

[SQUEAKING]

[RUSTLING]

[CLICKING]

**NANCY**

**KANWISHER:**

All right, it's 11:05. Let's get started. So the agenda for today, we're doing this whole thing on the methods in human cognitive neuroscience. And I'm illustrating those methods with the case of face perception. Not just because I'm into face perception, but it's a particularly rich domain of research where there's lots to say about it from all these different methods.

And so last time, we talked a bit about applying Marr's computational theory level to face perception. We talked a teeny bit about some behavioral data and a little bit about functional MRI. What I'm going to do today is quickly zoom through a speeded-up review of those things, and then we're going to get to some of these other methods. And there's a quiz at the end. All right?

OK, so methods in any field of science are just there to enable us to answer scientific questions. They're not to impress our friends with all the fancy things we know how to do or our colleagues. They're just to answer questions. And so you always have to start with the questions.

And so last time, I listed a bunch of questions. Not all of them, but a bunch of questions one would really want to know about face perception if we were to understand how it works in the brain. And last time, we focused on these first three. So let me just do a super quick review.

The questions at the level of Marr's computational theory, we ask, what is the problem that's being solved and why is that important to the organism? What is the input, what is the output? How much you get from that input to that output, right?

So for the case of face perception, here's a very simple version of it. Here's an example of the input. It goes in, hits the retina. The stuff that we want to understand happens in here, and you have an output. OK, so just even thinking about it that way, we can already just see, with common sense, that one of the big challenges in solving this problem is that faces look different every time you see them.

The lighting changes, the orientation of the face changes, the hair changes, the mood changes, all this stuff happens. People put on makeup, they shave off their facial hair, they do all these things to make it a big challenge to recognize faces. And yet, we manage really well.

So how do we do that? Well, our field has many methods to address this question. Last time, I talked about one little example of a behavioral study-- simple, cognitive psychology study measuring behavior-- where we showed that the way people solve this problem is fundamentally different with people they know well and people they don't know well. So I showed an example that all of you presumably would have no trouble determining that those are all pictures of the same person, even though at the pixel level they're wildly different.

And yet, you have a hell of a time saying which of those images are of the same people and which aren't. And so the point is that our ability to extract this invariant representation, that is to figure out abstractly who is that, is really-- well, to figure out that any of these images are the same as each other is much better for familiar than unfamiliar faces. And that means we don't have a perfectly general ability to take any face and abstract out this completely image-independent version of it. That's what invariant representation is. Yeah?

**AUDIENCE:** For the case of the Dutch politicians, did they ever do the study on people who were super recognizers?

**NANCY** I don't know about that, but they did do it on people who are professional TSA-type people.

**KANWISHER:**

**AUDIENCE:** OK.

**NANCY** Right? And I'll tell you guys about that later. But you could think about whether you think it might work better

**KANWISHER:** with those people or not. OK, everybody get this general point here? All right.

So I skipped over another simple behavioral finding last time that I want to mention now. And that is an extremely low tech-- charmingly low tech, and yet, I think very powerful-- discovery about face perception. One of the most important original bits of evidence that face perception might be a different thing in the brain came from a PhD thesis in this department by a guy named Robert Yin. And he used the extremely high tech equipment of a stopwatch and paper.

OK, so what did he do? He presented faces to people upright. And he said, study these 20 faces. And then he tested them later. Did you see this face? Did you see this face? Did you see this face?

And then he did the exact same experiment on a different set of faces, but they were all upside down. Studied upside down and tested upside down. And what did he find?

He found what's known as the face inversion effect. Namely, people do much worse at this task when the faces are upside down. Here's errors for inverted upside down, errors for upright at this task. Even though, importantly, they were studied and tested upside down or studied and tested inverted-- upright.

OK, everybody got what this shows? OK, so that's cool that this face inversion-- but the further cool thing is he showed that this face inversion effect is greater for faces than for other kinds of stimuli. So he tested lots of other things, including houses and stick figures. And he showed that that cost, when you turn the stimuli upside down, is greater, that difference is greater for faces than for other classes of stimuli.

So what that suggests is that face recognition may just work differently in some deep way from recognition of other classes of stimuli. And Robert Yin actually inferred in his PhD thesis-- way, way back before any imaging method-- that maybe there are special parts of the brain for face recognition. And maybe face recognition is just a totally different thing, that's why it is more affected by inversion than recognition of other kinds of things. Was there a question back there? Yeah.

**AUDIENCE:** I was going to ask, could that just be because faces are much more complex than houses or stick figures and that--

**NANCY** Good question. Hang on to--

**KANWISHER:**

**AUDIENCE:** --backwards.

**NANCY** Good question. That's a very good question. And many people have tried to grapple with that. And actually, about 10 years ago, the idea that this disproportionate effect for faces was standard textbook, completely accepted. And now there's another round of people doubting it with other kinds of stimuli. So it's kind of ongoing.

**KANWISHER:**

It's a very robust difference, but to say exactly what it is about face stimuli versus other kinds of things that is responsible for that difference, you can imagine it's subtle. For the purposes of this course, I'm trying to not quite lie to you guys, but give you the most standard view without freighting you with every possible objection to every little thing. Because pretty much every finding, there's somebody who has a beef with it. We'll tell you, that's not really true because blah-di-blah. OK?

So yes, there's a little bit of debate going on about this right now. But for the purposes of this course, it's pretty damn rock solid, at least as an empirical result. All right. So there's in fact lots of versions of the face inversion effect. One you may have seen before but which is very amusing. If you look at faces like this that are upside down, they look sort of normal.

But then if you rotate them, you realize there's something deeply weird going on. So the point is, you're much more sensitive to those grotesquely distorted faces when you see them right side up than when you see them upside down. So that's another version of the face inversion effect, and there are many, many incarnations of this effect. You'll see another one later in the lecture.

So where did we get last time with these questions? We got that one of the major, if not the major central challenge in face recognition at a computational level, is the fact that we deal with huge image variation each time we see a face. And yet, somehow we're able to grapple with it.

So to understand how face recognition works will be to understand, what is the code, ultimately-- nobody knows right now-- but what is a code running in our heads that enables us to do that? What is our mental representation of a face that enables us to deal with this problem? By looking at behavioral data, we got some evidence from the Dutch politician study that whatever that representation is that we extract from faces, it's not independent of the particular image.

It's not that we have some platonic ideal of the face that we can extract from any face that lands on our retina, platonic ideal of that person's face, right? So whatever we're doing, it's not completely invariant, because we can't do that so well with unfamiliar faces. Also, as I just showed you-- related, but not exactly the same point-- our mental representations of faces are very sensitive to the orientation of the face more than our mental representations of other classes of stimuli.

