

[SQUEAKING]

[RUSTLING]

[CLICKING]

NANCY All right. So I'm going to finish up some of the things that I talked about with experimental design last time. And
KANWISHER: then we're going to get on and talk about category-selective regions in the cortex, which, of course, we've been talking about in various ways all along.

But I'll raise some general controversies about that, some alternative views from the kind of one that I've been foisting on you, and what I consider to be some of the strongest, most important evidence against the view that I've been putting forth here. And then we'll talk about decoding signals from brains. That's the agenda.

Here we go. So last time, I had you guys work in groups to think about experimental design because, really, most decisions about experimental design, once you know the bare basics of the measurement methods, they're just applying common sense, thinking about what it's like for the subject. How are you going to get the data you need?

So in terms of what exact conditions to run in any experiment, I talked about the idea of a minimal pair, this kind of platonic ideal of the perfect contrast, which never exists in reality but that you aspire toward. So ideally, you want two conditions that are identical, except for the one little thing that you're interested in. And you don't want to have other things that co-vary with that thing you're manipulating, other than the thing you're interested in. And that's the crux of the matter in experimental design.

You guys talked about what kind of tasks to have subjects do in the scanner. There's a trade-off between kind of doing the most natural thing, which is they're just lying there and stimuli come-- visual, auditory, whatever-- versus the fact that subjects might fall asleep if they have nothing to do. And if they fall asleep, you won't know. And that's not good.

So it's sometimes better to have a task to keep them awake and to tell you that they're awake. Key important point-- don't have one task for one stimulus condition and a different task for a different stimulus condition. If you did, that would be a--

AUDIENCE: Confounder.

NANCY Sorry?

KANWISHER:

AUDIENCE: Confounder.

NANCY Confound. Exactly. That would be a confound. Don't do that.

KANWISHER:

We talked about baseline conditions. So for example, in a vision experiment, staring at a dot or a cross is kind of as far as you can go in turning off your visual system. Why would you want to bother with that?

Well, it's sometimes useful to have that kind of baseline because we sometimes want to look not just at a difference between two conditions-- remember, one condition alone in MRI tells you not a damn thing. All we can see is differences. But even just two conditions showing you a difference, that's something.

But it can be ambiguous. So for example, if you had a situation like this where there was a response in some brain region to the red condition here and the green condition there-- they're just two numbers. That's all you have-- that is different. That's kind of meh. There's a difference, but it's meh.

But if you have a good baseline and you really know that zero is zero or as close to zero as you can get-- and now imagine if zero was here. That would be like, wow, that's a really strong effect and especially in neuroscience, where we care a lot, as you may have noticed, about selectivity, about how much more we get a response in one condition than another. And selectivities are usually more interesting as a ratio than as a difference, as I'm illustrating here. And so you can't get a ratio unless you have a third condition, usually a baseline.

All right. A few other things-- we talked about how you allocate subjects to conditions. You could have all your subjects do-- one half of the subjects do the face condition for an hour in the scanner. Another half of your subjects do the object condition for an hour in the scanner. That's no good. We don't want to do that.

We want a within-subjects design. We want all the conditions within a subject whenever we can do that. Why? Well, my best analogy to this is suppose we decided to grade your assignments as follows. A third of the class is going to be graded only by Heather. This third of the class is going to be graded only by Dana.

Across the whole semester, you guys are Heather people. You guys are Dana people. You guys are Anya people. Is that fair? No, that's dumb.

What if Heather's a hard-ass? And she is kind of a hard-ass, not that you guys aren't. They're all a pretty tough crew there. I stand here just waiting for the gong to go, wrong. And you guys should do that. I'm sure I've already said wrong things and you knew it. So next time, sound the gong and correct me.

Anyway, that wouldn't be fair in grading exams. And neither is it good in experimental design. So for all the same reasons that you guys can hopefully get an intuition here, you want to have all the conditions within a person, because maybe one person's brain just activates more than another person's brain.

Maybe this person had more coffee. Coffee increases your bold response. We give away free chocolate espresso beans before scans in my lab to increase the MRI response. All of that-- do designs within subjects whenever possible.

How do you allocate conditions to runs, these kind of subsets of a whole hour-long experiment where you scan people for maybe five minutes at a time and give them a break and another five minutes? Well, the same logic applies.

Imagine you're in a scanner for an hour. You're getting sleepy. You're getting bored. You're thinking about other things. You're kind of not on the ball. Those things change over slow periods of time. And so you want to get all those conditions together within a run, just as you want to get conditions together within a subject whenever possible.

And so then we didn't really get into this. And I think you did in your groups. But how do we stick all these conditions together within a run? Do we clump them together in a batch? Or do we interleave them? And I think most of you guys realize that there's this deep set of trade-offs there.

And so here's what's sometimes called the block design, where you clump a condition-- a whole bunch of trials with one condition, then a whole bunch of trials of another with, in this case, some kind of baseline in between, versus a mixed interleaved condition, which is called event-related for uninteresting historical reasons. And if it's event-related, you can have it slow or fast.

So why wouldn't you-- what are the reasons to do this rather than that? Many of you guys came up with this last time. So nothing--

AUDIENCE: Biases.

NANCY Sorry?

KANWISHER:

AUDIENCE: Minimize the biases.

NANCY Yeah. What kind of biases?

KANWISHER:

AUDIENCE: In a blocked experiment, they might be biased by what they're looking for.

NANCY Yeah, all kinds of biases. Consider this trial here in a yellow condition. Well, you just did a bunch of yellow trials.

KANWISHER: So maybe your yellow system is adapted out or something or biased somehow. But you also know that the next one's going to be yellow.

And so there's all that previous stuff and anticipatory stuff, all on top of the actual effect of a single yellow trial. Yeah. Was that what you were going to say, as well?

AUDIENCE: I was going to say that you also have effects from previous trials.

NANCY Yeah, all of those things, the effects of recent history doing the same thing and anticipation of the future, all on

KANWISHER: top of what actually happens in this trial. So those are not deal killers. But they're things to be aware of. So those are reasons why you might want to go with this condition or this condition.

