Computational Neuroscience: Neuronal Dynamics of Cognition

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Neuronal Populations

Wulfram Gerstner
EPFL, Lausanne, Switzerland

Reading:
NEURONAL DYNAMICS
- Ch. 12.1 – 12.4.3
(except Section 12.3.7)

Cambridge Univ. Press

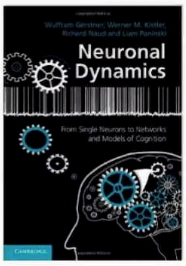

1. **Population activity**
   - definition and aims
2. **Cortical Populations**
   - columns and receptive fields
3. **Connectivity**
   - cortical connectivity
   - model connectivity schemes
4. **Mean-field argument**
   - input to one neuron
5. **Stationary mean-field**
   - asynchronous state: predict activity
6. **Random Networks**
   - Balanced state

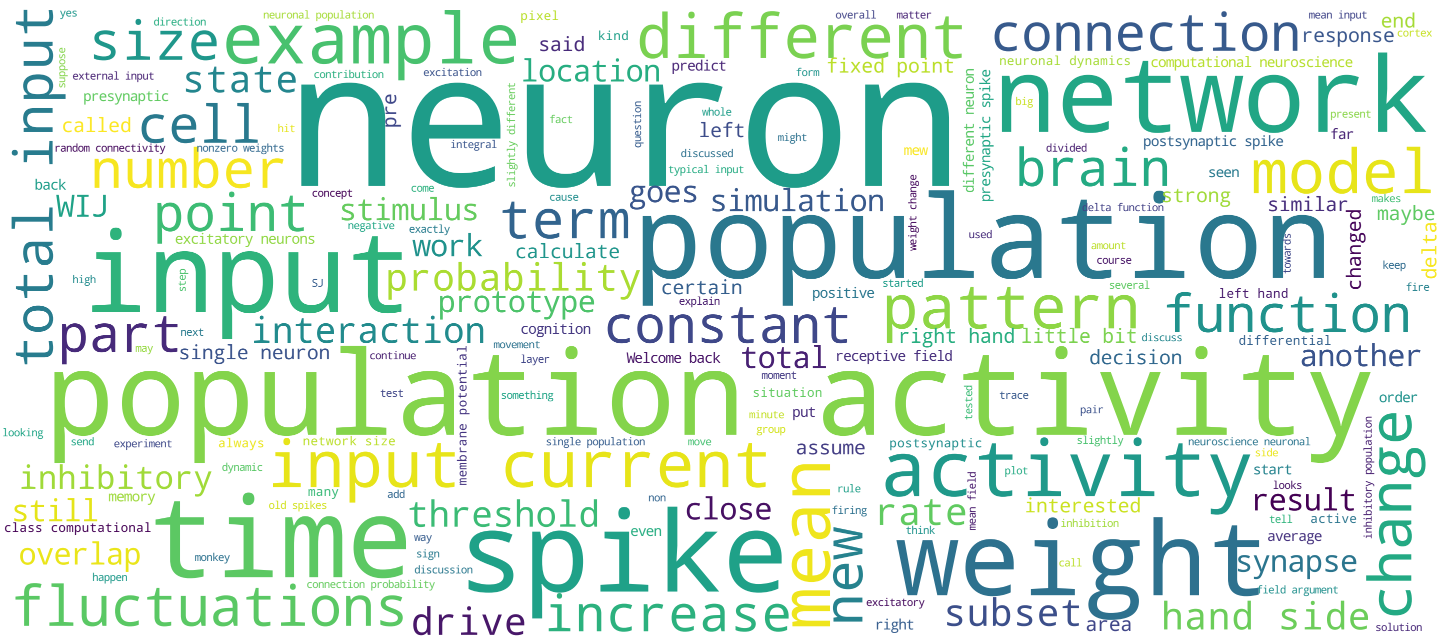

Search MOOC

Video

EPFL

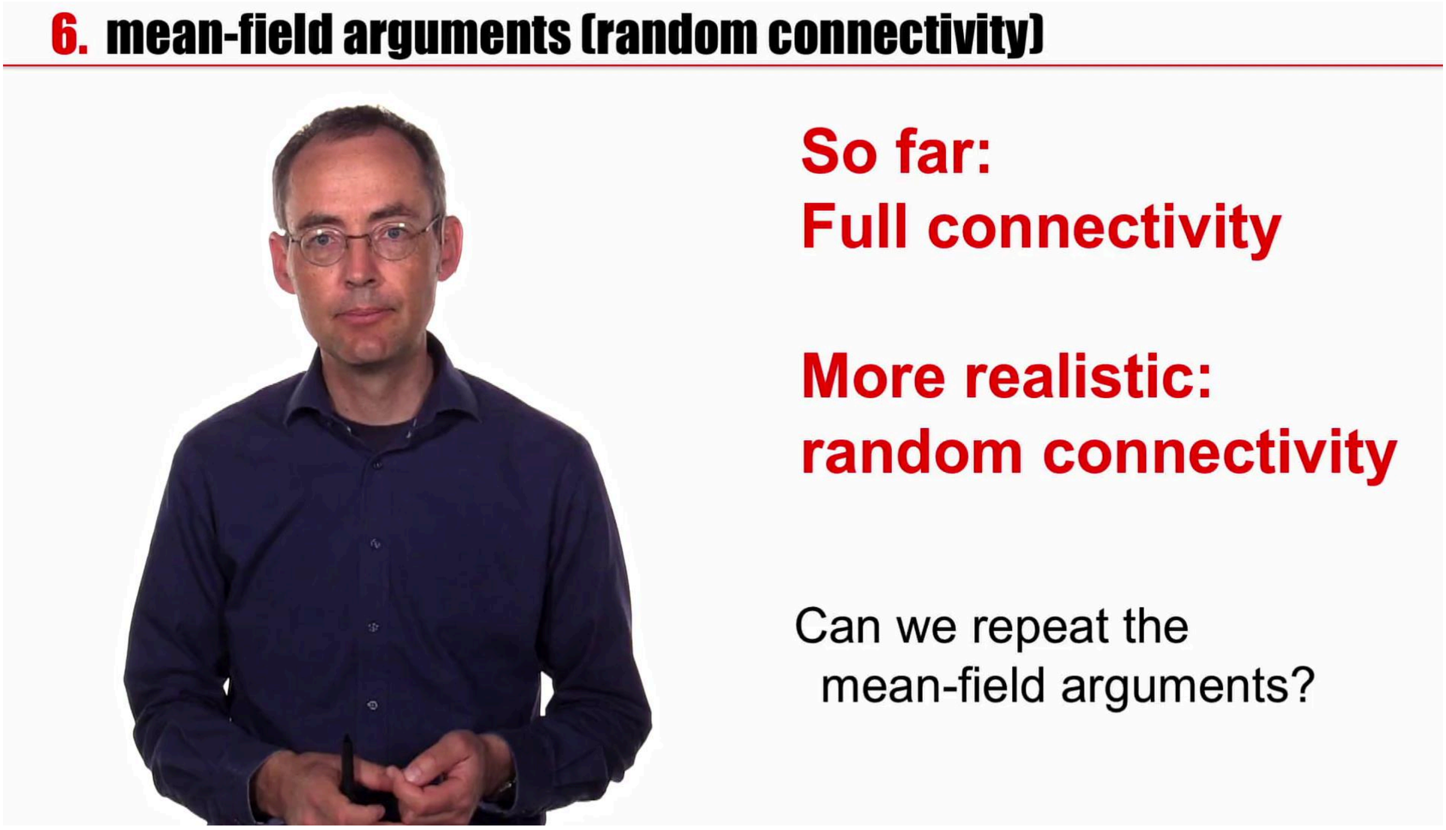

Welcome back to the class, computational neuroscience, neuronal dynamics of cognition. We'll continue our discussion of population activity. And so far, we have been restricted to full connectivity. That's not very realistic in biology and network it's not fully connected, we have seen that there are probabilities of connection so we have random connectivity. And so the question arises, can be generalize or mean field arguments to the case of random connectivity?

