Pr\{S_i(t+1)=+1\,|\,h_i=h^+\}=g\left[m^{17}(t)\right]$$

$$\Pr\{S_i(t+1)=+1\,|\,h_i=h^-\}=g\left[-m^{17}(t)\right]$$

for $p_i^{17}=-1$

Overlap (definition) $$m^{17}(t+1)=\sum_j p_j^{17}S_j$$

So I have g of minus 1-time m 17 of t. So this is the probability that a neuron and which is supposed to be of minus 1 in pattern number 17 will actually be will actually be active, they'll have an output SI equal plus 1. Now, what's the probability that 1 of these neurons staying the same group, not plus 1 plus or minus 1 is I of T plus 1 equals minus 1 for the same input, sees this same group of neurons? Well, this is 1 minus the probability to be active so its 1 minus g of minus m 17 of t. And the same argument could be made for the other group, neurons that should be active have a plus 1 in this pattern in our pattern number 17 can happen to be inactive, and there probably would be 1 minus g. 1 minus 2 probably to be active is the privilege to be inactive to be at minus 1. So we have these 2 groups of neurons. And whenever I write here Pi mew, I mean its mew equals 17. This is my current pattern, I've assumed that there's only overlap with pattern 17 no overlap with other patterns. So for those 2 groups of neurons, we have the driving force of plus m 17. If they belong to these plus pixels, plus 1 pixel in pattern 17 this is for PI 17 equal plus 1. And this is for PI 17 equal minus 1.

Notes

Summary

5m 25s

## 2. Stochastic Hopfield model

### Dynamics (2)

$$\Pr\{S_i(t+1)=+1\,|\,h_i\}=g[h_i]=g\left[\sum_j w_{ij}S_j(t)\right]$$

$$\Pr\{S_i(t+1)=+1\,|\,h_i\}=g\left[\sum_\mu p_i^\mu m^\mu(t)\right]$$

Assume that there is only overlap with pattern 17:
two groups of neurons: those that should be 'on' and 'off'

for $p_i^{17}=+1$

$$* \begin{cases} \Pr\{S_i(t+1)=+1\,|\,h_i=h^+\}=g\left[m^{17}(t)\right] \\ \Pr\{S_i(t+1)=+1\,|\,h_i=h^-\}=g\left[-m^{17}(t)\right] \end{cases}$$

for $p_i^{17}=-1$

Overlap (definition) $$m^{17}(t+1)=\sum_j p_j^{17}S_j(t+1)$$

(*

Now let's work on this little bit further, I see that on the right-hand side, I have the overlap at time t with pattern number 17. On the left-hand side of the state in the next time step and have 2 different groups of neurons, those that should be on and those that should be off should be off because the target pixel impact on 17 is minus 1. Now the idea is that I would like to take this information, plug it in here, time step t plus 1 and calculate the overlap in the next time.

Notes

Summary

7m 21s

## 2. Stochastic Hopfield model

Overlap (definition) $$m^{17}(t+1) = \frac{1}{N}\sum_{i=1}^{N} p_j^{17} S_j(t+1)$$

**Suppose initial overlap with pattern 17 is 0.4;**
**Find equation for overlap at time *(t+1)*,**
**given overlap at time *(t)*.**
**Assume overlap with other patterns stays zero.**

Hint: Use result from previous slide and consider 4 groups of neurons

- Those that should be ON and are ON
- Those that should be ON and are OFF
- Those that should be OFF and are ON
- Those that should be OFF and are OFF

So let's do this. You're interested in the overlap in the next time step, we assume that the initial overlap with this panel number 17 is non-negligible is, for example, 0.4. And do you want to find the equation for the overlap at the next time step? Now I want that there's only overlap with pedal number 17. And the overlap with other patterns, Stays zero are all the time. So I can really focus on pattern number 17. So now, given what you have seen on the previous slide there, in fact, four groups, they are those neurons that should be active. That means PI 17 equal plus 1, and happens to be active in the next time set. But they also those that should be active and have to be off in the next time step, then they are those that should be off, which means PI 17 equal minus 1, but happened to be active in the next time instead. And then others that should be open actually our so let's put this together.

