

MICHALE FEE: OK, good morning. So far in class, we have been developing an equivalent circuit model of a neuron, and we have extended that model to understanding how action potentials are generated. And, more recently, we extended the model to understanding the propagation of signals in dendrites. So today we are going to consider how we can record activity, record electrical signals related to neural activity in the brain, and we're going to understand a little bit about how we can, in particular, record extracellular signals. So that will be the focus of today's lecture.

So, so far in class, we have been analyzing measurements of electrical signals recorded inside of neurons. For example in the voltage clamp experiment, we imagined that we were placing electrodes inside of cells so that we could measure the voltage inside of these cells. But it's actually quite difficult, in general, to record membrane potentials of neurons in behaving animals, although it's certainly possible. It's much easier to record electrical signals outside of neurons. In this case, we can actually record action potentials.

And so those kinds of recordings are made by placing metal electrodes that are insulated everywhere along the shank of the electrode except right near the tip. And if we place a metal electrode in the brain, we can record voltage changes in the brain right near cells of interest. So in this case, we can record from action potentials of individual neurons in behaving animals and various aspects of either sensory stimuli or behavior. So this kind of recording is called extracellular recording, and we're going to go through a very simple analysis of how to think about extracellular recordings.

So recall, of course, that when we measure voltages in the brain, we're always measuring voltage differences. So when we place an electrode in the brain near a cell, we are-- we connect that electrode to an amplifier. Usually, we use a differential amplifier that gives us-- that provides us with a measurement of the voltage difference between two terminals, the plus terminal and the minus terminal. So we place-- we connect the electrode to the plus terminal, and we connect another electrode, called the ground electrode, that can be placed in the brain some distance away from the brain area that we're recording from, or the surface

of the skull. So we're measuring the voltage that's near a cell relative to the voltage that's someplace further away.

Now, the voltage that we measure in the brain, voltage changes, are usually or always associated with current flow through the extracellular space. So we can analyze this in terms of Ohm's law. So, basically, the voltage changes that we are measuring are going to be associated with some current through extracellular space times some effective resistance in the extracellular space. And you remember from previous lectures that the effective resistance of extracellular space is proportional to the resistivity in extracellular space times a length scale divided by an area scale.

So how do we think about what kind of voltage changes we might expect if we placed an electrode near a cell that is generating an action potential? So let's start with a spherical neuron, and let's give this spherical neuron sodium channels and potassium channels so that it can generate an action potential. So during an action potential, we have an influx of sodium, followed by an outflux of potassium, And that influx of sodium produces a large positive-going change in the voltage inside the cell. So you can see that during the rising phase of the action potential, we have sodium flowing in. But at the same time, we have current flowing outward through the membrane in the form of capacitive current.

Now, these two currents, the sodium ions flowing through the membrane and the capacitive current flowing outwards through the membrane, are co-localized on the same piece of membrane. And so there's no spatial separation of the currents flowing through the membrane. And as a result of that, there's actually no current flow in extracellular space, and so there's no extracellular voltage changes. The first lesson from this is that, if we were to record an extracellular space from a spherical neuron with no dendrites and no axon, we would actually not be able to measure any extracellular voltage change.

Now let's consider what happens if we have a neuron with a dendrite. So in this case, when the sodium current flows into the soma, that current, part of it will flow out through the membrane through capacitive current, but part of that current will flow down the dendrite, then out through the membrane through capacitive current back to the soma. So we have a closed circuit of current, current flowing into the

soma out through the dendrite, and then back to the soma through extracellular space.

So in this case, if we write down the equivalent circuit model of what this looks like-- this is the somatic compartment, this is the dendritic compartment-- we have-- in our earlier calculations of current flow through processes, through dendrites, we were neglecting the extracellular resistance, but in this case we are going to include that, because that extracellular resistance is what provides-- is what produces a voltage drop in extracellular [? space. ?] But during an action potential, we have current flowing in through the soma, out through the inside of the dendrite, back out through the membrane of the dendrite, and through extracellular space back to the soma. The voltage drop across this region of extracellular space will be proportional to this extracellular current times changes in extracellular voltage.

So now, we have a simple view that we have current flowing into the soma in a region-- from a region around the soma. So we have effectively what is known as a current sink. So we have charges flowing into the soma in the region around the soma. That current then flows out through the dendrite and appears in extracellular space in the region of the dendrite, and we call that a current source. So we have a combination of a current sink and a current source.