So those are just very simple insights about whatever our representations of faces are in our heads, just from simple behavioral data. OK, so let me just review some of the strengths and weaknesses of simple behavioral methods. Strengths are, they're good for characterizing the internal representation, right?

Not with huge computational precision, they're more like with gisty kind of ideas. They're not very invariant, they depend on the orientation, right? That's not very precise, but it's a whole lot better than nothing. That's what I mean by at least qualitatively.

They're good for disassociating mental phenomena. So you've already seen that when the inversion effect, it happens more for faces than other things. So that already starts to tell us, OK, maybe whatever the code in our head is that we use for face recognition, maybe it's pretty different than the code that we use in our head to recognize objects.

OK, it's also cheap. It's really cheap. Much cheaper than all the other methods. OK, weaknesses-- behavioral methods alone don't have any relationship to the brain, at least without doing extra work.

And it's not that they're useless until you link them to the, brain it's just that the brain is a whole source of other data. And it's nice to link them, because then you can connect with all those other data. Also, behavioral data are pretty sparse. For the most part, you have accuracy and reaction time, and that's it. And that's just not a whole lot of data to work with.

You have to actually be much smarter to be a behavioral cognitive psychologist, than you have to be a cognitive neuroscientist, where you have much richer data to reason from. Cognitive psychologists really have very, very clever designs because they're taking this extremely limited data and trying to pull out interesting insights about mental function. Another way of looking at that is, here's an eyeball and a bunch of processing going over stages and a response, right? With behavioral data, all you have is that response.

But presumably, for most of the mental processes that go on in our heads, there are many different stages of processing where different things are going on. Computations tend to have multiple stages and unfold over time. And all we have is the output. So really, what we want to be able to do is characterize the whole sequence of processes. And it's not that you can't get insights about some of those intermediates from behavioral data, it's just much more challenging. So if we had a way to look at those things independently, wouldn't that be awesome?

OK, so there's lots of ways to do that. And a particularly good one is functional MRI. So as I mentioned before-- I mentioned this very briefly-- this very early experiment that I did way back asking whether there is a region of the brain that's selectively involved in processing faces.

And I'm going to put a slightly different spin on it from what I put before. It's the same experiment, same data, but I want to emphasize more the logic of the experimental design because you guys will be designing an experiment on a different topic due Monday night. That we're going to discuss Sunday night-- that we're going to discuss in class on Monday. So we start with a hypothesis that there's a region of the brain that's selectively responsive to faces. That's the hypothesis.

The way we test it is to pop people in a scanner and show them faces and objects. The data that I showed you before is that this little patch of the brain-- remember, this is a horizontal slice, back of the head, left and right are flipped. So that little region in me is right about in there. Everybody oriented? OK.

That region responds much more to faces than objects. Is that clear to everybody what that is? OK. So yes, you see that in most subjects. So yes, there's a bit that responds more to faces than objects. But now, let's consider the hypothesis that that region is really selective to faces per se.

And the way you evaluate whether these data, how strongly these data support that hypothesis-- they're certainly consistent with it, but do they nail that hypothesis fully-- is to consider, are there any other alternative accounts we can think of that are consistent with these data and different from that hypothesis? Is that clear? It's really important. That's just the whole kernel of scientific thinking and evaluating evidence is asking yourself that question.

Is there any other way we could get those data where that hypothesis wasn't true? And if so, you've got to grapple with it. So what you do next is you think up alternative hypotheses to the one you started with, that is different accounts of the same data. And so in our case, you guys suggested a whole bunch, I suggested a bunch. And then the next thing I showed you is that we can test those alternative hypotheses, at least these ones here, by first-- what we did was, I didn't really emphasize this before-- but we reran that experiment in a new bunch of subjects, each subject individually.

We found in each subject the little bit that does this. We write down exactly where that is in that person's brain. Now that we found that region-- that's called a localizer run, because we're finding that region in each subject individually-- now we can ask it new questions. And so the new questions we asked it last time was to present faces and hands. And we found, oh, that region right there responds like this.

So the key ideas here is that we can identify that region in each subject individually with a functional scan. The reason that's important-- which I'll carry on about in more detail later-- is that the exact location of that region varies from one subject to the next. So if we just grab the whole fusiform gyrus or the whole lateral side of the fusiform gyrus in each subject, we'll get lots of stuff that is that region and lots of cortical neighbors that's something else.

And if we took the exact location of that region in my brain and registered it to any of your brains and said, OK, let's take the part of your brain that registers spatially as well as we can with mine, we're not going to exactly get the right bit. So to study that thing, we've got to first find it functionally. And then we can ask it new questions. Does that make sense?

OK, if anybody's unclear about that, I have actually online talks that go through the whole logic of this in painful detail. And I'm happy to answer other questions about it later. OK, so I put the word conditions in red because somebody asked one of the TAs what a condition was. And that's not stupid, I should have made that clear.

This is just experimental design gobbledygook that means any-- OK, what is the definition of condition? In an experimental design, you have things that you are manipulating and measuring. So in this case, we're manipulating the stimulus. And we're measuring the magnitude of response in the fusiform face area or in the brain.

So what we're manipulating, in this case, is the stimulus condition. So that would be one condition, that's another condition, that's another condition. Does that make sense? OK, so for your experimental design assignment for Monday night, you will be designing one or more experiments. And you will be describing exactly what conditions you are going to test. Everybody clear on that? OK.

All right, so these data enable us to rule out those hypotheses. And now what you could ask, OK, once you get more data like this, have you completely nailed that hypothesis? Is there just no way that hypothesis could be wrong now given these data and those data? And I'll let you percolate on that. There are ways it could be wrong, but you have to work harder to come up with them.

OK, so skipping ahead, just to give you the gist. This field has been going on for a long time. And there are now many, many studies-- 100 maybe even, I don't know, God, maybe even thousands, I don't know-- studies of this region in which-- and so this is sort of a summary statement from a long time ago. In my lab, we've tested the response of this region to lots of different kinds of stimuli. With that same method, localize it in each subject, measure its response when people look at that kind of stimulus,

And so what we know now was that this region is found in roughly the same location in pretty much every normal subject. It responds more to faces than to any other kind of stimuli anyone has ever tested. Let me just give you one example here.

If you haven't seen this stimulus before, raise your hand if you can tell what it is. Raise your hand if you can tell what that is. OK, some of you didn't quite get it yet. If you don't see it, don't worry. There's nothing wrong with you. It's a little subtle.

