

Physics Seminar

April 1, 2019

>>: All right. Can I have everyone's attention? I want to introduce our speaker for today. This is Greg Morrison, coming to us from University of Houston. This is his first time giving a talk here for our seminar series, so hopefully this is going to be a very good experience for him and he'll be happy to come back in the future.

He is an assistant professor in physics at University of Houston. And he has a Ph.D. from University of Maryland, where he studied biophysics and statistical mechanics. He did postdoctoral work at Harvard University where he focused on information theory and biological networks. And he also spent some time in Italy.

And so he's had a lot of different research, mostly spanning the work of single molecule biophysics and, you know, as you saw last week, there was a talk about how the brain works and how the individual neurons work within the brain. So I think this is going to be very complementary to that.

Also I do want to let you guys know, I am going to pass around a sign-in sheet. So expect that and also I do need to hand back papers. So with that I'll turn it over to Dr. Morrison.

>>: Well, thank you very much. Thank you for having me. For those who are watching at home, I'm impressed that at 7:00 p.m. you're watching someone give a talk on neuroscience. So thank you for tuning in.

I do have a background in single molecule biophysics. I'm not going to be talking about single molecule biophysics, though. I'm interested in biophysics in a lot of different contexts, and this talk is mostly going to cover network science and how one can understand the global properties of a highly densely connected network like the brain, given some very simple rules.

The organization of the talk, just to give you a sense of where we're going. So the first thing I'm going to talk about is describe a bit about how neurons work, how the brain works, and some important features in neurodynamics. I'm not going to assume that you know very much about how neurons work. And it is the case, though, that a lot of this is complementary to what I believe Dr. Cheung talked about last week. She talked about calmodulin, right? And so she's working at a very high resolution. I'm going to be working at a very low resolution.

Also going to talk about how we model these neurons mathematically and how essentially heterogeneity in the context between these things can lead to physically relevant correlations in the firing patterns of these neurons. Okay?

Those communities, though, I'm going to show, can -- if you choose them poorly, can lead to what's called a breakdown of the balance condition and lead to unrealistic firing. And then I'll talk about how we can recover that balance. And so I'll show that we're able to recover that balance and also that recovering that balance is actually not what we want to do. The brain can't be balanced in order to have the kind of behavior that we want. And so I'll talk about how we can break that in a controllable fashion to recover physically meaningful behavior.

And then finally I'll talk about hierarchy in these networks and how we can send signals from one side of the brain to another which is really a fundamentally important function of neural networks.

Also, I said the word neural networks. How many people have heard the word neural networks before now? Have you heard it in the context of brains or in the context of machine learning? Both. Yeah. So I'm going to condense the machine learning side of things down to 45 seconds or so. And so hopefully I will mention that somewhere around here, hopefully it's clear and of interest.

So let's get started on biologically motivated networks. So first I'm going to try to motivate why it is that we might care about these things and what sort of data is available. So this is a picture of a person with a micro EEG on their head. Essentially this device is able to track the electrical -- an electrical signal at all of the different locations along this brain.

Each of these lines correlates to one of these sensors, and this displays the electric signal at any of these points in this person's brain. This person has epilepsy. It causes one to have seizures randomly.

This is likely, though not certainly, likely not a full seizure that you're seeing, but simply a heightened activity within the brain where you see essentially this relatively quiescent behavior here that suddenly starts sparking wildly. This black line is a measure of overall activity in the brain. It says there's a lot more activity happening. And each of these lines talks about the behavior on a millimeter scale within the brain.

There's a couple of features that I want you to take away from this. First of all some of these lines are completely flat. If you're having a seizure, your entire brain doesn't light up. Only certain parts of it.

Second of all, you will you can see there are differences between these things. So here -- let's see, this blue one is lit up very brightly and is completely missing here. There are some other features that you can see, these guys are much lower here than they are here. And so you can see that the individual behavior, the behavior of individual neurons, isn't what drives this, isn't what drives this epileptic behavior. The significance trial to trial variation in the behavior of individual neurons or individual regions in this case. But we still see the same qualitative picture. We still see that the brain is behaving in a similar way.

So it's going to be important from this picture to understand the global properties of this thing, the global properties of the brain without focusing solely on a single neuron. We need to look at the statistics of this thing.

There are other mental health conditions -- schizophrenia, a few others -- that have been shown to have this sort of behavior as well where there is seemingly a structural deficiency that can lead to these mental -- these brain related issues.

Oh, and before I go on, please interrupt me if you have questions. I'm happy to answer questions at any time.

>>: (Inaudible).

>>: I cannot, but hopefully someone can.

>>: Oh, I can do that.

>>: I don't want to -- I'm worried that I will lift up the projector screen. So I will -- okay.

So, this is -- EEG is great. It tells us something, but it only tells us about the outside of the brain or something along these lines. A much more three dimensional representation of the dynamics of the brain can be measured using what's called fMRI, which is functional magnetic resonance imaging.

And in this particular set of pictures, which I'm sure all of us have seen something like this on a TV show where you're able to see the behavior of a patient's brain or something along those lines. This is real technology used to track the blood oxygen level at different regions of the brain in realtime. So this means that if there is a region of the brain that is active at a time, it needs more oxygen to function. It's a cell. It still needs blood to come and go and deliver it the oxygen that it needs. And so these things are correlated with activity.

