[SQUEAKING]

[RUSTLING]

[CLICKING]

JOSH MCDERMOTT: So today, we're going to wrap up talking about attention and then talk about touch. So last time, we got started talking about effects of attention that you can see in the brain. So we went over this general finding that when you're paying attention to something, it can be a part of space. It could be a type of feature, a type of object. You usually get an increase in neural activity in the region that we think is processing whatever it is that you're attending.

So this is a case where attention is directed to particular parts of space. And the way in which it's directed to those parts of space is people are asked to perform tasks that on stimuli that can appear at those parts of space. So they know that there's stuff happening in those regions that matters. And then you can measure with fMRI activity in the visual cortex and see that there's increases in activity in those regions of visual cortex that represent the parts of space that attention is directed to.

We also saw how you can attend to particular types of objects. So this is an experiment where images of objects are superimposed. One of them is moving. And you're directed to attend to either the moving one or the static one. And in this particular case, these are images of faces or houses which preferentially activate these two distinct regions of the ventral visual cortex, the fusiform face area and the parahippocampal place area.

And so the bottom line is that when you're attending to the face image, the activity in the FFA is higher. When you're attending to the house image, the activity in the PPA is higher-- so same sort of idea as back here but now with respect to types of objects. So that's where we left off last time.

So what I want to do now is talk a little bit about what happens at the level of single neurons. So fMRI is a very, very coarse measure of brain activity. You're measuring the signal that is indirectly related to blood flow and attributed to small cubes of tissue known as voxels that each contain hundreds of thousands of neurons.

If you want to measure the activity of single neurons, you usually have to do this in non-human animals. And so this is a very influential study that was conducted by our own Bob Desimone, who's now a professor here in BCS and director of the McGovern Institute. So he did this when he was at the NIH in the early '80s.

So it was a very early study of the effects of attention on neural responses. This was done with macaque monkeys. So again, the methodology here is you've got an electrode in the monkey brain in a particular visual region. You isolate a particular neuron. And you can measure its response while the monkey is looking at stuff and doing things.

And so in these experiments, the monkey was trained to make a discrimination on a particular stimulus. And so specifically, there were two stimuli-- a sample and a test presented sequentially. And the monkey responded differently depending on whether the two stimuli were the same or different. So they're doing a discrimination test on stimuli that are presented sequentially. And so the purpose of that is, again, to incentivize the monkey to pay attention to one of the stimuli because that's the one that they're going to have to make this discrimination on.

All right. And the way that they set this experiment up is that the stimuli are chosen to either be an effective stimulus for the neuron that they're recording from or an ineffective stimulus. So effective here means that it produces a big response in the neuron, and ineffective means that it doesn't.

So here's a picture from the paper that shows how this works. So the dashed line here represents the approximate boundaries of the receptive field of the neuron that's being recorded from.

So initially, the way the experiment works is you isolate a neuron. You map its receptive field. So you figure out where the boundaries are.

And then you set up your stimuli so that there are two stimuli that are appearing within the receptive field. One of them is chosen so that it produces a big response on its own. And the other one is chosen so that it doesn't. But they're both inside the receptive field.

And then the key manipulation here-- so there's two key manipulations here. In this particular case shown here at the top, what they're manipulating is which of the two stimuli the animal is attending to. So on one set of trials, this stimulus, the effective stimulus, is the one that the animal is going to have to perform the discrimination on. And on the other set of trials, it's the other stimulus, the ineffective stimulus, that they'll have to perform the discrimination on.

So they're interested in what does this change in what you're paying attention to do to the response of the single neuron in the brain. And the responses for this particular example neuron instead of stimuli are shown down here. These are raster plots.

Anybody remember what raster plots are? What's this dimension represent in a raster plot? Time. What does this dimension represent?

AUDIENCE: Neurons firing.

JOSH MCDERMOTT: No, this dimension just represents different trials. So each row here is like-- is a different repeat of the trial. And the dot-- what do the dots represent? Spikes, yeah. So there's a dot every time there's an action potential. And so what you're supposed to be looking at is like whether there's a lot of dots or not very many dots.

And so the S here is the sample. So that's the first presentation. The T is the test. That's the second presentation. So let's just look at when the sample is presented.

And the point here is that the stimuli are identical in these two cases. But attention can be directed towards the effective stimulus or the ineffective stimulus. And the result is that when attention is directed towards the effective stimulus, you get a larger response. So there's more dots here than when it's directed towards the ineffective stimulus. So there's fewer dots here.

Interestingly, these effects are really most pronounced when both of the stimuli are in the same receptive field. So this is one of the findings from this paper. So here's another thing that was manipulated.

So in this situation, the two stimuli are positioned that one is inside the receptive field, and one is outside. So it's the same attentional manipulation. People are either attending to the affective stimulus or the ineffective stimulus. And the argument is that in this setting, there's much less of a difference between the responses. So that's the empirical finding.

So there is this effect of what you're attending to on the response of single neurons. And those effects are most pronounced when there are multiple stimuli that are within the receptive field. And so one way to think about this is that attention is causing the receptive field effectively to shrink-- to shrink around the attended stimulus.

And so in this particular case, it shrinks around the effective stimulus. And so you get a big response. In this case, it shrinks around the ineffective stimulus, and you get a smaller response. All right. So that's just one example.

This is a summary of that effect for three different visual areas. And so what they're doing here is computing this thing that they call an attenuation index, which is a ratio between these two responses. So it's going to tell you the extent to which attention is causing the response to change. And so if attention is not having any effect, this ratio is going to tend to be one. And if it is having an effect, it will tend to be less than one.

All right. And so here's that ratio plotted for area V4. Here it is for TE, and here it is for V1. Now for V4, they did two versions of the experiment. One in which both of the stimuli, the one that's attended and the one that's ignored, are inside the receptive field, and that's what's shown here. And what you're supposed to take away from this is that the histogram is peaked down here near zero.

So most of the neurons are modulated by attention. Here is the variant of the experiment where one stimulus is inside the receptive field, and one stimulus is outside. And so here the result looks fairly different. Now you get a histogram. It's kind of peaked near one-- so much less attentional modulation.

Now they looked at the same thing in area TE, which is part of inferotemporal cortex, one or two stages beyond V4. Now, when you get to inferotemporal cortex, you probably remember that the receptive fields are very large. So it's actually not easy to do the condition where one stimulus is in the receptive field, and one is outside. So all they've got is a situation where both stimuli are inside the receptive field. And again, you've got a histogram where most of the mass is down here at small numbers, so you're getting lots of attentional modulation.