It's a face in profile, eyes, nose, mouth. Everyone got it? OK, so here's the thing. That's the same stimulus, it's just upside down. Another version of the face inversion effect.

In this case, you can't even make yourself see the face when it's upside down. If you think you see the upside down version of the face, you probably have the wrong bits. The thing you think is a nose probably isn't, et cetera.

OK, so this is an extreme version of the face inversion effect. And it's a gift to an experimental psychologist. Why is that such a gift? Because it's the same damn stimulus. But in one case you see a face, in another case you don't. All we did was tip it upside down.

And the response of the fusiform face area is much stronger to the upright version when you see the face than to the inverted version when you don't. So that enables us to stifle a whole line of attack from all of these hard core vision people who early on said, Kanwisher, your face area isn't really selective for faces. It's selective for these spatial frequencies or those, that kind of contrast, or this kind of shading information.

It's like, no, same stimulus. It's just upside down. It makes all the difference. It's really whether you see a face or not. Yeah?

**AUDIENCE:** When you were measuring the response of that example, did you have it so that at first, when people that the first time didn't recognize it and then you told them?

**NANCY** We did that later. Not in this experiment, but we did that later.

**KANWISHER:**

**AUDIENCE:** Are they looked see like what changed [INAUDIBLE]?

**NANCY**

OK, so it's a great question. And there's a lot you could do with that. And actually, I think other people have published studies like that since. I can't quite remember who all has done it.

**KANWISHER:**

But what we did was most of our subjects, especially in the context of a whole experiment, we chose stimuli so that most people could see the face in most of the upright stimuli and most people could not see the face in most of the inverted stimuli. It wasn't perfect at all. They didn't see faces in all of the upright ones and they didn't fail to see them in all of the inverted ones. And that's probably why this difference in response is not 2 to 1, but it's close.

But you could do lots of other experiments like that, and you should think about what kinds of designs would be good ones to do and what it would enable you to test exactly. All right. So OK, I'm, as usual taking too long to do things so I'm just going to throw out some questions for you to percolate on and we will come back to them later in the course. Do these data-- the fact that you can see this so robustly in all subjects and that all this evidence suggests it's really very selective for faces-- does that tell us that this region is innate?

It's in the same place, more or less, in pretty much everyone. Does that mean it's innate? Think about it, OK? It's not immediately obvious. Another question, does the fact that this thing responds so selectively to faces in pretty much everyone mean that it's necessary for face recognition? What do you guys think about that?

In the sense of, does that necessarily mean that if you lost that thing, you wouldn't be able to recognize faces? Isabelle. Is that Isabelle?

**AUDIENCE:**

Yes. Well, I would think to really test that hypothesis, you'd have to find someone that [INAUDIBLE] in that specific area.

**NANCY**

Exactly. Exactly. Exactly, and we'll talk more about that in a moment. The critical thing is that it's fabulous and powerful and cool to be able to find this thing in everybody, measure its response. It's taken us very far. But just the fact that people have that thing doesn't tell us that you need it for face recognition. It just tells you it turns on when you recognize faces.

**KANWISHER:**

This is really important. We'll keep coming around to this. Does this tell us how face recognition actually works in the human brain? No. I mean, it's important, but it's barely step zero. Unfortunately, the field is kind of still at step zero for most things. Step zero's better than I guess, I don't know, maybe I should call it step one. Anyway, it's something, but doesn't tell us how it works.

OK. All right, so advantages and disadvantages of functional MRI. Advantages, it is, as I mentioned last time, the best spatial resolution available for studies on normal subjects without opening their heads. That's what it means to say noninvasive. Disadvantages, as I just said, we don't know-- just because we see a response there doesn't mean that that region is causally involved in perception or cognition or experience.

We don't know exactly what is going on at a neural level underlying that bold response, that blood flow change. It could be any metabolic change, not necessarily neuronal spiking. So it's a little bit-- it's very indirect and a little imprecise.

Spatial resolution is much better than anything else in humans, but it's appallingly bad compared to anything that people who work on animals can do or they routinely record from individual neurons or even dendrites on a neuron. We are summing over hundreds of thousands of neurons in each pixel or voxel that we measure with functional MRI. It's very expensive. It's a little cheaper than that here, but in most places it's more than \$600 an hour. That is a lot.

There are other-- there are parts of the brain where it's really hard to get any signal for various physics-y reasons. And it makes a loud noise, which is not always a problem, but it's a problem for some things like scanning infants or like doing auditory experiments. The temporal resolution is not even close to the time scale on which vision happens.

So vision is really fast and functional MRI is really slow. Right? It's slow, why is it slow? Yeah.

**AUDIENCE:** Blood levels take time to change.

**NANCY** Yeah. Just takes a long time for blood flow to change after the increase in neural activity. All right. OK, so back to  
**KANWISHER:** our questions that we're asking about face perception. Where do we get with functional MRI? Well, actually from both behavior and functional MRI, it kind of looks like we have a distinct system for recognizing faces than for recognizing everything else. I don't think we've totally nailed it. Yes.

**AUDIENCE:** So quick question regarding the fMRI. So the resolution is field of a couple of seconds? [INAUDIBLE]?

**NANCY** Yeah, some people would say you could get it down to a couple hundred milliseconds but that's debated. You  
**KANWISHER:** have to go to great lengths to do that. Normal functional MRI, a couple of seconds at best. Yeah.

All right. So let's consider this next question. How fast does face recognition happen? Now, that may seem like a completely arbitrary question to ask, but it's not. Remember, we're trying to understand the computations that are running in your head when you recognize faces.

And you might imagine some computations that are iterative-- that involve multiple repeated testing of hypotheses, generative models, whatever-- things that involve lots of iterated feedback versus things where you just have a feed forward sweep up the visual system. And so there might be very different time scales for those different kinds of mental processes. So we just went through this.

Functional MRI is not going to answer this question. It's just not. It's a bummer, but that's life. We're adults, we're going to just move on and use a different method.

OK, so there's a bunch of different methods. One is kind of been around forever. You glue electrodes on the head, right? Sometimes you push the hair apart or try to find bald people and glue electrodes right on there. And you can use, in the old days, about 10 electrodes, or you can use in more modern devices these nets with a few hundred electrodes that you settle onto the head.

And so then you just measure directly electrical potentials right on the scalp. So what's cool about that is it's totally non-invasive. And it gives you a beautiful online temporal measure of underlying neural activity. What's not so cool about it is that electrical potentials blur all over the scalp and the spatial resolution is really awful.

So the analogy has been made that it would be like sticking a microphone on the inside of the top of a football stadium and collecting audio there. You would know when a touchdown was scored. There's a lot of noise all over. It's like, OK, there's an event, we detected that event, right?