Now, one thing I want to point out is these aren't telling you anything about individual neurons. They're telling you about big regions, millimeter, centimeter regions of the brain and how they're behaving collectively.

Alternatively, we can look at what's called the Human Connectome Project, begun in 2009 and it's unbelievable the things they're able to do now. But by combining dissection, which is obviously not good for an individual but good for science as a whole, we're able to determine the connections between individual regions of the brain to other regions of the brain. And these are much closer to an individual neuron level. These are micron scale.

So cells -- these cells aren't necessarily going to be micron scale. These are still long range, somewhat larger connections but we can draw these connections between different regions of the brain at incredibly high precision. This data is available for humans, mice, and a few other organisms.

And so this really opens up an enormous amount of ability to study the global behavior of the brain. And one feature that I want to point out here, so you can sort of see that these

bundles of filaments between different regions of the brain, they sort of clump together. You can see that some of them are bound together and then there are gaps. This is because regions talk to regions. You don't just have a uniform conduction across the entire network.

We can do a bit better than this, though, by designing systems that allow us to measure the activity of an individual neurons within an individual's brain. So unfortunately you have to cut into the person's head which is why this is not too often used in humans, but it is for those that have these brain conditions. More often we use model organisms.

This is -- in this particular case, this is -- what is this? This is a salamander that was put on a petri dish on top of one of these things and the electrical activity of every single neuron is monitored. That's what all these dots are. And then this image is the image of another salamander that they're showing the visual cortex of this first salamander in the petri dish.

So that's what this is. What you're seeing in this next picture is essentially where these arrows are point to the -- point to the particular location that's being monitored. This is as a function of time. Once you turn this image on, what is the response of the neurons along this -- along -- in between these two arrows?

And what you see is after about 50 to 75 milliseconds, nothing is happening. There's a lag time, which we've all experienced. And then what you see is there's an extremely highly correlated firing that indicates that something has been seen. If you look at the lag between the time of first firing for all of these things, this is what the visual cortex reconstructs.

The lagged -- the latency of this, that's what this latency is, zero means that it fired at the average time. Negative, which is dark, means it fired earlier than average, and white means that it fired later than average. This latency encodes the information for what it is that this visual cortex is looking at. And we can monitor this in realtime on a neuron level.

Let's see. A couple of other things to say. Let's see. So one thing to point out, essentially you show this visual cortex a picture, that picture is going to be registered essentially simultaneously. So this is there variation that gives you the information, but this is all very synchronized. It's not really random.

It's possible to have a asynchronous information, an asynchronous transfer of information or transfer of stimulation, where the firing rate is increased in certain regions of the brain but it's not synchronous, essentially. So there is a relationship between firing, there's a lot of firing going on even if it's not timed perfectly.

So we're going to actually in this talk mostly be interested in asynchronous correlation, which means that we stimulate something and downstream something else gets really excited. But it doesn't have to occur at the same time. Okay.

So those are some motivations. What is it -- so I said the word neuron a couple of times and I'm not sure how deeply Dr. Cheung got into this picture but it's definitely worth

talking about. This is a picture of a neuron. This is what's called a presynaptic neuron because it has a large axon here that is connected to a postsynaptic neuron. This neuron is going to send a signal along the axon to the neuron on the other side. That's the difference between pre and post.

Whatever -- so the way the signal is sent is that there is -- the cell is composed of a membrane. This membrane has ion channels. These are either potassium or sodium. And essentially if there is -- if the cell reaches a certain potential difference across this, which is greater than its resting potential, so at rest there's going to be some equilibrium difference across the membrane.

As positive charge flows from the outside to the inside, it's going to take on an additional positive charge in here and if the potential difference is -- if the change in the potential difference is large enough, it causes a cascade where suddenly the sodium -- so the sodium channel opens up and sodium from outside rushes in driving the potential difference to become more and more positive.

That signal is -- once that becomes high enough, then the potassium channels open up and all the potassium is released and so that is -- the entry of -- the rapid entry of sodium is called a depolarization and then repolarization is when the potassium is pumped out to recover the resting potential difference.

Once that happens, that process will propagate along this axon. Along this axon there are all these different channels and if we have a depolarization event here -- let me switch my pointer. Whenever we have a depolarization event here, essentially it'll repolarize.

But that is -- because this was depolarized, there's going to be a positive potential difference across the next section of the membrane, which means it's going to depolarize and then try to repolarize. But it's going to pass that signal of depolarization on. And so there's this cascade, this sort of avalanche of depolarization and repolarization that sends a signal from the presynaptic neuron to the postsynaptic through this axon. Okay.

This is an all or nothing event. Either this -- let's see. So either this depolarization wave propagates or it doesn't. So you can have little fluctuations in the potential here; but if there's not this depolarization wave, there's not a signal sent to the other neuron.

How does this work? So this is a schematic of the depolarization wave, you can imagine these are the depolarization events where calcium is flowing. And this is whatever the -- let's see. Whenever the wave hits the end of the axon, neurotransmitters are released and they can be detected by other neurons. And so this is a chemical signal that's driven by this potential difference in the axon. Okay?

Those neurotransmitters are things sort of like -- so glutamate is one of the more common ones in -- glutamate is commonly used to send an excitatory signal from one neuron to another and GABA is another neurotransmitter that sends an inhibitory signal. But there are lots of these neurotransmitters -- dopamine, adrenaline, a lot of things that you've heard of before.

