

Good morning. All right. So, it was a lot of fun to tell you last time about the work that's going on here at MIT in genomics.

It's, it's just fun. I mean it's fun to share what's going on and it's fun that you guys can understand what's going on already with simply one semester of biology, which I think is really cool.

What I would like to do today is talk about another exciting area, one I don't work in, but one that I know many people are interested in.

And I would like to lay the foundations for neurobiology.

Neurobiology is an incredibly interesting area.

I confess it's the subject that got me interested in biology in the first place, even though I don't work on it itself.

And, I mean, who can't be interested in the brain.

But if you want to take on the brain, you really have to understand its working components in some retail detail.

And so what we're going to focus on in this course for the next three lectures is the specific molecular mechanism by which nerve cells are able to transmit signals down their length, how they're able to transmit signals from one cell to the next, and then how they're able to change their properties over time, that is learn and remember.

So, the fundamental unit of neurological processing is a nerve cell which goes by the name a neuron. You have approximately ten to the twelfth neurons.

How many bases are in the human genome?

About three times ten to the ninth.

So, you have a lot more nerve cells than bases in the genome.

So, it's unlikely that all of the wiring diagram is completely specified in the sequence of the genome simply because we have a lot of them here. That's just a minor side point. And what does a typical neuron look like? So, neurons and connections.

So, a neuron. Well, a typical neuron, your ten to the twelfth neurons makes contacts with other neurons. It might receive contacts.

It might get contacts from, oh, ten to the third other neurons.

And sends signals to ten to the third other neurons.

So, that's a lot of connections if you figure ten to the third times ten to the twelfth. Ten to the fifteenth connections in this circuit diagram. Here's sort of a, kind of a picture of a nerve cell -- -- making connections to another nerve cell. Making connections to maybe a muscle.

Oh, and what activates this nerve cell? Well, maybe in your eye we have a photoreceptor cell.

And that photo receptor cell will synapse upon a first neuron which will synapse on a second neuron which will synapse on your muscle.

So maybe you'll see something, send a signal, send a signal, activate your muscle to pick it up. Needless to say, it's pretty much more complicated than that because it will take more than those two neurons to figure out this is a piece of chalk and how to coordinate that whole motion and all, but you get the idea. So, let's just take a look at some of these pieces here. What kinds of receptors might we have? We might have receptor neurons that receive light -- -- called photoreceptors.

And these photoreceptors in your eye are an extraordinary piece of engineering. Do you know how sensitive a photoreceptor can be?

What's the absolutely minimum possible detectable unit of light?

One photon. It turns out your photoreceptors can, under appropriate circumstances, detect a single photon. Not in the bright light but in dark adapted conditions, you actually have one photon sensitivity. Very impressive.

Under appropriate conditions, mind you. Sound receptors. You've got sound receptors in your ear and they are beautiful.

We're not going to talk about them at any length, but there's little flappy, these little spiky things going along in your ear and they can translate vibrational energy coming from your ear, hurting your eardrum, being translated into a vibration into the fluid in your ear into a physical motion of these little receptors there into an electrical motion, into an electrical signal that goes into your ear.

So, all of that, all of that's pretty impressive stuff. We're not going to talk about the details of it, but I invite some of you who want to learn more about this, particularly MIT students I think find receptors really quite remarkable kinds of devices. We're going to focus, though, more on this intermediate step here of a generic neuron.

So, here's my generic neuron. And the first thing that has to be said is there are no generic neurons. Neurons are all very specific, but I'm just going to use a generic neuron for the moment.

And the important features of a generic neuron here, it has these funny little processes on its cell body that are called dendrites. That's where it's going to receive signals.

Processes are the name for things that stick out.

So, these processes are called dendrites.

Here's our nucleus of our cell. Here's our cell body. We will label this your cell body with a nucleus in it.

Then we have this very long process here called an axon.

The axon is the wire here. This region here is called the axon hillock and it looks like a little hill. That's what hillock means, it's a little sort of hill-like region here.

And it's where all the electrical signals from the cell body are integrated and a, quote, decision is made about whether to fire. And the neuron will then fire a signal down the length of its axon and then it will get to the end here where we have these terminal processes that'll be synapsing on other cells here. These are called nerve terminals.