You might be able to tell a touchdown from something else. I don't know about football so I can't tell you what else. Anyway, something else, some other event that could happen.

OK, so that will be useful for some things, but kind of crude. But you'd have a hell of a time telling anything else, like what one person is saying to another person in the bleachers. So that's the old analogy. This is changing slightly, and we'll get to that later.

But first, I want to briefly mention one of the assigned readings that I just hoped you guys could figure out on your own. But just in case you were confused about it, the point I wanted you to get from the Thorpe reading is he's asking how quickly can we tell if an image contains an animal or not? It's a kind of way to say, how fast is object recognition?

So what does he do? He has people look at a bunch of images and they press this button if it has an animal in this button if it doesn't. Really simple task. So first question is, why not just use those reaction times? We can measure how fast it takes for people to press a button after the image comes on. Why not just use that?

Does that tell us how fast object recognition occurs? Yeah, Jimmy.

**AUDIENCE:** It doesn't because if you perceive that and then it also activates the motor neurons and it takes time to respond.

**NANCY** Yeah, you have to take all that time to figure out, OK, I see the animal. OK, which button is that? And then which  
**KANWISHER:** finger do I push? And then you've got to send a signal all the way down here, conduction velocity all the way down to your finger, that takes a long time. And so it includes all that motor stuff in with the perceptual stuff.

We could make some guesses about how long that motor stuff takes, but it's still not very precise. So the point of the Thorpe paper is they're basically trying to collect a reaction time out of the neurons in the head, right? So they're trying to actually collect-- it's essentially what they're collecting in this case is more of the motor response because they're collecting responses over frontal lobes, right? And we haven't talked about this much.

But all of the visual stuff we've been talking about all happens in the back of the head. More motor planning stuff mostly happens in the front of the head. And so they're collecting responses out of here, averaging over a bunch of frontal responses. And they see the average response when there's an animal-- this is just potential average over those frontal electrodes-- is like this. And when there's no animal it's like that.

And so what does that tell us about how fast people can distinguish whether an image has an animal or not? Yes? Yeah.

**AUDIENCE:** It's less than that number.

**NANCY** Less than?

**KANWISHER:**

**AUDIENCE:** 150, 160.

**NANCY** OK, why less than 150?

**KANWISHER:**

**AUDIENCE:** I've read the paper so it's kind of cheating, so.

**NANCY** That's OK. That's good. That's fine. Go ahead.

**KANWISHER:**

**AUDIENCE:** It gives you around-- the 150 second is giving you a [INAUDIBLE] saying some process has been registered and now you're trying to do something else in the case of non-animals.

**NANCY** Right.

**KANWISHER:**

**AUDIENCE:** So the deviation starts getting you that OK, two different actions have started taking place.

**NANCY** Yep.

**KANWISHER:**

**AUDIENCE:** So by that time, the image ought to have been sort of fully processed. So that should be something less than that number.

**NANCY** Yeah. Yeah, did everybody get that? It's actually quite subtle. So the key thing is, these curves diverge right there

**KANWISHER:** at 150. So that tells you that by 150 milliseconds, something in your brain is happening different if there's an animal and not an animal. That's the key question.

But what is that something? It may be your motor preparation of the response. In that case, the actual visual part happened before, because you wouldn't know which button to press if you hadn't already recognized it. So it's an upper bound for when that process happened, because maybe it happened before and we're looking at a later stage, OK? Does that make sense?

But also, it's an upper bound for the beginning of that process. Because the fact that those electrode responses have diverged doesn't mean you've finished processing whether it's an animal or not. So it's kind of a subtle business reasoning from this. OK, so that's all that.

So that's a case with detecting animals. What about faces, to get back to our theme for today? Yes, you can learn about the speed of face detection at least with the ERPs. And so here's the first paper that did that back in 1996.

They had electrodes where are these? Just right around here and here. I actually have those electrode locations tattooed on my scalp, color-coded anyway. Yes?

**AUDIENCE:** Is ERP just the same as an EEG, just in a specific plan e?

**NANCY** Yes, exactly. It's the same as an EEG except what you do is you time lock the data collection to stimulus onset.

**KANWISHER:** So it actually stands for Event-Related Potential. And the reason it's event-related is you collect all those trials and you time lock to stimulus onset, and then you signal average. I had a slide on that but I took it out. It was too detailed. But that's exactly the idea, yeah.

So here, stimulus onset is right around here. This is time going this way. And what you see-- it's hard to see here, but the faces are right there. And at 170 milliseconds after stimulus onset, there's a bigger bump for faces at an electrode approximately here. And even more so-- actually, even more so over the right hemisphere right there. Compared to cars and scrambled faces and stuff like that. Yeah?

**AUDIENCE:** What is ERP exactly measuring? Is it just activity?

**NANCY**  
**KANWISHER:** Yeah. So again, it's electrodes glued on your scalp or just stuck there with some kind of icky gel. And so they're just measuring potentials. And so the idea is that's neural activity somewhere underneath those electrodes, but maybe anywhere within inches. Like a long-- probably average is over much of the whole lobe underneath.

So it's very spatially blurry, but it's giving you summed idea of activity under that electrode. Make sense? Electrical activity, because it's the direct electrical consequence of neural activity, it's very precisely time locked, unlike functional MRI, which is going by way of blood flow.

OK, so that tells us that we have a face-specific response at 170 milliseconds. And that's sort of more evidence that there might be something special in the brain for face recognition. That's useful. It tells us that faces are discriminated from non-faces, or they've begun to be discriminated from non-faces by 170 milliseconds after the stimulus comes on. Make sense?

OK, now do we know whether the signals coming from the fusiform face area? No, we have no idea. It's probably somewhere in the back of the head, because you get it better with electrodes back here than electrodes up here. But that's about it. That's all you can tell.

So can we do a little bit better localizing the source of that signal? Well, maybe a hair better using a very similar method called magnetoencephalography. So this is a picture that Chris Brewer took of Leyla Isik postdoc in my lab, and me and the MEG system. This is in on the other side of the building.

So MEG is a lot like EEG and ERPs except that it detects magnetic fields, not electric fields. And it does this by having these several hundred devices that are placed right next to your head in this big hairdryer thing. There's 300 devices in there that measure teeny tiny magnetic field changes that happen with neural activity. And the crux of the idea is this is a cross-section through the brain.

So remember in Graybiel's dissection, this is cortex here and this is underlying. What is this stuff underneath it? Sorry?