They also did the experiment measuring responses from neurons in V1. That's panel D. Now in V1, the receptive fields are very small. And so you can't really fit two stimuli inside the receptive field very easily.

So they only did the condition where one stimulus was inside and one was outside. And there again, the histogram is peaked kind of at one, indicating that you can't really get much attentional modulation in V1.

So take-home message is that you can see attentional modulation in single neurons. The effects seem to be biggest when multiple stimuli are in the receptive field. And it's kind of as though attention causes receptive fields to shrink around the attended stimulus. And so I'll just note that this is loosely consistent with that effect that we talked about last time of attentional queuing on texture discrimination.

So remember, we looked at those results with texture discrimination where you put up a queue, and the pattern of results is consistent with the idea that the queue causes the receptive fields to shrink, which when the target is presented at the fovea makes you worse, and it causes the eccentricity at which you get the best performance to shift towards greater eccentricities. All right. Questions about that? Yeah.

AUDIENCE: We've done texture discrimination. [INAUDIBLE] Can you describe this effect again?

JOSH MCDERMOTT: The thing that I'm referring to here? Oh, yeah. I'll just go back quickly to that slide.

That's this experiment with this result where people have to detect the distinct texture region. And when you cue people to the location and you look at performance as a function of eccentricity, you find that cuing shifts the location on the screen where you get the best performance to greater eccentricities. And this is the task where there's reason to think that the variation in performance with eccentricity is really due to the scale of the filters at that eccentricity.

Make sense? Yeah. Let's talk a little bit about attention and eye movements.

So normally in vision, attention and eye movements are very closely linked. So you typically are kind of paying attention to whatever you're looking at. And that's in part because usually, if you're attending to something, it's because you want to see it well. And normally, to see something well, you move your fovea to that location because that's where you got the highest density of retinal sampling.

So attention often accompanies eye movements, but there are some important and interesting differences between attention and eye movements. So one is that you can attend to different spatial scales. So you could direct your attention to a particular location, but I could indicate to you that whatever is going to be happening is going to be happening in a very small spatial region, or it could also indicate that it could be happening in a spatial region. And you can show that there will be a benefit that would occur at different scales depending on those instructions.

As we saw earlier in the lecture, you can also attend to more than one thing at once, whereas you can only really look at one thing at a time. And of course, you can attend to things in other modalities where eye movements are not relevant, possibly via some central mechanism.

So these differences imply that the brain mechanisms of attention must in part be distinct from those for eye movements, but it remains likely that attentional mechanisms overlap to some extent with those for eye movements. And maybe actually attention evolved as an extension of the mechanisms for eye movement. So that's one reasonable hypothesis.

So just to give you a taste of some of what we know about this. This is a pretty interesting study by Tirin Moore, who was once a post-doc here in BCS and is now a professor at Stanford. And in this study, they looked at the effect of stimulating in a part of the brain that controls eye movements on visual responses in visual cortex.

So we haven't really talked about the brain circuits that control eye movements. It's a little bit beyond the scope of the course. But there's a few different parts of your brain that control your eyes, and one of them is called the frontal eye fields, or FEF. So that's like up there. So it exists in the macaque. It exists in humans.

And one of the reasons why we believe that that brain region controls eye movements is that if you have an electrode positioned in that brain region and you stimulate-- so you put some current into the electrode, which causes neurons at the tip of the electrode to fire-- that has the consequence of causing the animal to move its eyes to a particular location. So you stimulate the particular site in the frontal eye fields. Boom, the eyes go to a particular retinal location. And that's what's shown here.

So this is an experiment where there's an electrode positioned in the frontal eye fields. There's also an electrode positioned in V4 that's recording responses, which we'll talk about in a second. But when you stimulate the electrode in the frontal eye fields, this is a plot of what happens to the eyes. So each of these different lines is like a different trial. And what you're supposed to derive from this-- so this is the fixation point, and the animal's eyes move kind of down and to the left by about the same amount.

And so the idea is that normally, like when you're moving your eyes around, what's happening is there's this sequence of activity in the frontal eye fields. It's like planning all your eye movements and then sending these motor commands to your eye muscles to cause them to move in particular ways.

So in this experiment, as it says on the slides, there's electrodes in both the frontal eye fields and V4. And they're recording from pairs of neurons. And they move the electrodes around until they find pairs of neurons that are matched. And they're matched in the following sense.

So they're finding a pair of pairs of neurons, one of which is in V4, that has a receptive field at a particular location, and that receptive field location is approximately matched to where the eye movements go to when you stimulate at that location in the frontal eye fields. You're at this locus in the frontal eye fields where the neurons are coding for eye movements to a particular location. And then you're recording from the same-- the corresponding place in V4 whose neurons are representing that location in space.

All right. So that's the setup. And this is a hard experiment to do because you've got to get each of the electrodes in the right place in order to find these neurons. But that's what they did.

And I've just told you how if you stimulate in the frontal eye field above a certain threshold causes saccades. And so what they did in this study is to stimulate but below the threshold that evoked a saccade. So they stimulate the frontal eye fields but not so much as to cause an eye movement. And what they found is that subthreshold stimulation of the frontal eye field at that particular location causes attention-like enhancements of the response in V4. And so that's what's shown here.

So they've got stimuli that are being presented in the receptive field of the V4 neuron. And they're measuring the responses with and without frontal eye field microstimulation.

And so the red response is what happens following the microstimulation. And the black response is what happens without the microstimulation. And the eye traces-- so the H and the V represent the horizontal and vertical position of the eye. And so what that's supposed to tell you is that the stimulation is not causing the eyes to move. So the effect is not because the eyes are moving. It's something else.

And the claim is that that kind of enhancement of the response is a lot like what you would see if you were paying attention to that location. And this is showing that if there's no stimulus in the receptive field, then you don't really get much of an effect.

And so the inference here is that the circuitry that controls eye movements could mediate some attentional effects. So you might say, well, why is that? Well, possibly because normally, attention does move with the eyes, so they should be coordinated in some way. The eye movement circuitry also has some of the characteristics that you might want attention to have, in the sense that it can be spatially directed. Because you can look at-- move your eyes in all these different directions.

So it's an interesting hypothesis. So far, we've been talking about neural correlates of attention. So what changes in the brain when you're paying attention to something? And what we just discussed about eye movements kind of is one potential proposal for where attentional effects might originate from. But there's another part of the brain, the parietal lobe, that also seems to be really important.