Notes

Summary

8m 01s

# 2. Stochastic Hopfield model

Overlap $$m^{17}(t+1) = \frac{1}{N}\sum_{i=1}^{N} p_j^{17} S_j(t+1)$$

$\pm 1$

$P(+|+) = g(+m^{17})$

$p_i^{17} = +1$ $\qquad$ $p_i^{17} = -1$

$$= \underbrace{(+1)\cdot(+1)\cdot g(+m^{17}) + (+1)(-1)\cdot(1-g(+m^{17}))} + \underbrace{(-1)(+1)}$$

So in my sum of all different neurons, I have the 2 groups of neurons, I have those which have PI 17 equal plus 1, and then I those that have PI 17 equal minus 1. So for all those that have PI 17 equal plus 1, I will have in this sum a plus 1, and then I will have this state, but the state is SJ. Now SJ is also plus or minus 1. So SJ can be plus 1. And he does it will be plus 1 with a probability of being plus 1, which he calculated on the previous slide for this probability of being plus 1, given that the pixel is plus 1 is just g of plus m 17. Okay, g of plus, m 17. And then we have the second subgroup, it's still pixel plus 1, but now, it could be. So this is the pixel value. But now it could be that the neuron is actually inactive. So it could be that the neuron, despite the fact that it should be on is off for minus 1. And both the probability for this Well, it's 1 minus the probability that is on its 1 minus g plus m 17. So all this was the first group. Now let's work on the second group, put it here, I have a plus additional contributions, it could be that this pixel is negative, that's a minus 1. And then it could be that this neuron is currently active, it has a value of plus 1.

Notes

Summary

9m 05s

## 2. Stochastic Hopfield model

**Overlap** $m^{17}(t+1) = \frac{1}{N}\sum_{i=1}^{N} p_j^{17} S_j(t+1)$ with $S_j = \pm 1$

$P(+|+) = g(+m^{17})$

$p_i^{17} = +1$ $\qquad$ $p_i^{17} = -1$

$$= \frac{1}{N}\frac{N}{2}\left[(+1)\cdot(+1)\cdot g(+m^{17}) + (+1)(-1)\cdot(1-g(+m^{17}))\right] + \frac{1}{N}\frac{N}{2}\left[(-1)(+1)\,g(-m^{17}) + (-1)(-1)\left[1-g(-m^{17})\right]\right]$$

$$= \frac{1}{2}\left[g(m^{17}) - (1-g(m^{17}))\right] + \frac{1}{2}\left[(-1)\cdot g(-m^{17}) + (1-g(-m^{17}))\right]$$

$$= \frac{1}{2}\left[2\,g(m^{17}) - 1\right] + \frac{1}{2}\left[(-2)\cdot g(-m^{17}) + 1\right]$$

$$= g(m^{17})$$

And what's the probability for this; it's g of minus m 17. And then I have you on should be off PJ 17 is negative, I have the situation that the new and actually is off. And it is so with probably 1 minus g of minus m 17. Okay, these are four different cases, plus, plus, plus-minus minus plus-minus minus. Now, how many neurons do we have in each of these four subgroups? Well, we suppose that we have stored random patterns. So suppose that exactly 50% of the pixels are on a plaque plus 1. So I would have 1 half N and 1 half N of all these possibilities, and then the remaining is 1 half N all the other possibilities. But then I have an additional factor 1 over N in front, which will just cancel the N. So let's, let's sum this up, I plus G of M 17. I have minus 1 minus G of M 17 with the factor, 1 half since the end goes in also goes here, plus 1 half of a minus 1, g of minus 17 minus 1 plus 1 minus g of minus 17. Let's collect the terms, I have G of M 17. Let's put the brackets in here. That's what I meant minus 1 minus G of M 17. This gives 1 half 2 g of M 17 minus 1 plus 1 half minus 2 g of minus m 17 minus 1. So I have 1 half times 2 its G of M 17, then this 1 half gives a minus 1 half.