Why wouldn't you always do this? Alternate the order-- not alternate. Randomize the order of conditions over time and bunch them in together. Why is that not always a great idea?

People do that sometimes. It's not a terrible idea. But there are things to keep in mind here. What's the challenge with that? Yeah.

AUDIENCE: One possible challenge is that the bold response has a 10-second window. So it's scraggly.

NANCY Exactly. So the bold responses here are going to be massively on top of each other. That's why people sometimes

KANWISHER: do this. It's like, OK, we'll have a random order. And we'll put a big chunk of time in between.

But if you have to stick 10 seconds in between trials, your subject is going to fall asleep. And you're spending all that expensive scan time not collecting enough trials. So none of these is right or wrong. They're right or wrong in different conditions.

So as Eke-- am I saying it right? As Eke mentioned, the challenge here-- let me just give you my crude depiction of this. So let's suppose this is time, a series of trials with a house, a dot, a face, a dot, a dot-- I don't know where that dot went-- a face, a house, et cetera. And each of those trials is one second long.

Well, let's imagine the response in the fusiform face area to that first house. You get some kind of middling low response that's going to take many seconds to peak. Let's look at the response to this face. Well, it's going to be higher. And it's going to peak out there.

And so then you can look at the response of each of these things. And so you get this whole series of bold responses from each of those different trials.

But now, here's the problem. What we observe when we measure the response of a little voxel, a little three-dimensional pixel in the brain, is the sum of all of that, something like this. It'll be higher than that, but some big blurry sum of all that. So now, we want to go backwards from observing this to seeing the difference between that and that. And that's a problem.

So that's not great. But here's the crazy thing. It's not impossible. It's not impossible because by weird, mysterious, to me still kind of unfathomable physiological mechanisms, these things add up approximately linearly. It's really counterintuitive.

Who would think a big sloppy biological system with many different causal steps could produce something that is approximately linear? But it does. And because they add up linearly, if you have enough trials, you can take this thing and recover that and that.

We're not going to go through the math of it. It's just basically addition. It's solving for multiple equations because you have all these different time points. Did everybody get the gist of the idea that even if you're observing something really slowly varying and weakly varying because it's massively blurred, you could, in principle, with enough trials, go backwards and solve for that and that? Everybody get that idea?

So what that means is it's a bit of an uphill battle to do a fast event-related thing. You can't just look at the response. You have to actually do a lot of math. And you may or may not have enough trials to pull it out. But under some circumstances where you really need things to be interleaved, you can pull that off.

All right. So that's what I just said. Blah, blah, blah. A few other design things that I didn't really talk about in detail-- one I've mentioned glancingly, but I want it to be more explicit about it, this whole idea that we've talked about a few times of defining a region of the brain that we're going to look at with a localizer scan with functional MRI.

We talked about that with the case of characterizing face areas. Go run a face versus object scan. Find the face area in that subject. And then do new experiments and testing.

Or when you guys proposed your snake experiments, you said, first, localize a candidate snake-specific region with snakes versus non-snakes. And then do repeated tests in that region that you found in each subject.

Why do we have to do all that within each subject? You don't technically have to. Lots of people don't. But the reason I think it's important and the reason we do it in my lab and all of my intellectual descendants do that and lots of other people do, too, the reason we do that is that region is not in exactly the same place in each subject.

So I have a dopey analogy. Brains are physically different from one person to the next. If we scan you guys just anatomically and look at the structure of your brains, your brains are as different from each other as your faces are.

That is, you all have the same basic structure, the same major lobes and sulci, just as you all have eyes and nose and mouth. But they're in slightly different positions. And that's just the anatomy. The function on top of that is even more variable.

So it's like trying to align faces. So if you have a bunch of photographs of faces and you try to align them on top of each other and superimpose them, even if you have a few degrees of stretch, you can't do it perfectly. You'll get some kind of mess like this. Just they're different. So they don't perfectly superimpose.

Well, it's the same deal with brains. You try to align them perfectly from one person to the next, but they're physically different. They do not perfectly superimpose.

So now, imagine that-- this is a totally crazy analogy. But it's the best I could come up with. Suppose you're a dermatologist. And you're interested in skin cancers that arise in the upper lip. Well, it could happen. There's more sunlight hitting the upper lip, whatever.

So now, you could take a whole-- and you're studying photographs to try to see how many people have it or something like that. So you could take a whole bunch of photographs. And you could just say, OK, I'm going to look right there. It's usually going to be the upper lip.

But it's not always going to be the upper lip. And so you're really throwing away a lot of information by choosing the wrong location. For this person down here, you missed it. You're looking at the wrong thing.

So in the same way, if you want to study that region, you've got to find it on each individual photograph. And similarly, if you want to study the fusiform face area or the snake area-- which doesn't exist, but whatever-- you've got to go find that thing in that person individually. Otherwise, you're really blurring your data just as those data are blurred there. Make sense?

All right. Good. Different topic about design-- these are just kind of different topics. I couldn't find a good segue.

So far, we have been talking about the most rudimentary simple possible experimental design. That means two conditions-- faces and objects, snakes and non-snakes, moving or stationary, whatever-- two conditions where you contrast and you look in the brain. Is there a higher response to A than B? Nothing wrong with that-- you can get pretty far with this.

But first of all, of course, we can have more than two conditions. So you can have one factor-- in this case, stimulus category-- with many different conditions-- faces, bodies, objects, scenes, whatever. So that's not rocket science. We've just added a few more conditions of the same factor.

Here, factor is the dimension you're varying. In this case, it's stimulus type. But we could get fancy. And we could have four conditions that are two factors varied orthogonally, like this. This is sometimes called a 2 by 2 design. We're going to vary one thing on this axis and another thing on this axis.

Why would we want to do that? Well, let's look at an example. Now, let's suppose that you were going to compare faces to objects-- in this case, chairs. But beyond just those two conditions of comparing the response in the brain when people are looking at faces versus objects, we could now ask, does a response in the brain to faces and objects depend on whether you're paying attention to the faces and objects?