So the parietal lobe is typically active when people are asked to move their attention around. And when the parietal lobe is damaged, you see disorders that are often interpreted as disorders of attention. So remember, think of the visual system as having the ventral and dorsal stream organization. The parietal lobe is up here. So it's where the dorsal stream kind of ends up.

And lesions to the parietal lobe result in a phenomenon that's typically referred to as neglect. And it's called neglect because one of the key symptoms is that people neglect one side of space. It's like they ignore it.

So this is an excerpt of a case study. It says Mr. P neglected the left side of his body and of the world. When asked to lift up his arms, he failed to lift his left arm, but could do so if one took his arm and asked him to lift it. When asked to draw a clock face, he crowded all the numbers onto the right side of the clock.

When asked to read compound words such as ice cream and football, he read cream and ball. When he dressed, he did not attempt to put on the left side of his clothing, and when he shaved, he shaved only the right side of his face. He ignored tactile sensation on the left side of his body. Finally, he appeared unaware that anything was wrong with him and was uncertain as to what all the fuss was about.

So this is interpreted as an attentional disorder because it's not the case that this individual is blind on the left side of space. If you do things to direct their attention over to that side of space, they can see things. But it's like they don't on their own pay attention to it and don't notice it. So here's an example of what happens if you ask them to do certain things.

So this is a line cancelation task-- so part of a standard assessment. So you give them a sheet-- a piece of paper that's got all these line segments. And you ask them to put a mark through every line segment. And this is what somebody with neglect would do. So they X out a lot of the line segments on the right, but none of the ones that are on the left.

You can ask them also to do a line bisection task. So put a marker through the middle of the line segment, and they're way off the center. You can ask them to copy pictures. So here's the model. It's a clock. This is an actual drawing by a patient with neglect.

Here's another example. This is a house. They leave out a lot of the stuff that's on the left side. Flower-- and again leaving out some of the stuff that's on the left side.

So neglect typically and the parietal lobe lesions that result from it typically would result from a stroke. So remember what happens when people have a stroke? The brain is deprived of blood flow and oxygen and dies pretty quickly. And so depending on what type of stroke you have, there'll be different parts of the brain that are damaged. And so parietal lobe lesions cause this phenomenon of neglect.

Now one of the things that often characterizes brain damage is that there's plasticity. So there's often some process of recovery. So this is a really interesting example of somebody who is an artist who suffered right parietal damage.

And this is a series of self-portraits over the course of their recovery. And so initially, you can see that they're entirely just drawing stuff that's on the right side. And then as they recover, they're able to get more and more of the stuff that's on the left.

All right. So you might ask, so what is it? What's special about the left? I think that's not totally well understood. So in general, neglect tends to be much more a left visual hemifield thing. But it's just a phenomenon at this point. Yeah.

AUDIENCE: Like the line bisection task-- if you move the paper over to the right, they will do the whole thing?

JOSH MCDERMOTT: So OK, so one of the really interesting things about neglect is that it's often, at least to some extent, object-based. So it's not just about whether the stuff is in the left hemifield-- so to the left of fixation. So they'll neglect the left part of objects.

So you can have an object that's in the good hemifield. And they'll still tend to neglect the left side of it. Yeah. And so it's yeah, one of many pieces of evidence in support of the idea that attention is kind of object-based, that you tend to direct your attention to objects. But yeah, you can neglect sides of objects too. Any other questions about that?

So we talked about a lot of stuff because attention is kind of rich and diverse and not very theoretically unified at this point. So this is just a quick recap.

So we talked about how attention has been studied with lots of different paradigms. We just discussed queuing and the distinction between exogenous and endogenous queuing. Talked about how queuing results gave rise to a spotlight metaphor of attention. Talked about exceptions to the spotlight metaphor, like multiple object tracking, where you seem to be able to direct your attention to multiple locations at once. We talked about the paradigm of visual search, where there's an array of elements, and you have to find the one that's kind of different from the others.

Talked about differences between feature and conjunction searches, where feature searches tend to be fast, and the reaction time tends not to scale with the number of items in the display. Whereas conjunction searches tend to be slow, and reaction time increases as the number of elements in the display increases. Talked about the idea of feature integration theory, which was traditionally used to explain the differences between feature and conjunction searches and illusory conjunctions as evidence that you need attention to bind features.

Talked about evidence that attention improves spatial resolution and increases apparent contrast. Talked about evidence that attention seems to dilate time. Saw that it speeds responses, integrates features. Talked about how there is an attentional blink. So if you select something with your attention, there's a little bit of a refractory period before you can grab the next thing.

Talked about change blindness. These are instances where you don't see changes without attention. So you can miss really salient things in the image. Talked about selective listening and evidence that there-- for things that you're not paying attention to, there's pretty strict limits on the extent to which it gets processed.

And we talked about the neural basis of attention. So some of the key findings here are that attention to something generally increases the response to that something. Attention often acts as though it shrinks receptive fields around the attended stimulus. Eye movement circuits seem to be able to mediate attention-like neural effects. The phenomenon of neglect results from lesions to parietal cortex presumably, which is the locus of some critical attentional mechanism.

And so I'll emphasize that we're using this one word attention to talk about all this different stuff. But it's probably not just one thing, but rather a loose assortment of related phenomena. We mostly lack computational models of attention at present. I think this is a really important next frontier. And we're actually well-positioned as a field to really be able to make progress on that. And that's where we'll leave things.

We got 1 and 1/2 classes left, and we've covered two of the five senses. So we spent about the first third of the class on audition. Since then, we've been talking about vision.

So you may ask, well, why is that? That doesn't seem fair. But our coverage is, I would say, arguably proportional to the density of research on the different senses.

So vision is by far the best-studied sense. I mean, audition is pretty well-studied. Touch, taste, and smell are much less well studied. And I mean, it's mostly because they're harder, I would say. The stimulus is like a lot harder to manipulate.

So if you want to study vision, you just make a picture, and you show it to people. And we have computer screens and stuff that make that really easy. And if you want to study sound, give people headphones, and you can make a sound waveform and play it to them. You all have done that.

If you want to study touch, you have to generate tactile stimulation. And that's hard to do with a cheap device that would be easily available-- same with taste and smell. You got to generate this chemical or something and give it to people. So everything is a little bit more complicated.