Notes

Summary

0m 05s

6. mean-field arguments (random connectivity)

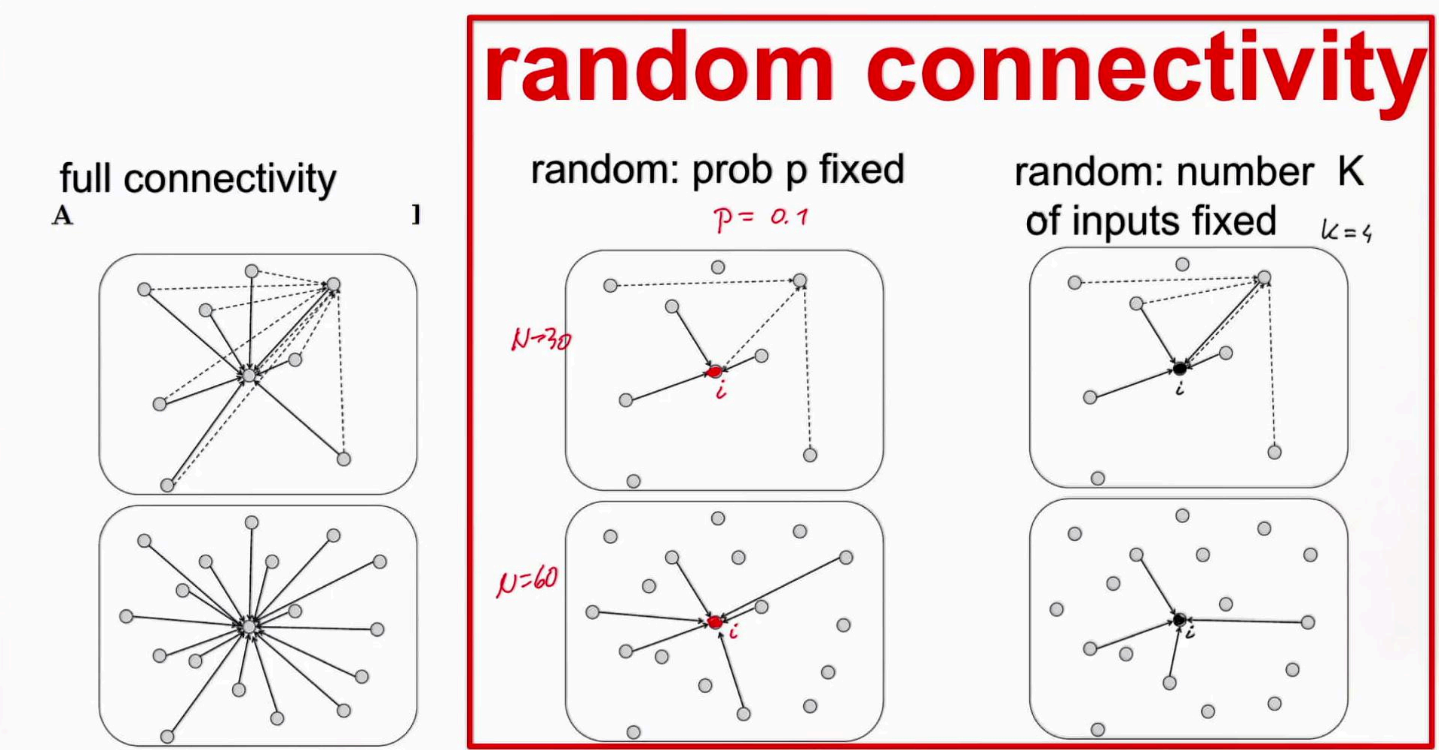

To random connectivity comes, in two different flavors. And we'll discuss both. The first one is, you're on the right hand side, for giving you an I, in a network of say, 5000 neurons you pick exactly K connections. In this case, for example K equal four and this neuron gets exactly four inputs. If we now go to a network which is twice as big and you get again at a newer I, you will still receive exactly four connections. Now the other possibility is that you save a fixed probability. For example, p equal 0.1 10% connection probability. Then, if I'm in a network of, say, 13 neuron, then this neuron I would receive three connections. The same neuron I, in a network of 16 neurons would receive six connections, so it would receive more inputs. Now for both cases will ask whether we can calculate the mean input to neuron I, and whether we can calculate the population activity.

Notes

Summary

0m 33s

6. Review - Random Connectivity: fixed p

Can we mathematically predict the population activity? A_0

given

- **connection probability *p* and weight *w*_ij**
- **properties of individual neurons**
- **large population**

$N \to \infty$

asynchronous activity

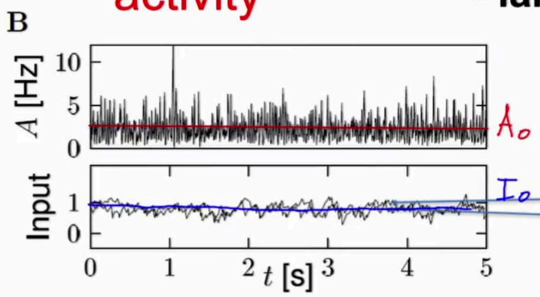

Input is nearly identical for different neurons

So our aim is as before, to predict the population activity. And what we can use is just the connection probability, and the values of the weights. And of course we will as before use the probability of neurons, and then assume a large population. So N is really large and goes to infinity, maybe N is 10,000. And here is a simulation of integrating fine neurons. I'll tell you in a minute how these works. And what have you observed is that we have this state of asynchronous activity with some mean value, a zero. This is the value we would like to calculate. And in order to calculate this, we will see that we will need the typical input into one of the neurons, which will call (I zero). And this graph shows that in my network of say 5000 neurons, two different neurons have a very similar meaning. Let's try to understand how this happens in order.