And you can see that the current in extracellular space is flowing from the current source to the current sink. So in this case, you can see in our simple equivalent circuit model that the voltages are more positive in regions corresponding to current sources, and the voltages are more negative in regions of extracellular space corresponding to current sink. Current is flowing from-- in extracellular space, this current-- current is flowing from the region of the dendrite to the region of the soma. The voltage here is more positive, the voltage here is more negative.

Now, let's take a look at understanding the relationship between the extracellular voltage change and the intracellular voltage change. So let's just write down the equation for the voltage drop across this extracellular space. That voltage drop is just the external current times some effective external resistance. The external current is just the sum of a capacitive current and a resistive current through the membrane of the dendrite.

So we can write down that voltage drop is now proportional to an external resistance-- extracellular resistance, I should say. So the-- so we can now write down an expression for these two currents as a function of membrane potential. And you recall from earlier lectures that the capacitive current is just given by $C \, dV \, dt$, and the membrane ionic current is given by some conductance, membrane ionic conductance, times the driving potential.

Now, in an action potential, voltage change is very rapid. So $dV \, dt$ is large. And, in fact, this term generally dominates over this term in an action potential. And so what we find is that the voltage change in extracellular space is proportional to the derivative of the membrane potential. And we can see that this is the case. In earlier experiments from Gyorgy Buzsaki's lab, they were able to record simultaneously from a cell intracellularly-- that's shown in this trace here-- and the extracellular voltage recorded from a microwire electrode placed near the soma. Extracellular recording is shown here, that the extracellular signal is actually quite close to the derivative of the intracellular signal.

Now, why is this voltage negative? Because the membrane potential near the soma here, the voltage here goes negative because during the rising phase of the action potential, we have sodium ions flowing into the soma from extracellular space. A current is flowing out of the dendrite out here and traveling back through extracellular space to the soma. So, again, the soma is acting like a voltage sink, and so the voltage is going negative near the soma.

So now let's take a look at what happens when we have synaptic input onto a neuron. So it turns out you can also observe extracellular signals that result not from action potentials but from synaptic inputs. So let's take our example neuron and attach an excitatory synapse to the dendrite. In this case, when the synapse is activated, if the cell is hyperpolarized, when a neurotransmitter is released onto the postsynaptic compartment of the synapse, it turns on conductances which allow current to flow into the dendrite. That current flows down the dendrite to the soma, flows back capacitively out through the soma and back to the synapse through extracellular space.

So you can see that in this case, the region near the synapse looks like a current sink as charges are flowing into the cell. The region near the soma looks like a

current source as current flows out of the soma and back to the current sink. And see that, in this case, you have charges flowing into the cell, and that should correspond to a decrease in the extracellular voltage. For the soma, you have positive charges flowing into extracellular space from the inside of the cell, and that corresponds to an increase.

OK, things look different when you have an inhibitory synapse. In the case of an inhibitory synapse, for example if GABA is released onto this GABAergic synapse, it opens chloride channels. Chloride is a negative ion that flows into the cell, but that corresponds to an outward-going current. So now we have the region around the inhibitory synapse looking like a current source. Current flows through extracellular space to the soma, where the soma now looks like a current sink. And so in the presence of activation of inhibitory inputs in the-- near the dendrites, you actually have an increase in voltage of extracellular space. While near the soma, you have a decrease in the voltage.

One of the important things to consider is that, in the discussion we've just had, we've been thinking about an individual neuron, how current sources and current sinks appear as a result of synaptic activity and action potentials around a single neuron. But in the brain, neurons are not isolated but are in the tissue next to many other neurons, that are also receiving inputs and spiking. So it turns out that the types of electrical extracellular signals that you see in the tissue depends very much on how different cells are organized spatially.

So in some types of tissue, for example in the hippocampus or in the cortex, cell bodies are lumped together in a layer, and the dendrites are collinear and are organized in a different layer. So this is called a laminar morphology. In this case, many of the synaptic inputs arrive onto the dendrites, and currents then flow into the somata. And in this case, these extracellular currents are reinforced. They reinforce each other, and they sum together to produce very large extracellular voltage changes.