**AUDIENCE:** White matter.

**NANCY**  
**KANWISHER:** White matter, yeah. Well, those are all the fibers. OK, so the activity that underlies perception and cognition mostly happens in the gray matter where the cell bodies are. And so a lot of that activity goes in a direction perpendicular to the cortical orientation with these cells that cross the cortical surface like that.

So if you remember 8.02-- if you have activity that's going through the cortex like this, right hand rule, the magnetic field here is going to be a consequence of that electrical activity in this direction. It's going to mostly stay within the cortex. Everybody see how that's true? That's not so great, because our detectors are out there, outside the cortex.

However, consider the activity that's in the sulcus in here, in this fold of the brain. Electrical activity in this direction, right hand rule, will stick outside the brain. And we can detect it with our magnetic sensors. Does that make sense?

So you can sort of see most cortical activity better if it's in a sulcus, or at least in part of the cortical surface that's perpendicular to the scalp where the detectors are just because of the orientation in the right hand rule. OK, so it primarily sees activity in the folds or sulci, not in the outer bumps gyri. Field strengths are minuscule as a consequence of neural activity. So the fields we measure are 10 to the minus 13th Tesla, a million times weaker than the Earth's magnetic field.

So you can imagine that if you set up an MEG system you need a lot of shielding. We had a whole rigmarole when the MEG system was set up in this building because it's right near the subway and the train. And so there are many, many layers of copper shielding to protect it. So we can detect these teeny tiny magnetic fields from the brain's activity separated from the noise of the outside world, which is much greater in magnitude.

OK, so-- all right. So actually, MEG was invented here at MIT by this guy, David Cohen. And this is the first MEG device ever built, very cool, way back in 1968. And what can it tell us about face perception? Well, a lot. I'll give you just one rudimentary example.

That M170 that you can detect with scalp electrodes, you can also detect with magnetic sensors on the head. So here's some of our data from a long time ago. This is the strength of the magnetic field at sites right about out here. And you can see a face-selective response also at 170 milliseconds, just like you can with scalp electrodes.

So that tells us that at least you've started to detect faces by 170 milliseconds. That's pretty fast. And again, it's more evidence that there's specialized machinery. These data don't yet go beyond the EEG data, the ERP data from electrical potentials. But they might, in principle, and there's lots of ongoing work trying to do that.

OK, overview, advantages of these methods, both EEG and MEG. They're non-invasive-- that means you don't need to open the head. A very good thing, especially if you're the subject. They have very good temporal resolution.

And if we want to see computations unfolding over time in the brain, this is a good way. I just said why we'd care about that. OK, so far-- well, never mind, I'm going to skip this point. Not that important. We will get back and do more sophisticated things with EEG and MEG in subsequent lectures.

Disadvantages-- spatial resolution is terrible. And this is another kind of ill-posed problem. So just as the brain is facing lots of ill-posed problems in perception and cognition, we scientists are facing ill-posed problems when we collect electrical or magnetic activity at the scalp and try to infer the exact location in the brain where it's coming from. It's a similar problem to the problem of invariant object recognition.

There are many possible configurations of sources in the brain that could give rise to the same set of electrical and magnetic fields out of the scalp. And that means it's ill-posed. We don't have a way to get a unique solution.

So all that to say we can't figure out the exact sources. We can make some guesses, but it's not very good. So what do we do? Just give up? No, we use another method.

So here's an amazing method. This is the one method in humans that gives us high resolution in both space and time. And that's when we have the very rare opportunity to record directly from inside the human brain. This happens only in the context of neurosurgery.

So neurosurgical patients-- like this guy here, who you'll meet in a little bit-- this guy had intractable epilepsy. And most people with epilepsy are treated well by drugs that suppress seizures. But some people are just not responsive to drugs.

And if the seizures are bad enough, they can be totally life disrupting. If they happen dozens of times a day, you just can't live a normal life. And under those rather extreme circumstances, sometimes the best option is neurosurgery. That is, trying to find the source of those seizures and trying to remove it surgically.

OK, so you hope you never have to go through this or anyone you care about has to go through it. It's no picnic. But actually, this surgical treatment is often very effective.

So when neurosurgeons decide to do this, they have to remove a whole piece of skull bone to get access to the brain. They have to go through what structure that Ann Graybiel showed you in her dissection the other day. What do you have to-- after you take off the a skull patch? Yes.

**AUDIENCE:** Dura mater.

**NANCY**  
**KANWISHER:** Dura mater, exactly. That nice big piece of white, leathery stuff that was sitting over the surface of the brain. So you to take off a piece of skull, then you need to cut through and push apart the dura. And then what they sometimes do is stick electrodes straight on the surface of the brain.

And they do that for two reasons. One, if they have enough of them sampled far enough apart, they can kind of triangulate and figure out where is the source of the seizure. So the patient hangs out in the hospital for a week or so with these electrodes in their head waiting to have seizures. And then when they have a seizure, the clinicians can figure out where the source is so they know what bit to cut out.

The other reason to do this is to map functions. Because once the surgeons decide they have to go in and cut, they want to try to not cut out any of the most important parts. I don't know what it means to have unimportant parts of the brain, but they try to avoid language regions and stuff like that because then patients really notice if they lose those things or motor control regions. OK, so they map out functions where they might be planning their route. OK, make sense?

Now, some of these patients are very kind and generous to the world and say, yes, you scientists can measure responses in my brain while I look at your damn stimuli. And so whenever we can, we ask them please, please, please, can we show you some pictures or play you some tones or have you read some sentences while we record from your brain. And some of those patients very kindly let us do that.

And that gives us the most amazing data you can get from human brains. So for example, I had a rare opportunity to do this a few years ago from this lovely guy who was undergoing neurosurgery in Japan. And while he had electrodes in his brain, a colleague of mine was there and emailed me and said, look where these electrodes are-- right near regions I care about-- do you want to show us some stimuli and we'll record responses from those electrodes? And I said, damn straight I want to send you some stimuli.

So my students and I stayed up for a couple of days and made some stimuli and shot them to Japan and got some responses from those very electrodes. And here they are. So this is a strip of two parallel strips of electrodes right along the fusiform gyrus, right where the fusiform face area should be in most people. And here are the responses of each of those electrodes. 174 is here, what's 173 and so forth.

And what you see is this batch of electrodes right here-- this is a response when the patient was looking at faces. And these are the responses when they looked at a whole bunch of different kinds of stimuli. Objects, and this guy is Japanese so we showed him Kana and Kanji and digit strings and other kinds of stuff. Very low response to those other things.

This is an extremely selective response. It's much more selective than you see with functional MRI because we were recording directly from the surface of the brain. Further, we have time information.