And each will make connections either with other dendrites or maybe with a muscle, although you wouldn't see both of those occurring, so I'm just drawing this for illustration purposes here. And it will send a signal. And these points of contact are called synapses. So, nerve or muscle. OK? So, that's your general picture there. You have a synapse which transmits.

This is an electrical signal. This is typically a chemical signal.

And, as you might imagine, the pre-synaptic cell is this guy, the post-synaptic cell is that guy. So, I'll say pre-synaptic and post-synaptic. And what neurobiologists want to do is understand how does this all work? How does the initial signal impinge upon the dendrite from a synapse? How does that signal get collected and integrated into a decision about whether to fire an action potential that runs along the nerve? And then how does that get transmitted from an electrical signal back into a chemical signal which then restarts an electrical signal in the next cell, etc., etc., etc? Those are some of the kinds of questions we want to ask. Now, I said generic neuron. Of course there's really nothing generic about it. Some neurons are very tiny.

Some neurons are very big. What's the typical size of an ordinary eukaryotic cell in a liver cell? Ballpark. Ten microns.

So, a typical eukaryotic cell might be -- -- and ten microns.

But a neuron can be anywhere from that size, ten microns, all the way up to three meters. Where's the longest neuron known?

Squid giant axon. That's possible, actually. I'm not sure how big squid, I mean there are these giant squids and they may qualify.

The longest one that I know of is in a similar sort of setting.

It's in the giraffe. It's a motor neuron in the neck of the giraffe which runs three meters, a single cell. But these monstrous squids that, at least in stories, glom onto boats might actually get up there. Will somebody check the squid, the largest squid on record here? I might need to revise that. But the largest one I know about definitively is the giraffe that has single neurons three meters long. Ten foot long single cells. So, it's really very impressive what a cell can do there.

So, I'll put down the giraffe but we'll check out on the squid there.

Giraffe neck. OK. So, what are the kinds of questions that we might want to ask? Well, the kinds of questions we might want to ask about this process, and I'm not going to be able to answer all of them but I'll get them up here. So, how do receptors transduce signals -- -- from the outside world?

Light, sounds, touch, etc.

That's a good question. People know a lot about it.

How do electrical signals propagate along the neuron?

Along an axon?

People know a lot about that one, and that's what we're going to discuss today.

How do signals transmit across a synapse?

We'll talk about that, some today and some next time. And how do they transmit them to effector cells? So how do signals transmit specifically to effector cells like a muscle?

We'll talk about that.

Then how does the pattern of connections -- -- give rise to a computation?

That's pretty tricky.

We'll talk just a teeny little bit about the simplest kind of computation, but the complex computations of recognizing that that's somebody's Volkswagen or that that's your grandmother or something are still beyond us today. But we know things in between that.

How does this pattern of connections arise during development?

That's a fascinating question. And a lot is known about it, but we're not going to talk about it in this course.

How does this pattern of connections get modified by experience?

That's something that we'll mention briefly on Friday that a colleague of mine who's very knowledgeable about this is going to give.

And that's learning and memory. And then how does this all give rise to consciousness?

And we haven't got the first clue. We have no idea. It's fascinating.

I digress for a second. A famous biologist, J.B.

. Haldane in the mid-century, middle of the 20th century wrote a final exam to be given in the year 2000 or so. And on the final exam he had about 20 odd questions. And if you go through the exam virtually all of those questions could indeed be answered by a student in the year 2000, except for the question that says consciousness arises on embryonic day 18, example.

And we have no progress toward that particular question.

That was one that he completely was way off in terms of being able to predict that we'd make any progress on. Maybe he meant it as a joke.

Anyway. So, let's now dive into the specifics.

So, electrical signals in axons. Here is an axon. We're going to give a very close up view. Here's the lipid bilayer. I'm just going to take a cross-sectional view of my axon here.

That's the lipid bilayer. It turns out if I stick an electrode into an axon and I measure the voltage gradient from the outside of the cell to the inside of the cell, what I'm going to find is that the inside of the cell is negative compared to the outside of the cell. There is an electrical potential across the plasma membrane of this axon. This plasma membrane you remember, you know, you all remember about this.

This is, what, three nanometers wide. There's an electrical membrane, electrical potential across this equal to

about minus 70 millivolts.