What if you're paying attention to something else? What if we have little colored letters right in the middle of the display and they're changing rapidly over time and your task is to monitor for a repetition of a letter, a one-back task? And it's going really fast. So it's very demanding.

You're just looking at those letters. They're flashing up. Oh, two B's. You hit a button. It's very demanding. The information hitting your retina is still coming in from the face, because the little letter is tiny. It's not hiding much of the face.

What do you think? If you're doing the letter task, do you think you'll still get a response in the fusiform face area when the face comes up? And will it be higher than when the chairs come up? Any intuitions? Yes? Talk to me about that.

AUDIENCE: It should be because I think what people are doing is the signal winds and hits the retina. There'll be one in the process of it.

NANCY It's still coming in. Yeah. Will it be just as high? What do you think?

KANWISHER:

AUDIENCE: No, I don't think it'd be as high. But I think there'd be some response that's higher than that for chairs.

NANCY Does everybody see how this is kind of an interesting question? The machinery is the same. All the feed-forward stuff is the same. You can't-- when I tell you just now you're doing the letter task, now you're doing the face task, when you switch to the letter task, the wiring in your brain doesn't change.

KANWISHER:

All the same wiring is there. The stimulus is still hitting your retina. It's still going up the system. So it becomes interesting to ask how could it be different. Would it be different?

All right. I just wanted you all in the grip of this is a question that we might ask. So how could we ask this question? Well, we can do-- as I just said, we can have subjects in one case do their standard object task.

Look for consecutive repetitions of a face or of a chair. We can have all different kinds of chairs. But every once in a while, two in a row are the same. Or we could have this other task where they're monitoring for letter repetitions.

So does everybody get this 2 by 2 design? On one factor, we're varying the stimulus. Is it faces or objects? Faces or objects, those are the two conditions. That's just terminology.

And on the other factor, we're varying task. Are you doing the face-object task or the letter task? Yeah. Ben, is it?

AUDIENCE: So what is it that this task-- what conclusions does it allow you to draw that the simpler task won't?

NANCY
KANWISHER: Good question. Good question. Anybody have an intuition here? You mean other than just doing that, never mind the letters?

AUDIENCE: Right.

NANCY
KANWISHER: Yes, exactly. Exactly the right question. You guys, what do you think? Is there any reason to do this?

Does anybody care about this? What would it tell us? Yeah. I forget your name.

LAUREN: Lauren.

NANCY
KANWISHER: Lauren. Yeah.

LAUREN: The effect of attention on perception.

NANCY
KANWISHER: Yeah. So if we want to know not just is there some bit that responds more to faces than objects-- we've been doing that for weeks. Enough already. We know there is.

Now, we want to know does it matter what you're paying attention to. Is that thing just like a little machine that's going to do its thing no matter what? Or do you, the perceiver, have any control over it?

Here's another version of that question. You guys can all sit there and look bright-eyed and bushy-tailed and look at me and smile and nod and think about whatever you want to think about. And I won't know. You could be bored out of your mind, thinking about what you did last night, whatever. And I won't know. And that's great.

Isn't that nice, that we human beings are not trapped by the stimulus that's in front of us at any moment? Instead, we can control our mental processes to some degree. And if you choose to think about something else, you go for it. You have good judgment. That is fine. It happens to me all the time.

You have that ability. I have that ability-- not really when I'm lecturing. I kind of have to stay on task. That's why it's exhausting.

But anyway, we are not trapped. We are not completely controlled by the sensory world impinging on us. And that's a good thing. And so if you wanted to find out about how that works and study how well we can control our own mental processes, you would do something like this. Make sense?

All right. So this design enables us to ask a whole bunch of things. One, does a response in some region or voxel or wherever we're looking depend on stimulus category? This is what we've been talking about for a couple weeks now.

To do that, you could just say, OK, is there an overall higher response to these two conditions than those two conditions? You wouldn't worry about task. You just say, overall, is there a bit that likes faces more than objects? Everybody got that?

That's one thing. That's sort of what we've been doing so far is just comparing two levels of one factor. That's called a main effect-- in this case, a main effect of the factor stimulus type.

Or we could ask a different question. Does the response of a region of the brain depend on attention? So overall, never mind whether it's faces or objects. There are photographs flashing up there. Does it matter if you're paying attention to those photographs or paying attention to something else?

So for that, we compare the average of these two versus the average of those two. That would be a main effect of task. Make sense? It's just terminology. But it's important to see that we can ask these different questions of a 2 by 2 design.

Everybody with me? Anybody want to ask me something? This main effect isn't very interesting. It's kind of a weird one, but you could do it. So that's main effect of attention or task.

Now, we could ask, as someone else said a moment ago-- was that you, Lauren? Yes-- if we want to know, does the effect of stimulus category depend on attention? That's what a 2 by 2 design is-- it's that kind of question that a 2 by 2 design enables you to ask.

So to ask that question, really, what we would do is essentially look at this row and then that row. And then we compare them. So we might say how much higher of a response you get for faces and objects. And we get some number in that cell-- sorry, when you're paying attention to them.

And how much would you get when you're not paying attention to them? You're paying attention to the letters. And then we could say, oh, how selective is the face response when it's attended versus unattended? In other words, how do the response to stimuli depend on task?

It's not rocket science. But it's important to see how this humble little 2 by 2 enables you to ask these very different questions. So this question of how the effect of one factor depends on the level you're at with the other factor is called an interaction. And it's often the most interesting kind of question to ask of any kind of data, whether it's MRI or anything else.

You could think of it as a difference of differences or, more directly, how the effect of one factor depends on the level of the other factor. In this case, the terminology would be we would be looking at an interaction of stimulus category by task. Make sense? Everybody with the program about how this question is different than the two different main effect questions?

To get some practice with this, I'm going to have you guys come up here and draw some data. So we're going to consider a-- just to get experience with main effects and interactions, we're going to consider a main effect of factor X, which is an overall effect of X, the difference between condition one versus condition two within X. And we're going to consider interactions of factor X and factor Y. That is how the effect of X depends on Y and vice versa.