Notes

Summary

1m 57s

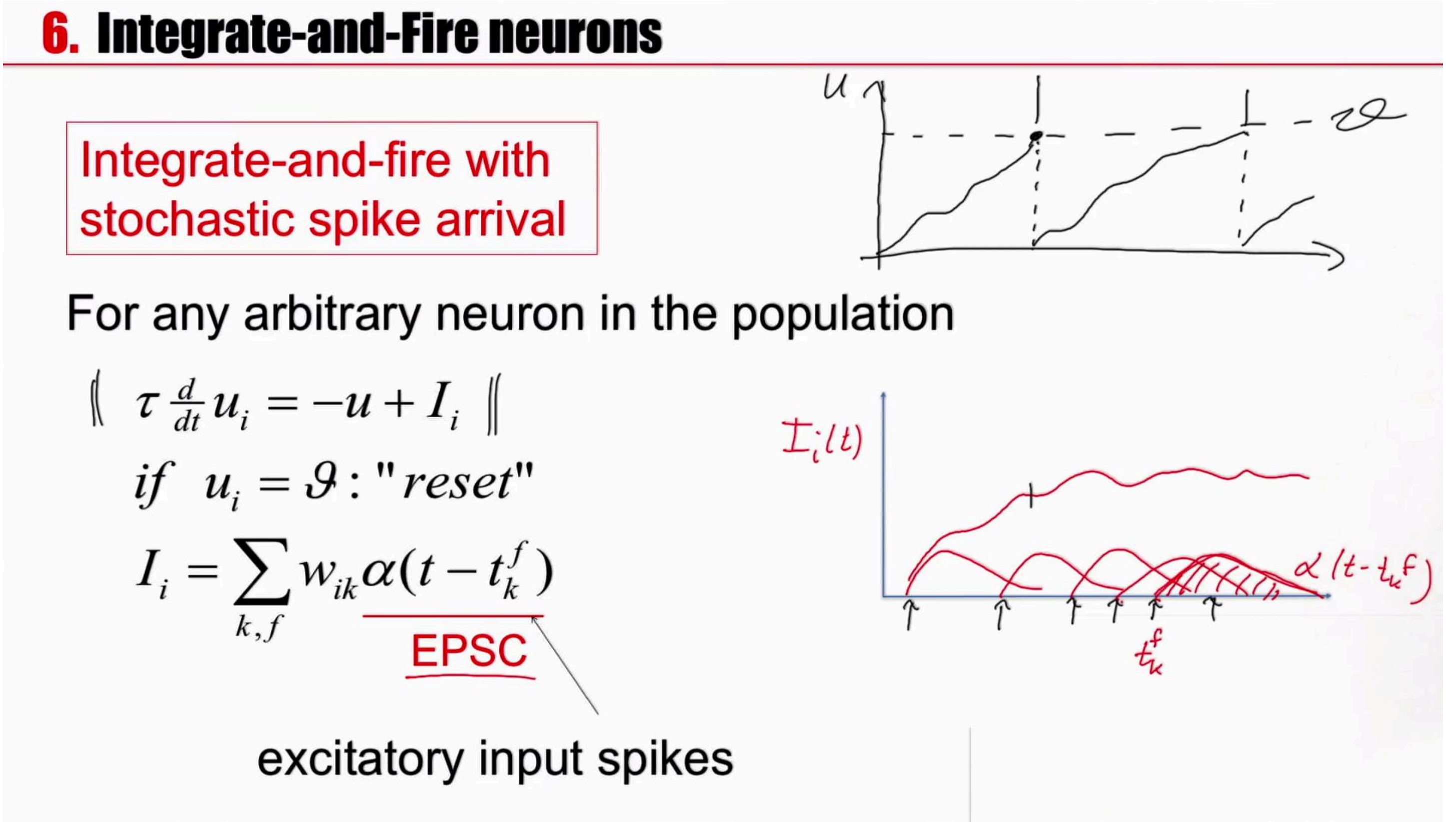

In order to explain the formalism, let me make a little detour and introduce the integrated fire neurons, to those who are not familiar with these. So I have a neuron here, and I can plot the input current and the input current would be for each spike that arise. Either include input current pulse, where one of these spikes, for example this one here is my spike emitted by neuron K and spike number F, and so this part here, this would be a synaptic current pulse, or excitatory post synaptic current alpha of T minus TKF. So this is the input current for neuron I. Now I take this input current and put it into a differential equation for the membrane potential. If I take this total input current, and as before, but what matters is the sum so this is my input current. Now I take this input current, put it into this differential equation for the membrane potential. You then this will start to increase. And at some point it will hit a threshold theta in it hits the threshold I say okay this is the moment when the neuron fires its spikes. Now afterwards, we reset, and then we use the remaining current to integrate again up, and maybe we hit the threshold again here that we said, and be integrating again up.

Notes

Summary

3m 09s

6. Network of integrate-and-fire neurons (random connectivity)

Integrate-and-fire neurons with stochastic spike arrival

For any arbitrary neuron in the population

$$\tau \frac{d}{dt} u_i = -u + I_i$$

$$if \quad u_i = \vartheta : \text{"reset"}$$

$$I_i = \sum_{k,f} w_{ik} \underbrace{\alpha(t - t_k^f)}_{\text{EPSC}}$$

excitatory input spikes

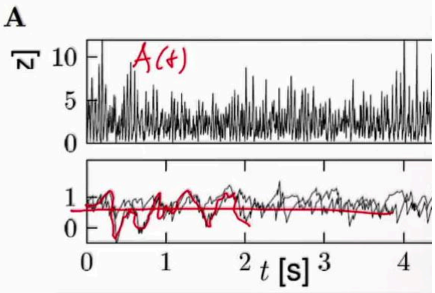

Can we predict the mean current?

So this is an integrated and fire model neuron. And it's driven by spikes, to derive stochastically, they arrive, because this new one is embedded in a network. Now this neuron is embedded in a network. And I'm interested in the total current that this new and receives. And here is the current of new and I. This is one of these traces goes up and down like this. And then I can have another new on K, and it would be the remaining trace. But I'm interested not in the fluctuations, then just interested in the mean current. And this mean current will drive the other neurons, and as a result they will generate this population activity A of T as we have here. Okay, so let's work on this current expression.

Notes

Summary

4m 57s

6. mean-field argument: random connectivity

p that w is nonzero

$$w_{ij} = \frac{w_0}{pN} = c$$

$$I_i = \sum_j \sum_f w_{ij}\, \alpha(t - t_j^f)$$

$$= \sum_j \sum_f \frac{w_0}{p \cdot N} \int \alpha(s)\, \delta(t - t_j^f - s)\, ds$$

$w_{ij} \neq 0$

$$A_0 = \nu = g(I_0) = g(J_0 w_0 A_0 + I_0^{ext})$$

What is my current Ii is the sum of all neurons J, all spike times W IJ alpha of T minus TJF. And that's it. That was the current that had shown before. Now let's do the same trick that we had in a previous section. Let’s write this slightly more complicated. We used to delta functions, we write alpha of S, delta of T minus JFT minus SES, which is the same, and then I have my sum. My sum over F, my WIJ. For WIJ, let's assume that the weights that are present. The not all the weights are present. But most are zero. But there's a probability P, that the weight is non zero. In for those that aren't zero. I'll take a weight W 0 divided P times N. So, this is just a constant. But I've chosen the constant in the particular form, which makes it easy to do the following calculations. Okay, now let's collect terms. I've here, my spikes. This is some of all New Years Old spikes. And also the one over P times N. I have N neurons in my network. But the probability that there is a nonzero connection is just P, so that each neuron, receives, P times N inputs. And it's just over those nonzero weights that I take this sum. Okay. This is limited to the case that WIJ is non zero.