Now, come on, minus 70 millivolts is pretty trivial, right? You go to the store, you buy a battery, what has it, what's it got like one and a half volts or something?

It's minus 70 millivolts like who cares, right? Except it's minus 70 millivolts across three nanometers. What's the electrical field of minus 70 millivolts across the incredibly tiny distance of three nanometers?

Well, the electrical field strength minus 70 millivolts over three nanometers is about 200, 00 volts per centimeter.

That's a very impressive number, 200,000 volts per centimeter.

Suppose I could arrange to change the electrical potential of a cell from minus 70 millivolts inside to, I don't know, minus 70 millivolts outside. That would be a swing of 400,000 volts per centimeter.

Do you think that a protein which had an alpha helix that had a dipole moment on it could feel a change of 400,000 volts per centimeter?

You bet. Some alpha helix that had some dipole moment on it would swing wildly in the presence of a change of 400,000 volts per centimeter.

That's the key to how things work down there is this tiny little minus 70 millivolts. It's a huge potential gradient.

And if we can change that we can swing the shapes of proteins quite dramatically. OK? That's the principle.

Now, it turns out that if you take this electrode and use it to change the electrical gradient, what you will get is the following bizarre and fascinating behavior. So, take my axon here, I use this, and what I do is I start off, here's zero, at minus 70, minus 70, this is the electrical potential.

If I use the electrode to change this to about minus 50 or so then all by itself, with no further input on our part, the cell wildly shoots up to plus 50 before rapidly coming back down and reestablishing itself at minus 70. So, this depolarization slightly shifting it away from being as polar as it is, it has a polarity of minus 70, if we depolarize it, make it less polar, make it less negative, all by itself the cell executes something called an action potential. This action potential involves rapidly changing to being positive inside the cell instead of negative and then restoring itself.

And, as you would imagine, this massive change in the field has a huge effect on proteins in the membrane. This is called depolarization phase. Not shockingly this is the repolarization that occurs there afterwards. So, it goes from minus 70 millivolts up to about plus 50 millivolts there.

All right. How does this happen? That's my job to explain right now.

Well, the way this happens, what kind of a beautiful piece of engineering explains this? Well, the first thing you need to know is that there are some concentration gradients set up in the cell. So, the concentration gradients in the cell are as follows. There are certain ions, sodium. Sodium happens to be low inside the cell and high outside the cell. The concentration of sodium inside the cell is about 12 millimolar. Whereas, outside the cell is about 145 millimolar. By contrast, potassium ions are high inside the cell, 139 millimolar. Whereas, only four millimolar on the outside of the cell. Calcium ions also have a gradient across the cell. They are virtually rare.

0.1 micro, not milli, micromolar inside the cell and 2.

millimolar outside the cell. So there are very big differences in the concentrations. Let me try to write these concentration gradients. If this is the outside and this is the inside, sodium has a concentration gradient higher on the outside than the inside. Potassium has a concentration gradient that's higher on the inside than the outside.

Calcium has a concentration gradient higher on the outside than the inside. This turns out to be the force that drives this action potential, is that energy has been stored up in these concentration gradients. A concentration gradient is a form of energy.

Why is that? Well, if there was no membrane there, what would be the concentration on the inside and the outside?

The same. There wouldn't be an inside and outside but, you know, so let's imagine. Suppose I just drilled holes in the membrane. I get in there vrm, vrm, vrm, drill a lot of holes, it will equilibrate and, you know, because just by diffusion there, we'll go to the entropy, you know, entropy will set in and it will be equal concentrations. If I want to establish a concentration gradient then, I have to work against entropy.

That takes energy here. And that is a form of energy then.

I had to put work into moving ions around in order to establish a concentration gradient. Well, who's in charge of carrying out that work? Who is it that sets up these concentration gradients in the cell? Membrane proteins. Certain membrane proteins. In particular, membrane transporters are involved. So, we're going to talk about transporters and then we'll talk about some channels.

All right. The first one.

There is an ATP-driven sodium potassium pump.

Any ideas what an ATP-driven sodium potassium pump does? It sets up, how does it set up a concentration

gradient? So, what does it pump? It pumps a sodium. Does it pump a sodium in or out?