So I'm going to have you guys draw data. I need my first volunteer. This is not hard. How do I put this thing up? I forgot to check if I have red and black pens. Hopefully, I do.

If you don't volunteer, I'm going to pick randomly. And that could be worse. It's not too awful. Is it Carrie? Unfortunately, I remember your name. So come on up here.

So you got an easy one. This doesn't write very well, but it will do. This is your red pen. That's your black pen. So we have here the response in red or orange will be the attended case.

We're looking at a response in the fusiform face area, a possible response-- in this case, a pretty unlikely one. But never mind-- so the attended one. And there's an unattended one. And there's a response to objects and faces.

And what I want you to draw is a pattern of data in which there's no main effect of stimulus type, no main effect of attention, and no interaction of stimulus type by attention. So you're going to draw four dots. You can do X's and O's or whatever.

CARRIE: So no effect of the stimulus-- so that means it doesn't matter if it was a face or an object.

NANCY Uh-huh. Exactly.

KANWISHER:

CARRIE: So I guess--

NANCY You could do that for one. Go do that for the attended task first. Do that in red. Yeah.

KANWISHER:

CARRIE: So about midway.

NANCY You have to really lean on it. Oh, it worked for me. Sorry. We'll have the extremely counterintuitive thing of this is

KANWISHER: attention. There we go. Here you go.

Perfect, no main effect of stimulus type. Good. Now, we've got no main effect of attention. So take the blue pen.

CARRIE: So this is like--

NANCY And no interaction--

KANWISHER:

CARRIE: No effect of attention relative to attention?

NANCY That's right. No main effect of attention means no difference for attended and unattended.

KANWISHER:

CARRIE: OK. But stimulus type is important now.

NANCY No, no, no. We're still-- this is all the same. We're drawing all the same situation. Yeah, exactly. It's a little bit of

KANWISHER: a-- there you go. Beautiful.

Nicely done, Carrie. So that's kind of a dopey case. Well done. You can sit down. Yeah.

CARRIE: Hopefully, I seemed right.

NANCY So we're just starting basic here. That's what it looks like if you have no main effects and no interactions.

KANWISHER: Everything's the same.

All right. That's not going to happen if you're in the fusiform face area. If you get that, there's something wrong with your scan. Or something went way wrong. But we're just fleshing out the logical possibilities.

I need the next volunteer. Who's going to do a main effect of stimulus type, no main effect of attention, and no interaction of stimulus type by attention? Yes. Come on up here. Is it-- what's your name?

AKWILE: Akwile.

NANCY Sorry? Akwile. Yes. Right. Great. So go ahead and draw that for me.

KANWISHER:

I'm just going to clarify that this-- what did we do? This is unattended-- wait a minute. Yeah, unattended here. Here you go.

AKWILE: Oh, static. So you start with the--

NANCY There's main effect of stimulus type. It's probably easiest if you start-- yeah, start with attended. There's a main

KANWISHER: effect of stimulus type.

Great. You're in the FFA. The faces are going to be higher than the objects. And main effect of stimulus type says you're going to get a difference. Good.

AKWILE: Is that good?

NANCY Beautiful. Well done. Make sense to everyone? Thank you, Akwile. Does that make sense, everyone?

KANWISHER:

So what would this mean if you got this? OK, Akwile. You're not quite done. So you get that. What's that telling you?

AKWILE: It tells you that it responds to the stimulus. But the attention doesn't make any difference to the energy.

NANCY Yeah. The selectivity you get doesn't depend on attention in this case. Again, these are all-- we're just making up

KANWISHER: data. We're just considering the different ways the data could come out and what they would tell us. Everybody got that?

Now, the plot is going to thicken a little bit. Now, we're going to have a main effect of stimulus and main effect of attention and no interaction of stimulus by attention. Let's go ahead up here.

Is it Talia? Come on up here. Here. Mm-hmm.

Beautiful. Thank you. Everybody see how this is a main effect of stimulus type. Faces are higher than objects, a main effect of attention. Attended is higher than unattended, but no interaction. The effect of stimulus type is the same at each level. Yeah? Ben.

BEN: Just something that maybe was unclear for me-- does attention usually affect the selectivity or the average response?

NANCY These are great questions. Right now, we're just considering the logical possibilities. We'll talk about that later.

KANWISHER: Yeah, it's a good-- you should be wondering. Yeah.

So Talia, tell us. If you found that, what would that mean?

TALIA: So because the difference in effect between attended and unattended objects and faces is the same, that does show that attention plays an effect and the stimulus plays an effect. But there's no interaction between them, because the difference is the same.

NANCY It's like there are these two different things. There's face selectivity. And then there's just a big overall if you're
KANWISHER: looking at stuff, you get higher responses than if you're looking at the letters. Yeah, exactly.

All right. One more-- I need a volunteer. David. That's not a volunteer, I realize. It's different than a volunteer.

So draw me a case where you have a main effect of stimulus, a main effect of attention, and an interaction of stimulus by attention.

DAVID: This one. And then it goes like this.

NANCY Yeah. Beautiful. So here, we have-- oh, wait. Actually, hang on. Hang on. Hang on.
KANWISHER:

Wait a second. You got a main effect of stimulus. Actually, you don't have a main effect of stimulus here.

DAVID: Did I get rid of that?

NANCY Yeah, you got rid of that.
KANWISHER:

DAVID: I come back.

NANCY Ah. Now, you have a main effect of attention.
KANWISHER:

DAVID: Wait.

NANCY Wait. Oh, maybe I said it wrong. You didn't have-- OK. Wait a second.
KANWISHER:

DAVID: Oh, yeah. You're right.

AUDIENCE: No, he has a main effect.

NANCY I think I screwed you up here. We want a main effect of stimulus. Yeah.
KANWISHER:

DAVID: Yeah. And then let's just we move it a little bit like that. And we get--

NANCY Exactly. So don't go away. Does everybody see how this is a main effect of stimulus? Those guys are higher than
KANWISHER: those guys. A main effect of attention-- the green guys are higher than the blue guys, but an interaction like that difference is bigger than that difference.