So now let's turn to the question of how one actually records neural activity in the brain. So let's go back to our experimental setup now where we have an electrode placed near the soma of a neuron. And that electrode is connected to a differential amplifier or an instrumentation amplifier, giving us the voltage difference between

the extracellular space near the soma and the extracellular space somewhere far away. So we-- so this amplifier then measures that voltage difference and multiplies it by a gain of typically, let's say, a couple hundred, or 1,000, or even 10,000. We then take the output of that amplifier and put it into an analog to digital converter, that then measures the voltage at regularly spaced samples and stores those voltage digitally in computer memory.

So in analog to digital converters, the voltage is sampled at regular intervals, Δt , corresponding to some sampling rate, which is a frequency or a rate that's given by $1/\Delta t$. So the rate at which the samples are acquired is referred to as the sampling rate or sampling frequency. So if we were to record the extracellular voltage in a region of the hippocampus, we might see a signal that looks very much like this. These are data from Matt Wilson's lab.

The signal has a number of features. So you can see that there is a slow modulation of the signal that actually corresponds to the theta rhythm in a rat. That slow modulation of the voltage is actually caused by synaptic currents in the hippocampus you can also see that there are very fast deflections of the voltage corresponding to action potentials. So, once again, this is about a second's worth of data of intracellular-- sorry-- of extracellular recording from rat hippocampus, and we can see both slow and very fast components of this signal. The extracellular-- the action potentials, you can see, are very brief. They typically last about a millisecond.

If we were to look at the amount of power at different frequencies in this signal using a technique called measuring the power spectrum, which we will cover in more detail later in class, you can see that there is a lot of power at low frequencies and much less power at high frequencies. So this is a representation of the amount of power at different frequencies. So we can actually extract the fast and slow components of the signal using a technique called high pass and low pass filtering. So what we're going to do is we are going to develop this technique that allows us to remove high frequency components from the original signal to reveal just low-frequency structure within the signals. So how do we do that?

Well, basically, what we do is we're going to start by using a technique of low pass filtering that works basically by convolving this signal with a kernel. What that kernel does is it locally averages the signal over short periods of time. So what we

do is we take that signal, and we're going to place this kernel over the signal at different points in time, average together. We're going to multiply the kernel by this signal at different points of time-- in time, and plot the result down here. So let me explain what that looks like.

So let's say this is the original signal. You can see that it's a little bit noisy. It's fluctuating between 1 and 3, 1, 3, 1, 3. And then, at a certain point here, it jumps to a higher value, 5, 3, 5, 3, 5. So, intuitively, we would expect low pass filtered version of this signal to be low here, and then jump up to a higher value here. Now, here's our kernel. This is a representation of a kernel that looks like this. The kernel is 0, 0.5, 0.5, and 0.

And, basically, what we're going to do is place the kernel at some point in time over our signal, multiply the kernel by the signal time element by time element. So you can see that the product of the kernel with the signal is 0 here. 0 times 1 is 0, 0.5 times 3 is 1.5, 0.5 times 1 is 0.5, and 0 times 3 is 0. So that's the product of the kernel and the signal within that time window. Then sum up that-- the elements of that product. 0 plus 1.5 plus 0.5 plus 0 is 2. I'm going to write down that sum at this point in the filtered output.

Now, what we're going to do is just slide the kernel over by one element and repeat that. And I also added the earlier values of the output here. So now we slide the kernel over by one, we repeat. We get 0, 0.5, 1.5, and 0. Sum that up, we get a 2. And write down that filtered output down here. So you can see that the low pass filtered result of this signal is 2 everywhere up to here.

Now, if we slide the kernel over one more and multiply, you can see we get a 0, a 1.5. Here's 0.5 times 5. That's 2.5. 0 times 3 is 0. And the sum of that, the sum of those four elements is 4. Write down a 4 here. And if we keep doing that, all of the rest of those values is 4, which you can verify. So you can see that the low pass filtered version of this signal, filtered by this kernel, is 2 up to this point, and then it jumps up to 4, which was consistent with our intuition about what low pass filtering would do.

So that was low pass filtering. Now, how would we high pass filter? How would we extract out these high frequency components from the data, throw away the low

frequency components? Well, one way to think about this is that we actually-- we can get rid of the low frequency elements simply by subtracting off the low pass signal that we just calculated. So how do we do that?

So we're going to use this kernel here to do our high pass filtering. And notice that this kernel has two components. It has a square component that's negative, and it has a delta function at 0. So you can see that this negative component of this high pass filter looks a lot like the negative of our-- looks like the negative of our low pass filter. So if we were to take this kernel and instead use a kernel that was the negative of that kernel, like this, then what we would get is the negative of this low passed, of this filtered signal. So this part right here, this part of the high pass filter, is producing the negative of our low pass filter.