But they're all really interesting and really important. And the sense of touch is critical to all kinds of things that we do on a daily basis. So this is a slide that shows a bunch of things that we do that utilizes our sense of touch. So if you are holding a pen to write, you have to grip the thing with the appropriate grip strength so that you hang on to it, that you can still-- and move it around in the way that you want to.

And so there's feedback that you get from your sense of touch. It's kind of critical for holding the thing in the right way. So if you're playing instrument, like a guitar, in order to pluck those strings, you have to be able to feel like the position of the string with respect to your fingertip. The sense of touch is critical.

You want to throw a baseball? Again, you got to be able to grip it exactly the right way, position your fingers in exactly the right way if you want to throw a curveball or a slider or whatever. You're holding a tweezers. Got to again, have to grip that in the right way.

Manipulating a mouse-- getting the mouse kind of positioned in your hand correctly. That depends on a sense of touch. If you're typing on a keyboard, the sense of touch is what helps you get your fingers on the keys in the right way.

If you'd like to play video games, these controllers, you have to hold and get your fingers on the button. Playing golf-- sense of touch helps you figure out how to hold the golf club in exactly the right way. Holding a phone and so on and so forth. So we do lots of things on a daily basis that uses our sense of touch.

This has got some other examples of what kinds of information we get from touch. So take your fingers and move them over something. So I'm moving them over my shirt. You can get a sense of what the material is just by that lateral motion.

So you get a lot of information about texture. We were talking about texture and in vision and hearing. It's also really important for touch.

Take your finger and push it into something around you. So you get a sense of how hard something is just by kind of tapping it. So that's again giving you information about material.

Take your finger. And again, put it on something around you, and you'll be able to get a sense of the temperature. So the metal thing on the chair might be a little bit cold. Your body is probably pretty warm. A cup of coffee is probably pretty warm. So again, there are sensors in your skin that give you information about temperature.

Grab something and hold it in the palm of your hand. Just from holding it, you get a sense of how heavy it is. Again, that's a sense of touch in conjunction with information from your muscles. Just by grabbing something like this cup, you have a sense of how big it is, what the shape is.

We also get a lot of information about shape by kind of running our fingers around stuff. So these are all different things that we do with our body to give us different kinds of tactile stimulation and derive different kinds of information about the world.

So this is a graph that shows some surface depth maps for three different materials. So the y-axis here is depth. The x-axis is lateral position along the surface. Or well, and yeah, in this case, this is being swept along. So it's as though you're moving your finger across something-- so time and/or space.

The point is that for a surface that's very, very smooth-- I mean, there are variations in the surface depth, but they're kind of small. So glass and silk, you have these pretty small indentations. Or something that's a lot rougher, like sandpaper-- indeed, if you look at it a fine scale, you can see these variations in the surface depth map. And that's something that you can sense if you move your finger over it.

All right. So how does all this work? So big picture here is like you've got sensors that are in your skin. Those sensors send signals. They travel up very long nerve fibers to the spinal cord. The spinal cord relays that information to the thalamus. The thalamus projects to primary somatosensory cortex or secondary somatosensory cortex, and that is involved in the analysis of that touch information.

So a lot of what we know about touch, and a lot of what I'm going to talk about today-- and in fact, this is true for all three of the remaining senses. A lot of the action is at the level of receptors. So much less is known about what's happening cortically to analyze the information that comes from the receptors. So a lot of what we'll talk about is the receptors and what kinds of information they capture, but not exclusively.

So in the somatosensory system, like at least within the skin, skin simulation is measured by four main types of receptors. So they each have different names. You don't have to memorize the names but should be familiar with the conceptual distinctions between the different receptors. So we've got the Meissner corpuscle, the Pacinian corpuscle, the Merkel disk, and the Ruffini ending.

So is a picture of the skin. So skin's got a few different layers to it-- the epidermis and the dermis. So the touch receptors are kind of beneath the outer layer of skin. And then they're situated at different depth within the skin. And that will impart-- give them kind of different properties. And then each of these has axons that then travel up to the spinal cord.

So there's different kinds of mechanical stimulation that can result that these things will sense. And the color-coding here-- and then the color-coding of the receptors is supposed to indicate that the different types of receptors are differentially sensitive to different types of stimulation.

So you can have heavy pressure. You can have vibration. You can have light touch. You can have stretching of the skin. And these different types of receptors are maximally sensitive to each of those.

Here's another picture that more or less shows the same thing. So this is like a little chunk-- a zoom-in on a little chunk of skin and say, the fingertip, and you've got these different receptor types that are situated at different depths within the skin.

So you will remember how at various points in the class, we dealt with some kind of annoying nomenclature issues where you have things that are given different names by different people, and they're often not used very consistently. And so it's the same situation here. The nomenclature that has been used to refer to these things kind of has evolved over time. And with most of these receptors, there's a few different ways that they're referred to.

And so the Meissner corpuscle is often referred to as the RA receptors or RA fibers. That stands for rapidly adapting. Pacinian corpuscles are usually just called PCs, and then the Merkel disks are usually called the SA1 receptors because they slowly adapt. Sometimes they're referred to in functional terms, other times with the name of somebody who discovered them.

So tactile receptors are called mechanoreceptors because they respond to mechanical stimulation-- so pressure, vibration, or movement. And the four main types of mechanoreceptors, they kind of vary on two main dimensions. So that's the kind of conceptual thing that everybody should know. They vary in terms of how quickly they adapt to stimulation, and they also vary in the size of the receptive field. So you end up with a 2 by 2 grid, which in fact, we'll see on the next slide.

So the Meissner corpuscles exhibit fast adaptation and have small receptive fields. The Merkel complexes is a bit slow adaptation but with small receptive fields. And then the Pacinian corpuscles and Ruffini endings have fast and slow adaptation, respectively, and large receptive fields.

Now one complexity here is the Ruffini receptors supposedly don't exist in monkeys, and a lot of the research on touch, of the neurophysiology of touch, happens in monkeys. And so you're not going to see the Ruffini receptors on a lot of the slides. So with monkeys, you mostly study the three.

All right. So each of these receptors has a different range of responsiveness and functionality. And then, to further complicate things-- so sometimes they use the word fast to refer to adaptation, sometimes use rapid. So you'll have RA and FA. Those mean the same thing. But the main point here is that we've got a 2 by 2 grid. So we have different rates of adaptation and different sizes of receptive fields.