Notes

Summary

5m 51s

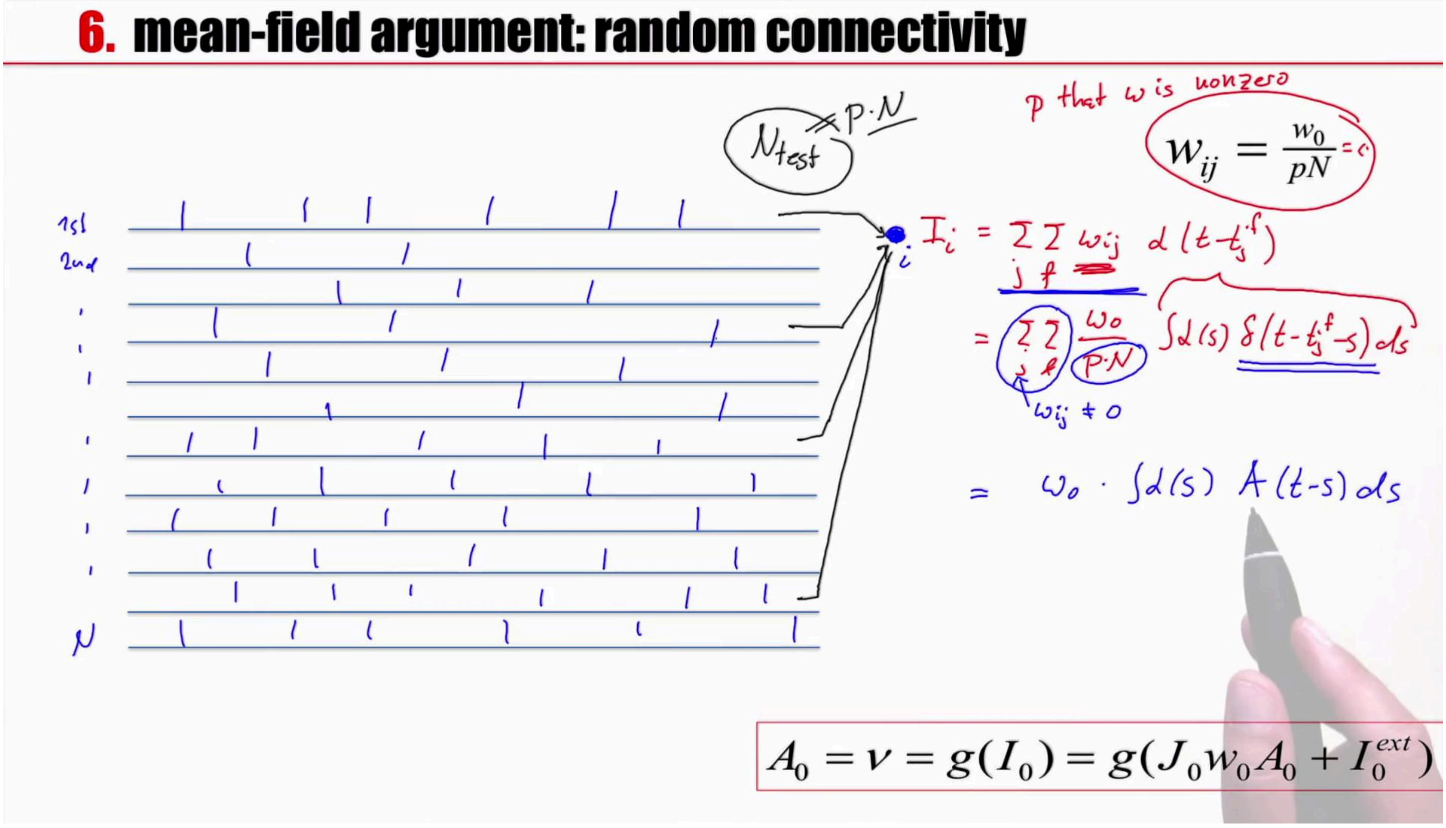

Now I will argue that this is SB for W zero times alpha of S, population activity of T minus S yes. Now, before we had some of the old spikes of neurons. Now we have something similar, and let me explain why this is still the population activity. So, these are my N neurons in the population. I've N neurons total. And as before. These neurons emit spikes. Now, what this formula here tells me that my neuron I, does not get input from the whole population. But this neuron I gets input from a subset of neurons. For example, from this neuron. From this neuron, from neuron, and from that neuron. So this neuron does not take the whole population activity, but it looks at a subset of neurons. This is very similar to the discussion we had earlier when I said, you can estimate the population activity by not looking at the whole network, but just looking at a tested a subset of neurons, at which you look at the average. And that's what this neuron is doing and how many neurons are in the test set of this, this neuron tested for the population activity. Well, it's just P times N. So this neuron, I actually some, some old spikes on neurons it’s tested of size P times N.

Notes

Summary

8m 01s

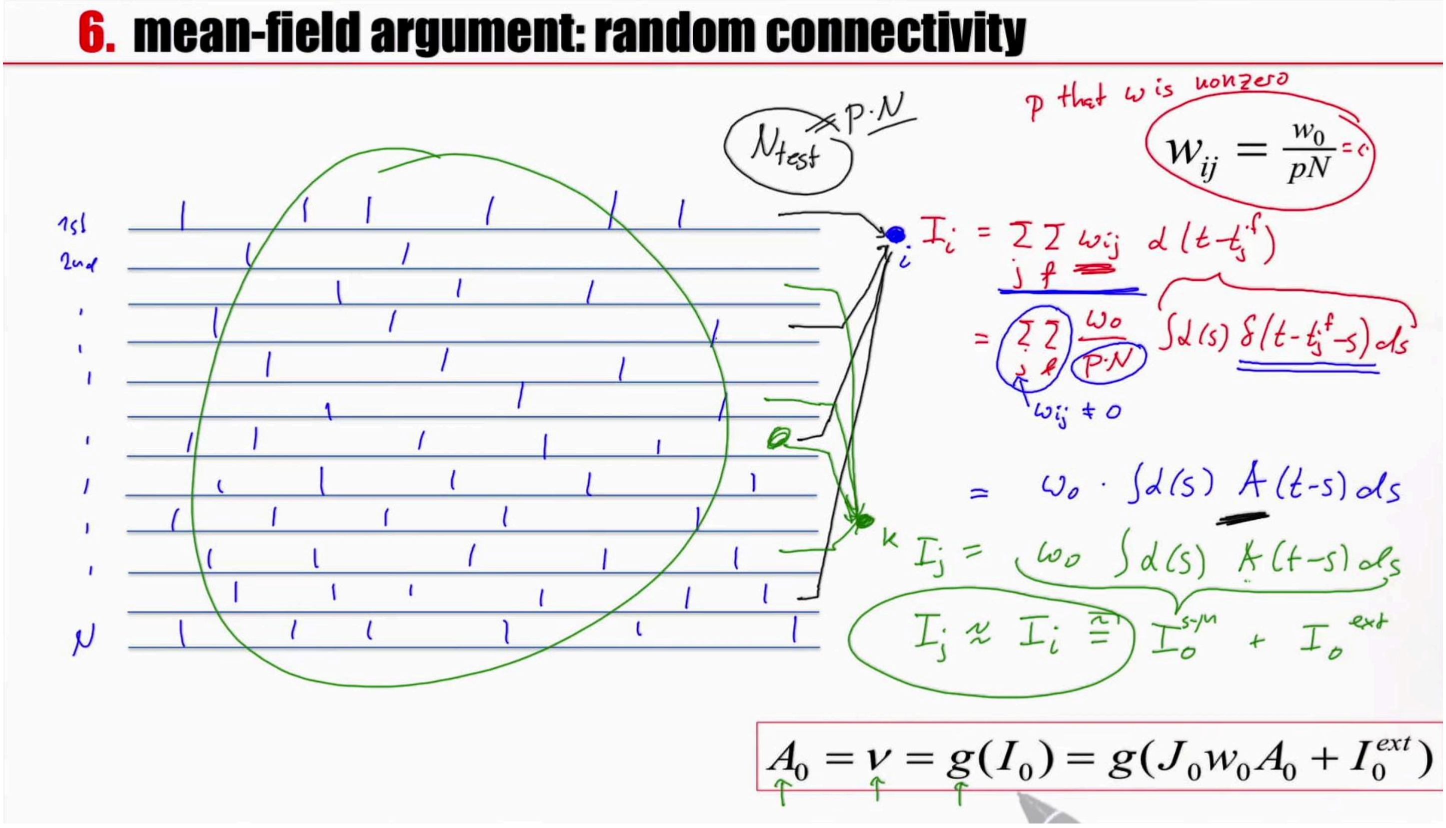

So this is a sample of the population activity for neuron I. Now there will be another neuron. K. And this other neuron will also take a subset of inputs. It's an independent subset it's an Independent sample of neurons that will drive this neuron I. And it's also a sample of size in test P times N, and this is worth the same input current drive, just that the activity is estimated from a slightly different subset, some new ones might be shared, others are different. But it's a slightly different subset, but overall, these are subsets that all look at the same global population activity. So in the end, the input drive to neuron J, will be very similar to that of neuron I. So in the end, it doesn't matter whether I look at neuron I = 10 or I = 25 on your number 3021. I will end up with the same nearly the same input current. So, the total input current is, I zero and I zero is this expression here. This is the synaptic part. And then I will also have the external part. And from this point on, I can exactly repeat the arguments that we had before. The population activity is equal to the single neuron activity, but for the single neuron I can use the gain function, the fi curve.