Now, don't go away, David. If you got that, what would you conclude about the fusiform face area if you got those data?

DAVID: Well, the FFA, if it was like this, it depends-- not only does it depend on attention. But it kind of depends on attention more than-- maybe the object attention doesn't depend on attention so much.

NANCY

KANWISHER:

That's right. That's what your data show is that the response to faces is more strongly affected by attention than the response to objects. But another way of saying the same thing is to say that the selectivity is greater when you're attending than when you're not attending. Make sense? Or the differential response is greater.

Great. Thank you. Everybody got these basic ideas? They're pretty rudimentary. I don't want to insult your intelligence. But I really found that people often don't get main effects and interactions.

And often, really, the crux of an interesting design is an interaction. And keeping it straight from the main effects sometimes takes a little doing. So let's consider what is the key sign of an interaction.

Oh, well, we already have a case where, often, people draw an interaction where the lines cross. But they don't need to cross. David just showed you a nice interaction where the lines don't cross.

All right. Moving on-- that was all leftovers. That's bad planning.

Oh, sorry. What? Oh, yes, put the thing up. Good point-- or down. Thank you, Chris.

Let's talk about category-selective regions of the visual cortex. We have been talking about these all along. But it's time to get a little more critical. So first, I've been talking about how there's a patch in there that responds pretty selectively to faces.

There's a patch out there on the lateral surface that responds pretty selectively to bodies. And we haven't mentioned it much. But next week, you'll hear more than you want to hear about a patch smack in the middle there that responds selectively to images of scenes.

So you just look at that. And it's really damn near impossible not to wonder what else is lurking in there. What else is in there? And of course, we wondered that many years ago, me and Paul Downing who did the body area paper. He was my postdoc at the time.

And we said, well, let's just scan people looking at 20 different categories of objects. And we put in all kinds of silly stuff in there. I'm phobic about snakes, so I wanted snakes. He's phobic about spiders. We compromised in our creepies condition-- threw them both in there. It was kind of sloppy.

But we had food and plants, because we figured those are biologically important. We had weapons because those are-- and tools because those are important in other ways.

We had flowers because Steve Pinker has this line in one of his books saying that "a flower is a veritable microfiche of biologically relevant information." And he hypothesized based on that that people might have special-purpose neural machinery for flowers. It sounded like a crock to me, but it's an empirical question. So we threw flowers in there for Steve Pinker.

And so then we scanned people looking at all of these things. And we replicated in every subject the existence of selective regions for places, faces, and bodies. And we didn't find anything else. None of these other categories produced clear whopping selectivities in systematic regions of the kind that you see in every subject for faces, places, and bodies.

Now, I hasten to say that there are lots of ways with any method to not see something that's actually there. You might not have enough statistical power to see it. It might be that there's a whole bunch of neurons that do that, but they're scattered all over the brain. And so they're spatially interleaved with neurons that do other things, in which case MRI will never see it.

There are big black holes in MRI images where there are artifacts. And you can't see anything. And if the soul was right there, we wouldn't have discovered it yet because we can't see it in our MRI images, not that I know what the contrast is for the soul. You could work on that. Did you have a question?

AUDIENCE: Yeah, I'm just curious. Did you try it on text?

NANCY Yes. And we will get to that later. And there is absolutely a specialized region for text. And we'll talk about that in
KANWISHER: a few weeks. Yeah. We didn't in this experiment, but we and lots of others have in other cases.

So don't take this too seriously. My main point is just that you don't find a little patch of brain for any damn thing you test. Mostly, you don't find it.

There is some disagreement in the field about the case of tools and hands. There are many reports that if you look at pictures of tools or look at pictures of hands, you can get a nice little selective blob. I have looked at both of those many times. I don't see it. I don't know what everyone else is on about.

I'm confused about that. I just leave that as in play. I don't know. But with that exception, there's good agreement that faces, places, and bodies, everyone replicates. And most of these others, no one replicates.

And so in particular, nobody reports selective patches of brain that respond selectively to cars, chairs, food, or lots of other things. We have tested snakes, by the way, and not found anything, at least in the cortex.

So what does that mean? That implies, kinda sorta, that some categories are special in the brain, at least at this crude grain that we can see with functional MRI. And that seems pretty interesting and important. Yes.

AUDIENCE: I had a question about places. Did you distinguish between human-made places and naturals?

NANCY We'll get into all of that in excruciating detail next week. Yeah. It doesn't really make much of a difference. It
KANWISHER: likes all of those things. Yeah.

So I've been going around for 20 years saying, see, these categories are really special in the brain and the mind. And that's what we're getting from this. And that's deep and fundamental. It's telling us something about who we are as human beings or whatever.

Sometimes, I go off the deep end with huge claims. But not everybody buys this. And so what I want to do is allude briefly to the general kinds of ways you could argue against this and then talk in some detail about one main one.

So ongoing controversies-- this view here is highly caricaturized. And this is actually not right. The brain doesn't have completely discrete little regions. It's a mucky biological system.

If you actually look at the part, the face-selective regions, they have ratty edges and little kind of archipelagos of sub-blobs and stuff. It's kind of a bit of a mess. There's a general cluster in that vicinity in most subjects, but it isn't always a discrete blob unless you blur your data. You take any data and blur it enough, it looks nice and clean.

But if you want to know the actual native form in the brain unblurred, it's kind of mucky. So one could react to that in different ways. My reaction to that is like, what do you expect? It's a biological system. Does it really need to be perfectly oval-shaped with a perfectly sharp edge?

I don't really care if it's interleaved with other stuff around the edges. But people react different ways. And one kind of important alternative view is, look, how do we know that these are really things in the brain? I'm talking about them as things, pieces, parts of brain and mind.

And maybe they're just kind of peaks in a broader landscape of responses across the cortex that are fluctuating. And empirically, that's true. There isn't just one butte and then nothing else in the cortex around it. There's some kind of profile.

So it's a bit of a judgment call how excited you want to be about a big peak in a fluctuating background. And so there's much discussion about that. Is it really just a peak in a broader spatial organization? And if so, what is that broader spatial organization all about?