Essentially the process of receiving this signal is a complicated chemical process, as well. There's diffusion of these neurotransmitters to reach the receptors at the postsynaptic neuron, and that process is strongly related to what Dr. Cheung was talking about before. So this is the scale that y'all heard about last week and I'm going to be interested at the much more global scale of the kind of data that we can get from the Human Connectome. Okay?

All right. So how do we model this?

So in order to mathematically model these things, there are many, many steps. From the calcium and sodium pumps all the way down to -- I'm sorry, potassium and sodium pumps all the way to the propagation of this wave and then the chemical process that underlies it. All of that is very complicated and difficult to model simultaneously.

And so what we're going to do is use a much simpler model, which is called a Leaky integrate and fire model. That means that we're going to pretend like a neuron is just a ball in space, and it's got a connection to other balls in space. We're even going to ignore that right now.

We've got just some sphere and it has a potential associated with it. That potential has a rest value that it likes to sit at and it reaches that exponentially. So it exponentially decays either up and down. If its potential is too high, it decreases its potential. If it's too low, it increases the potential.

So this is the dynamics of the potential of the Leaky integrate and fire with one exception. If the potential is high enough, it fires. What it's going to do whenever it fires is drop -- so this is what the dynamics is going to look like. The potential drops to zero and stays at zero for some amount of time, which is called the refractory period. And then after that refractory period, it's going to return to its usual dynamics of exponentially, in this case, growing from zero to whatever its rest potential is going to be. Okay?

So this means that if we just have a neuron sitting around and as long as it's -- the rest threshold, the potential it likes to rest at is greater than the threshold potential at which it fires, we're going to have regular, like clockwork, firing of this neuron. Right? It's going to continually increase up to the threshold then drop, wait for a little while, increase, then drop and so on.

So this is the sort of dynamics that we suppose, but that's for a neuron that's in space living by itself. In order to account for any external signal or interactions with other neurons, we're going to add a synaptic current. Which is just going to be -- this is the -- this DBDN is the normal dynamics and then there's going to be something that comes from outside.

What's going to be coming from outside is due to -- we're going to focus primarily on the synaptic current being driven by other neurons in the network. So we imagine that we've got a whole bunch of these idealized spherical neurons. And that's what's represented here.

Each of these dots represents a neuron and each of the lines represents a connection between them. In this particular diagram, they're completely randomly connected. There are in this picture 4,000 neurons or supposed to be. This is -- this may be scaled down to make it visible.

But there are 4,000 of these exciter neurons and 800 connections between every neuron. So every neuron is connected to 800 others in the network and it's completely random and this is the sort of picture that you should see from this -- for this network is essentially this ball of hair where you have lots of connections and there's no real structure. Okay?

So how do these things talk to each other is one of the last things that we need to talk about. So this is the synaptic current at neuron N. It's going to be composed of two pieces. One is an exciter piece, which is going to be multiplied by this KP that says anytime some other neuron I'm connected to fires, I get a signal. And so I'm going to add to my potential -- the time derivative of my potential. I'm going to get a jump, potentially and this size of that jump is going to be this exciter signal. Okay?

But there's also inhibitors in the network that are connected. Anytime one of them fires, I reduce my potential. There's going to be a drop in my potential. So like I said, these are the number of neurons, the number of exciter neurons, the number of inhibitor neurons, and then this is (inaudible). Okay?

And this is the sort of dynamics that you see. So this is a plot of the potential of a particular neuron. So we're going to imagine just one neuron, tracking it. And what we see are these firing neurons where the potential suddenly drops. And then it builds back up and then floats around for a little while and then drops again.

We can visualize this, these individual drops. There's a lot of noise here. We can visualize this more efficiently by drawing what's called a raster diagram. This is the time of the -- this is the time running from zero to two seconds and this is the neuron index. So neuron one is down here. Neuron 4,000 is up here.

And every one of these dots corresponds to one of these firing events. So every time it fires, we draw a dot. So if one was firing -- if they were all firing synchronously, what you would see would be straight marks. Everyone fires at whatever time this is and then this time and so on.

A couple of features to see here. It looks random. A few things is there are a couple of pictures I'll show you where you might be able to convince yourself it's not random, but it is. You might actually convince yourself -- does anyone see it looks a little lighter here? Random. Dumb luck.

So all of these are going to have weird features, and your brain is really clever and tries to tell you that there's something here, but there just isn't. It's always fun to look at these things because they can be random, and you will see something in them.

But this picture says this model brain is firing randomly and there's no structure whatsoever to it.

Before I go on, I just really briefly want to talk about neural networks in the machine learning sense and tell you why they're different than what we're talking about. So a neural network in a machine learning sense is a machine that tells you whether or not a picture is of a cat or a dog, basically.

There are hundreds of applications of these things, but at its most basic, it takes in data at what's called the input layer. And then it does something to that data. Whatever it does to that data, it sends it downstream to what's called a hidden layer. And that hidden layer processes it somehow.

Essentially it takes the pixels here and averages them or swaps them around or whatever. But essentially it sends a signal from the early input layer down to one of these hidden layers, and that can pass through another hidden layer and another hidden layer. And then at the end it has an output layer that says this is a cat or a dog or a plane or a bus or whatever. And it's unbelievable the stuff they can do. It's incredibly accurate in some cases. It's a really remarkable system.

The important thing about this is if you take data, you take a hundred thousand pictures of cats and a hundred thousand pictures of dogs, you say one is a cat and one is a dog, and then the computer teaches itself what is the correct way to connect each of these layers to recover cat or dog. That's essentially what's going on.