Notes

Summary

9m 57s

6. mean-field argument: random connectivity

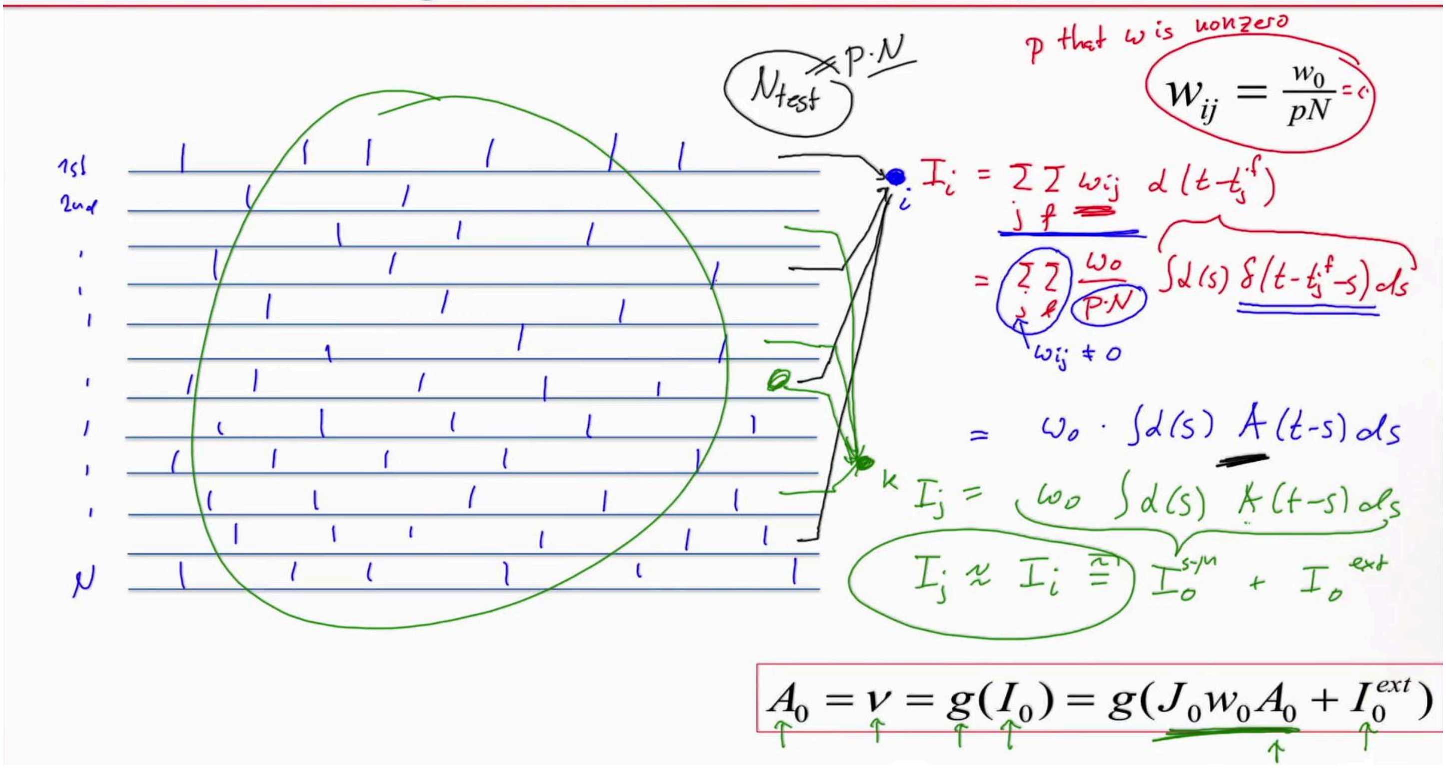

Again function depends on the total input, the total input has two contribution. The contribution from the external input and the synaptic drive, and the synaptic drive depends on a zero. And now you see that construction of the mean field argument, and the construction that you do graphically for solution is going to be exactly the same. So yes, we have extended the mean field argument to the case of stochastic connectivity.

Notes

Summary

11m 38s

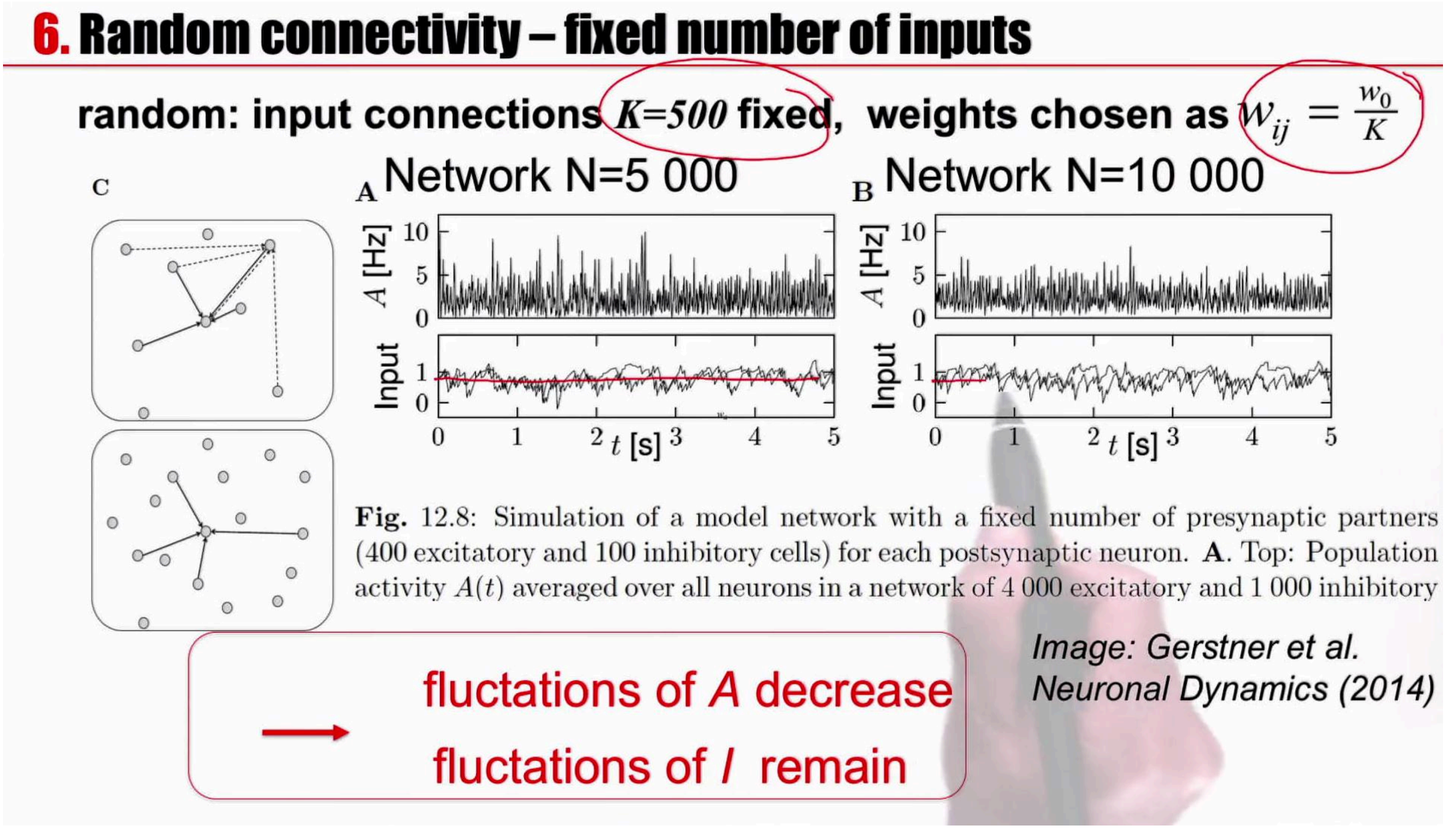

Now let's look at a couple of simulations, with integrating file neurons. So here it's fixed connection probability of 0.1 I have a network of 5000 neurons. So typically each neuron receives 500 inputs. You see, the input current have two different neurons in the network. And you see the population activity. Now I increase the network size. And when I increase the network size. I pick the nonzero weights, with this value, which has changed because now it has increased. So, weights has become twice, smaller, but since that's raw weights with probability p, I've twice as many nonzero weights. And now you see that the typical input is the same. What has changed its just the amount of fluctuations. The typical population activity is the same. What has changed it's just the size of the fluctuations in the population activity. So this was the situation for fixed probability, we compared to different network sizes. Now let's do the same for the case of a probabilistic network, were are each new and receive exactly 500 inputs. In which case, the weights, those weights that are non-zero all have the same value. W zero divided by k. And now you see again. I have a typical input which is the same for both neurons.