Now let's take a look at this component here. This component is a delta function with a value of 1 at the peak. You can see that if you convolve that kernel with your original signal, it just gives you back the same original signal. So what we've done is this component simply gives us back the original signal. This component subtracts off the low pass version of the signal. What we're left with is the high pass version.

Let's look at what that does to the spectrum of our signal. You can see that, the high pass, we've gotten rid of all of these low frequencies, all of the power at low frequencies, and we're left with just the high frequency part of our signal. And if we go back and look at the spectrum of what the low passed signal looks like, we can see that the low passed output retains all of the power at low frequencies and gets rid of these high frequency components [AUDIO OUT]. Low passed signal has no power at high frequencies. So once we have extracted high pass filtered version of the signal, you can see that what you're left with are action potentials.

So what we're going to talk about now is how you actually extract these action potentials and figure out when action potentials occurred during this behavior. Next thing we're going to do is detect-- we're going to do spike detection. So, basically, the best way to detect spikes is to look at the signal, plot the signal, and figure out what amplitude the spikes are. So look at the voltage of the peak, the peak of the spikes. And then set a threshold that consistently is crossed by that peak in the spike waveform.

So here are the individual samples associated with one action potential. You can see that a voltage right about here will reliably detect these spikes. And then, basically, you write one line of MATLAB code that detects where this voltage crossed from being below your threshold on one sample to being above your threshold on the next. Now what we can do is, once we detect the time at which that threshold crossing occurred, we can write down that threshold crossing time for each spike in our wave form. The spike here, we write down that time. That's t_1 . If you have another threshold crossing here, you write that time down, t_2 . And you collect all of these spike times into an array.

So we can now represent the spike train as a list of spike times. So we're going to represent this as variable t with an index i , where i goes from 1 to N . Now, we can also think of spike trains as a delta function. So you may remember that a delta function is 0 everywhere except at the time where the argument is 0. So this delta function is 0 everywhere except when t is equal to t_{signal} . That's when t is equal to the time of the signal. At that time, the delta function has a non-zero value.

So we can write down this spike train as a sum of delta functions. So the spike train as a function of time is delta of t minus t_1 , corresponding to the first spike, plus another delta function at time t_2 corresponding to that spike, and so on. So we can now write down mathematically our spike train as a sum of delta functions, one at each spike time.

We can also think of a spike train as being the derivative of a spike count function. So a spike count function will reflect the number of spikes that have occurred at times less than the time in the argument. So if this is our spike train, then prior to the first spike, there will be zero spikes in our spike count function. After the first spike, then we'll have a spike count of 1. And you can see that you get this staircase increasing to the right for each spike in the spike train. Since the integral of this spike train has units-- the integral over time has units of spikes, you can see that the spike train, this ρ of t , has units of spikes per second.

So now let's turn to the question of what we can extract from neural activity by measuring spike trains. So one of the simplest properties that we find about cells in the brain is that they usually have a firing rate that depends on something, either a motor behavior or a sensory stimulus. So, for example, simple cells in primary visual

cortex, of the cat in this case, are responsive to the orientation of visual-- of stimuli in the visual field. So this shows a bar of light, but represents a bright bar of light on a black background, actually.

And you can see that if you move this bar in space, at this orientation this neuron doesn't respond. But if we rotate the-- if we rotate this bar of light and move it in this direction, now the cell responds with a high firing rate. So if we quantify the firing rate as a function of orientation of this bar of light, you can see that the neuron has a higher firing rate for some orientations than for others. That property of being tuned for particular stimulus features is called tuning, and the measurement of that firing rate as a function of some parameter is called a tuning curve. So you can see that in primary visual cortex, neurons are tuned to orientation. And the tuning curves of neurons in primary visual cortex often have this characteristic of being highly responsive at particular orientations, and then smoothly dropping off to being unresponsive at other orientations.

So in a similar-- similar to the way that neurons in visual cortex are tuned to orientation, neurons in auditory cortex are tuned to different frequencies. So, for example, in the auditory system, when sound impinges on the ear, it transmits vibrations into the cochlea. Those vibrations enter the cochlea and propagate along a basilar membrane, where vibrations of particular frequencies are amplified at particular locations, and the membrane is unresponsive to vibrations of other frequencies.