This axis here is time, and you can see that that response-- well, you can't see the axis, but that response starts up at around 130 milliseconds and peaks up there at around 170. Everybody clear what we're seeing here and why this is so vastly better than either functional MRI or MEG or ERPs or anything else? Make sense? OK, so these are very, very precious data.

OK, nonetheless, the electrodes in this case are about 2 millimeters across, each electrode. And that is about the size of a functional MRI pixel or voxel, a little bit smaller. It has less blurring because functional MRI blurs spatially because it's looking at blood flow.

So this is a more precise spatial measurement than functional MRI, but it is still averaging over probably tens of thousands of neurons, down from hundreds of thousands of neurons with functional MRI. So can we ever get responses from individual neurons in the human brain? Yes, occasionally.

In fact, a paper came out on the bioRxiv a couple of months ago. I was on this guy's PhD thesis defense. And this is a guy who works with a neurosurgeon on Long Island. And this neurosurgeon specializes in epilepsy neurosurgery.

And he's very interested in not damaging people's ability to recognize faces. And so he sticks electrodes to map out neural activity and to discover seizure foci. Before the neurosurgery, he sticks electrodes in parts of the brain near the fusiform face area.

So this is a slice like this through the brain. I showed you before horizontal slices, OK, so left and right are flipped, that region is right in there, everybody oriented with this picture here?

So this is an MRI image of this person. It was scanned with functional MRI before the electrodes were put in. And that shows you their fusiform face area right there. So now, the neurosurgeons put in electrodes for clinical reasons, but the electrodes this surgeon uses have these little tiny micro wires that come out of the tip of the electrode that enable him to record from individual neurons.

And so these guys, for the first time, have recorded from individual neurons in the fusiform face area in humans. And here's an example of one of these neurons. So here are the different stimuli here. A bunch of different face stimuli, body stimuli, houses, patterns, and tools. And this shows you time across here. Each one of those dots is-- this is all the response of a single neuron that's been identified in a human brain.

Each dot is an action potential, is a spike out of that neuron. So you can see them happening over time here to all the faces. And this is an average amount of activity to all of the faces and average amount of activity to all the other stimuli. Make sense?

So that's pretty breathtaking to me because I've been using these very indirect methods for a long time, inferring that they must result from the average across a lot of neurons doing that, but it's pretty awesome to actually see individual neurons doing that. Yeah? OK. Here's the time course of responses just averaging over this raster over time, showing you a similar time course to what I've shown before. And in this guy's thesis, he found three other face selective neurons in the FFA, but the electrodes are so rarely in the right location that they only have a few in this whole thesis, and there they are. Yeah?

**AUDIENCE:** Even if we could measure individual neurons, we don't really know which neuron it is, right? If I wanted to go back and find the same neuron Again, That's pretty much impossible.

**NANCY**  
**KANWISHER:** Forget it. Yep. Yep. So people like me who almost never get to see responses from individual neurons in human brains have kind of neuron envy. It's like everyone else in this building has-- they can measure stuff from dendrites or ion channels or individual neurons. They can do all this amazing stuff.

But actually, there are a lot of limitations in those methods too. And you just put your finger on one of them. So they're like, OK they found those neurons, there are four neurons. We can't go back and find those neurons again. That's that, right? And they're probably subtly different in different brains, right? So it's cool and powerful but has still has many limitations.

OK, does this tell us that these neurons are involved in discriminating one face from another or just detecting faces? Can we tell from these data? Are they just saying, here's a face or are they saying, that's Joe?

**AUDIENCE:** Did they have different conditions for different people?

**NANCY**  
**KANWISHER:** These are different faces here. What do you think? What are these neurons doing? Yeah?

**AUDIENCE:** They're just recognizing faces [INAUDIBLE].

**NANCY**  
**KANWISHER:** You mean just detecting? No, just say more. What do you think they're doing?

**AUDIENCE:** They're just selecting for faces. There's no evidence to show that they distinguished different faces.

**NANCY**  
**KANWISHER:** Well, how about this? These are different faces here. These are different faces here.

**AUDIENCE:** But one could ask, if it does involve them sort of acknowledging what faces, did they have to put a name to the face?

**NANCY**  
**KANWISHER:** Nope, they're just sitting there looking at stuff. So bottom line is, we don't know from this. It could be just responding and saying essentially, there's a face. But the fact that there's different responses to different faces suggests that maybe there's some information in there.

If you ran some machine learning code on this, you could tell a little bit, which face was being presented. Because those neurons are responding differently to different faces. Yeah?

**AUDIENCE:** Is it really like if they just showed the same face repeatedly, wouldn't it just be like [INAUDIBLE]?

**NANCY**  
**KANWISHER:** OK, very good question. Very good question. That's why I said suggest, right? You're absolutely right. That could be just noise. It could be that if you presented the same face every time you'd get that same distribution. You're exactly right.

And so we will talk-- not next time, I think Wednesday next week. But anyways, very soon we'll talk about methods that enable us to exactly deal with that question and ask, is there actually information in this pattern of response across neurons or voxels or whatever it is? Or is that just the noise of variation? Yeah? OK.

**AUDIENCE:** But how many neurons are in [INAUDIBLE]?

**NANCY**  
**KANWISHER:** Oh good question. Let's see. I would say, I think a few million. So let's think about it. Each voxel is about a half a million, and they are typically maybe like 30 voxels, something like that. Somewhere on the order of 20 million, something like that. I mean, with huge error bars.

OK, so this is cool and tantalizing, but it doesn't even tell us what these neurons-- what exactly they're participating in. It doesn't tell us if those neurons are telling that person which face is there or maybe what facial expression the person has or how old they are or whether they're male or female or God knows what else, right? And it certainly doesn't tell us how those neurons get that information. Still, it's cool.

OK, so intracranial recording, both with the grids that I showed you and the single unit version. Advantages are, this is the only method in humans that has both pretty good spatial resolution and temporal resolution at the same time. Disadvantage-- well, you need to have a craniotomy, which is no picnic, to put it mildly. You need to have a huge piece of your skull removed and neurosurgery.

And that means that the only times we get to do this are when it's required clinically and everything is under control of the doctors, as it should be. So the doctors make all the choices about where the electrodes go, and we just get to sit in the background and say, please, please, please, look at these stimuli, but try not to hassle the patients too much. Right now there's a patient in Albany, New York who has electrodes right over a really exciting part of the brain to us that I'll talk about in a few months.

This patient has electrodes that respond specifically to music. We will talk about that later. It's pretty amazing. And for the last couple of days, Dana and I have been-- mostly Dana has been collecting stimuli because we really want to ask questions about the response of those electrodes. And this patient is not too thrilled listening to our stimuli.