Notes

Summary

12m 08s

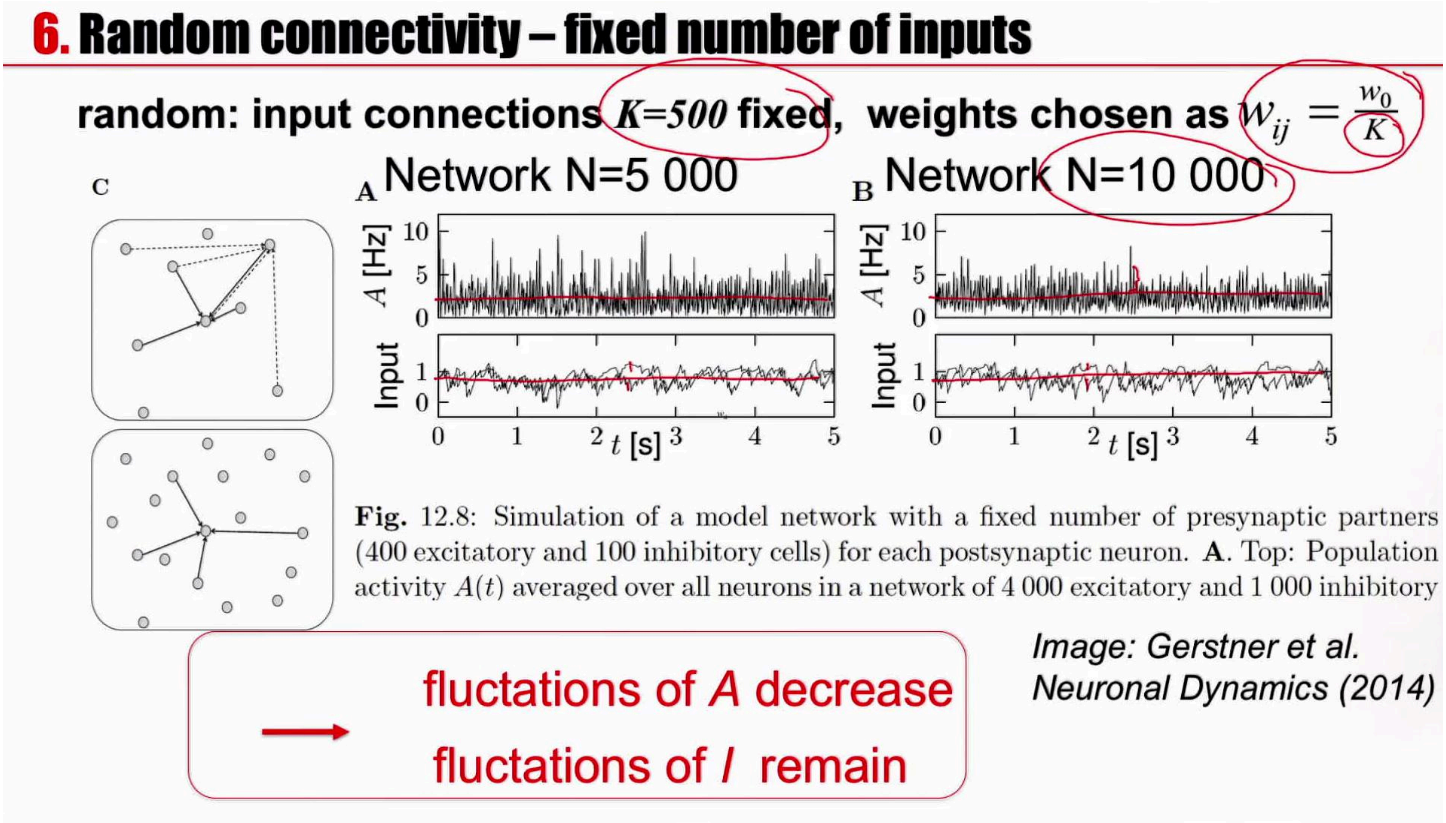

And it's the same but it’s a network of 5000 or 10,000 neurons. Now the size of the fluctuations. Also stays the same. Why? Because we said that neuron draws a sample of the population activity. The sample size is K and K does not change between the small network and the large network. The population activity fluctuates around a constant value easier, which we can predict which is the same in the left hand side is on the right hand side, and the fluctuations have decreased a little bit. Why, because to block this I average over 10,000 neurons. So just by this larger sample size in the plot. I've slightly reduced fluctuations.

Notes

Summary

13m 50s

6. Connectivity schemes – random, fixed p, but balanced

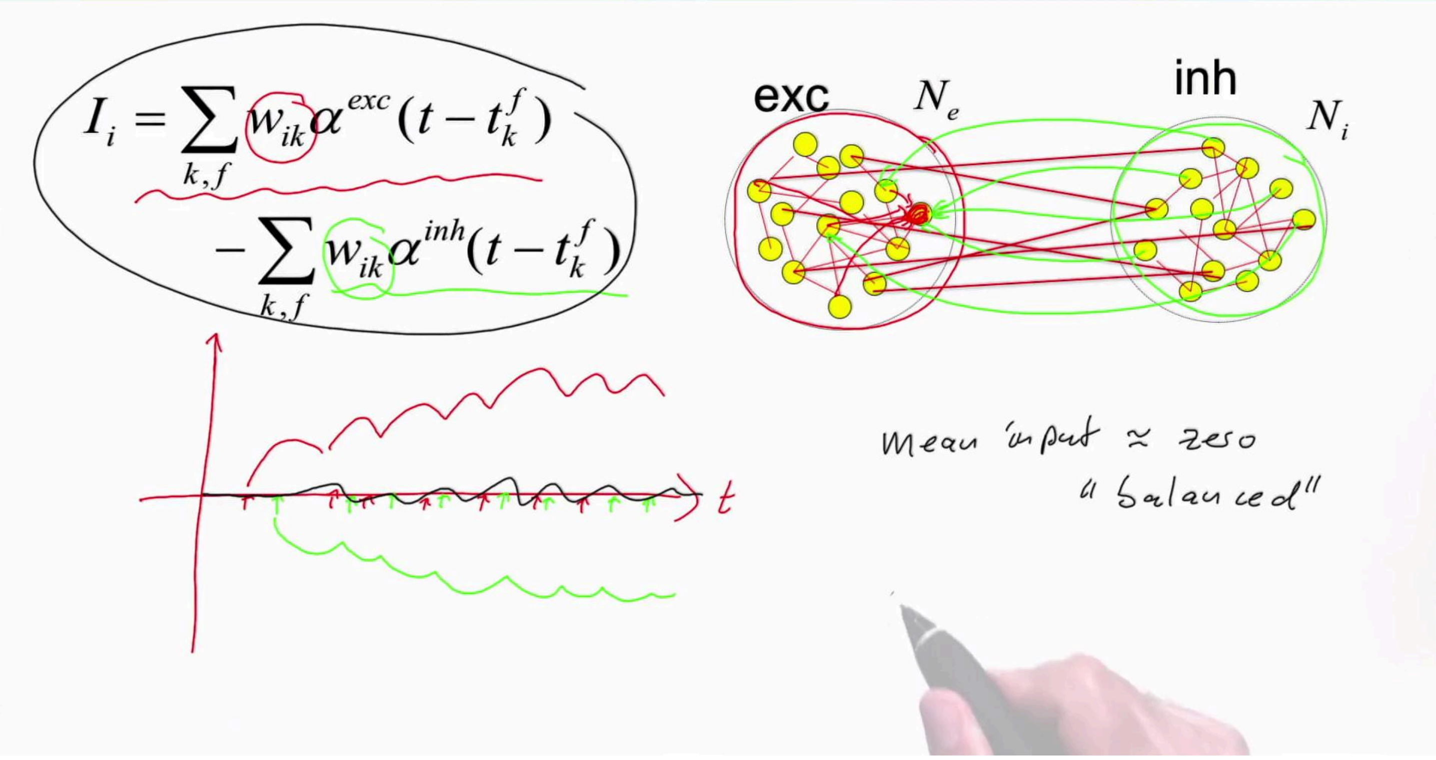

So far I have focused in the discussions on one single population. Now new ones coming two principal types. Excitatory neurons, and there are inhibitory neurons. Now for the excitatory neurons. I already discussed the total input current as a function of time if the excitatory spikes arrive here so I look at one neuron gets excited through input from subset of the other neurons, and the total input, totally excited for input will look like this. But this neuron, will also get inhibitory inputs from a subset of the inhibitory population of neuron. As a result, there are spikes coming from the other population from the green population, which also cause some response, and which will also add up and adroit downward because it's inhibitory. Now if you adjust the wait parameters for the inhibitory connections, and for the excited hurry connections appropriately. Then you will see that the total input is always very close to zero fluctuates around zero. So what drives a neuron and then indicate fire and what’s threshold is not the mean input, which is very close to zero. But it's the fluctuations of the input. So the mean input is zero, which means it's balanced between excitation and inhibition.

