MICHALE FEE: So today, we're going to introduce a new topic, which is related to the idea of fine-tuning curves, and that is the notion of receptive fields. So most of you have probably been, at least those of you who've taken 9.01 or 9.00 maybe, have been exposed to the idea of what a receptive field is. The idea is basically that in sensory systems neurons receive input from the sensory periphery, and neurons generally have some kind of sensory stimulus that causes them to spike.

And so one of the classic examples of how to find receptive fields comes from the work of Huble and Wiesel. So I'll show you some movies made from early experiments of Huble-Wiesel where they are recording in the visual cortex of the cat. So they place a fine metal electrode into a primary visual cortex, and they present.

So then they anesthetize the cat so the cat can't move. They open the eye, and the cat's now looking at a screen that looks like this, where they play a visual stimulus. And they actually did this with essentially a slide projector that they could put a card in front of that had a little hole in it, for example, that allowed a spot of light to project onto the screen. And then they can move that spot of light around while they record from neurons in visual cortex and present different visual stimuli to the retina.

So here's what one of those movies looks like. So you're hearing the actual potential of a neuron visual cortex. So you can see the neuron generates lots of spikes when you turn a spot of light on a particular part of the visual field.

So they will basically play around with spots of light or bars of light and see where the neuron spikes a light, and then they would draw on the screen-- I think they're going to draw in a moment here-- what they would call the receptive field of the neuron. So you can see that this neuron responds with high firing rate when you turn on a stimulus in that small region there.

Notice that the cell also responds when you turn off a stimulus that is right in the area surrounding that small receptor field. So the neuron has parts of its receptive field that respond with increased firing when you apply light, and they also have parts of the receptive field that respond with higher firing rate when you remove light. That was actually a cell, I should have said, that's in the thalamus that projects to visual cortex. So that was a thalamic neuron. Here's what a neuron in cortex might look like.

So they started recording in the thalamus. They saw that those neurons responded to spots of light in small parts of the visual field. They were actually recording from neurons in the visual cortex. They got kind of-- they couldn't really figure out what the neurons were doing, and they pulled the slide out of the projector, which made an edge of light moving across the visual field.

And the neuron they were recording from at that moment responded robustly when they pulled the slide out. And they realized, oh, maybe it's an edge that the neuron is responding to. And so then they started doing experiments with bars of light.

Here's an example. So you can see the neuron responds when you turn a light on in this area here. But is responds when you turn light off in this area here. And so you can see they're marking a symbol for positive responses, positive responses to light on here, and negative responses or increased firing when you turn the light off.

So there's different parts of the receptive field that have positive and negative components. But you can see that the general picture here is that the process of finding receptive fields at this early stage was kind of random. You just tried different things and hoped to make the neurons spike. And we're going to come back to this idea of finding receptive fields by trying random things, but in a more systematic way, at the end of the lecture today.

So here's what we're going to be talking about. So you can see that Hubel and Wiesel were able to describe that receptive field by finding positive and negative parts and writing symbols down on a screen. We're going to take a more mathematical approach and think about what that means in a quantitative model of how neurons respond to stimuli.

And the basic model that we'll be talking about is called an LN model, linear/nonlinear model. And we're going to describe neural responses as a linear filter that acts on the sensory stimulus followed by a nonlinear function that just says neurons can only fire at positive rates. So we're going to have our neurons spike when that filter output is positive, but not when the filter output is negative.

And we're going to describe spatial receptive fields as a correlation of the receptive field with the stimulus. And we're also going to talk about the idea of temporal receptive fields, which will be a convolution of a temporal receptive field with the stimulus. So the firing rate will be a convolution of a receptive field with the temporal structure of the stimulus.

We're going to then turn to the-- combine these things into the concept of a spatial temporal receptive field that simultaneously describes the spatial sensitivity and the temporal sensitivity of a neuron, as an STRF, as it's called. And we'll talk about the concept of separability. And finally, we're going to talk about the idea of using random noise to try to drive neurons, to drive activity in neurons, and using what's called a spike-triggered average to extract the stimulus that makes a neuron spike. And we're going to use that to compute-- we're going to see how to use that to compute a spatial temporal receptive field in the visual system or a spectral temporal receptive field in the auditory system.

So let's start with this. What are spatial and temporal receptive fields? So we just saw how you can think of a region of the visual space that makes a neuron spike when you turn light on or makes a neuron spike when you turn light off.

And at the simplest level, you can think of that in the visual system as just a part of the visual field that a neuron will respond to. So if you flash of light over here, the neuron might respond. If you flash of light over here, it won't respond.

So there's this region of the visual field where neurons respond, but it's more than just a region. There's actually a pattern of features within that area that will make a neuron spike, and other patterns will keep the neuron from spiking. And so we can think of a neuron as having some spatial filter that has positive parts, and I'll use green throughout my lecture today for positive parts of a receptive field, and negative parts. And this is a classic organization of receptive fields, let's say, in the retina or in the thalamus, where you have an excitatory central part of a receptive field and an inhibitory surround of the receptive field.

So we can think of this as a filter that acts on the sensory input. And the better the stimulus overlaps with that filter, the more the neuron will spike. So let's formalize this a little bit into a model.

So let's say we have some visual stimulus that is an intensity as a function of position x and y. We have sum filter, G, that filters that stimulus. So we put the stimulus into this filter. This filter, in this case, just looks like, in this case, an excitatory surround in an inhibitory center.