So we finally said, oh, OK, tell the patient they can just do Instagram on their phone and we'll play the stimuli in the background. So hopefully, we'll have cool data from that soon. OK, so to say that these data are limited and hard to control is an understatement. We basically can't control it at all. All we can control occasionally is the stimuli.

And it also, like functional MRI, just because we see those beautiful responses, it doesn't tell us how those responses are connected to behavior. So that's a real challenge. So that won't do.

We need to get beyond this problem. I keep saying this method is great, but it doesn't tell us the causal role of that neural phenomenon in cognition and behavior. As scientists, science is all about discovering causal mechanisms. We're not just interested in what is correlated with what, we want to know what's causing what. That's really of the essence, and so we need to do better here.

So what are we going to do? Somebody mentioned a while ago, maybe it was Isabelle, that one of the ways to do that and ask whether the face area is causally involved in face perception is to look at a case where the face area is altered. So there's a bunch of ways to do that.

And one of them-- OK, that's just a review. We said faces are recognized fast but we haven't learned much more. How do we test causality? OK, patients with focal brain damage. Here is a patient. These are vertical slices through the back of this patient's head.

OK, let me get oriented. The slice is maybe this here. And as you go rightward, you're marching back in the brain like that. Everybody oriented? What's this thing right there?

**AUDIENCE:** Cerebellum.

**NANCY** Yeah, cerebellum, right. That thing right there is this patient's lesion that spans several slices going back like  
**KANWISHER:** that. And this patient's lesion looks a whole lot like my FFA. There's my FFA, greater response to faces than objects, on similar slices.

We don't have functional MRI from this patient so we don't know exactly where this guy's FFA was. But there's a good bet that it was blitzed by that lesion, because it's right in the zone where it usually lands. And this patient can't recognize faces at all.

And importantly, the patient is absolutely normal at recognizing objects. No problem whatsoever at recognizing objects. How does this take us beyond functional MRI? Yeah?

**AUDIENCE:** It implies causation.

**NANCY** Speak up.

**KANWISHER:**

**AUDIENCE:** It implies causation.

**NANCY** Yeah, say more. What does it tell us?

**KANWISHER:**

**AUDIENCE:** So because of the fact that area's damaged and then it makes it to not be able to recognize faces and I can see that there's causality that oh, that area's [INAUDIBLE].

**NANCY** Exactly. Exactly. It says you need that bit to recognize faces. But also says something else. What else does it?

**KANWISHER:**

**AUDIENCE:** That you don't need it for recognizing objects.

**NANCY** You don't need it for recognizing objects. So this is actually really strong evidence that that bit of brain is very  
**KANWISHER:** specialized for face recognition. Specialized and necessary for face recognition.

OK, so--

**AUDIENCE:** Can that person still detect faces?

**NANCY**  
**KANWISHER:** Oh, yes. Good question, absolutely. OK, so let me just distinguish-- this person here has prosopagnosia-- that means a selective deficit in face recognition-- like Jacob Hodes, who I described yesterday, who has no brain damage whatsoever but has just never been able to recognize faces at any point in his life. So this syndrome can arise just from some weird developmental thing where you're atypical and you're just really bad at it, or it can result from damage to that part of the brain.

So now we're talking about the case of damage, but in both cases, people with prosopagnosia have no problem knowing that a face is a face. They just don't know who it is. Yeah?

**AUDIENCE:** Has there ever been a case of problems of people who can't recognize faces who [INAUDIBLE]?

**NANCY**  
**KANWISHER:** Indeed. Indeed. Jacob Hodes, who I talked about last time who is just absolutely awful at face recognition, including family members, close friends, can't do it, like not at all. He has a very normal looking fusiform face area. So after I told you that I had that conversation with him a dozen years ago or something like that, I scanned him. And he had a beautiful fusiform face area, like textbook. It looked-- well, looked like mine, which is a damn fine one if I do say so myself.

And I looked at that and I went, oh shit. I better publish this before someone else does. And I didn't get my act together, and then a whole bunch of papers came out saying, oh, people with developmental prosopagnosia have normal looking face areas. Take that, Kanwisher. What do you say about that?

And it was a little shocking. But upon further reflection, it's not really devastating, right? I mean, it's bracing, it's informative. But it tells you that having a face area that is a region that responds more to faces and objects isn't sufficient for normal face recognition, right? You need other stuff.

What might that other stuff be? Well, the circuits in there need to work right. It's not enough to just respond more to faces and objects. To recognize faces, they need to be able to distinguish faces from objects. We don't know if that's working right. What else do you need?

**AUDIENCE:** Memory.

**NANCY**  
**KANWISHER:** Memory, absolutely. Yes, you need to remember faces. What else?

**AUDIENCE:** [INAUDIBLE]

**NANCY**  
**KANWISHER:** Could be, but in Jacob's case, it was close friends he couldn't recognize. So what's another possible account of how could he have a normal face area and yeah? David?

**AUDIENCE:** It might be a gap between recognizing a face and connecting that to recognizing a person.

**NANCY** Yeah. Yeah. Or to put that neuroanatomically, you got to get the information out of there. maybe, for all we know,  
**KANWISHER:** that little face area is working perfectly. Maybe that face area knows who that person is, in a sense. But if the connection's out of that brain region to the rest of the brain are messed up, it doesn't do you any good. You need to be able to read that information out and act on the basis of it.

Anyway, that's a big sidebar. Point is, you can have prosopagnosia either as just a developmental disorder or as a result of brain damage. Oh, God, I knew this was going to happen. All right, so OK, so very briefly, it messes up ability to discriminate and recognize faces, not your ability to detect a face, right? So as [INAUDIBLE] had asked, it's not just you can't tell the thing is a face, they're fine with that.

Importantly, they are normal and voice recognition. So it's not that they're confused about distinguishing one person from another. They can do it fine from audition, just not from vision.

In the rare cases where the lesion is small, it can be very specific, leaving object recognition intact. More often, there's kind of a blurry mess. You have a big lesion and a bunch of things are affected. OK, so we've talked about that.

OK now, it's very important in neuropsychology reasoning-- like we want to say, OK, that's really powerful, the case of prosopagnosia. You lose that bit, you can't recognize faces. And that establishes a kind of causality that we didn't have before with just functional MRI. But is that sufficient to say that that region is specialized for face recognition only?

It's not. Whenever I ask this, the answer's no. Your task is to say, why? How could you have great difficulty at recognizing faces and be OK at object recognition? And yet, not have a deficit that's specific to faces? How might that arise? You guys have suggested this hypothesis in different context before.