We're doing things that are different. One is that we're going to have positive and negative connections, and that's true here. We have a positive connection here, a negative connection there. There's no problem with that.

We're going to be -- for real neurons, they either spit out exciter neurotransmitters or they spit out inhibitor neurotransmitters. They don't do both. Neurons have to be one kind or another.

We care a lot about dynamics. Whenever I have one of these networks in the machine learning sense, essentially I sense pixels from here to here and then from here to here. For neurons, all the signals can pass at different times. It makes it a big mess and hard to extract information.

And then finally, whenever you built this thing you handed it cat pictures and dog pictures. So if I show you an elephant, it's going to say 50 percent cat, 50 percent dog because it doesn't know what to do. Our brains aren't that. We're probably not all that optimal, as anyone who has gone into a room and said I have no idea why I went into this room, knows. Our brain doesn't do everything perfectly, and these things are designed to be as perfect as possible.

So neural networks are amazing. It's not quite what we're talking about here.

Okay. Okay. One other mathematical thing, so I mentioned that these things were random. What do I mean by random? A Poisson process is one where you have essentially an exponential distribution of the time between events. So sometimes things happen close together, like these two right here and sometimes they happen really far apart and that's just random chance.

So Poisson process is what we're thinking about whenever we think about random. And if we just had one of these idealized things sitting out without any connections, it would be a Poisson process. The firing would be random when distributed.

If the network matters, though, we kind of expect that it's not going to be a Poisson process. It's going to behave a little bit differently. And so we're going to monitor this network using two things, often. One is the firing rate, the activity of the neuron is going to matter. Also we're going to be interested in what's called the Fano factor, which is the variance divided by the mean.

And it's a nice feature of the exponential -- of the Poisson process is it has a Fano factor of one. If we see a Fano factor less than one, then our signal is more organized than exponential. And if it's greater than one, then our signal is more noisy, has more diversity in activity than you would expect by random chance. So F is -- I'm sorry. F is less than one means that our system is more predictable than you would expect by chance. Okay?

Okay. How are we doing on time? So, so.

Okay. So we drew -- in this picture -- sorry. In this picture we had the worst model for a brain possibly imaginable. We have a whole bunch of neurons that are this idealized thing, then we connect these things randomly, which is probably not how the brain is organized.

Neurons are organized, are connected to each other spatially, right? Where they are in the brain matters. They're connected between regions. And so the function of the region in your brain has an influence on who you're connected to.

If you draw essentially the connections between the -- the long-range connections between neurons, what you see is that you see a binodal structure of the brain. Essentially this is left side and right side or vice versa, side one and two, I suppose, where each of these cords corresponds to a connection between these things. What this says is that it depends -- there's a much denser set of connections here than here. Right? You don't have all that many connections that are really long range.

Another -- so this is -- I want to say this is macaque visual cortex, but I'm not 100 percent sure about this. This is a visual cortex of some creature that I can look up. What you see is that there's a much denser set of connections here than there are here. Some of these neurons are highly connected and some are not. So there's a whole lot of diversity in the structure of the connections between neurons in the brain.

So we can take a step forward and say, okay, why don't we include some of that diversity? Where essentially we're going to go from this really basic idea to a slightly more complex idea where now we have multiple regions of the brain that are more dense within themselves than between each other. So you can clearly see there's a denser set of connections here and here and here.

These are called communities. They're called clusters in the neural networks literature. And so I'm going to use both of these words interchangeably. And this is not very clear.

So there's lots of firing events that are occurring here but they're not pictured well. It looks basically like this. Right? The firing is still going on. But on top of that, there is this -- these brighter bands.

And what's happening here is that within -- so each of these corresponds to a single community of neurons that are clustered like this. And here what we've done to make these clusters, we've increased the probability of connecting to one of your friends or one of your same type of neurons. We've increased the probability of those connections by a factor of 2.5. You can increase it by less or more. You'll get roughly the same behavior, but 2.5 makes for a nice picture.

We also increase the strength, so you remember those J factors? So neurons of the same type are more exciting to neurons of the same type than -- in comparison to neurons of other type.

And what we see is that sometimes -- so this -- in this region which is difficult to see because of this projector, there is very little firing or normal firing sometimes. But other times there are these really bright spots, and that means that those communities are really lit up. There's a massive increase in the activity. And so this is essentially (mic cuts out) excitement of the activity within these communities. Okay?

We can also stimulate these. This is a little bit clearer. You can kind of see that there is actually activity outside of the bright spots, right? So under this band you see the gray area. These are communities that we've grouped together and decided to stimulate from one second to 1.5 seconds. And by stimulate, all I mean is that we add a constant value to that, a synaptic current that I talked about before.

While it's being stimulated, the Fano factors drop, which means that it becomes more predictable. We're seeing that this thing is behaving more predictably. (Inaudible) is more predictable than (inaudible). So that's the dynamics that we're seeing here. Okay?

So, an interesting feature of this is that this is the Fano factor for neurons that are directly stimulated. If they're not directly stimulated, so all the neurons up there, we also see that the Fano factor drops. And what's happening there is that for a somewhat unexpected reason, these neurons are being suppressed. The neurons that aren't being directly stimulated are being suppressed.