Notes

Summary

14m 42s

6. Connectivity schemes – random, fixed p, but balanced

$$I_i = \sum_{k,f} w_{ik} \alpha^{exc}(t - t_k^f) - \sum_{k,f} w_{ik} \alpha^{inh}(t - t_k^f)$$

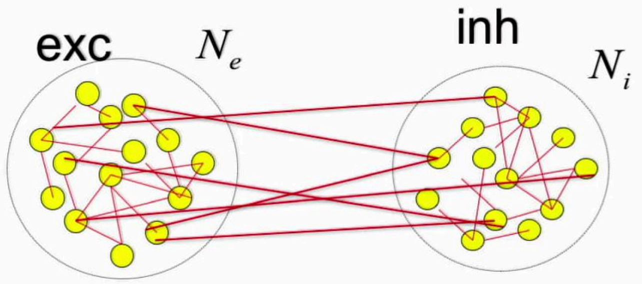

make network bigger, but

-keep mean input close to zero

$$p N_e J_e = -p N_i J_i$$

-keep variance of input

$$w_{ij} = \frac{J_e}{\sqrt{pN}}$$

$$w_{ij} = \frac{J_i}{\sqrt{pN}}$$

C. van Vreeswijk and H. Sompolinsky (1996) Science

But the fluctuations can still drive a neuron across the threshold. So the appropriate scaling for the weights. So that means zero, and the, the fluctuations are big enough to drive the neuron towards the threshold is in this form there it's the square root of the probability of connection times the number of neurons in the network. And this allows to keep the meaning put always close to zero. But if you increase the network size, the variance of the input which means the amount of fluctuations, roughly remains constant.

Notes

Summary

16m 17s

6. Connectivity schemes – random, fixed p, but balanced

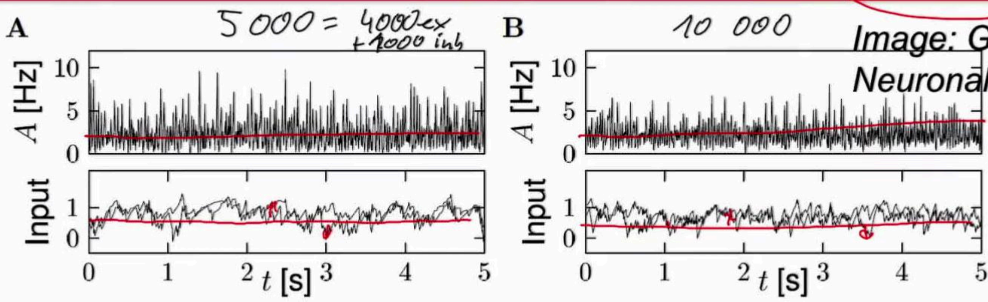

Image: Gerstner et al. Neuronal Dynamics (2014)

Fig. 12.9: Simulation of a model network with balanced excitation and inhibition and fixed connectivity $p = 0.1$ **A**. Top: Population activity $A(t)$ averaged over all neurons in a network of 4 000 excitatory and 1 000 inhibitory neurons. Bottom: Total input current $I_i(t)$ into two randomly chosen neurons. **B**. Same as A, but for a network with 8 000 excitatory and 2 000 inhibitory neurons. The synaptic weights have been rescaled by a factor $1/\sqrt{2}$ and the common constant input has been adjusted. All neurons are leaky integrate-and-fire units with identical parameters coupled interacting by short current pulses.

→ fluctations of *A* decrease

fluctations of *I* become 'smoother'

And this is simulation again with the integrated fire on neurons. On the left hand side, it's 5000 neurons. On the right hand side, its 10,000 neurons, total out of the 5000 neurons. There's 4000 the excitatory and 1000 inhibitory. And you see that the value of the population activity, and all the changes between the left and the right hand side and the input here is not really balanced at zero. It is slightly above zero, for certain mean, and the size of the fluctuations, also roughly remains the same. In the big network on the small network, what changes is that the trajectories become a little bit smoother. So we have these kind of highly volatile trajectories that go up and down, a very typical for neurons in vivo that are measured in the life brain. And this is an indication that this balanced regime, where excitation and inhibition roughly cancel each other, might be the regime in which the brain is operating.

Notes

Summary

16m 57s

6. Neuronal populations: outlook

One population
→ multiple populations

Application to visual cortex
→ visual processing

Application to decision making
→ competitive networks

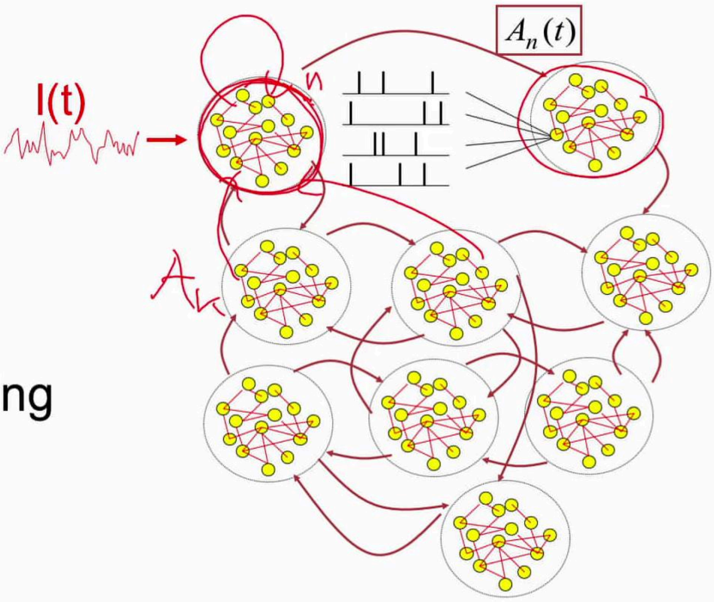

So what our next steps. So far, I focused on a single population. However, this single population is maybe just excitatory neurons, to one particular political column. This population of neurons will be connected to other populations of neurons and will receive input from other populations of neurons from other crowded columns in the same crowded area, and even from other areas. The framework that I've discussed this week is generally enough to extend to the situation with multiple populations. And we'll use this in the following sessions, for example, for applications in visual processing and applications in decision making. The drive of a single neuron, in this general case is the activity of the network onto itself. And it's also the activity that it received from other populations' k. The population N sends since, its activity, to other populations that receive input from this first population. So instead of interacting network of neurons, we now have an interacting network of neuronal populations. And that's the level at which it is useful to discuss cognitive processing in the brain. So let me summarize, we have gone a long way. You've started off this neuronal populations, population in cortex, for example, a set of neurons excited, in one specific layer for cortical columns.