It pumps a sodium out. It pumps potassium in. And it does it as an even exchange. Why would it like to do it as an even exchange, one sodium for one potassium?

Charge conservation, exactly. So, there's no electrical work to be done if it swaps them one for one. And is there work, though, to be done? Yes, there is, because if I'm going to pump sodium out of the cell, I'm doing it against the concentration gradient. Or at least once I get going I'm doing it against a concentration gradient. So, energy is needed to move a sodium out of the cell against its concentration gradient. Energy is also needed to move a potassium into the cell against its concentration gradient.

And where do we get the energy? ATP. So, an ATP is burned in order to do that. This guy gets called an antiporter sometimes because it's an anti-transporter or something like that. OK. The pump then is ATPase as it uses the energy from an ATP to exchange this. Good. Next. There is an ATP-driven calcium pump or transporter.

And, as you might imagine, what does it do? It transports a calcium ion out of the cell, and it happens not to do that in exchange for any other ion. And it, too, is driven by ATP.

So, we go pump, pump, pump, pump, pump, pump, pump, pump, keep going.

Will this thing be able to keep going forever, set up arbitrarily large concentration gradients? Why not?

Well, what is nothing leaked back in?

Could we, could we keep going? It gets harder and harder to put stuff out because you're working up against a bigger and bigger concentration gradient. And the ATP is only going to give you so much energy. So, at a certain point, the burning of that ATP or however many ATPs it uses for its specific mechanism won't suffice. So, there's going to be some natural upper bound to how far it could go in setting up a gradient because there's a nature amount of energy it can spend. OK.

It's useful to think about these gradients as, you know, work that you've got to do and the hill gets steeper and steeper.

All right. So that could, in principle, set the electrical gradient, the concentration change across the cell.

But there's one other important component that we have to think about. And that is something called a resting potassium channel. So, what is a resting potassium channel? The resting potassium channel is a protein that sits in the membrane and it's got a hole in it, a pore.

And that pore here is designed so that while a sodium can't escape through that pore, and you can imagine there's some little bit of clever molecular architecture to make the sodium atom not be able to, sodium ion not be able to get out but a potassium ion be able to get out, a potassium ion can get out.

This is a completely open door that allows potassium ions to escape.

Now, isn't this stupid? We just spent all this ATP getting potassium in and sodium out, and here I go opening the door for potassium saying you're free to leave despite all the work we put into bringing you into the cell. That makes no sense. All the potassium is just going to rush out.

Can all the potassium rush out?

The concentration gradient goes which way? It's more inside, more potassium inside, less outside, so the potassium is going to rush out. Yes. Ooh, electrical gradient.

If there's an electrical, so maybe at the beginning there's no electrical gradient, so potassium rushes out.

What does it do? It makes it more positive on the outside.

Now the next potassium wants to rush out.

It's going to, it's going down its concentration gradient but it's going up a teeny weenie little electrical gradient.

Now that gets out there, and what does it do to the outside?

Makes it more positive. So, the third potassium makes it even more positive. And as more and more potassiums rush out, the outside becomes more positive. And every potassium now has to do work to get up that electrical gradient.

But, of course, it's still got a concentration gradient pushing it out. So, will they keep going forever?

No, they're going to balance each other. There will come a point when the concentration gradient, the force driving potassium out due to the concentration gradient is offset by the electrical gradient pushing potassium in, or keeping potassium in.

So, what happens is potassium goes out, rushes out, or goes out, leaks out, potassium leaks out which now sets up an electrical gradient.

And what happens is the cell becomes more positive on the outside, more negative on the inside, and an equilibrium potential is reached, equilibrium is reached.

So, an equilibrium electrical potential -- -- is reached when the concentration gradient exactly offsets, when it balances the electrical gradient.

Any guesses to what that electrical concentration gradient is?

Minus 70 millivolts. That's where minus 70 millivolts comes from.

It comes because that's the concentration gradient at which potassium going up that gradient just offsets the, yes?

Well, because, oh, it doesn't. The electrical potential wants to keep it in. The concentration wants it out. The electrical wants it in.

That's why we have an equilibrium. Exactly. It's those two balancing forces. Very good. So, how much potassium, by the way, has to rush out to set this up? It turns out a trivial amount. It turns out that about ten to, one part in ten to the sixth of all the potassium ions leaving sets up a gradient of about minus, of about five millivolts, minus five millivolts. One part in a million of the potassium ions will suffice to set up a concentration gradient of minus five millivolts.