And the reason why is because these guys are more excited, so they send a signal to the inhibitor saying I'm excited and those inhibitors inhibit everyone else in the network. So there's some evidence that perhaps that plays a role in actual brains, but there's also evidence that it doesn't. And it would seem strange to light up a community and then turn off of everything else. If you see something, you don't want to be incapable of moving your hand. So it's interesting to see whether or not this is something that we can suppress. Okay.

So this is the pay dirt. We have homogeneous communities. The pictures I showed previously, every community is the same size. So what I got interested in -- we're finally to the part where this is my work -- and the question I got interested in is in the heterogeneity of these things.

So this is the picture I showed before. These communities are clearly very different. This is another visual cortex of another creature whose name I don't recall. And what you can see is there's a community here. Every dot here is a connection between these things, like an anatomical connection. So these things are connected in a physical sense.

And what you can see is that there's big communities and also small. And this is true for most of these brains -- most of the brains that have been studied at this level looking for connections between neurons shows that there's a lot of heterogeneity in this.

So I was curious, what does the heterogeneity do? So what we wanted to do, let's have some big communities and some small communities and see what happens. This is the model we chose. We chose the size distribution to be exponential. So with some probability, we have a big community. With other probability, we have a small community. And we just randomly generate networks.

And what we see is something bad. So a couple of features to point out. (Inaudible) blue then red, blue then red, blue then red. The blue corresponds to the first largest community. The next red corresponds to the second community. The next blue corresponds to the next community and so on. It's alternating in color so you can visualize it.

We also see that everyone is sort of happy and doing their own thing up until about 500 milliseconds, at which point the largest community freaks out, becomes completely hyperactive, and because its activity affects the inhibitors, it completely turns off everything else in the network.

This is not good news. If the brain were wired like this, it would be very bad because this happened spontaneously. This is a topology that causes extreme hyperactivity that leads to behavior that would be medically very bad. Right? So this was a big surprise. The project that we were starting to work on had already sort of gone off the rails because this was not the expected thing.

So why does this happen? What's happening is we're breaking a principle that is called balance. And it has a particular definition which means that on average, the amount of excitement that a single neuron receives should be about the same as the amount of inhibition it receives. If it gets too much excitement, it's just going to keep firing. If it receives too much inhibition, it's going to turn itself off.

For a balanced network, this idea of balance, each neuron is going to experience roughly equal of these things, roughly equal excitation and inhibition. And it's been shown -- this is something that is well studied -- that for a balanced random network, these two conditions should be satisfied. Okay?

And before -- some of y'all might be thinking, what's wrong with this guy? Why didn't he balance the networks to begin with, if this is so well understood. We actually did impose these two conditions. We didn't screw up that bad. So these are insufficient to balance a clustered network. The reason that they're insufficient is because not only do we want the balanced condition to exist, we want it to be stable.

So if we imagine that we've got some set of exciters and inhibitors that are evolving under a mean field potential -- so here r is the firing rate. τ is some time scale we don't know. This is a mean field approximation, so, it's some number.

Imagine that the firing rate forgets about what its current state after some amount of time, this time τ , but interacts with other neurons in the network with some interaction matrix. With some function that we don't know. Okay? So this is all that this equation is saying is that there is some relationship, some relationship with the current firing rate and its dynamics.

Now, if we assume that we can linearize this function about its -- about the mean field solution for this, which is just going to be these F_{es} , essentially we see that this differential equation vanishes whenever r_e is equal to (inaudible). And so long as these two things vanish. That's only going to be stable if these things are negative.

So the solution exists and it's going to be stable, which means that small perturbations are going to let it stick to that solution if this has negative eigenvalues. If this has positive eigenvalues, it's going to lead us somewhere else. Who knows where it's going to lead us, but it's going to go somewhere else. So this is the basic idea with balance is not only do we need to impose balance, but we need to make sure it's the same state that we end up with.

For a clustered network, this argument still works. This is just r_e and r_i . If we have a bunch of different communities, we've got to break them in -- essentially we have -- instead of having two equations, we have C plus one equations. But we can -- all we need to do is make sure that we have all the negative eigenvalues. Okay?

So sorry for the matrix math. I'm just going to wave my hands a little bit. Don't get too stressed about the number parameters here. If we have homogeneous community sizes, we have a very regular matrix. The matrix is going to be some stuff on the diagonal, which is the self-interaction. Some stuff off diagonal, which is going to be the interaction between exciters. And then these bottom and side rows and columns are going to be essentially the interaction with the inhibitors. Okay?

And we can actually diagonalize this thing. It's not hard to do. This is the condition -- this is C minus one conditions to balance a homogeneous network.

And if we do the same thing for heterogeneous networks, essentially the only thing that changes is that some are bigger and some are smaller so we have to include these fractions. So this says that if I increase the connection strength within A by some fraction that's proportional to the number of neurons within that community. Okay?

So that's nice. But a problem is that if we pick bad values of F , everything falls apart. The eigenvalues can be negative and we end up with really badly behaved systems.

So what we're going to do is we're going to say okay, well, the only reason why that's happening is because we had to scale everything up for interactions due to the fact that there are more neurons. Why don't we scale down the interaction strength so that we balance that out?

This is not the only solution to create a balanced network, but it's a solution. Turns out it's very easy to do. And so we decided to work this out. These are the conditions that are required. And so it's absolutely possible to construct a very simple weight matrix that allows us to look at balanced in neural networks.