That filter has an output, L, which is the response of the filter. Then we have some nonlinearity. So we take the response of the filter, L, we add it to some spontaneous firing rate, and we take the positive part of that sum and call that our firing rate. So that would be a typical output nonlinearity. It's called "a threshold nonlinearity," where if the sum of the filter output and the spontaneous firing rate is greater than 0, then that corresponds to the firing rate of the neuron.

So in this case, you can see that as a function of L-- I should have labeled that axis L-- as a function of L, when L is 0, you can see the neuron has some spontaneous firing rate, R naught. And if L is positive, the rate goes up. If L is negative, the rate goes down until the rate hits 0. And then once the neuron stops firing, it can't go negative. So the neuron firing rate stays at 0.

And then once you have this firing rate of the neuron, what is it that actually determines whether the neuron will spike? So in most models like this, there is a probabilistic spike generator that is a function of the rate output of this nonlinear output. It's basically a random process that generates spikes at a rate corresponding to this R.

And in the next lecture, we're going to come back and talk a lot more about what spike trains look like, how to characterize their randomness, and what different kinds of random processes you actually see in neurons. A very common one is the Poisson process, where there's an equal probability per unit time of a neuron generating a spike, and that probability is controlled by the firing rate. We'll come back to that and discuss it more. Any questions? Yes, [INAUDIBLE]

AUDIENCE: Is something biologically [INAUDIBLE] something like if the overlap [INAUDIBLE] it's just it's more excitatory?

MICHALE FEE: Yeah. So we're going to come back in, I think, a couple lectures where we're going to talk about exactly how you would build a filter like this in a simple feed forward network. So at the simplest level, you can just imagine you have a sensory periphery that has neurons in it that detect, let's say, light at different positions.

Those neurons send axons that then impinge on, let's say, the neuron that we're modeling right here. And the pattern of those projections, both excitatory and inhibitory projections from the periphery, would give you this linear filter. And then this nonlinearity would be a property of this neuron that we're modeling. Yes, Jasmine?

AUDIENCE: [INAUDIBLE]

MICHALE FEE: Great, exactly. So it's going to turn out that we're going to treat this as a linear filter. The output of this filter will be calculated for a spatial receptive field as the correlation of this filter with the stimulus.

But in the time domain, when we calculate a temporal receptive field, we're going to use a convolution. And we'll get to that in a minute. That's the very next thing. Great question. Anything else?

So that's called a LN model. I should have put that on the slide-- linear/nonlinear model. So let's describe that mathematically. So let's say we have a two-dimensional receptive field. We're going to call that G of x and y.

So remember, we had intensity as a function of x and y. There's our stimulus input. And we're going to ask, how well does that stimulus overlap with this receptive field? And we're going to describe the receptive field as a function on this space, x and y.

And our linear model is going to be how well the stimulus matches or overlaps with the receptive field. And we do that just by multiplying the receptive field times the stimulus and integrating over x and y. x and y, just think of it as a position on the retina. So let's look at this in one dimension.

So remember, this was a receptive field that has a positive central region and an inhibitory surround. So if we just take a slice through that and plot G as a function of x, you can see that there is a positive central lobe and inhibitory surround, inhibitory side load. That's a very, very common receptive field early in the visual system, in the retina and in the lateral geniculate nucleus.

So in one dimension, we just take this receptive field, G, multiply it by the stimulus pattern, and integrate over position. That is L. That's the output of the linear filter. We're going to add that to a spontaneous firing rate, and that gives us the firing rate of our neuron.

And you can see that that's like-- that this product, an integral over x, is just like a correlation-- G of I times intensity of I summed over I. So let's walk through what that looks like. So here's G of x, the receptive field. Let's say that's our intensity profile.

So we're going to have a bright spot of light surrounded by a darker side lobe. So the way to think about this is, in visual neuroscience experiments, usually the background is kind of gray. And you'll have bright spots, like here, and dark spots, like there. And the rest will just be gray. So that's how you get positive and negative intensities here, because they're relative to some kind of gray background.

And so now we can just multiply those two together. And you can see that when you multiply positive times positive, you get positive. And when you multiply negative times negative, you get positive. And when you integrate over position x, you get a big number. You get some positive number.

So that neuron, that stimulus, would make this a neuron with this receptive field likely to spike. Let's consider this case. Now, instead of a small spot of light centered over the excitatory lobe of the receptive field, you have a broad spot of light that covers both the excitatory and inhibitory lobes of the receptive field. What's going to happen? Yeah?

AUDIENCE: [INAUDIBLE]

MICHALE FEE: Yeah. You're going to get a positive times of positive here, and a negative times a positive. And so you're going to get negative contribution from the side lobes, and those things can exactly cancel out when you integrate over position. And so you can get a very small response.

If you have-- I'm not going to go through this example, but it's pretty obvious. If you were to have light here, and then dark here, and then light there, you would have only these negative side lobes activated. You would have no contribution from this excitatory lobe, and the integral would actually be negative. And so the firing rate of the neuron would go down.

Do you remember seeing those different points in that first movie where you saw the donut of light turning on? The neuron kind of shuts off when you turn that light on. Yes?

AUDIENCE: So I'm just curious what we're signing up 0 [INAUDIBLE].

MICHALE FEE: Yeah. So 0 is just this gray background. It's some intermediate level of light intensity-- pretty straightforward. So that's a spatial receptive field right there.

We refer to this correlation process, this linear filter as linear, because you can see that if you put in half the light intensity, let's say, you get half the product. And when you integrate, you get half the neural response. If you take the stimulus and you cut it in half so that you only apply light and dark to half the receptive field, then you'll also get half the response of the neuron. Because this will contribute to the integral and this won't, and so you'll get a neural response that's half as big.