So if you [? learn ?] from neurons in visual cortex and you-- sorry-- in auditory cortex and you play a tone, so this shows a phonograph of a auditory stimulus. At some frequency, you can see that this neuron spiked robustly in response to that. So individual neurons are tuned to respond robustly at particular frequencies but not other frequencies. And you can see that different neurons are selective to different frequencies. So you can see that this curve representing one neuron, that neuron is most active for frequencies around a little bit above 5 kilohertz, whereas other neurons are most responsive to frequencies around 6 kilohertz, and so on.

So we saw an example now of how firing rates of neurons in sensory cortex are sensitive to particular parameters of the sensory stimulus. This property of tuning applies not only to sensory neurons but also to neurons in motor cortex. This shows

the results of a classic experiment on analyzing the spiking activity of neurons in motor cortex during the movements, during our movements in different directions. So this shows a manipulandum. This is basically a little arm that the monkey can grab the handle here and move that arm around. The monkey's task is to hold that-- hold this arm at a central location. And then a light comes on, and the monkey has to move this manipulandum from the center out to the location of the light that turned on.

And so the experiment is repeated multiple times where the monkey has to move to different directions. You can see here the trajectories that the monkey went through as it moves from the center location out to these eight different target locations. And in this experiment, neurons were recorded at different regions of motor cortex, and the results-- the resulting spike trains were plotted. In this figure, you can see what are called raster plots, where-- for example, for movements in this direction, five trials were collected together. So each row here corresponds to the spikes that a neuron generated during movements from the center to this direction. You can see that the neuron became active just after the [INAUDIBLE], indicating which direction was turned on, and prior to the onset of the movement, which is indicated [AUDIO OUT].

You can see that the neuron responded robustly on every single trial to movements in this direction. But the neuron responded quite differently in some other direction. So you can see that the response to downward movement was quite weak. There was essentially no change in the firing rate. And you can see the to movements in other directions, for example to the right, were associated with an actual suppression of the spiking activity. So I should just point out briefly that these spikes here and here, before and after the onset of the trial, are spontaneous spikes that occurred continuously even when the monkey wasn't engaged in moving the handle. So we have a spontaneous firing rate, the trial initiates an increase in firing rate, and then a recovery to the baseline position.

So you can see that these motor cortical neurons exhibit tuning for particular movement directions. And we can quantify this now by counting the number of spikes in this movement interval, and plotting that as a function of the angle of the movement. And when you that, you can see that movements in particular

directions, in this case [AUDIO OUT] by degree direction, resulted in high firing rates, whereas movements in other directions resulted in lower firing rates.

So in order to do this kind of quantification, we need to be able to-- we need to understand a few different methods for how we can actually quantify firing rates. The simplest thing that I just described here is to just count the number of spikes in some particular interval that you decide is relevant for your experiment. So, for example, let's say we have a stimulus that turns on at a particular time, is-- it stays on, and then turns off at some later time. On different trials of that-- or different presentations of that stimulus, you can see that the spike train-- that the neuron spikes in a somewhat different way. But you can see that this sample neuron here, that I just made up, you can see that, in general, there is a vague increase in the firing rate after the onset of the stimulus.

So we can quantify that by simply setting a relevant time window and counting the number of spikes in that time window. So we're going to set a time window T from the onset of the stimulus to the offset of the stimulus, and simply count the number of spikes on each trial. So N is the number of spikes that occurred on the i th trial. The brackets here represent the average over that quantity i , which is trial number. So we're going to count the number of spikes on the i th trial and average that over trials.

So then, once we have that average count, we simply divide by the interval T that we're counting over, and that gives us a rate. Now, you can see that the number of spikes, the firing rate is not constant. So you can see in this little toy example here that, the way I've drawn it, that the spike rate increases at stimulus onset, and then it decayed away, which is very typical. So you're throwing away a lot of information. Now, the way that we just quantified that firing rate, we just counted spikes over the whole stimulus presentation. But if we want to get more temporal resolution in how we quantify the firing rate, we can actually just break the period of interest up into smaller pieces, count the number of spikes, the average number of spikes in each one of these smaller bins, and divide by the interval of these smaller bins.

So, for example, we can have-- we can count the number of spikes on trial i and then j . And we can average the number of spikes in the j th bin, overall trials. So that's just the average number of spikes in this first bin, for example, and divide by

the interval ΔT . And so now you can see that you have a different rate, finer temporal resolution.