It just pushes that question back. It says, we're wrong to think about discrete things. But that still leaves many mysteries about what that continuous gradient is. So that's kind of one line of response, which I think is completely legitimate.

Any sort of version of that kind of blurs into this next view, which we've talked about a little bit. And that is to what extent can these things, if I'm calling them things, be accounted for just by their perceptual features. So we've grappled with that in a number of ways so far.

One of the first things we asked about the face area is, is it just responding to curvy stuff or round things or whatever? And so there are many lines of work where me and many other people have asked that question. And for the most part, the answer seems to be there are some featural selectivities in these regions, but probably not enough to account for their category selectivity. But that one, too, is still in play.

And there's this dude in England who publishes several papers a year saying, no, this thing isn't category-selective. It's just that. I'm going to try to assign one of his papers to you because I want you to expose you to alternate views. But I haven't yet taught you the key methods you need for that paper.

Anyway, so there's room for debate in that question, as well. Then there's just a continuum of OK, exactly how selective are these regions. I'm excited if a face area responds like this to faces and like that to objects. But hey, it responds like that to objects. Is that selective enough?

So there's a lot of debate about what that means. So there's a lot of room to push back on the simpleminded story I've been serving up to you guys. But what I want to do next is talk about what I take to be the most smart and serious challenge, which is somewhat different from all of these.

And this comes from a guy up at Dartmouth named Jim Haxby, who published the paper that was assigned for today. And I intended for you to struggle with it a little bit and try to understand it. But if you didn't understand it fully, I'm going to talk about it here. And hopefully, that'll make it more intelligible.

So here's the big idea that Haxby-- there are many ideas in that paper. But the part of it that's most relevant to us for now is the following. Even if the fusiform face area responds weakly to chairs and cars, in contrast with strong response to faces, that doesn't mean that it doesn't hold information about chairs and cars.

So all along, I've been just talking about one dimension. Does it respond like this or like that? And that's gotten us pretty far. But the essence of Haxby's idea is that we should care not just about the overall mean response. We should ask if there's information present in the pattern of response across voxels.

And his point is that even if there's a low mean response, you could still have information in the pattern across voxels, even if it averages to some low number. And that pattern of information could enable you to distinguish different categories.

So let's get very particular. So how exactly would you tell? So here's what Haxby did, essentially. Or here's the subset of the assigned paper that's relevant to the current question.

If we want to know, does the fusiform face area hold information about cars and chairs, thereby arguing against its selectivity for faces-- we should care about information in the brain, not just magnitude of response. If the brain is an information processing system, we care what information the parts contain, not just how much the neurons are firing.

All right. So if we want to know this, here's what you can do. Here's a version of what Haxby did. You scan subjects while they're looking at chairs and cars. You've localized the fusiform face area so you know where it is.

So now, you get the response. This is highly schematic. This is an idealized version of the cortical surface. Remember, the cortex is a surface. So we can mathematically unfold it and look at the magnitude of response of each voxel in the FFA.

FFA isn't square, but we're idealizing it here. Everybody get how that could be a pattern of response across voxels in the FFA when the subject looks at chairs? And maybe you have some other pattern when the subject is looking at cars.

Now, certainly, the pattern when they're looking at faces, all of these bars would be much higher. But our point is that even if these are low, they're different across voxels. So that's step one.

So then what Haxby says is you do the same thing in the same subject. You do it again, hopefully in the same scanning session. And you get another pattern, like this and this.

Now, here's the key question. If those patterns are systematic for chairs and systematically different for cars, then there is information in that region about the difference between chairs and cars. And chairs and cars aren't faces. So that's an important challenge to my story about how that region only does faces.

So how do you measure that? Well, there's lots of ways. Haxby's is the lowest tech and most intuitive. He just says, let's look at the similarity of this pattern to that pattern, repeated measures on cars-- I'm sorry-- chairs, same subject-- chairs on the even runs and chairs on the odd runs.

By the way, why do you split your data like this rather than like this? He does eight runs. We could take the first half of the runs, put those data here, and the second half and put them over there.

Or we could take the data like this and take even runs and odd runs. Why is even and odd better than first half, second half? Yeah.

AUDIENCE: I guess it doesn't allow the subjects to get used to one particular thing one after the other.

NANCY
KANWISHER: Well, they're doing the same thing. It's all the same data. It's just how you analyze it. Yes. What's your name?

BAYLA: Bayla.

NANCY
KANWISHER: Bayla. Yeah.

BAYLA: I'm not sure. I think it's probably easier to compare between one face and the other, I guess.

NANCY
KANWISHER: You can actually do it either way. It's like you scan these eight runs. Here they are. You can do that.

I don't know if you can see what I'm doing here with my whole crew. Or you can do this. Why is this better than this? Yeah. Isabel.

ISABEL: It could be a subject was really tired.

NANCY
KANWISHER: Yeah. Maybe they fell asleep halfway through the scan. Then if you do like this, the odd and even are going to be better compared to each other than first half, second half. Make sense?

It's another version of why you do things within subjects. It's the same kind of argument. Yeah. So he splits into even and odd.

And so you ask, how similar are they within a category, within chairs and within cars? You get two different correlation values. Just how similar are those patterns? You get an r value. And we compare that with how similar the patterns are between chairs, even, to cars, odd, and cars, even, to chairs, odd.

And so the key question you ask-- if there's information about chairs and cars in this pattern of responses, then the correlations will be higher within-category than between-category. In other words, two different times you scan looking at chairs, those patterns are more similar than chairs are to cars.

Make sense? It's pretty basic. But it's one of these things that's simple and yet subtle at the same time. Does everybody get this?

So you just do these repeated measures. And you look at these patterns of correlations. And if the patterns are more similar or more correlated within a category than between categories, then you have information in that pattern that enables you to distinguish those categories.

Yeah? So that's what Haxby did. Yes?

AUDIENCE: So if this information doubles, so would it be difficult to look at correlation?

NANCY Nothing, really. Well, wait a second. Oh, that's the same. That's essentially like this.