So in this model, the response varies linearly with this overlap of the receptive field and the intensity. So that's where the term linear comes from. Any question about that? Correlation is a linear operation.

So the next thing we're going to talk about is temporal receptive field. So we just talked about spatial receptive fields. Neurons are also very sensitive to how things vary in time. So we're going to take the same concept.

Instead of a stimulus that's a function of position on the retina, let's say, we're going to take a stimulus that's a function of time. And we're going to operate on that temporal stimulus with a filter that [INAUDIBLE] a temporal sensitivity. We're going to get the output of that filter, add it to a spontaneous firing rate, and we're going to get a time-dependent firing rate.

So let me just show you what this looks like. So let's say that you have a stimulus that's fluctuating in time. So imagine that you have a neuron that has a spatial receptive field that's just a blob, and you shine just a positive bump.

And you shine light on it, and the intensity of that light varies. And this is the intensity that you apply to that spatial receptive field as a function of time. And again, 0 is some kind of average gray level. And so you can go dark, dark, or bright around that.

So now, neurons generally have receptive fields in time, and this is what a typical receptive field might look like, a temporal receptive field might look like, for a neuron. Neurons are often particularly driven by stimuli that go dark briefly and then go bright very suddenly, and that causes a neuron to spike. So we can imagine that temporal receptive field, and sliding it across the stimulus, and measuring the overlap of that temporal receptive field with the stimulus at each time. Does that makes sense?

And so you can see that most of the time that negative bump and positive bump are just going to be overlapping with just lots of random wiggles. But let's say that the stimulus has a negative part and then a positive part, dark, then bright. You can see that at this point that filter will have a strong overlap with the stimulus.

Why? Because the negative overlaps with the negative, and that product is positive. Positive overlaps with positive, that product is positive.

And so when you integrate over time you get peak. Does that make sense? This is, what I'm plotting here, is the product of these two functions at each different time step as I slide this temporal receptive field over the stimulus.

Does that make sense? Any questions about that? I'm going to go through that. We're going to work on that idea little bit more, but if you have any questions, now's a good time.

AUDIENCE: [INAUDIBLE]

MICHALE FEE: So you're asking, why is it correlation in the spatial domain? Because-- well, let me answer that question after we define what this is mathematically. So what is this mathematically?

AUDIENCE: A convolution.

MICHALE FEE: It's a convolution. It's exactly what we talked-- it's a lot like what we talked about when we talked synapses. In that case, we had some delta functions here corresponding to spikes coming in, and the synaptic response was like some decaying exponential. And we slid that over the stimulus. In this case, we have this fluctuating sensory input, this light intensity, and we're sliding that linear response of the neuron over and measuring the overlap as a function of position, and that's a convolution.

Mathematically, what we're doing is we're taking this linear kernel, this linear filter, sliding it over the stimulus, using this variable t, and we're integrating over this variable tau. So we have a kernel, D, multiplied by the stimulus at different times shifts. We integrate over tau, and that's the output of our temporal receptive field. That's the linear output of our receptive field. And we add that to spontaneous firing right, and that gives us a time-dependent firing rate of the neuron.

Yes?

AUDIENCE: So is tau how much we [INAUDIBLE]?

MICHALE FEE: Great question. t is the location of this kernel as we're sliding it along. Tau is the variable that we're integrating over after we multiply them. Does that make sense?

So we're going to pick a t, place the kernel down at that time, multiply this. And remember, this is 0 everywhere outside of here. And so we're going to multiply the stimulus by this kernel.

It's going to be-- that product is going to be 0 everywhere except right in here. You're going to get a positive bump when you multiply these two, a positive bump when you multiply those two. And the integral over tau gives us this positive peak here.

If we picked a slightly different t so that this thing was lined up with this positive peak here, then you'd see that you'd have positive times negative. That gives you a negative. When you integrate over tau, that gives you this negative peak here. Does that makes sense? So let's just go back to the math.

So you can see that we're integrating over tau, but we're sliding the relative position of D and S with this variable t. Yes?

AUDIENCE: Is that the [INAUDIBLE]?

MICHALE FEE: Yes. S is the stimulus.

AUDIENCE: Oh. And D is the kernel?

MICHALE FEE: D is the linear kernel. Yes?

AUDIENCE: [INAUDIBLE]

MICHALE FEE: So nature chooses the shape of the kernel for us. So that is the receptive field of neurons. Now, I just made this up to demonstrate what this process looks like. But in real life, this is the property of a neuron, and we're going to figure out how to extract this property from neurons using a technique called spike-triggered average, which we'll get to later.

But for now, what I'm trying to convey is, if we knew this temporal receptive field of a neuron, then we could predict the firing rate of the neuron to a time varying stimulus. That was a very important question. Does everyone understand that? Because it's one of those cases where once you see it it's pretty obvious, but sometimes I don't explain it well enough. Yes?

AUDIENCE: [INAUDIBLE]

MICHALE FEE: Yes. So I've already flipped it, and sometimes you'll see. So this is all going this way-- positive tau.

I've flipped it for you already. Sometimes you'll see it plotted the other way with tau going positive to the right, but I've plotted it this way already. Any questions?

Oh, and so that was actually the very next question. You might normally-- you might sometimes see temporal receptive fields plotted this way with positive tau going to the right. And kind of meant-- I always just flip it back over. Because in this view, you see that what the neuron responds to is dark followed by light, and then right there is when you have a peak spiking probability. Peak firing rate happens right here.

Yes?

AUDIENCE: [INAUDIBLE]

MICHALE FEE: So we're going to get to that. But typically, neurons in the retina-- I'll show you an example in the retina. A typical time scale here might be tens to 100 milliseconds, so pretty fast.