So, for example, if you look at the analysis that Georgopoulos did in that 1982 paper for the arm movements of the monkey, they broke the trials up into small bins of about 10 or 20 milliseconds each, counted the number of spikes in each one of those bins, divided by that bin width, and computed the firing rate, spikes per second, in each one of those bins. You can see that they did one other thing here. The average firing rate during each bin in which we've subtracted off the firing rate in the pre-stimulus period. So that is, very typically, what you'll see in a neuroscience paper describing the behavior of neurons to a stimulus or during a motor task.

And so we can use similar tricks to estimate the firing rate of neurons in just a continuous spike train. So not all neuroscience experiments are done in trials like this, where we present a stimulus and then turn it off. Some trials-- some experiments are done, for example, where a movie may-- a monkey or an animal might be in a movie where stimuli are presented continuously. So you don't have this clear trial structure.

So we can also quantify firing rates in cases where we just have a continuous spike train without trials. And we can do that, again, by just taking that continuous spike train and breaking it up into intervals. So there will be N_j spikes in bin j , and then the bin has some width ΔT . Now, one problem that you can see immediately from this is that the answer you get will depend on where you place the boundaries of the bin. So if you take all these bins and you shift them over by, let's say, half of-- shift them over by $\Delta T / 2$, you can see that you can get a completely different set of-- a different set of firing rates for this same spike train. So there's not a unique answer.

So another way to do this is to quantify firing rates in bins that are shifted to all possible times. So, for example, we can take a window where we have 0 everywhere except within this window. We can now multiply the spike train. Well, one way to do this would simply be to count the number of spikes within that window, and then shift the window over and count the number of spikes, and shift the window over and count the number of spikes. And now you get a count of number of spikes in

each of those windows for all the windows that have width ΔT , but we've shifted it in small time steps.

But how can we describe that mathematically? Well, you may recall that this, in fact, looks a lot like a convolution. We're going to take this square kernel, and we're going to multiply it by the spike train and count the number of spikes, take the integral over that product. And you can see that that's basically going to give you the number of spikes within that window from t_1 to t_2 . So, for example, in this case we're going to use a square window. The firing rate is just going to be given by the number of spikes divided by the width of the window. And that's just $1/\Delta T$ times the integral $t - \Delta T/2$ to $t + \Delta T/2$, sliding that gradually over the spike train. So we're effectively convolving our spike train with this rectangular kernel.

And that's what it looks like mathematically. The firing rate is the convolution of the spike train with this smoothing kernel. So, as I mentioned, that's just a convolution, and we're convolving our spike train with this square kernel of this width. So that is the mathematical expression for a convolution. The firing rate is just the convolution of this kernel with the spike train. And, again, the kernel is 0 everywhere except within this window, $t - \Delta T/2$ to $t + \Delta T/2$, and it has a height of $1/\Delta T$ such that we make the area 1. Notationally, we often write that as a star, $\rho \star K$.

Now, in this case we were convolving our spike train with a square kernel. The problem with a square kernel is that it changes abruptly every time a spike comes into the window or drops out of the window. A more common way of quantifying firing rates is to convolve the spike train with a Gaussian kernel. So instead of using a square kernel here, we're going to use a Gaussian kernel that looks like this. The kernel is just defined as this Gaussian function. And it's normalized by $1/\sigma\sqrt{2\pi}$, and this normalization gives area under that kernel, 1.

So it's still essentially counting the number of spikes within that window. It's again just a weighted average of the number of spikes divided by the width of the kernel, but there's less weight at the edges. It has smoother edges, and that gives you a less steppey-looking result. So let me show you what that looks like. So here's a spike train. If we take fixed bins and compute the firing rate as a function of time, it looks

like this. If we take that spike train and we convolve it, and that estimate of firing it would depend very much on where exactly the windows are placed.

On the other hand, this shows the firing rate estimated with a square kernel of the same width. And you can see that it shows a much smoother estimate, a smoother representation of the firing rate varying over time. And if you take that same spike train and you convolve it with a Gaussian, then you get a function that looks like this. And we think of this as perhaps better representing the underlying process in the brain that produces this time-varying firing rate, this time-varying spike train, than either of these two. You don't really think that the firing rate of this neuron is jumping around in rectangular steps. We think of it as being some kind of smooth, continuous underlying function that represents maybe the input to that neuron.