So, all I've got to do is set up about, let's see, I don't know, less than, I don't know, maybe about one part in ten to the fifth of all the potassium ions leaving.

A tiny contribution of the concentration.

Changing the concentration by one part in ten to the fifth will get me minus 50 millivolts already. So, what's very interesting is teeny amounts of ions set up a very big electrical gradient and have no real effect on the concentration, but they have a big effect on the electrical gradient. So, when I say potassium leaks out, it only takes a teeny bit of potassium to leak out in order to accomplish this. OK, guys. Now let's get ready to make an action potential. Here it goes. So, the action potential mechanism.

We've set up our resting potential by transporters, by this open channel. Now, the first thing we have is a new kind of membrane channel. We have a voltage gated sodium channel.

Check this out. This guy here is a channel that's closed. It's closed. Except, so at minus 70 millivolts it's closed. But if I can transiently depolarize the cell to minus 50 millivolts it opens and it admits sodium.

So, when I get to minus 50 it opens. Now, what will sodium do when that door opens? Let's see. Electrically what would sodium, what would sodium do?

So what, well, how about concentration?

From the point of view of concentration what would sodium, oh, what did I just draw here? Ooh, forget my error.

Electrically what would sodium like to do? It would like to come in, wouldn't it? Because it's negative inside and sodium is positive.

But from a concentration gradient what would it like to do?

Come in also. So what's stopping it? Nothing. We've got the ion that would like to come in electrically and would like to come in from the point of view of concentration.

So what happens? It comes in. And it comes rushing in. Now, what happens as it comes rushing in? What will it do to our electrical potential?

We went from minus 70, we got it transiently up to minus 50, and as it comes in what does it do to our electrical gradient?

Brings in positive charge. The electrical gradient goes towards zero. Does it stop at zero? No, because now, even as the cell becomes positive in the interior, sodium still wants to keep coming in because of its concentration gradient.

It now has to do some work against the electrical gradient that's set up, but it keeps going and going and going. Does it go on forever?

No, because eventually it reaches a point where the electrical gradient, positive now inside the cell, offsets the concentration gradient and it'll stop and it'll reach a new equilibrium potential which turns out to be at about plus 50. Amazing. Opens the door.

Sodium comes rushing in down its concentration electrical gradient, shifts the electrical gradient from being negative to positive, and eventually it slows itself down because it's now going against an electrical gradient. But that's not the end of the story because what there also is, there are two other interesting functions. One, after a certain amount of time, after a very brief time interval -- -- this channel also has the property that it closes.

So now when it closes, what's going to happen? The pumps will start working again and reestablish our negative 50. It's going to take too long, though.

I mean it'll happen. You'd get back to minus 50, but it's going to take a long time. So it's good that the channels, that the voltage gated sodium channel closed, but even better nature has arranged to have a voltage gated potassium channel.

And what happens to this voltage gated potassium channel is around plus 30 millivolts it opens and admits potassium.

Now, instead of having to wait for the relatively slow ATP-driven pumps, what happens when I admit potassium? Well, if the cell is positive on the inside, negative on the outside, what happens to our potassium?

Potassium starts coming out because there's more potassium on the inside and there's also a favorable electrical gradient.

So, potassium explosively starts rushing out and rapidly reestablishes the resting potential. All this happening in a millisecond.

It's very impressive. So, in something like one millisecond we open up the potassium channels. Now, you have to ask how did it happen that the channels got open in the first place? How did we get to minus 50? That's the work of these dendrites. The dendrites integrating a signal from all the things impacting on it will get the cell to minus 50, but once the cell gets to minus 50 at its axon hillock, bingo, the action potential fires, the cell zips up to plus 50 in a big rush of sodiums and then zips down to minus 70 again in a big rush of potassiums. Yes?

Yup. Remember I said it was only like one part and ten to the fifth of the ions? It's a trivial actual, the number of ions necessary to set up these electrical things is such a tiny fraction of the concentration that through this whole thing those concentrations don't noticeably change. That's what's so cool, is there are different regimes, right? And they do change. They change by one part and ten to the fifth of those numbers.