So that's called the temporal kernel or the temporal receptive field. And again, it's linear in the sense that if you, for example, had a stimulus intensity that just had this positive bump without the negative bump, then the response would be lower just by the ratio of areas. So if you got rid of this big negative bump here, then the response would be, I don't know, a third as big. It would be linear in the area. Let's push on.

So now, let's extend this. So we've been talking about spatial receptive fields and temporal receptive fields. But in reality, you can combine those things together into a single concept, called a "spatial temporal receptive field," and that's usually referred to as an STRF. If you're working in the auditory system, STRF, it's the same acronym, but it just means spectral temporal receptive field, because it's sensitive to the spectral content of the sounds, not the spatial structure of the visual stimulus.

So in general, when you have a visual stimulus, it actually depends on x- and y-coordinates in the retina and time. So just I of x and y, which would be like a still image presented to you. I of x, y, and t is-- any movie can be written like that. Your favorite movie is just some function of I of x, y, and t.

And so we're going to now present to our retina, and we're going to simplify this by considering just one spatial dimension. So we're going to take your favorite movie and just collapse it into intensity as a function of position. It's probably not nearly as interesting, but it's much easier to analyze.

So we're going to write the firing rate as a function of time as a spontaneous firing rate plus a filter, D, which is a spatial temporal receptive field acting on that intensity. And you can see that we're doing stuff in here that looks like a convolution integrating over tau, and we're also doing stuff that looks like a correlation when we integrate over x. So there's the convolution integrating over tau.

What I've done is I've pulled out the D tau, because we can consider-- I've just written this as two separate integrals. So we have an integral over tau that looks like a convolution. And we have an integral over x that looks like a correlation.

So what is separability mean? So separability is just a particularly-- if a receptive field is separable, it means that you can write down a spatial receptive field and a temporal receptive field separately. And that looks like this.

So I imagine that if you have a spatial temporal receptive field, D, that's a function of position and time. But you can see that you can just write it as a product of the spatial part and the temporal part. So here, you have a temporal receptive field that looks like this, a positive lobe here and a negative lobe there, a spatial receptive field that looks like this, just a positive lobe.

And if you multiply this function of x by this function of t, you can see that you get a function of x and t that looks like this, where at any position the function of time just looks like this-- scaled. And at any time, the spatial receptive field just looks like this. Does that make sense?

Other receptive fields are not separable. You can see that you can't write this receptive field as a product of a temporal receptive field and a spatial receptive field. Does that make sense? Is that clear why that is?

So basically, you can see that if you take a slice here at a particular position, you can see that the temporal pattern here looks very different than the temporal pattern here. And so you can't write this simply as a product of a spatial and a temporal receptive field-- separable, inseparable. So let's take a look at what happens when you have a separable receptive field.

Things kind of become very simple. We can now write our spatial temporal receptive field as a spatial receptive field, which is a function of position, times a temporal receptive field that's a function of time. And when you put that into this integral, what you find is that you can pull that spatial part of the receptive field out of the temporal integral.

So basically, the way you think about this is that you find the correlation of the spatial receptive field with the stimulus, and that gives you a time-dependent stimulus, a stimulus that's just a function of time. Then you can convolve the temporal receptive field with that time-dependent stimulus. So you can really just treat it as two separate processes, which can be kind of convenient just for thinking about how a neuron will respond to different stimuli.

So let's just think about, develop some intuition about, how neurons with a particular receptive field will respond to a particular stimulus. So here's what I've done. I've taken a spatial temporal receptive field here. This is a function of position and time, and we're going to figure out how that neuron responds to this stimulus.

So this stimulus is also a function of space and time. It's one-dimensional in space. So what does this look like?

This looks like a bar of light that extends from position 2, let's say, 2 millimeters to 4 millimeters on our screen. And it turns on at time point 1, stays on, and turns off at time point 6. Let's say 1 second to 6 seconds. Does that make sense?

So just imagine we have a 1D screen, just a bar, and we turn on light that's a bar between 2 and 4. So we turn on a bar of light. We turn it on at time 1, and we turn it off at time 6.

It's just a very simple case. We flash of light on at a particular position, and then we turn it off. So let's see how this neuron responds.

So what we're going to do is we're going to slide-- remember, in the 1D case where we had the temporal receptive field, we just slid it across the stimulus. So we're going to do the same thing here. We're going to take that spatial temporal receptive field, and we're going to slide it across the stimulus.

And we're going to integrate, we're going to take the product, and we're going to integrate. And the integral plus the spontaneous rate is going to be the firing rate of our neuron. So what is the integral right there? The product is--

AUDIENCE: 0.

MICHALE FEE: --0. where? integrate, it's 0. So we're going to add a spontaneous firing rate, which will be right there. So that will be our firing rate. Now, let's slide the stimulus a little further.

That means that this time we're asking, what is the firing rate of that neuron? So what is it going to look like?

AUDIENCE: Go up a bit.

MICHALE FEE: It's going to go up a little bit, because we have a positive part of the receptive field. Green is positive in our pictures here. It's going to overlap with this bar of light, because that neuron is sensitive to light between, let's say, 1 and 4, positions 1 and 4 on the screen. And so the light is falling within the positive part of that receptive field, and so the neuron's going to increase its firing rate.

So now what's going to happen?

AUDIENCE: [INAUDIBLE]

MICHALE FEE: It's going to cancel. You're going to get a positive contribution to the firing rate here-- whoops-- and a negative contribution here. And those two are going to add up.

You're going to multiply that times that. That gives you a plus. That times that gives you-- sorry.