Yes? You look like you know. No?

**AUDIENCE:** But I just you could do other things. It doesn't have to only be for facial recognition. Because it response to animals [INAUDIBLE], right?

**NANCY** Sort of. But the question here is-- OK, let's just start bare bones. You have a lesion, you get around fine in the  
**KANWISHER:** world, you can do everything else but you have a real problem recognizing faces. Does that mean that the region lesioned is specialized for face recognition per se?

**AUDIENCE:** There might be other [INAUDIBLE].

**NANCY** That's true. There could be other things going on, absolutely. But let's suppose they're not. Let's suppose you  
**KANWISHER:** had good reason to think there weren't. Yeah?

**AUDIENCE:** It could be-- it'd be a path-- it could be one point in a pathway.

**NANCY** That's true. It could be totally a point in a pathway. Absolutely, that's another account. What else?  
**KANWISHER:**

**AUDIENCE:** Well, I couldn't get last comment, so.

**NANCY** He said maybe you damage a pathway. Yeah?

**KANWISHER:**

**AUDIENCE:** Maybe there's some other function we haven't tested in that person.

**NANCY** All these are very good alternative hypotheses. You guys are very good at this. The one I'm fishing for is, maybe  
**KANWISHER:** face recognition is just harder than object recognition. Maybe the part that's damaged is just generically involved in object recognition, but you damage part of the object recognition system and face recognition takes a bigger hit because it's harder. Right? Does that make sense?

Do you see how the case of prosopagnosia is consistent with that? So that means we cannot infer from these data alone that that region's specialized for face recognition. Now, we can do various things like test them on really hard versions of object recognition. And people have done that.

But there's another kind of data that are really powerful here. And that's when we have the opposite syndrome. So there's only a couple of cases of this. The best one is called CK, published in a paper in 1997. You don't need to remember that.

The point about this is that this guy has the opposite syndrome. He's severely impaired at object recognition. He can't tell a chair from a table from a car from a toaster, but he's 100% normal at face recognition. Totally normal at face recognition. In fact, better than average.

Do you see how that's in some ways even more powerful evidence that face recognition goes on in specialized brain machinery than the case of prosopagnosia? Face recognition isn't even a special thing that sits on top of normal object recognition. It's a totally different pathway. You can have no ability to recognize objects and your OK a face recognition.

Does everybody see how that's really powerful? And how those two kinds of evidence together are vastly more powerful than either one alone. Well, that's called a double association. We'll skip all of that for now.

Doubled associations are particularly powerful examples-- powerful forms of evidence in cognitive neuroscience where we have opposite syndromes that collectively make it really hard to wiggle out and come up with alternative accounts other than that there's a bit of brain that's really specialized for face recognition. It's not just that face recognition is harder, or else you'd never get this syndrome.

All right, I just wanted to finish that point. OK now, how much time do I have until the quiz? 15 minutes, OK good.

**AUDIENCE:** 13.

**NANCY** OK, good. We're going to skip over TMS. I'm sorry about that, guys. Someday I'll learn to time things in a lecture.

**KANWISHER:** Actually, I knew this was going to happen, I just-- we'll get back to TMS later.

And we will skip to the most amazing method in all of cognitive neuroscience for which-- we're going to come back to this dude who you met before who has the face selective responses in that part of his brain. Remember how I said that even though these data are gorgeous and spectacular and the only way we can get high spatial and temporal resolution together, but they don't tell us causality? Right? That's true here?

Resolution doesn't get you causality. To test the causal role of something, you need to mess with it. So it turns out that sometimes the neurosurgeons electrically stimulate through those same electrodes. And they do that to test the function of those regions causally. They also do it to test their hypotheses about the location of the seizure foci.

So in those rare cases, where you have a patient like this with selective electrodes like that where the clinicians decide that they are going to electrically stimulate through some of those electrodes, then we're in a position to kind of have it all scientifically, right? I don't mean to be so crude. This is a horrible situation for that lovely guy to be in, but scientifically, it's extremely powerful.

So I'm going to show you-- we did in fact have an opportunity. The same guys in Japan emailed me and said, OK, we're going to be stimulating that electrode. What do we do? And I said, OK, have him look at faces and have them look at other objects and ask him if anything changes.

And I'm going to show you a video of what happens when that goes on. OK, here we go. Oh, and I need to turn on the audio. OK, he's getting stimulated right there and he says--

[VIDEO PLAYBACK]

- [NON-ENGLISH SPEECH]

**NANCY** He's such a good subject, this guy.

**KANWISHER:**

- One more time.

- [NON-ENGLISH SPEECH]

- His eyes.

- [NON-ENGLISH SPEECH]

**NANCY** OK, that tells us that that region is causally involved in face perception. Is it causally involved in perception of things that aren't faces? He's getting stimulated in the same electrode. He doesn't know that there's a face area.

**KANWISHER:**

- [NON-ENGLISH SPEECH]

**NANCY** He doesn't know which electrode is being stimulated.

**KANWISHER:**

- [NON-ENGLISH SPEECH]

**NANCY** This is a Kanji character on a card here.

**KANWISHER:**

- [NON-ENGLISH SPEECH]

- One more time.

- [NON-ENGLISH SPEECH]

[END PLAYBACK]

**NANCY** Awesome, huh? What did we just learn?

**KANWISHER:**

**AUDIENCE:** You can trigger it.

**NANCY** You can trigger it, yeah. Yeah. So what does that tell us about the function of that region? Why is this-- I mean,

**KANWISHER:** it's amazing to see, no question, but what does it tell us scientifically?

**AUDIENCE:** It's specific.

**NANCY** Yeah. How does it tell us that it's specific?

**KANWISHER:**

**AUDIENCE:** Because when you stimulate it, it particularly sees a face.

**NANCY** Yeah. And what happens when he's looking at things that aren't faces?

**KANWISHER:**

**AUDIENCE:** [INAUDIBLE]

**NANCY** Yeah. So if that region was causally involved in perception of things that aren't faces, you might think that it would distort-- the box would look different or the ball would look different or the Kanji would look different. It doesn't, there's just a face on top.

So I think that's very strong evidence that that region is not only causally involved in face perception, but very specifically causally involved in face perception only. Everybody get that? Do I have to stop? OK. OK, I have another video.

Consider-- and we'll get back to this later-- consider other alternative hypotheses to this. This is pretty powerful. This is more powerful than most of the other things I showed you, but there's always ways to come up with alternative hypotheses, and that's the business we're in here. So be percolating on what other control conditions you'd want from this guy to really believe these data.