It's really cool, isn't it? We haven't screwed up our concentration gradient? We moved tiny numbers of ions to accomplish all this. This is very clever engineering.

It's very clever engineering. Now, how do we manage to transmit this signal down a cell? Well, let's talk about how we propagate. I think this is really cool.

This is one of the truly great mechanisms that was invented.

Transmitting an action potential.

Suppose I transmit my action. Suppose I have a neuron here and over here, at the axon hillock, I transiently depolarize to minus 50.

Then I fire and fire.

Well, what happens is this part of the cell was originally minus, minus, minus, plus, plus, plus, it becomes positive

temporarily, right? It now becomes -- -- plus, plus, plus. So, let's go minus along the whole cell here. Plus, plus, plus, plus, plus, plus. Now, some little patch of membrane at the beginning of the axon has become positive inside, right? Because whatever local affect here caused this to flip. If this becomes positive at this part of the membrane, what happens to the negative charges over here a little bit further down the axon? Some of them will be pulled over to the positive. Ho-ho. Some negative charges get pulled over. Well, what happens when some of those negative charges get pulled over to my minus 70 millivolts?

Is it as negative as it was before? No. It becomes minus 50 millivolts.

Oh. What happens when this becomes minus 50 millivolts?

Fires an action potential. So, what happens is, here's my axon, I fire an action, I have an action potential here, and in the course of that I pull over some negative charge.

That, of course, transiently depolarizes here and causes an action potential to fire. Then, of course, this becomes positive which pulls over some negative charge which causes the action potential to fire here, etc., etc., etc. So, if I can manage to get the thing started with my dendrites causing a transient depolarization from minus 70 to minus 50 then the action potential itself will draw over some charge and start the process at the next patch of membrane, the next patch of the next membrane, the next patch of the next membrane and the next patch of membrane.

And it gets all the way down to the end. It's a brilliant mechanism.

You can calculate, based on the diffusion coefficient of ions, how long it will take to transmit that signal. And it's OK but not good enough. I'd like it to go faster.

How can I make it go faster?

How about put an insulator around it? It turns out that with one little trick, I can speed this up dramatically. Oh, sorry. Sorry about that. The trick is suppose I were to wrap some insulator around the axon that would make this such that it could not, that it was not an electrical contact outside, that it was only an electric contact with the outside solution here, here, here. Then, when the action potential fires here, it's going to pull charges over, not from the tiny little patch of membrane here, but the charges are going to have to come from here, it turns out, because that's where I'm going to next feel my action potential. And so what'll happen is that the action potential, if this stuff is electrically insulated, will only be happening at this little gaps between the insulation. And I can dramatically speed up the electrical transmission if I'm willing to insulate the wire.

Because then there's much less leakage along the way.

So, it turns out that there are special cells that wrap themselves around the axon that are called Schwann cells.

The Schwann cell wraps itself around and around and around and squeezes out all of its cytoplasm leaving only a lipid bilayer.

Lipid bilayers wrapped around and around and around make great insulators. This is called the myelin sheath of a nerve. These myelin sheaths allow the transmission of this action potential to proceed about a hundred times faster. Very clever.

Now you ask me how do I know this matters? I mean you can calculate that it should go about 100 times faster.

But how would you really know that it mattered. Try taking them off.

Unfortunately, there is a disease that takes them off.

Some human patients make autoimmune reactions against their own myelin.

It attacks their own myelin and leads to the demyelination of the nerves. This disease is called multiple sclerosis. Multiple sclerosis involves an autoimmune attack on your very own myelin sheaths, and the effects, the very serious effects that multiple sclerosis can have on an individual come from the greatly diminished, hundred-fold diminished conduction velocity of electrical signals down motor neurons along these large distances. So, this is, in fact, the basic electrical setup. We have taken a situation where we don't have electrical wires at all but have concocted, through chemistry, a powerful way to create signals.

First, by setting up a resting potential using pumps and an open potassium channel. Then by setting up this explosive mechanism where a little change in voltage which, of course, has a huge change in field, swings open a sodium channel, shuts it, swings open a potassium channel, and then cleverly recruits charges from down the neuron and sends the signal.

Next time we shall talk about what happens when the signal gets all the way to the other end and it has to talk to the next neuron.

Till next time.