All right. So what do we mean by this? Well, the receptive field is just what it sounds like. So this is like the area of the skin that would evoke a response on a receptor. And some of the receptive fields are small, and some are larger. And that's mostly a function of how deeply the receptor is positioned in the skin. So the ones that are more deeply positioned in the skin will be sensitive to stimulation over a wider area. If they're kind of close to the surface of the skin, the receptive field will typically be kind of small.

So this picture is showing some example receptive fields for the four main types of receptors. So we've got the Merkel receptors, which are also known as SA1. The Meissner corpuscles, which are known as RA. And the point is that those receptive fields are small. And then we have the Ruffini endings and the Pacinian corpuscles, or the PC receptors, and those are large receptive fields.

And this is a plot that shows the density of these receptors across the hand. And we're going to return to this in a moment because the density varies quite enormously over your body. So there's some areas of skin.

And so it's a bit like the fovea in the retina. So there's the fovea is this particular part of the retina, which has got a ton of receptors. So you get really high-resolution sensing there. The fingertips are sort of like the fovea for touch. It's got tons of receptors in your fingertips.

But there are other areas of the body that are also pretty densely innervated. And the hands in general are pretty sensitive, but the fingertips are very sensitive. And so in particular, these two receptor types, the ones that have small receptive fields, are very dense in the fingertips. So they're doing kind of high-resolution analysis of touch. Whereas the Ruffini and the Pacinian are not really as densely concentrated in the fingertips because their receptor fields are-- receptive fields are bigger, and I guess there's less point to having them be very dense.

All right. So the adaptation that we were talking about that also differentiates these receptors, that's shown in the top right, panel C. So there's a stimulus. That red line indicates pressure that's being exerted on the skin. So you have a little rod, and it gives you a poke. So it presses into the skin. That's when the red thing goes up.

The two graphs below that are showing the response of a slowly adapting receptor and a rapidly adapting receptor. And the vertical lines there indicate spikes. And so the point is that if you have this fixed pressure stimulus, the slowly adapting receptor responds a lot, and the response decreases. But it kind of continues constantly. Whereas the rapidly adapting ones, the response kind of goes away very, very quickly. So mostly our signaling changes in pressure.

All right. So the different receptors are maximally sensitive to different features of touch stimulation. And so this is a table that is showing the type of stimulation that the four different types of mechanoreceptors are maximally sensitive to-- so sustained pressure, very low frequency, sustained downward pressure, temporal changes in skin deformation. And so essentially, the ones that are rapidly adapting are very sensitive to temperature changes but in different frequency ranges.

So the receptor sensitivities constrain perceptual sensitivity. So this is a graph that I like to show because it's got some interesting analogies to stuff that we have seen elsewhere. So this is kind of analogous to the contrast sensitivity function.

So what's being plotted on the y-axis is threshold. So it's a little different. The contrast sensitivity function is like plotting the inverse of threshold. So this is just the threshold. So now, down is better.

And it's applying the threshold as a function of the frequency of stimulation. So you've got a little rod that's making contact with the skin. And it's vibrating back and forth at some frequency. And then you can control the amplitude of the vibration and measure what that amplitude needs to be in order for you to detect it. And so the dots here are plotting human thresholds.

And so what this is showing is that peak sensitivity is actually occurring for fairly high frequencies. And then it gets a bit worse when the frequencies get too high. And it's worse when the frequencies are low.

Superimposed on this graph of human thresholds, which are the solid dots, are the dashed lines, which are the sensitivities of the individual types of receptors-- SA1, FA1, and FA2. And what you're supposed to take away from this is that at any given frequency, there is a receptor that's kind of maximally sensitive to that frequency, and that's what ends up dictating your perceptual sensitivity. So it's like the contrast sensitivity function being determined by a combination of spatial frequency channels. Is there anything else that this kind of is vaguely evocative of?

AUDIENCE: The psychometric curve but like two of them are connected.

JOSH MCDERMOTT: Which psychometric curve?

AUDIENCE: The one that we always--

JOSH MCDERMOTT: The classic.

AUDIENCE: Where you [INAUDIBLE]

JOSH MCDERMOTT: Two sigmoids. Yeah, OK, a little bit, yeah. There's something else that this is kind of reminiscent of. Yeah.

AUDIENCE: Rods and cones.

JOSH MCDERMOTT: Exactly-- yeah. So it's a little bit like the dark adaptation curve. So remember, if you walk into a dark room and then you measure your sensitivity to light as a function of the amount of time you spent in the room, the thresholds kind of drop a little bit, and then they drop again. And so that curve is explained by the sensitivity of the cones and the rods. And so initially, your sensitivity is determined by the cones. And then when the rods become more sensitive, the rods take over. So it's a little bit like that.

So the point is that these different receptor types are sensitive to different kinds of stimulation and collectively determine your perceptual sensitivity. So now is the time for an experiment. So I mentioned to you that there is variation in the density of receptors over the body.

So we saw some examples of this in the fingertips. But there's variation across the entire body. And this causes very large variations in how-- in the acuity of your touch perception. And so this is standardly measured with a two-point sensitivity test.

So the idea is that you stimulate the person with either kind of one poke or two pokes that can be some distance apart. And you're trying to measure the threshold, the threshold for the distance between the two points at which you can tell that you're being poked twice instead of being poked once. So clearly, eventually, if you move the things far enough apart, you'll be able to tell that there are two pokes. But then it also clearly, if they get really close together, it would be indistinguishable from one poke.

And so this graph is plotting the threshold for distinguishing 1 and 2 pokes at all these different places on the body. And you can see that the fingers and the thumbs, the threshold is pretty good. So this is 5 millimeters. So that's saying that if things are 3 millimeters apart, you can tell whether you got two pokes or one.

By contrast, the arms, lower arm and upper arm, that's like 40 millimeters, so 4 centimeters apart. If things are 3 centimeters apart, in general, you won't be able to tell whether you got one or two pokes. Seems kind of crazy. And the back and the thigh and the calf are also pretty bad.

So the other place where sensitivity is really good is the face-- so in particular, the lips. So the lips have to have very good touch acuity, presumably because in order to eat and make sure the food goes in your mouth and stuff, you have to know where food is relative to stuff. And I guess the rest of your face is very sensitive, just because there's a lot of important stuff on the face, and you want to know if something's on you, I guess.

So all right, let's do a demonstration. So would anyone like to volunteer to be tested for touch acuity? It'll be a lot of fun. Do you want to volunteer? Yeah, come on.