Notes

Summary

11m 05s

# 2. Stochastic Hopfield model

**Overlap** $m^{17}(t+1) = \frac{1}{N}\sum_{i=1}^{N} p_j^{17} S_j(t+1)$

$\pm 1$

$P(+1|+) = g(+m^{17})$

$p_i^{17} = +1$ $\qquad$ $p_i^{17} = -1$

$$= \frac{1}{2}\left[(+1)\cdot(+1)\cdot g(+m^{17}) + (+1)(-1)\cdot(1-g(+m^{17}))\right] + \frac{1}{2}\left[(-1)(+1)\,g(-m^{17}) + (-1)(-1)\left[1-g(-m^{17})\right]\right]$$

$$= \frac{1}{2}\left[g(m^{17}) - (1-g(m^{17}))\right] + \frac{1}{2}\left[(-1)\cdot g(-m^{17}) + (1-g(-m^{17}))\right]$$

$$= \frac{1}{2}\left[2\,g(m^{17}) - 1\right] + \frac{1}{2}\left[(-2)\cdot g(-m^{17}) + 1\right]$$

$$= g(m^{17}) - g(-m^{17})$$

$$m^{17}(t+1) = \tilde{F}\left(m^{17}(t)\right)$$

And I made a mistake here, okay, this is this is a plus 1 that comes in here. So as minus 1 half, plus 1 half, they go 1 half minus 2 gives minus 1, T of minus and 17. So this is the overlap M 17 at time t, And on the left-hand side of the overlap time T plus, 1 and the combination of these 2 G's level just defines as a function F tilde. So I take the overlap at time t, I pass it through this function F tilde; I get the overlap at time t plus 1.