That times the light that's shining on it is negative. Add those up, and it's going to cancel. And the firing rate's going to go back to baseline.

Now, the light in this receptive field, we're continuing to slide it in time over our stimulus. What happens here?

AUDIENCE: Same thing.

MICHALE FEE: Same, good. How about here? It's going to go--

AUDIENCE: Down.

MICHALE FEE: --down. It's going to dip down, that's right. And then?

AUDIENCE: [INAUDIBLE]

MICHALE FEE: Yeah. By 0, you mean the spontaneous firing. Yeah, exactly. Cool. Yes?

AUDIENCE: [INAUDIBLE] the rate of response because the slope of the line [INAUDIBLE]?

MICHALE FEE: So you should think about this thing sliding over the stimulus in real time. So if these are units of seconds, then this thing is sliding across the stimulus at 1 second per second sliding across. And so that is firing rate as a function of time in those units. Does that make sense?

AUDIENCE: Yeah. But why [INAUDIBLE]?

MICHALE FEE: Oh, OK. Like why doesn't this go up to here? So what's the answer to that?

AUDIENCE: [INAUDIBLE]

MICHALE FEE: Yeah. So how would I make that steeper? How would I make that go up to here?

AUDIENCE: [INAUDIBLE] the light.

MICHALE FEE: What's that?

AUDIENCE: You'd turn the light up.

MICHALE FEE: Yeah, you'd turn the light up, that's right. This is the receptive field of a neuron, so we generally can't control that. So if we wanted to make this neuron respond more, we'd turn the light up to a higher intensity. Any other questions?

So neurons often have more complex receptive fields. So here's an example. What is this going to do as we slide this across the stimulus? What's that?

AUDIENCE: [INAUDIBLE]

MICHALE FEE: Yeah. It's not going to-- the neuron isn't going to respond at all. Because as soon as it overlaps, it has a positive contribution. The light activates these lobes of the receptive field, but inhibits these lobes of the receptive field. And the net result is that when you integrate over the product, you're going to get 0.

Does anyone have any idea what kind of stimulus might make this neuron respond? This is a very special kind of receptive field. Yes, [INAUDIBLE]

AUDIENCE: The light goes from [INAUDIBLE]

MICHALE FEE: Yeah. What is that called?

AUDIENCE: I'm not sure.

MICHALE FEE: It's called a stimulus that moves.

AUDIENCE: [INAUDIBLE]

MICHALE FEE: Moves-- a moving stimulus. Good. So that's a receptive field that response to a moving stimulus. So let's take a look at that.

So here we go. Anybody want to take a guess at what this stimulus will do to this neuron? Can you visualize sliding it across?

AUDIENCE: [INAUDIBLE]

MICHALE FEE: And then what?

AUDIENCE: [INAUDIBLE]

MICHALE FEE: Yeah. You can see that this-- so let's describe what this [AUDIO OUT]. So we've turned a bar of light on here between 0 and 2, and then we slide it up over the course of a few seconds. So we've turned a spot of light on, and then we move it up-- off.

So it's a spot of light that turns on, moves, and then disappears. So let's walk through it. So there's a little bit of overlap there, so the neuron's firing rate is going to start going up.

But then as it goes further, this light is now activating those inhibitory lobes, which is going to have a negative contribution. So when you take the product, you're going to get lots of negatives there, very little contribution from the positive lobes, and so the firing rate's going to go down. And what happens is it goes down, and once the firing rate hits 0, it can't go any more negative. So the firing rate is just going to sit at 0 until this stimulus moves out of the temporal receptor of this neuron. And then what?

AUDIENCE: Back up.

MICHALE FEE: It's going to go back up to the spontaneous rate. So what kind of stimulus will activate this neuron? A stimulus that moves from top down. So let's take a look at that.

So here's our stimulus. You see that it's going to just hit that inhibitory lobe, go down a little bit. And then the excitatory lobes of the receptive field are going to overlap with the stimulus. You're going to get a big positive peak, and then the stimulus will move out of the receptive field, and the firing rate will go back down to baseline.

Any questions about that? So that's very common, in both the visual system and in the auditory system, to have neurons that are responsive to moving stimuli. What does moving stimulus mean in the auditory system?

AUDIENCE: Changing pitch.

MICHALE FEE: Right, changing pitch. So [WHISTLE], like that. That activated a gazillion neurons in your brain that are sensitive to upward-going pitches that you have structure like this. Isn't that crazy?

[WHISTLE]

I can control all the neurons in your brain--

[WHISTLE]

[LAUGHS]

--at least the ones that respond to whistles.

So now that we've seen mathematically how to think about what a receptive field is and how it interacts with a sensory stimulus, how do you actually discover what the receptive field of a neuron is? That turns out to actually be a very challenging problem. So in very early parts of the visual and the auditory system, like in the retina and the LGN, and as far as, let's say, V1 in visual cortex, it's been possible to find receptive fields of neurons by basically just randomly flashing bars and dots of light and just hoping to get lucky and find what the response is.

It turns out that that's generally a very-- it can be a very time-consuming process. And so people have worked out ways of discovering the receptive fields of neurons in a much more systematic way. And that's what we're going to talk about next-- the idea of a spike-triggered average. So here's the idea.

We're going to take a stimulus, and we're going to basically-- we're basically just going to make noise, just a very noisy stimulus. So we're going to take, let's say, an intensity, a light, a spot of light, and we're going to fluctuate the intensity of that light very rapidly. And we're going to do that basically with a computer.