Oh, the only thing is you're going to have to have a bare arm. Yeah, that'll work. Yeah, that'll work, yeah. Come on up here.

Thanks for volunteering. This will be painless. OK, yeah, I don't want to draw on her. So just make sure this doesn't hurt. Is that OK? OK, good.

All right, so we have two pencils here-- mechanical pencils. So the tips are pretty thin. So we're going to measure two-point discrimination. So turn your hand upside-down. And move it back just a little bit. There we go.

So close your eyes. And we're going to do this first on the fingertip. So you're going to tell me if there's one poke or two poke.

STUDENT: One

JOSH MCDERMOTT: Yep, very good.

STUDENT: Two.

JOSH MCDERMOTT: Very good.

STUDENT: One. I can't really tell-- maybe 2, 2, 1, 2, 2.

JOSH MCDERMOTT: So she's able to tell them apart when they're about this far apart, a few millimeters-- so pretty consistent with the graph up there. So now we're going to repeat the experiment on your forearm. Got to close your eyes.

STUDENT: Oh, I don't know. One or two, I can't really tell.

JOSH MCDERMOTT: OK, that's fine.

STUDENT: Two.

JOSH MCDERMOTT: OK.

STUDENT: 2, 2, 2, 1, 1, 1, 2, 1, 2, 1.

JOSH MCDERMOTT: So the threshold is about that.

STUDENT: Oh, interesting.

JOSH MCDERMOTT: Yeah, maybe two or two or three centimeters. All right. Yeah, you're a little bit better than the participants in the textbook. So yeah, but you get the point.

So there are these huge differences in perceptual sensitivity. And it's a function of the touch receptors. Thanks very much for volunteering. Any questions about that? Yeah, round of applause. All right. Any questions about that? Yeah.

AUDIENCE: Is this related to the amount of muscular control? Do you think fingers have a lot of fine muscles, and the face has a lot of fine muscles? But then, like the legs, I guess they have fine muscles, but you don't do as many dexterous things with them.

JOSH MCDERMOTT: I mean, they're not related in the sense that what's mediating your sensitivity are mechanoreceptors in the skin. I mean, you do have receptors in your muscles, as we'll talk about in a second, but that's not what's determining your discrimination in this particular case. So they're not proximally causally related in that sense. But they might be related just in the sense that the places where you need to have a lot of fine motor control are probably also places where you want a lot of information about what you're touching.

So those things probably go hand in hand. And the places that are like the phobias for touch are also places where you want to have a lot of musculature because you want to be able to move them around to be able to touch things and so forth. Yeah. Yeah, any other questions? OK.

So your skin is covered with these receptors. Actually, some of them have pretty small receptive fields. And the set of all of the receptors on your skin is, in a sense, going to give you an image.

And so there's been a lot of experiments showing that the receptive field-- the receptors accurately transmit spatial structure to some extent. And the way that these experiments typically work-- so you might imagine that what you'd actually want to do is be able to simultaneously record from all of the receptors in the skin, and that's not technically feasible.

So typically, the way these experiments will work is there's an electrode that will be in the arm of the animal subject. And you isolate the nerve that contains all of the nerve fibers that are carrying signals up to the spinal cord. And you hook into one of them. So you're typically recording from one receptor at a time.

But what you can do is you can measure the response of that one receptor as it is swept across, like a tactile image, if you will. This is a classic experimental paradigm for doing just that.

So the way that this works, this is the monkey's hand. And this is a cylinder that is embossed. In this case, it's embossed with letters. All right.

And so the cylinder rotates around, sweeping over the monkey's finger. And then you move the cylinder back and forth such that the finger stays in one place, and with different sweeps of-- rotations of the cylinder, the finger comes in contact with different points on the image. And each of those kind of gets swept around.

And so the consequence is that you can effectively measure the response of the receptor like at every point in this image. And the inference is that would capture what the full set of receptors would be capturing when the thing gets swept over the fingertips once. Everybody get that?

So what happens here when you do this? And so this is a schematic of how this works. So the green circle is the receptive field of the receptor that's being recorded from. And these different sweeps are kind of different rotations of the cylinder positioned at different locations with respect to the finger.

And so the idea is that when there's a ridge on the surface that comes in contact with the receptive field, you're going to get a response. And so you can display those responses spatially to get a visual sense of the extent to which the action potentials from the receptor are actually accurately capturing the spatial structure of the tactile stimulus. And so in this case, the point is that this kind of looks like a letter G.

So this is a summary figure that shows the responses of the three main types of receptors. So remember how I told you in monkeys, they don't have the Ruffini receptor as far as people know? So there's really three main types of receptors-- SA1, RA, and PC. So it's the response of those three different types of receptors to the first few letters of the alphabet.

And what you're supposed to take away from this is that the receptor types that have small receptive fields, so the SA1 and the RA, are doing a pretty good job of capturing the spatial structure of the letters. But you can see that the responses are qualitatively different. So they're kind of responding to different things.

You can also see that the PC receptors-- those are the ones that are deeper in the skin. They've got bigger receptive fields. Those are actually not doing such a great job of capturing the letters. They're responding to larger-scale spatial structure rather than fine-scale spatial structure.

All right. So this is for letters. And the main reason to use letters is that they're familiar visually to humans. And so you can kind of just glance at the picture and easily understand it.

A more typical tactile stimulus would be a texture. So we use our sense of touch to identify material. And a big part of that is texture.

So these are different textures. So these are like Braille dot patterns on one of these drums. So it's the same kind of experiment. You've got this rotating cylinder with an embossed surface on it. But now, the embossing is these dot patterns, and they vary in their density. And again, we've got the three main receptor types that you see in monkeys.

And what do you see? Well, when the dots are very sparse, you get spatially localized responses to the individual dots in the SA and RA receptors, the ones with the smaller receptive fields. You don't really see much of anything in the PC receptors. But you can also see that as the dots become more dense, you stop being able to see of individual responses to the individual dots. It's like the texture kind of gets too fine for there to be a spatial representation of the texture in the receptor responses.

So that's for some artificial textures. These are a few examples of real-world textures. So well, this one is artificial embossed dots.

This is something called Huck towel. I guess it's a type of fabric. Got nylon and chiffon. So these are all things that if you felt them, they would feel different.

But you can see that really it's only the very kind of coarse texture features that end up showing up in this spatial pattern of the responses. And so this fine-grained stuff that you can see in nylon and chiffon just doesn't really make it into the receptor responses as analyzed spatially.