You'll see that in these estimates of firing rate, we had to actually choose a width for our window. We had to choose a bin size. We had to choose the width of our square kernel here, and we had to choose the width of our Gaussian kernel here. And you can see that the answer, actually, that you get for firing rate as a function of time depends very strongly on the size or the width of that kernel that you choose to use. And now we smooth it with a Gaussian kernel that has a width, a sigma, of 4 milliseconds. So that's the standard deviation of the Gaussian that we use to smooth the spike train.

And you can see that what you get is a very peaky estimate of the firing rate. So that spike produces this little peak. But the question is, what's the right answer? How should you actually choose the size of the kernel to use to estimate firing rate for your experiment? The answer is that it really depends on the experiment. So neural spike trains have widely different temporal structure. So, first of all, neuronal responses aren't constant. Firing rates are not a constant thing. Neurons-- neural firing rates are constantly changing depending on the type of stimulus that you use. And different types of neurons have different temporal structure in response to the same stimuli.

So, for example, this shows the response of four different neurons to stimuli, to-- four different neurons in rat vibrissa cortex in response to whisker deflections. So this shows a raster, a histogram of firing rate as a function of time for one neuron in response to a deflection of one of the whiskers right here. So the whisker is at rest,

deflected, and then relaxed back to its original position. You can see that this neuron shows a little burst of activity at the onset of the deflection, and it's fairly persistent throughout the deflection. Here's another neuron during a deflection. Increased activity at the onset of the deflection, followed by a fairly persistent increased spiking rate throughout the deflection.

Now, a different neuron was quite a different behavior. So this neuron shows a brief increase in firing rate just at the time of the deflection, a little increase of firing rate when the deflection is removed, and [INAUDIBLE] no activity that persists during this constant part of the deflection. Here's another neuron. This neuron was primarily active at the onset of the deflection. Here's another neuron that was silent at the onset of the deflection, and then gave a robust response. Neurons, the changes in firing rate are not of a particular timescale. Different neurons have different time scales on which the changes in their firing rate are important.

But let's come back to our example of our auditory neuron. Here's the spiking of an auditory neuron during the presentation of an auditory stimulus [INAUDIBLE] right here. This neuron-- in fact, you can't see it here, but this neuron shows spiking at a particular phase of the auditory stimulus, much like the auditory neurons that we discussed for sound localization in the owl. So if you plot the firing rate of this neuron as a function of time during the presentation of this stimulus, you can see that the firing rates are rapidly modulated in time. You can see that, at particular phases of this stimulus, the firing rate is very high. And then just a millisecond later, the firing rate is very low.

So in this case, you can see that the spikes are locked temporally at particular phases of the stimulus. That's reflected when you make plots of firing rate as a function of time. It's reflected in various modulations. The neurons are firing at a particular time. So this corresponds to a case in which we would say that the spike timing is precisely controlled by the stimulus. It's, in many ways, more natural here to think about spike times as being controlled rather than spike firing rate being modulated.

So you can see here that sensory neurons, neurons can spike more in response to some stimuli than others. Motor neurons spike more during some behaviors more than others. We can think about information being carried in the number of spikes

that are generated by a stimulus or during a motor act.

Now, all neurons exhibit temporal modulation in their firing rates. They fire more during a movement or after the presentation of a stimulus. Sometimes that information is carried by slow modulations in the firing rate. For example, the response to different oriented bars is carried in the average firing rate of the neuron during that particular stimulus. And we can refer to that kind of code and that kind of representation of information as rate coding. We can say that information about the stimulus is carried in the rate, the firing rate of the neurons.

But in other cases, like the auditory neuron that we just saw, we can see that information is carried by rapid, by fast modulations, rapid changes in spike probability. And in that case, we often say that information is coded by spike timing. So a common question that you often hear about neurons is whether they're coding information using firing rate or temporal coding, rate coding versus temporal coding. Really, this is a false dichotomy. You shouldn't think about neurons coding information one way or the other. These are just-- really just two limits of a continuum. The brain uses information at fast time scales as well as slow time scales.

And how do we determine, how do we know what's important for the brain? What time scales are important? And the answer to that question really comes from understanding the way spike trains are read out by the neurons that they project to. What time scale is relevant for the computation that's being done in the system that you're studying? What are the biophysical mechanisms that those spikes act on. To understand these questions, then that's the appropriate level of analysis to think about how spike trains are important for our sensory coding and for motor behavior.