KANWISHER:

It's just that since we're going from even to odd in the within case, we're going to go even to odd in the between case. You could have done it this way, but then-- yeah. So that's the method.

What does Haxby find? Well, you guys could all look at the paper some more so you get a sense of it, because it's actually really nicely written, even though it's dense. Those science papers are very dense. But basically, here's what happened.

So in that paper, he says, yes, he can distinguish between cars and chairs in the FFA. And therefore, to quote from his paper, "regions such as the 'FFA'"-- notice the scare quotes he's putting there to diss me. I hear you. I hear you, Jim.

"Regions such as the 'FFA' are not dedicated to representing only human faces. Rather, they're part of a more extended representation for all objects." Them's fighting words. Everybody see how this is a serious challenge with a very elegant method?

So when I first read that paper, I was like, huh. All right. I'm paying attention.

But he didn't do everything right. I didn't like the way he defined the FFA. I found a million reasons to diss it. And I ran my own version.

And in my paper that we published, we could not discriminate those. So we said, ha. You can. We can't. You did it wrong. We did it right.

Then a few years later, Jim publishes a paper with a collaborator in which they re-analyzed their old data and said, actually, you really can't discriminate it very well. It was significantly above chance, but really lousy.

And so they concluded, "Preferred regions for faces and houses"-- that is, regions that respond preferentially to faces or houses-- "are not well-suited to object classifications that do not involve faces and houses, respectively." But I didn't get to gloat because right about the same time, we were redoing our experiments at higher resolution. And actually, we could distinguish two different non-faces in the fusiform face area.

So that was the little drama that unfolded. And so the current status is yes, you really can discriminate two different non-face categories within the fusiform face area, even if you do it right. Even if I do it right and I don't want that result and I do it right, I can get that result.

So that's true empirically. The ability to discriminate is feeble. It's not very strong, but it's significantly greater than chance. So does that mean I'm toast and I wasted the last few weeks telling you guys a bunch of BS that has been disproven and that I should not have been telling you? Yeah. David.

DAVID: Isn't it kind of like saying that you could use a vending machine like a clock and then ask the question, then what is this thing for? Well, it's obviously the office clock.

NANCY That's a great analogy. I love that. Absolutely. Absolutely.

KANWISHER:

So now, to me, the central question-- and here's another example that I think is exactly like that, but even more on point. And that is that there are deep nets that people have trained on faces.

VGG Face, it's really good at face recognition. It has only ever seen faces. It has only been trained on faces. That is all it's about. And if you feed it chairs or-- what do I have-- chairs and cars, it can discriminate between chairs and cars.

So even if you have this perfect representation that's only been trained on faces, that has only evolved-- if it evolved. We'll get to that later-- to deal with faces, it can still give you a somewhat different response to chairs and cars. And that doesn't mean that that's what it's doing.

So I think this is a really important challenge. But I think centrally, crucially, what we really need to be thinking about-- maybe Akwile has a contribution. Yeah.

AKWILE: So if it's only been trained on faces and you feed it a chair, what's the outcome? What's it say?

NANCY So it's just a bunch of feed-forward layers that are connected with boatloads of units at each thing and
KANWISHER: connected in the systematic pattern. And once you train it up, you can feed it any stimulus. And you can collect a response out the top.

So even though it was designed for and only been trained on faces, you can feed it in non-faces and get the response out the top and see. That is not the category, not the top layer where it says that's Joe or that's Bob. But just before that layer, there's a whole bunch of units that have some representation distributed across units.

You can take that and try to read it out and ask if there's information there. I'm not giving you all the details of how you do that. But hopefully, you can get at least the gist. And later in the semester, Katharina Dobs is going to tell you more about how you do all this kind of stuff.

I spent a lot of time in the last few weeks talking about a key difference between two different kinds of methods, one set of methods that allows this kind of inference and another set of methods that allows that kind of inference. I'm trying to give you guys a clue here.

Actually, what I'm going to do is let you percolate on this. I don't think this is obvious. I worried about this for years. I think there are many answers to it. It's not cut and dried.

I will say I have already presented to you at least two different lines of work that provide an important counterargument to this. One of the people who gave me crappy teaching evaluations last year said, she told us about counterarguments and then made us tell her how they could in fact, after all, be consistent with her data. I thought that was weird.

I was just trying to teach people to think about data. But anyway, I won't make you do that because somebody didn't like that before. But you can think about it. And we'll talk later.

And it's actually good to think about. And we will come back to it. But I want to get on with the rest. I mention all this because it is an important challenge. Yeah.

AUDIENCE: I'm wondering if objects are not processed in FFA, they must be processed somewhere else.

NANCY Totally.

KANWISHER:

AUDIENCE: Somewhere else.

NANCY Totally. I had a whole piece of this lecture on that. And then I thought, for once, I'm not going to go over my time.

KANWISHER: So I'm not going to talk about that.

But remember, there's all those other bits of cortex. I've just identified a few particular ones. There's lots of cortex in between. And the simple statement is there's a lot of nearby cortex near the FFA and the PPA that seems to respond generically to object shape. And the first-pass guess is that there's a general-purpose visual machine in there, in addition to some more specialized ones.

But I'm going to not say more at the moment. And I'll just say, actually, you may read it in papers. It's sometimes called LO or LOC. That's kind of a shape-selective region which is arguably the kind of generic, let's process everything else system.

Only if it's a clarification question. Ask it.

AUDIENCE: No. I was just wondering which came first, this work or the transcranial stuff.

NANCY Sorry. This work or--

KANWISHER:

AUDIENCE: Or the transcranial.

NANCY Ah, good question. The transcranial stuff has actually been going on for a long time. But the relevant kind that I
KANWISHER: talked to you about is more recent.

And you're right. It is one of the very strong answers to this kind of critique. There's several. Actually, I've told you about three, so far, answers to this. But think about it.

So what we're going to do now is talk about not just this particular use of this method to ask a serious question about the selectivity of regions in the ventral visual pathway. Now, what I'm going to do is argue that, actually-- I think I just said all of this-- what Haxby has given us is also a method to ask what information is present in this little patch of the brain. And that's an awesome thing.