Notes

Summary

13m 27s

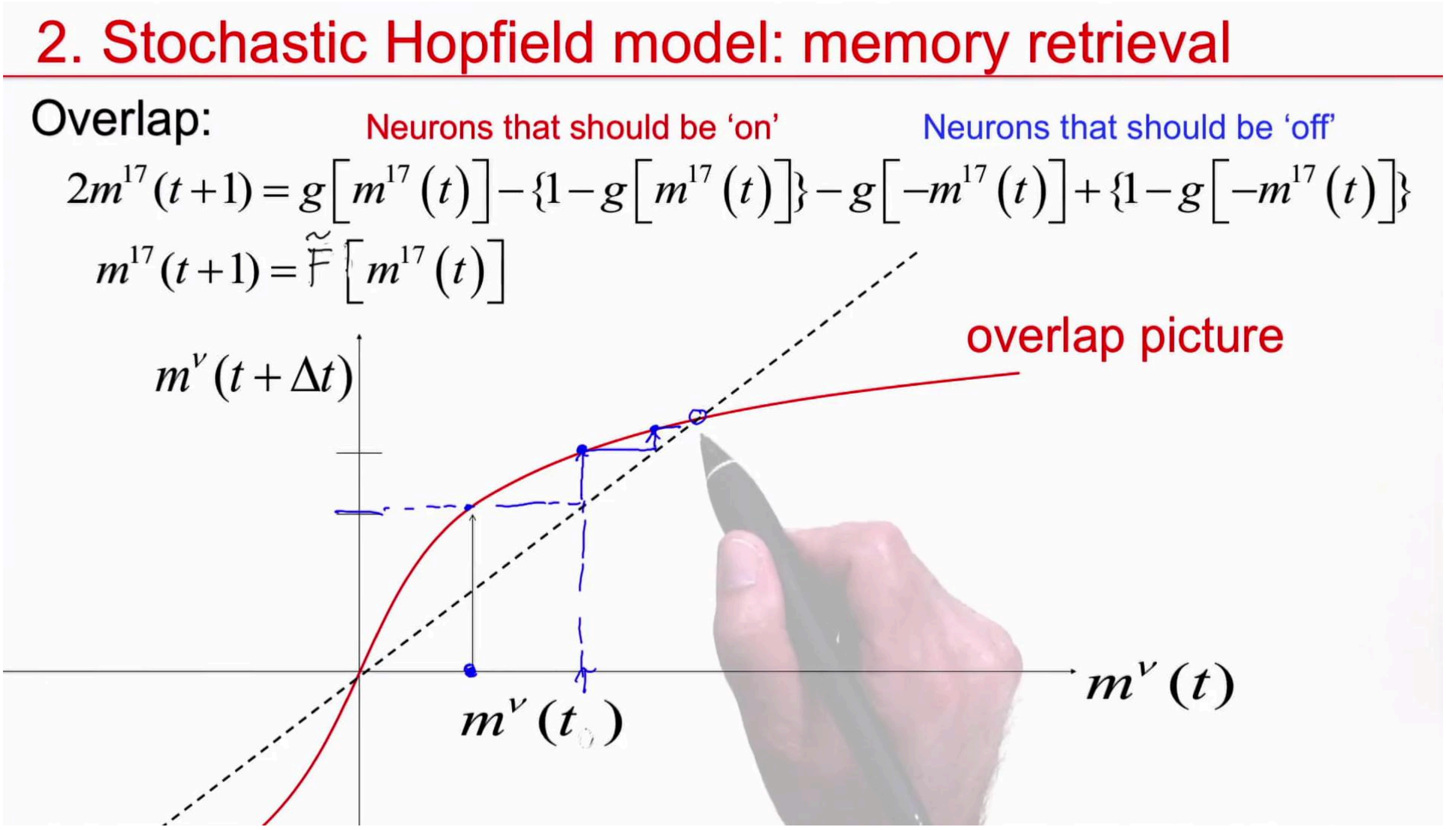


So let me summarize this, the result of the calculation is that I have some function which I call F tilde, I take the overlap at time t, passes through this f tilde and get the overlap of time t plus 1. And of course, pattern number 17 is completely arbitrary could be any pattern. So let's call it pattern mew the state of pattern new at time t plus delta t or t plus 1 the state in the previous time step. So we go from t to t plus delta t. And suppose we have an initial overlap of some value, some small value may be 0.4or 0.6, then we just read off the overlap in the next time step and this would be the overlap the next time step. Okay. But then this is the overlap times t plus delta t, which will inject so we will injected here now if you think about it, to find this reinjection point, we can just go to the diagonal, and this would be the reinfection point. Okay, and then we go up here. And this would be the overlapping the next state, we reject, we are here, overlapping the next step, and so forth. And you see, it will converge to this fixed point. And this 6th point is stable under this dynamics, if I started with a bigger overlap, I would go downwards. So this is a stable fixed point.