We just take a computer, make a random number generator, hook that up to, let's say, a light source that we can control the brightness of with a voltage. And then have the computer generate-- put out that random number sequence, control the light level, and then play that to our, let's say, our visual neuron. And that neuron is going to spike.

And now, what we can do is take the times of those spikes and go back figure out basically what made the neuron fire post-hoc. So if we do that here, you can basically take the spike times. Now, you know that whatever made the neuron spike happened before the spike. It didn't happen after the spike.

So you can basically ignore whatever happened after the spike and just consider the stimulus that came in prior to the spike. So we're just going to take a little block of the stimulus prior to the spike, and we're going to do that for every spike that the neuron generates, and we're going to pile those up and take an average-- spike-triggered. That's it.

And that is going to be-- what's really cool is that you can show that under some conditions that spike-triggered average is actually just the receptive field of the neuron. It's the linear receptive field of the neuron. So let's think-- and you can write that down as follows.

We're going to add a stimulus. We're going to write down the times at which all these spikes occur, t sub i, or the times in the stimulus at which the spikes occur. We're going to take the stimulus at those times minus some tau, and we're going to average them over all the spikes, all the n spikes, that we've measured. And that K of tau is going to be the spike-triggered average, and in many cases, it's actually the linear kernel.

Now, let's think for a moment about what the conditions are. What kind of stimulus do you have to use in order to get the spike-triggered average to actually be the linear kernel of the neuron, that receptive field of the neuron? Any guesses? Let me give you a hint.

What happens if I take a stimulus that varies very slowly? So instead of having these wiggles, it just goes like this. It has very slow, random wiggles.

Will that be a good stimulus for extracting the receptive field? Why is that? Yes?

AUDIENCE: [INAUDIBLE]

MICHALE FEE: Yeah. So I think what you're saying is that that stimulus is very slow, and it doesn't actually have the fast fluctuations in it that makes the neuron spike. If the stimulus varies very slowly, then it-- see, this neuron likes to have this very fast wiggle, this negative followed by a positive. But if the stimulus you put in just varies slowly, then that stimulus doesn't actually have the kind of signal that's needed to activate this neuron. Yes?

AUDIENCE: [INAUDIBLE] the stimulus [INAUDIBLE] smaller than tau?

MICHALE FEE: Well, tau is just the variable that describes the temporal receptive field. But I think what you're saying is that the stimulus varies more slowly than the temporal structure in the receptive field, in the temporal receptive field. That's right. Tau is just this variable that we define the receptive field on. Yes, [INAUDIBLE]

AUDIENCE: So when we add up an average, are we actually adding up from everything before the spike, everything before the spike?

MICHALE FEE: So great question. So how far back do you think you would need to average?

AUDIENCE: [INAUDIBLE]

MICHALE FEE: Maybe. I mean, in principle, you could have spikes happening very fast, and you could have signal that affects the response of a neuron from even before the last spike. But in general, what do you think the answer to that question is? Brainstorm some more ideas.

So let's say that you were recording in the retina, and you knew that neurons tend to respond to visual stimuli only for-- that temporal receptive fields in the retina never extend back more than 100 milliseconds. Then how would you choose that window? You would just choose this window to be 100 milliseconds, and that would be it.

If you're recording in a brain area that you really have no idea, then you have to actually try different things. So you can try a window that goes back 100 milliseconds. And if when you do the spike-triggered average, it hasn't gone to 0 yet, then you need to take more window. So you can figure it out. You can create a short window, and it only takes-- like you change one number in your Matlab code, and hit Run again, and do it again. It's pretty simple. Yes?

AUDIENCE: So when you've got like [INAUDIBLE].

MICHALE FEE: Yes.

AUDIENCE: Wouldn't that depend [INAUDIBLE] what kind of filter [INAUDIBLE]

MICHALE FEE: Yeah. So you're saying that the stimulus that you choose actually depends on the kinds of filters that the neurons are, right? Actually, the right answer, it depends.

AUDIENCE: And so we have-- [INAUDIBLE]

MICHALE FEE: Yeah. So generally, the statement is that the stimulus you use has to have fluctuations in it that are faster than the fluctuations in the kernel that you're trying to measure. And so most people choose what's called a "white noise stimulus." And white noise stimulus comes from the idea that when you take the spectrum-- and we're going to get into spectra next week. But when you look at the spectrum of the stimulus, and you take the Fourier transform of it and look at how much power there is as a function of frequency, the spectrum is flat.

And just like in colors, white light has a flat spectrum. And so the term evolved to a noise that has a flat spectrum white noise. And that's what people generally refer to when they do spike-triggered averages. They use noise that has a flat spectrum. And so you'll often refer to people saying that they've used a white noise stimulus to extract a spike-triggered average.

Now, of course, you can't ever make a noise that truly has a flat spectrum. You have to-- you can only make things fluctuate as fast as your experimental setup can make them fluctuate. So things eventually fall off.

Fortunately, neurons tend to have receptive fields that only have fluctuations that are on the scale of 10 milliseconds, or maybe a millisecond in extreme cases. Maybe in the auditory system you might get a little fluctuation for a millisecond, and that's generally pretty slow for an experimental setup. So you choose a white noise stimulus, where white noise means it's got fluctuations faster than the fastest fluctuations in the temporal receptive field, and that for real neurons tends to be, even in early sensory areas, tends to be on the scale of millisecond fluctuations. And in higher brain areas, they would have even slower fluctuations.

Now, let me just say one word about spike-triggered average. It works really well in lower sensor areas. But once you get up out of primary sensory areas, this method doesn't work any more. Neurons don't actually have simple receptive fields. And this method starts not working so well outside of primary sensory areas.