And what we see is global synchronization. So I mentioned earlier, we have these straight lines going up and down. That means that almost everyone fires at a particular time. We got very excited about this, and it turns out that this is garbage. The reason why this is garbage is --

This, by the way, whenever you're doing original research and you get excited about something and nothing is working right, it looks like it's ridiculous, figure out why it's ridiculous because it might lead you to something interesting.

The reason why this is happening is (inaudible) divided by F_i , but F_i is really small. But I just increased the strength of some of these connections by a lot. And so that means that if I have a community of just one, right, just one neuron sitting by himself, one lonely guy, the interaction -- let's make that two, two neurons in the community, and that's the only size of the community. Well, they're going to have an enormous interaction between each other. That's going to suppress everyone else and we get this synchronization. Everyone is inhibited until these two guys aren't. And that's ridiculous.

The problem here is we've included these tiny, tiny communities in our -- whenever we rebalanced the network. And so what we're going to do is throw them away and we have some mathematical formulas to choose what to throw away. But basically small communities are out, we keep the big communities. And this is finally the firing pattern for this balanced network. Essentially this is the largest community, second largest, and so on and so on.

And it behaves reasonably. Nothing is going crazy, and we're actually able to see reasonable dynamics here. This is the Fano factor saying that if we have these balanced communities, all of these communities are balanced and they have a slightly lower than one Fano factor, meaning they are a little bit lower. There are correlations inside of these communities, more than one would expect by random chance.

And then we have these other communities. These are small. All of these guys way up here, they're in small communities. Who cares what they do? But we prevent them from destroying the dynamics of the brain. So that's fantastic. We have a project that makes sense.

So what we're going to try to do now is stimulate one of these and we're going to stimulate a subset of neurons within a community and see what happens to the network. And what we see is absolutely nothing. So we light up this section of the community, which is what this set of points here is. And then we look to see what's the behavior in the other half of that community, which is what that line is, and there's absolutely no response whatsoever.

Well, that's not great because being able to -- being able to send a signal, being able to respond to stimulus is one of the basic functions of a brain. And so it's a very important feature to have.

So what have we done wrong? Well, the thing that we did wrong was balance the network. So all the stuff that I said about the brain needing to be balanced is absolutely true if we want -- if we want perfectly stable dynamics.

If we were to have a -- because we have a perfectly balanced network, we stimulate this part of the community. That sends a signal to the inhibitors. The inhibitors say, oh, I'm going to inhibit this other part of the community. And because we balanced it correctly, it prevents any signal from going to any other neuron. It prevents this coordination. Networks can't respond to stimulus in this fashion.

So what do we do? We go back to this stability equation and say we still need negative eigenvalues. In order for balance to be perfect, we need negative eigenvalues. Why don't we perturb that just a little? What we mean by just a little is that we're going to imagine we perturbed these rates and to do so we perturbed the strength matrix.

And so to first order, we can work out what the perturbation in the firing rate will be in terms of the perturbation of this weight matrix. Right? And we can make that small. It doesn't have to be zero. Just something that's manageable. It's got to be a reasonable number, whatever it is.

So we do that. This particular set of values, in particular we choose ΔW to just be along the diagonal. So we're increasing the positive feedback within each community whenever we do this. And I think we chose ϵ was .1. Small number. You can do it for other numbers and you'll get roughly similar behaviors so long as it's reasonably small.

And this is what we get. What we see it we've increased the Fano factor now. This is what we saw in the homogeneous cluster communities. There's a bit more higher density firing here. Maybe you can convince yourself, once again, that might be your brain tricking you, but it might also be real. I'm not sure.

But it is the case that this is firing in a less coordinated fashion, which is essentially driven by changes in the firing rates. That mean firing rate here is something that fluctuates such that sometimes it's firing slowly, sometimes it's firing quickly.

So this is reasonable. But do we actually see a stimulation? The answer is yes. We stimulate the bottom half and we get a response in the top half. This is something that we were surprised by. The way that we've chosen to balance this prevents us from turning off other parts of the brain. You don't see any dip in the firing rate anywhere else, right?

This means that not only have we increased -- not only have we moved ourselves towards balance and then a little bit away from balance enough that we see reasonable behavior, but we prevent some of this nonphysical (inaudible). And we do this to the first and second community. You see that it works reasonably well. If you stimulate some part, you get a response in the other half. Okay?

For this particular case, we don't see a one to one correspondence, though. Okay? So the response is always going to be weaker than the stimulus.

We can imagine changing the fraction that we stimulate and changing the strength, right? So if I were to grab one neuron and just drive it crazy, that shouldn't necessarily turn on the other 500 neurons. Right? And alternatively, if we were to try to stimulate everyone but with a really, really weak stimulation, nobody would really feel anything. There wouldn't be a lot of change.

That's basically what we see. As long as we're above say 25 to 30 percent and as long as we're in a -- 25 to 30 percent, we're able to -- we can look at the response of the unstimulated region whenever the stimulus is applied -- I'm sorry. This is the response of the stimulated -- the indirectly stimulated region in comparison to this behavior whenever the stimulus is turned off basically. So it is this value divided by that value.

And what we see is that if you stimulate about 30 percent or so, it's possible to get a very clear response from the other half of the community. And this is true for the first largest and the second largest community. This is the relative response of the directly versus indirectly stimulated, and what we see is that as long as we're past that 20 or 30 percent, we can get something on the order of 50 percent relative response rate.