Notes

Summary

14m 11s

# 2. Stochastic Hopfield model = attractor model

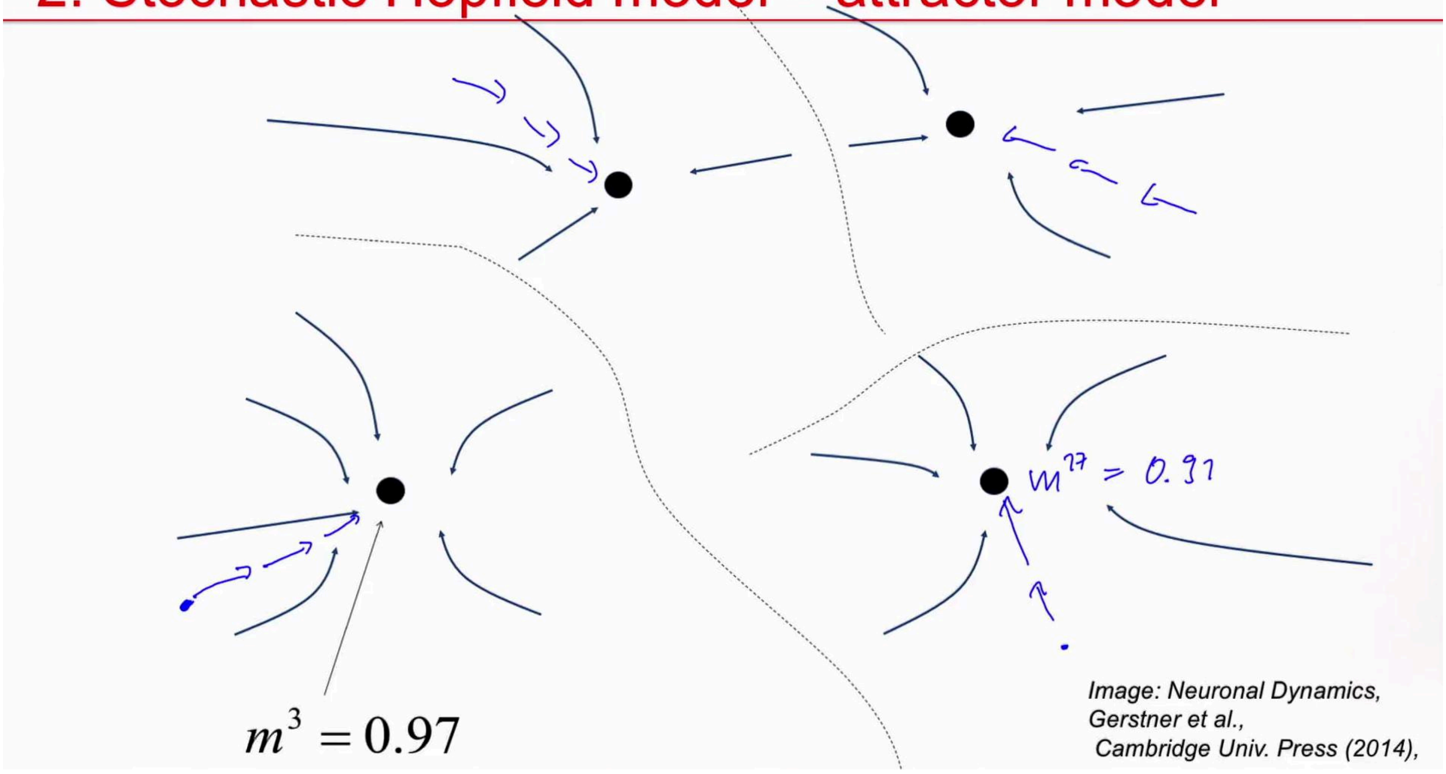


So what I've shown is that if you start in a certain pattern, for example, pattern number three, which has an initial overlap of all point four, then you may take a couple of steps and then arrive at the 6th point, the 6.6 very close to the perfect pattern, he doesn't see that overnight 1 for example, at 0.97, it may sit for pattern 17 the fixed point may sit at 0.91. An important fact here is we've converted to a fixed point go the pattern is nearly perfectly retrieved. If I start over here, it will converge to some other prototype maybe pet on number 12 if I start over there, it will yet progress to a different prototype. So depending on the initial condition, I will move towards the fixed point the ne2rk is attracted towards fixed point and that's why it's called the attractor memory. Now, you know that this is true here in the setting of the stochastic Hopfield model.

Notes

Summary

15m 35s

## 2. Stochastic Hopfield model: memory retrieval

- Memory retrieval possible with stochastic dynamics
- Fixed point at value with large overlap (e.g., 0.95)
- Need to check that overlap of other patterns remains small
- Random patterns: nearly orthogonal but 'noise' term

So what have you seen, memory retrieval is possible just the stochastic dynamics, we don't need the deterministic dynamics of the Hopfield model. If we converge to the fixed point, the fixed point will be very close to the pattern but he did not have a perfect overlap of 1 but a large overlap. For example, 0.954 or 0.97 now, for the calculations I did in this section, I sort of assumed that if I have a large overlap with 1 pattern, I can keep other patterns as your overlap. Now, this is implicitly the assumption of orthogonal patterns. Now note that random patterns are nearly or orthogonal. Therefore, the argument should also work for random patterns. So we have generalized to Hopfield model to the setting of stochastic dynamics. And we have shown that the concept of attractive dynamics is much more general than the original Hopfield model now to wrap this up, have a look at the quizz and see you for the next section.

Notes

Summary

16m 44s

# Quiz 2: Stochastic networks and overlap equations

[ ] The update of the overlap leads always to a fixed point with overlap m=1
[ ] The update equation as derived here implicitly assumed **orthogonal** patterns because otherwise we would have to analyze overlaps with several patterns in **parallel**
[ ] The update equation as derived here requires a function

$$g(h_i) = 0.5\left[1 + \tanh(2h)\right]$$

Notes

Summary

17m 44s