But before I get too much into the limitations of this, let me just show you some examples where it works really beautifully well. So this is-- I'll show you some slides that I got from Marcus Meister, who studies-- he's at Caltech. He used to be at Harvard, and he studies the retina. And so he developed this setup for extracting receptive fields of retinal neurons.

And here's the idea. So here's a piece of retina, thats a representation of the circuitry in the cells within a piece of retina. You extract the retina, and you place it on a dish, a special dish, that has electrodes embedded, metal electrodes embedded, in the glass, sort of on the surface of the glass.

You take the retina out. You press it down onto the glass. So now the electrodes are sensing the spiking activity of these neurons down here in the retinal ganglion cell layer. These are the photoreceptors up here.

And then what he does is he has a computer monitor that's generating random patterns of visual stimuli, and you project that using a lens down onto the photoreceptors of the retina. And those neurons now make lots of spikes, and you can extract those spikes using the methods that you saw in the video from Tuesday. So here's what those signals look like.

This just shows the signal on four different electrodes that happen to be right near each other, like four adjacent electrodes on this electrode array. And you can see that you get spike wave forms on all these different electrodes. You can see that you get-- that you see what looks like lots of cells on these four electrodes. One really interesting thing to note is that these electrodes are actually placed close enough together that multiple electrodes detect the spike signal from a single cell.

So you can see right here, here's a spike. It exactly lines up with the spike on this other electrode. So here is a spike on one electrode that lines up with a spike on another electrode. You can see there's a little blip there and a little blip there. All of those spikes are actually from a single cell whose electrical activity is picked up on four adjacent electrodes. David?

AUDIENCE: Is this the raw data?

MICHALE FEE: Yeah. This is the raw voltage data coming out of those electrodes. And you can see here's at a different cell right here. You can see that this cell has a peak of voltage fluctuation on electrode two. You see a little blip there, and a little blip there, and nothing there.

And here is yet another cell that has a big peak on electrode three, essentially nothing, maybe a small bump there. So you can actually extract many different cells looking at the patterns of activity that appear on nearby electrodes. And it turns out that this multi-electrode array system is actually very powerful for extracting many different cells, the spiking activity of many different cells, out of a piece of tissue.

So what you can do is you put this through what's called a "spike-sorting algorithm," which uses these different spike wave forms on these different electrodes to pull out a spike train. And the spike train is now going to be a delta function for each different neuron that you've identified in this data set. So even though different neurons appear on these different electrodes, you're eventually going to extract this now so that you have a spike train for one neuron, a spike train for another neuron, a spike train for a third neuron, and so on.

And then you can plot the firing rate of those different neurons. This is actually a histogram, a peristimulus histogram, of the activity of a bunch of different neurons to a movie being played to this piece of retina in the dish. And that's just literally a movie from a forest with trees swaying around in the breeze.

So you have all these different neurons. And you can see that each neuron responds to a different feature of that movie. And that's because each neuron has a receptive field that's in a slightly different location, has a slightly different spatial and temporal receptive field, that it allows it to pick out different features of the visual stimulus.

And there are about a million of these neurons that project on the back of the retina. Actually, I should be careful. It's actually the-- it's the back of the retina. Because the light goes through the ganglion cells through the photo receptors, which are actually-- sorry.

Photoreceptors are actually on the backside of the retina. Ganglion cells are on the front, and light goes through the ganglion cells to get to the photoreceptors. And there are million of those retinal ganglion cells that then project up through the optic nerve to the thalamus.

So how do we figure out what each of those neurons is actually responding to in this movie? So what we can do is-- you could imagine doing a spike-triggered average of these neurons to the movie that's playing the trees swaying in the breeze. But why would you not want to do that? Why would that be a bad idea?

What is it that we just decided is the best kind of stimulus to use to extract receptive fields? This is a highly structured stimulus that's got particular patterns in the stimulus and both in space and in time. So it's really not an optimal stimulus for finding the receptive fields in neurons.

What we want to do is to make a very noisy stimulus that we can play, and that's what they did. So then they make this, what they call in the visual system, a "random flicker stimulus." So it's basically a movie where you randomly choose that the pixel values, both in R, G, and B-- red, blue, and green-- for each stimulus at each time step. And here's what that looks like.

So now, you play that movie to the retina, and you record the spike trains. So there's the neurons spiking.

And now what you do is you-- because this is now a two-dimensional stimulus, what you do is you have to collect the samples of the movie at a bunch of time steps prior to the neurons spiking. Does that make sense? Now, you do that for each spike that occurs.

And now, you average those all together to produce a little movie of what happened on average before each spike of the neuron. And here's what that looks like. So this is for two different neurons.

So what is that? So this is time across the top. So it starts at minus half a second.

So what did that look like? What was it that made that neurons spike? What was it that happened right before that neuron spiked?

AUDIENCE: [INAUDIBLE]

MICHALE FEE: Yeah, a dark spot. So that neuron was excited by a spot of light, sorry, by a stimulus that looked like a darkness right in that location right there. So that neuron is essentially being inhibited by light at that location. And when the light at that location goes away, boom, the neuron is released from inhibition and spikes.

Here's another cell. So that neuron responded to a spot of light right there in that location. And that's because that neuron gets excitatory input from bipolar cells that are located in the retina at that location. And those bipolar cells respond to input from the photoreceptors at that location.

That's called an on cell. That's called an off cell. So there are many different kinds of neurons in the retina. There's something like-- I forget the latest count-- 40 or 50 different types of retinal ganglion cells that have very specific responses to visual stimuli.