So this is pretty good. We're actually able to see that if we stimulate a community, we get a response within that community.

The last thing we're working on, and this is still unfinished, is to try to understand how hierarchy can play a role. This is a picture of the visual cortex for humans -- is this humans? I don't remember. This is another creature whose name I don't remember.

Essentially down here is the eye. Up here is the brain. All these are related to the visual cortex. And this is a (inaudible) diagram of the stuff that receives the image and how it passes it upstream so that the brain can interpret. The most important feature here is that this is hierarchical. Right? The B3 doesn't send a signal everywhere. It sends a signal downstream mostly to this area. Not everything is connected uniformly.

So if we were to -- so a question is if we take our network but we have a hierarchy built in where some communities like each other more than the rest of the network, are we able to propagate stimulus through that network? And the answer is question is yes.

What we've done here is we've connected the first and the second communities a little more than the rest. Essentially we're breaking balance even worse than we were before. And what we see is if we stimulate 50 percent of the biggest, we not only get stimulation for the rest of the biggest but we also get stimulation for the second largest.

And likewise, if we stimulate 50 percent of the second largest, we not only light up all of the second largest, but we also stimulate the first largest. So it's not a function of size. We are able to share signals between these things and we're able to do so in a way that respects balance enough that we have physical behavior.

Okay. I am pretty close to on time because we started a little bit late.

So just a summarize what we talked about, the surprises today were that balance makes -- balance is an enormously important thing and is not easy to impose correctly for clustered networks. But if you do impose balance, if you figure out a way to do it, it can't send signal. A balanced network can't propagate stimulus in a way that is reasonable.

And so it's necessary to have a nearly balanced but not quite balanced network in order to actually propagate these things. And we can do so pretty reliably. Once we're past about 25 percent it's possible to send signals upstream and downstream or by connecting smaller communities to the larger one.

A few things to -- also I want to mention our solution isn't unique. All we did is divide some stuff by this additional normalizing factor. There are lots of other ways you could imagine to balance this thing. All of them that we've tried now are hard except for that first one, but they're not impossible. We're curious about that.

Neuroplasticity is something that's important. We learn things. The way that we learn things is by the brain rewiring itself effectively. So if you have this unbalanced state of hyperactivity, does that drive you to balance? Essentially because the neurons can rewire themselves, will the brain drive itself to try to go back to near balanced? There are lots of models of plasticity and we're still exploring which ones could do that.

And then finally I'm interested in if you hand me a network, what parts of the brain are important? Where do you stimulate optimally to actually get a response at a particular spot? So all of these things are things that we're still working on.

Okay and that's all I have. So thank you.

(Applause.)

>>: Questions?

>>: Yeah?

>>: So in one of your later slides where you had your partially balanced network and you're stimulating one and getting a response in another -- that one -- is it my imagination or am I seeing a small time delay?

>>: We're still trying to figure that out. That is a feature we see -- sorry, the question was, there seems to be a bit of a lag between, I think is what you're saying, when the onset of the blue and the onset of the red. And if you look closely enough, you can probably convince yourself that there's also a longer lifetime, although I haven't drawn clearly enough. It lives a little bit longer. We're still exploring that. That is something that happens in the homogeneously clustered networks and we have not shown that's something that happens in this case in a provable way.

Other questions? Yeah?

>>: Of a clarification question. Excuse my (inaudible). Whenever you said at the beginning (inaudible) more oxygen is going to that area?

>>: Yeah. So essentially what this -- I have to hop all the way back. What's being highlighted here is the response time of this on the order of seconds. And so it's looking at the average change in oxygen levels at a particular time.

And so I think that is close to answering your question?

>>: I just want to make sure, when a level (inaudible) if a level of the brain is highlighted, it's been stimulated, right?

>>: Yes.

>>: And so that means there are more oxygen in that specific area?

>>: Exactly, yeah. So the base-line, all the gray stuff is base-line where there's not extra flow there but all the reds, I believe -- so I believe, although -- yeah. So the reds correspond to a large amount of activity. The blues correspond to a small amount. So some is flowing from the red regions to the blue regions.

Other questions? Yeah?

>>: So, is there anything (inaudible)?

>>: Yeah. So there absolutely is. So certainly we see that this is the case. So it is the case. So just a little bit of detail about this. If I take the data that exists in these pictures, right, and I draw lines between these things and then I just try to lay the out in a reasonable way, so this doesn't any information about the actual structure of the brain. Just connections between neurons. And then I tried to draw it in a way that makes the lines as short as possible. This is the feature that I see.

Now, so this means that the bilobal structure of the brain is encoded into the brain. What it doesn't tell us is what any of these features o do. So all of the discussions of this is right brain and left brain activity, that can be built into this but we don't have any -- at this level, we don't have any idea how.

Yeah?

>>: You have the groupings, you have some large groupings and small groupings (inaudible). Is there, like, limits on -- have you noticed any kind of limits on how big these groupings get?

>>: So it depends -- it's an interesting (inaudible). There is, of course, theoretical (inaudible) neurons in the brain but that's going be pretty big. Largely the problem that we have is data like this at this resolution is hard to come by. It's hard to track the communication between neurons. All the -- like this picture -- let's see. This picture, so these are individual neurons. These are subregions of the brain. So connections on this scale are hard to come by on the single neuron level. And so it is -- there is no known upper bound, I would argue. And I argue that it's because we have no idea what the upper bound could possibly be because we simply don't have the ability to look at that many neurons simultaneously.