So how can we account for the fact that these things all feel differently? So what we believe is the key is that there is temporal information in the receptor responses, even in cases where things are not resolved spatially. So these are, again, kind of depth maps of, in this case, three types of cloth. So there's premiere velvet-- sounds like it might be very smooth-- drapery tape, and blizzard fleece, which I guess is a little bit rougher.

These are spectrograms. Remember spectrograms? We all love spectrograms. So we use those to analyze audio. So what happens when you're feeling a texture is that spatial depth map kind of gets turned into a time series because you're sweeping this thing across the receptor. So if you assume that it's moving at a certain rate, you can then plot-- turn it into a spectrogram where you have the frequency as a content as a function of time.

And the point of showing you it in this way is that in the frequency domain, these three materials kind of have pretty different signatures. So the premier velvet has got like a lot of power fairly consistently at this one particular frequency. The drapery tape has got some higher frequency and then some stuff at lower frequencies. And then the blizzard fleece is kind of a little bit more generally low-pass.

So this is just like in the physical stimulus. So these different materials, when you look at them in the frequency domain, they have different signatures.

So how would this be transduced? So remember phase locking? We all love phase locking. It's a characteristic of the auditory nerve where the action potentials are time locked to the stimulus. What's the upper limit of phase locking, roughly speaking, in the auditory nerve?

AUDIENCE: That's 4000.

JOSH MCDERMOTT: 4000 hertz-- yep, exactly, roughly. So this is phase locking in touch receptors. So we've got a sinusoidal stimulus.

So here you can see the stimulator here. It's like a little piston-- just goes poke, poke, poke, poke, poke, moves back and forth sinusoidally. You can vary the amplitude of stimulation and the frequency. So this is a sinusoid at 400 hertz at three different amplitudes measured in micrometers.

And this is a plot of action potentials. And what you're supposed to take away from this is that the action potentials here are time locked to the stimulus. So this is phase locking a lot like what you would see in the auditory nerve. All right. So there's some temporal information in the response.

So does this actually have relevance for real-world textures? So this is one of those textures we were looking at before, nylon, along with a couple others-- stretch denim and silk jacquard, different cloths. At least the top ones have pretty fine-grained structure. And so that wouldn't really show up in the spatial pattern very clearly.

These are raster plots. So remember, this is time. And the different rows are different trials. And then this is the frequency spectrum that you see here from those rasters.

And what you're supposed to take away from this. So this is the frequency spectrum of the neural response. And what you're supposed to take away from this is that there is information that differentiates these different textures in the temporal pattern of the response, specifically in the frequency at which the spikes are fired.

So the nylon gives you some peaks at these two points in the spectrum. The stretch denim has a peak at a different place, and it's a little bit more spread out. The silk jacquard has got some other kind of pattern to it. So even though the spatial pattern isn't really enough to tell apart these very fine-grained textures, the temporal pattern is.

And so the general idea here is that there are these two components by which we sense texture and probably just other kinds of patterns more generally. There are fingertip deformations that cause a spatial image in the receptors that have very small receptive fields. And then there are high-frequency skin oscillations that cause spike patterns in the rapidly adapting receptors, the ones that are very sensitive to rapid changes in stimulation. And those get combined in some unspecified way that to yield the percept.

So texture is important in touch. We've talked a lot about it in vision and audition. Now in vision and audition, there are actually pretty well-developed theories of texture. So we don't really have that yet for touch. All we're talking about here is this idea that there is information being conveyed by receptors that would probably be diagnostic of texture.

And it would be really cool to actually try to build a model of tactile texture perception, kind of in the spirit of the ones that we have in these other domains. But that hasn't happened yet. So that'd be a cool thing for the future.

So we got the mechanoreceptors in the skin. There's four main types that we talked about. But there's lots of other types of mechanoreceptors. Specifically within muscles, tendons, and joints, there are kinesthetic receptors that play an important role in our sense of where our limbs are and what kinds of movements are made. So that's part of what's called proprioception.

There are receptors in muscle spindles that sense the tension in the muscle. We talked a little bit about that in the context of eye movements. So the receptors in tendons that signal the tension in the muscles that are attached to tendons, receptors in joints that react when a joint is bent at an extreme angle, all these are kind of things that help you control your body and make sure you don't break things if you don't have to. So here's just a picture showing some of the complexity of muscle fibers and the sensors that you have in there.

So there are some interesting kind of neurological patient cases that illustrate some of the importance of these kinesthetic receptors. So there's a famous patient called Ian Waterman, who got this kind of virus that caused the cutaneous nerves connecting his kinesthetic receptors to the brain to be destroyed. So this guy was not getting any input from proprioception. So he lacked kinesthetic senses, became dependent on vision to tell his limb position. So it's a very strange way to be.

All right. So those are kinesthetic receptors. There's also thermoreceptors. So these are sensory receptors that signal information about changes in skin temperature. So there's two main populations of thermoreceptors. There are those that are sensitive to warmth and those that are sensitive to cold.

So your body is constantly regulating its internal temperature. And thermoreceptors respond when you make contact with an object that is warmer or colder than your skin. So this is a picture that shows thermoreceptive fields for these two types of fibers-- the cold fibers and the warmth fibers.

So skin temperature is normally kept between 30 and 36 degrees Celsius. So it's regulated pretty tightly. And the cold fibers start to respond if you encounter something that's significantly colder than your skin temperature. The warm fibers respond when you encounter something that's significantly warmer than the skin temperature, telling you when the skin gets heated or cooled out of normal range.

So that, of course, is part of material perception. It's also, again, a safety mechanism so that you don't burn yourself or freeze yourself, so forth.

There's also nociceptors. So these are receptors that signal pain. So they transmit information about noxious stimulation that could damage your skin. So the key thing to note here is that there are two main groups of these nociceptors. So they have two different names-- A delta fibers and C fibers.

But so the key thing to know about them is that one of them is myelinated, and the other one is unmyelinated. They both transmit information about pain and temperature. And so the consequence of that-- so what does myelin do? Anybody remember? Yeah.

AUDIENCE: It insulates the signal of axons to make it travel a further distance.

JOSH MCDERMOTT: Yeah, it also makes things faster. So myelinated fibers carry signals faster.

And many of you have probably had the experience of something painful. So imagine you're walking around barefoot early in the morning. You haven't had your cup of coffee, and you step on a LEGO-- something that happens to me fairly routinely.