So now let's break that down into a spatial and temporal receptive field. Most-- I probably shouldn't say most-- but many retinal ganglion cells are separable in the sense that they have a spatial receptive field and a temporal receptive field that are just a product of each other. The STRF is a product of a spatial and temporal component. So here you can see as a function of time before the spike. So this is the stimulus at the time of the spike.

This neuron responds with a spike to a spot of light that happened about 150 milliseconds earlier. And here's what that stimulus looked like as a function of space on the retina. So that's the spatial receptive field. Sorry, that's the spatial temporal receptive field-- a spatial stimulus as a function of time. You can write that as a product of a spatial receptive field and a temporal receptive field.

So here's what the spatial receptive field looks like, and here's the temporal receptive field. You can see that this neuron, just like the example that we talked about earlier, this neuron likes to respond when there's a darkening in the central area, followed by a bright spot. You can see that little bit of darkening right here. So the response when this goes dark and then bright.

So that's the visual system. Let's take a look at auditory - you use this same method for finding receptive fields in the auditory system. So we're going to talk briefly. We're going to come back to spectral analysis and spectral processing of signals in a couple lectures, but let me just introduce some of the basic ideas.

So we're going to talk about the idea of a spectral representation of a sound. So this is a microphone signal of-- let me see if you can guess what it is-- of a creature. There are parts of this stimulus that have high frequency.

So this is a microphone signal. It fluctuates due to fluctuations in air pressure when you hear something. Parts of that signal have high-frequency fluctuations. Parts of that signal have low-frequency fluctuations.

You can compute a Fourier transform-- which we'll talk about more later-- as a function of time stimulus and see what the spectral components are. So this is a spectrogram of the sound that I just-- right now. But you can see frequency as a function of time, and the intensity, or in this case the darkness on the plot, shows you how much energy there is at a particular frequency at a particular time, so frequency as a function of time.

And now, neurons respond to stimuli like this. It's a canary song. And neurons respond to different sounds.

And so you can discover what sounds activate neurons by doing the same trick. So I'll show you. This is from a paper from Michael Merzenich's lab. This was worked on by Christophe deCharms who was a post-doc in the Merzenich lab.

And basically, what you can do is-- OK. So this is for calculating a visual receptive field. For calculating an auditory receptive field, what you can do is you can basically play noisy stimuli in auditory space.

So what you can do is present random patterns of tones. So this is frequency, and this is time. And so what you can do is you can make a little chords of tone, tone, tone that last, let's say, 20 milliseconds. And then you make a different random combination of tones, and then a different random combination of tones. And this sounds like a very scrambled, noisy stimulus.

And you play this to the animal while you're recording a neuron in auditory cortex. The neuron spikes. And then what you can do is just do exactly the same trick. You can look at the stimulus that occurred before each spike, pile up those columns. There's a little spectrum temporal pattern of stimuli that the bird-- in this case, a monkey heard right before that neuron spiked.

And you can do the same thing. You can take that little snapshot of that sound and average them together. And here are the kinds of things you see.

So that's a spectro-temporal receptive field. You can see I plotted it. It's plotted in a way that this is the stimulus that occurs with the spike.

So this is like the D plotted already flipped. And you can see that-- how would you describe what this neuron responds to? How would you describe that?

So this is frequently. Sorry, this is frequency in kilohertz. And that's time in milliseconds. So what do you think this neuron responds to?

AUDIENCE: [INAUDIBLE]

MICHALE FEE: What's that?

AUDIENCE: It responds [INAUDIBLE].

MICHALE FEE: It responds maximally, actually, to a very short tone at 4 kilohertz. You see how it kind of has some inhibition there? See how it's kind of darker right there?

So this neuron actually will respond better to a tone that only lasts 20 milliseconds than it will to a tone that lasts a long time. So this is a neuron that responds to a short tone pulse. What happens if we play a stimulus to this neuron that's broad? That instead of just being a tone, [WHISTLE], is broad, like [STATIC]? The noise that's at 4 kilohertz here will tend to excite the neuron, but the noise that's over here at 5 kilohertz or 3 kilohertz will tend to inhibit the neuron.

So the best response, the best stimulus to make this neuron respond, is a pure tone at 4 kilohertz that lasts about 20-ish milliseconds. How about this? Let's take a look at this neuron right here. What about that neuron?

What does that neuron want to respond to? What does it like to hear? I'm anthropomorphizing shamelessly. You're not supposed to do that. What kind of stimulus drives this neuron?

AUDIENCE: [INAUDIBLE]

MICHALE FEE: Good, a downward sweep, tone sweep, that goes from about 4 kilohertz to 3 kilohertz. In how long?

AUDIENCE: 100 [INAUDIBLE].

MICHALE FEE: In about 100 milliseconds, that's right. Maybe 50 would do, [WHISTLE], like that. How about this? It's kind of messy, right?

So you can see that neurons have receptive fields that can be very complex in space, or in this case, in frequency and time. They are very selective to particular patterns in the stimulus. So we talked about a mathematical version of receptive fields, which are essentially describing patterns of sensory inputs that make a neuron spike. And we've talked about a very specific model, called a linear/nonlinear model, that describes how neurons respond to stimuli or become selective to stimuli.

We've talked about a spatial receptive field, described the response of a neuron as a correlation between the spatial receptive field and the stimulus. Temporal receptive fields, where we've used convolution to predict the response of a neuron to a temporal-- to a stimulus. We've talked about the idea of a spatio- or spectro-temporal receptive field, and we've talked about how to use a spike-triggered average to extract the spectro-temporal or spatio-temporal receptive field of a neuron using white noise or random stimuli.