Yeah?

>>: Does whether the body is in a sympathetic or parasympathetic state have anything to do with it?

>>: I don't know -- yeah. Essentially your emotional state has a huge impact. That is something that people can see in EEGs, I'm sorry, and FMRI. It's also -- there are lots of studies looking at the difference between resting sleep and active sleep. And you see radically different patterns. I didn't include the slide but I was thinking about doing. So in rem sleep some regions of the brain are completely silent and other are doing whatever they want. And then during active sleep, the other regions (inaudible) calm down. So you see very different behavior depending on your mental state.

>>: (inaudible).

>>: So the way that they connect to each other, it is not clear -- so basically the way that the brain rewires itself, the current thinking as far as I understand, although I read something two weeks that maybe this wasn't true I I can't say if it's right on-or wrong. The general thinking is the wiring happens whenever you're a child and you develop the wiring you develop. And the strength of the connection here is the (inaudible) the thing that strengthens or weakens building up memories.

>>: (inaudible).

>>: Exactly. Which is how learning stuff works, too. That's -- I believe the current understanding of how the strengthening and weakening happens. And so if you have a traumatic situation, particularly an injurious situation that can completely interrupt this individual neuron's behavior but every neuron that's a few steps away from it. And so you can have radically different behavior due to an injury or due to trauma. Yeah.

Yeah?

>>: Going back to (inaudible) but whenever you said the largest cluster (inaudible) whenever you highlighted (inaudible)?

>>: Yeah.

>>: And you said the second largest also (inaudible).

>>: So we designed it so that it wouldn't. What we designed is that -- let's see. What we did was if you -- imagine the picture. You have all of this stuff and then we cut all of the lines up here. So we're looking at these guys and they're talking to each other but they don't talk up stream anymore than they would by random chance. So the only reason that it speaks to the second -- the first and second largest are affecting each other is because we built it that way. That's a really good we. What we're working on is building up a hierarchy where essentially we have -- so for example these guys talk to each other and these guys talk to each other -- (mic cuts out).

>>: Of a followup question. In a true hierarchy like this, would you end up with a situation where stimulating region one causes region two to stimulate, but (inaudible).

>>: Yes, that is definitely possible. It breaks the symmetry so basically you would need to have arrows on this and that's certainly possible. I'm interested in -- so I've got this stuff. If you imagine that I also have say a different eyeball over here. This is sense of smell or something like that. How can I couple these different stimuli to light up some particular spot? It's a very hard problem because there's an infinity number of ways you can do this. But it's something you can play around with, what signals can propagate pathways?

>>: (inaudible) at Rice University, and a student did -- the son of the student (inaudible) so instead of (inaudible) she's doing the speaking. So when you tried to talk, like, right now I'm thinking and then I compose and then I speak out.

>>: Yeah.

>>: So it activate (inaudible) so actually, the input into the (inaudible) of your hierarchy. So it's a sequence of actions. And also, I have seen a show, like, (inaudible) it was talking about in our (inaudible) one is the video detection. They say well, I see you. You look like my brother but then I don't feel like you're my brother. You're pretending to be my brother. So there's an emotional part of the hierarchy that is negative association. So there was a guy, he had brain damage and it broke the emotional part. So he keep running away from the home. He said I see my mom, I see my house, they look like it but I don't feel it. So I don't think they are really my relative and my home. And he run away. So maybe in your hierarchy they are (inaudible) you have visual and then you have emotional --

>>: That's a hundred percent the case. This is salamander again and basically they're showing this under the assumption that salamanders recognize salamanders. But the input is the (inaudible) and the inference that that means something has to come from the inside of the network. And so if you were to break some wiring up stream, you'd still see the signal. You'd still see your parent's house or whatever. But if you traumatically rewire regions of the brain, you can completely destroy the ability to interpret what you're seeing in a physically meaningful way.

>>: (inaudible) one is the sequence of event and another one is (inaudible) the original input and (inaudible) how you progress with the hierarchy (inaudible).

Another thing S there was another (inaudible) so what they did is (inaudible) so, like, for example, you'll have some sort of distortion just because (inaudible) so I guess you can study that part with --

>>: Yeah. Yeah. So because you have -- yeah. Because you have these downstream structures that interpret the thing that you're shown, if you're shown something that isn't - - that your eye doesn't respond to well, so if the physical interaction between the light that you're seeing, the retina, if that -- if the downstream interpretation network isn't robust to what it is you're seeing, it's going to do something with it. In the same way that I was able to convince at least some of you that there's structure in this, there isn't. And the reason is your brain tries to interpret what it is that you're seeing. So, yeah, if you're shown optical I illusions, the brain does something with the thing you're seeing.

>>: One example would be (inaudible) and then everybody would say well that is (inaudible) exactly the same intensity. But the background is different, so it (inaudible).

>>: Yeah. Yeah. I'm color blind so those don't affect me, but -- (laughter) okay. Any other questions? Okay.

>>: Let's thank our speaker.

(Applause.)

>>: And thank you at home.

(End of class)

DISCLAIMER

THIS TRANSCRIPT IS A ROUGH DRAFT TRANSLATION FROM THE CART PROVIDER'S OUTPUT FILE. THIS FILE MAY CONTAIN ERRORS. THIS TRANSCRIPT MAY NOT BE COPIED OR DISSEMINATED TO ANYONE.