There will be two distinct stages to that experience of stepping on a LEGO. You get this kind of quick, sharp pain followed by a throbbing sensation. And that's believed to actually correspond to these two classes of fibers.

There's a very quick signal that's supposed to tell you to get your foot off the LEGO. So that's fast. And then this unmyelinated signal, that is, I guess, probably just telling you, is my foot damaged? Do I need to be careful for a while? And that one doesn't need to be as fast.

The difference is speed is due to myelination. So that's the key concept there. Two classes of these pain receptor fibers.

So again, there are these crazy cases of individuals that are born with insensitivity to some aspects of this. So a famous patient was born with insensitivity to pain. And this had major consequences for her life. She didn't sneeze or cough or gag or blink reflexively, suffered injuries, burning herself on radiators, biting her tongue when chewing food because these receptors are sending you signals to not do certain things. And this particular individual is from a long time ago, but died very young from infections that possibly might have been prevented if she had been able to sense pain.

So you've got this massive array of different types of receptors. Where do they end up? So they project ultimately to somatosensory cortex. So touch sensations are represented somatotopically in the brain. So this is another key concept.

And I should emphasize. So this is an idea that's been in all the textbooks forever. This has very recently kind of been challenged-- this idea that there is a full kind of map of the body that's laid out in somatosensory cortex and is kind of currently under debate. So it may well be that in another year or two, the story is going to change. But this is the classical story.

And so the classical idea is that there is a somatotopic map in somatosensory cortex that's analogous to retinotopic mapping that's found in vision, where adjacent areas on the skin connect to adjacent areas in the brain. And people will often refer to homunculus. It's like a map-like representation of regions of the body in the brain. And somatosensory cortex typically is divided into two regions-- S1 and S2. S1 is primary somatosensory cortex, kind of gets direct projections from the thalamus or most of them. And then, whereas S2 doesn't, or to a lesser extent.

So if you're interested in old-school neuroanatomy, S1 includes Brodmann areas 1, 2, 3a, and 3b. S2 is within the lateral sulcus. So remember, the central sulcus is this big sulcus here that divides the frontal lobe and the parietal lobe. And on one side of that is motor cortex. On the other side of that is somatosensory cortex.

So what this is showing is the regions of the body that are represented in different places within somatosensory cortex. And superimposed on this are the classical divisions of Brodmann areas 1, 2, and 3. And what you're supposed to take away from this is that kind of adjacent places on the body are adjacent on the cortex.

And so the kind of classic picture that it's so classic, many of you have already probably seen this. How many of you have seen this before? Yeah, this is shown in every introductory class. So you end up with this map of the body that's laid out.

Now the other thing that you'll see is that the amount of cortical real estate that is devoted to a body part is not proportional to the size of the body part. What is this kind of loosely analogous to that we've studied in another sensory system? Yeah.

AUDIENCE: Like the vision system [INAUDIBLE].

JOSH MCDERMOTT: Yeah. So we talked about cortical magnification in vision where there's much more area in the cortex representing the fovea than representing the periphery. It's the same idea.

And so again, in the fovea in the vision, there's lots and lots of receptors. They have to be packed together very densely because there's an image that is formed on the fovea. Similarly, in your fingertips or your lips, there's lots and lots of receptors that have to be packed very densely because they're supposed to be detecting stuff there. But then once you get to the cortex, kind of things spread out, and the density of the cells is equal. And you devote much more area to things like the face or the fingers than to the arm, for instance, or the back.

And so a lot of what we know about this originally derived from experiments by neurosurgeons who were stimulating with electrodes in these different regions. This would be done with epilepsy patients who they'd potentially be performing surgery on. So they'd have an electrode in somatosensory cortex, and they'd inject a little current.

And people would say they felt tickled. And so they'd move the electrode around, and like the location of the tickle would move around. And this led to this idea of a map. As I said, this has recently kind of been challenged. And we'll see where this kind of ends up.

Supposedly, you got this map of the body that exhibits something kind of like cortical magnification. One of the other things that you see in somatosensory cortex is orientation tuning. So this is very much analogous to the orientation tuning that you see in visual cortex.

So this is a receptive field that would be measured. The way that this is measured is with a device that looks like this. So you've got all these little pins. And so you can press the pins up or down. That will provide localized mechanical stimulation.

So the animal's got its fingertip on this thing. You can generate different kinds of stimuli and derive the receptive field in different ways. And this is an oriented receptive field that you see. So this is showing how this would work.

So you'd present like a little embossed line segment. And you can vary the orientation. And you can see here that there are orientations that give you big responses. And then as you rotate it off, the response gets less. The orientation tuning is present if you just press a thing into the skin versus scanning it over the skin, and it's independent of the pressure by which the line segment is pressed into the skin.

Two last things, and then we will call it quits. So one kind of other interesting analog with other systems-- we just talked about visual search where you have an array of elements and you have to detect something that's different. You can also perform haptic search where there will be elements that are presented to the different fingertips. So you have to represent the-- recognize the presence of material properties that are presented haptically to the fingers with this kind of device. The question is whether some properties kind of pop out.

So you could present a rough texture or a smooth texture. So remember, we analyze visual search results by plotting reaction time versus the number of elements. Here it's the number of fingers as opposed to the number of line segments or whatever.

And when you're searching for rough stimuli among smooth stimuli, you get pretty flat search results, indicating pop out. So rough among smooth pops out, hard among soft, cool among warm, edge surfaces among smooth surfaces. But unlike the visual system, you don't get pop-out for distinct orientations. So if you do horizontal lines among vertical lines, you don't get pop-out. So same kind of phenomena but the features that seem to be kind of fundamental may be very different across sensory systems.

Finally, a really cool kind of application of touch is Braille. So this is how someone without sight would read. So Braille consists of these patterns of embossed dots. And you might wonder, well, what's the point of having these dots? Why not just actually emboss the actual letters?

And what this is supposed to show is that at the kind of coarse spatial scales that the receptors in our fingertips are measuring things, those patterns of embossed dots are more discriminable than the actual letters would be. All right. So this is supposed to be what the pattern of stimulation would be for each of those different Braille letters and for each of the different actual letters. And the idea is that these things all look subtly different, and these all look the same. So Braille is an interesting application of the sense of touch.

All right. That's all I got for you about the sense of touch. What questions do you have? Again, a lot of the action is at the receptor level. We don't have much in the way of understanding of higher-order aspects of perception. Hopefully, that's to come in the future.

OK, have a great weekend. We got one more class to go. I will see you Tuesday.