Notes

Summary

18m 10s

6. Summary: Neuronal Populations

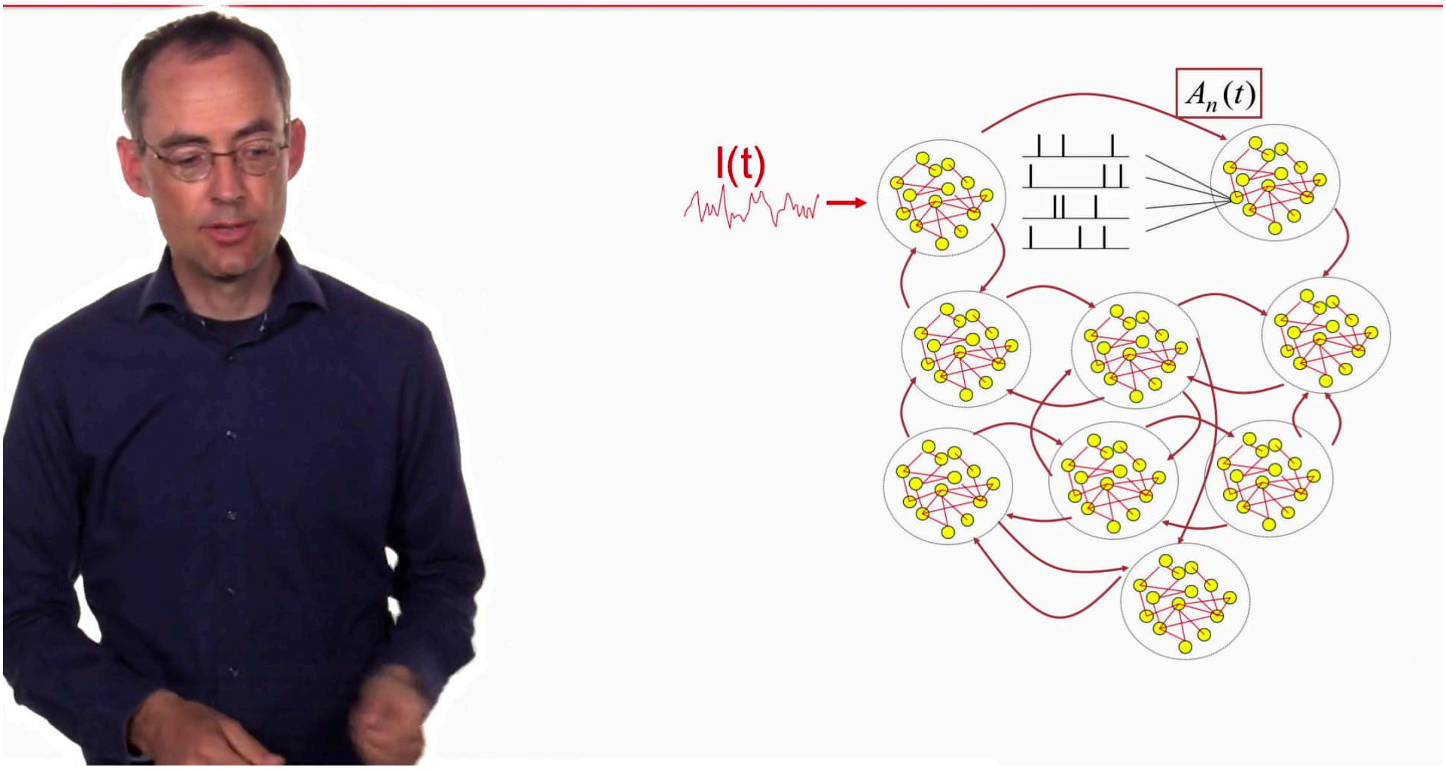

I've argued that these are fairly homogeneous populations. And now for such homogeneous populations, we can write down a mathematical framework that allows us to predict the population activity based on neuronal powers parades, based on connectivity parameters. Now, this was just one neuronal population. In the next weeks we will go on, and discuss the interaction between different cortical populations. And that allows us to look at specific perceptual, and cognitive phenomena. A lot of people have worked on the theory of population activity on normal populations, on the formalization of the mathematical tools.

Notes

Summary

19m 50s

6. Selected References: Neuronal Populations

Receptive fields, columns, and cortical connectivity

D. H. Hubel and T. N. Wiesel (1962) Receptive fields, binocular interaction and functional architecture in the cat's visual cortex.. J. Physiol. (London) 160, pp. 106–154.
T. Bonhoeffer and A. Grinvald (1991) Iso-orientation domains in cat visual cortex are arranged in pinwheel-like patterns.. Nature 353, pp. 429–431.
S. Lefort, C. Tomm, J.C.F. Sarria and C.C.H. Petersen (2009) The excitatory neuronal network of the c2 barrel column in mouse primary somatosensory cortex. neuron 61: 301-316.. Neuron 61, pp. 301–316.

Modeling populations

H. R. Wilson and J. D. Cowan (1972) Excitatory and inhibitory interactions in localized populations of model neurons.. Biophys. J. 12, pp. 1–24.
C. van Vreeswijk and H. Sompolinsky (1996) Chaos in neuronal networks with balanced excitatory and inhibitory activity. Science 274, pp. 1724–1726.
N. Brunel (2000) Dynamics of sparsely connected networls of excitatory and inhibitory neurons. Computational Neuroscience 8, pp. 183–208.
W. Gerstner (2000) Population dynamics of spiking neurons: fast transients, asynchronous states and locking. Neural Computation 12, pp. 43–89.

For those not familiar with the Dirac delta: https://www.youtube.com/watch?v=l3hvrx33lZc

More info on neuron models: http://lcn.epfl.ch/~gerstner/NeuronalDynamics-MOOC1.html

Notes

Summary

20m 40s

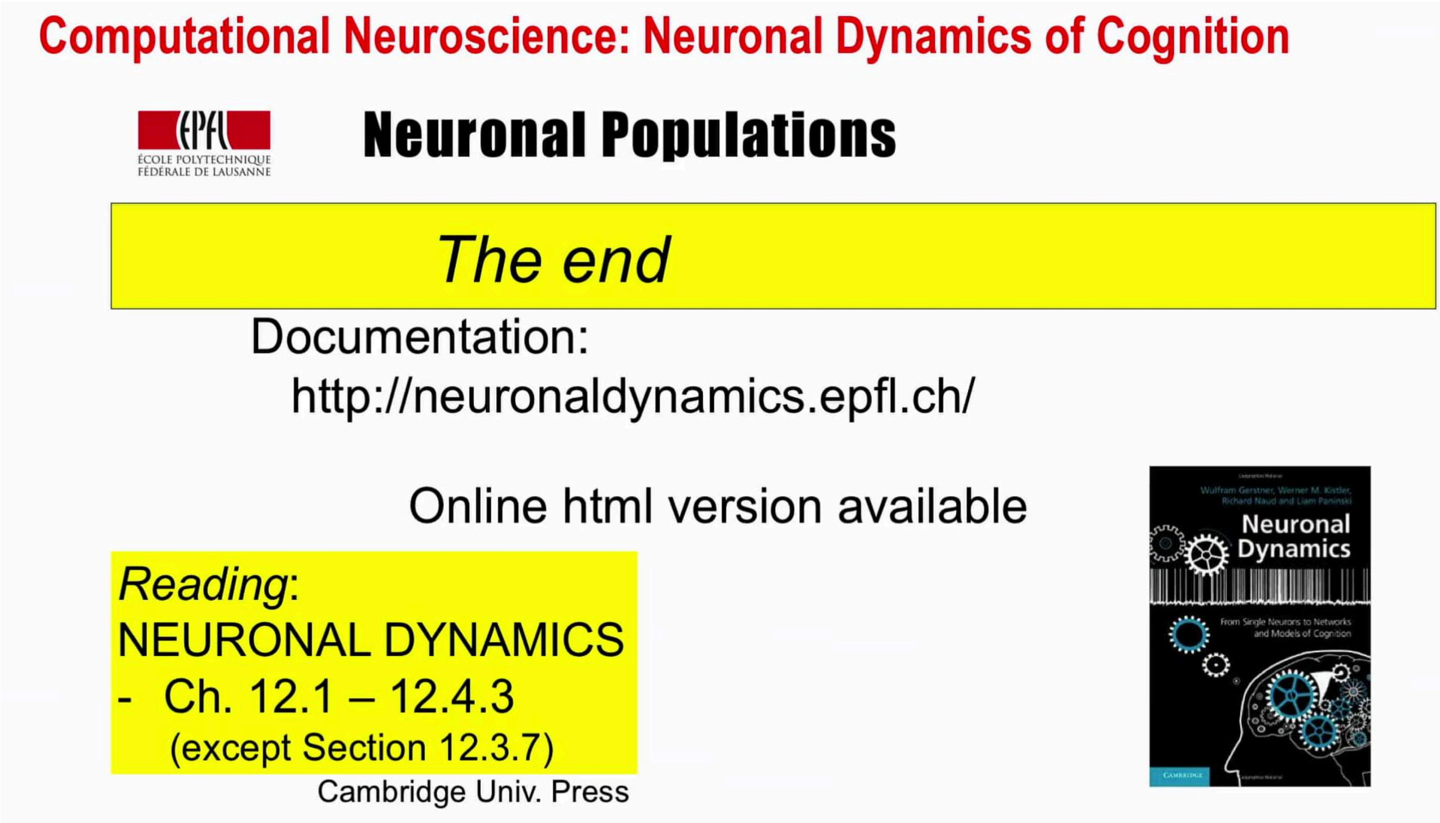

And please have a look at their references and as always, there's also the textbook available in its printed form, and in its online, for, for those who have not been exposed to DR delta function before they may also want to look at the appropriate chapters in the book, or at the videos that are part of the references. I hope you enjoyed it. See you again next week.

Notes

Summary

20m 46s