So let's go on and talk about that. Let's talk about neural decoding with functional MRI. So that was an instance of it, but I'm going to cash it out in another way more generally.

So let's take the case where there's a person with a patch of their brain and a pattern of response across voxels in that patch of their brain when they look at some stimulus. Let's suppose you're given this. And you want to know what was that person looking at to produce that pattern.

What was the stimulus out in the world that produced that pattern? Can you do that? So more generally, can you read the mind with functional MRI? Or maybe a little more honestly, can you at least tell what the person saw from their pattern of brain response? Everybody get the question here?

How can we try this? Well, they're all variations of that Haxby method that I just told you about. But let's walk through this.

So the first thing you need is you have this pattern. And you're trying to figure out what stimulus produced that pattern in that part of this person's brain. Well, you need a decoder. You need to know what those voxels respond like when they look at different things, where you know the answer.

So what you do is you scan the subject on a bunch of different conditions to get your decoder. And then you can take your unknown data and compare it to those decoder data. So in particular, you have to train your decoder.

So you scan the person looking at, say, shoes, and you get pattern. You scan them looking at cats, and you get a pattern. Maybe you scan them looking at five, 10, 100 other things-- probably not 100. You don't have enough scan time-- but some number of things.

And so now, you know. You know this is how those voxels respond when the person looks at shoes. And this is how those voxels respond when they look at cats. Now, you test your decoder with your mystery pattern.

Now, you have your mystery unknown pattern. And you want to know, was that shoes or cats? Well, you can just look. What is it most similar to?

All the methods are versions of that. They're just fancy mathematical versions of that. So what do you think that pattern-- what produced that pattern?

AUDIENCE: Shoes.

NANCY Shoes. It's more similar to the shoe pattern. Exactly. You guys just did neural decoding. So that's exactly how you
KANWISHER: do this.

There are all kinds of ways of doing this, from just saying, is this more correlated with that than that? That's Haxby's version. Or you can put a whole big fancy machine learning rigmarole in there to do pattern classification, because that is, after all, what machine learning is so awesome at is pattern classification.

And this is just a straightforward pattern classification task. Train on these. Test on that. Is that sort of intuitive, what we're doing here?

So that's the agenda. That's the logic of how we do this. And so does that work? Well, a little bit. But you don't have to worry, at least at the moment, because there are a million ways.

About 10 years ago, I was getting called up by legal types all the time because, are people going to detect lies with functional MRI? And I thought this was a total crock. And I was going around giving talks on all the reasons why nobody has to worry that they're going to be compelled to testify by being shoved in a scanner and have their brains read.

It's not a totally stupid thing to worry about. But lest anybody-- I don't think this will happen. But lest anybody try to read your mind against your will while you're in an MRI scanner, you can totally foil it in any number of ways.

One, move your head. Two, if they've got your head bolted down, move your tongue. You totally mess up your whole signal if you move your tongue. Three, do mental arithmetic. You can totally shut down whatever they're trying to do if you think about something else.

Anyway, so we don't need to worry about it. It's not good for insidious kind of legal efforts. But it is pretty good for science sometimes.

So there are lots of versions of neural decoding with functional MRI. So we've been talking so far about decoding functional MRI patterns of response across voxels. That's called MVPA, Multiple Voxel Pattern Analysis. You don't need to memorize that. But when you see MVPA in a paper, this is what it's talking about.

But you can also do it with lots of other kinds of neural data. Oh, sorry. Within MVPA, you can ask it of a particular ROI in the brain, Region Of Interest, like V1 or the face area or the body area or something else. But you can also apply it to the whole damn pile of data from the whole brain and say, can I tell what this person is thinking by looking at their whole brain?

Beyond functional MRI, you can apply it to lots of other kinds of data. So you can do monkey neurophysiology, as we discussed briefly last time, where you have actual firing rates from individual neurons. And you can look at the response of each stimulus class to each neuron in a region of the brain.

And you can do the same deal, running a pattern classifier or a simple correlation method on the pattern of response across neurons, rather than voxels. Everybody see how that's sort of the same deal, just better?

Or you can do magnetoencephalography, as we've talked about. Stick your head in the big expensive hairdryer. Collect magnetic signals from all around the head, 300 channels. And now, those magnetic signals are changing over time.

So the cool thing about neural decoding with MEG is you can say, OK, let's take the data from just exactly 80 milliseconds after the stimulus flashed on. And let's ask, what can you decode then? What can you decode at 100 milliseconds, 120 milliseconds?

You can see the growth of information over time as neural information processing proceeds by running the decoder separately at each time point. I'm going to try to squeeze into a future lecture more talk about that, because I think it's cool. And we're doing a lot of it in my lab right now. Does everybody get the gist of this, at least?

So that gives you the time course of information extraction. Similarly, there are lots of different decoding methods. You can use, as I mentioned, the kind of simple, low-tech Haxby-style correlations. Or you can use something called linear support vector machines or various other kinds of fancy machine learning math to do those classifiers.

Let's take-- do I have time to do this? I'm going to skip this. Yeah, we're going to skip that. It's cool, but we're going to cut to the-- oh, I don't know.

No, we'll do it all. We've got time. All right. So now that I've wasted all that time deciding whether we had time, we're going to compare how well this works when you do it on MRI versus how well it works when you do it on neurons in monkey brains.

So there was a beautiful paper a few years ago that looked at this. So the question is here are these face patches in monkeys that I told you about and that David Leopold will be talking about at 4 o'clock today. And so the question is this particular one, AM, one of the nice face patches up there, these guys wanted to know what information is represented up there in face patch AM.

Is there information about different individual face identities? Can you use it to decode which face the monkey saw? And so they did this experiment two ways.

One, they did monkey neurophysiology. They recorded from 167 different individual neurons in that region. And for each neuron, they measured its response to five different faces.

In another condition, they popped the very same monkeys in the scanner. And they scanned them with functional MRI. And they did the same experiment. And they measured the magnitude of response of each of 100 voxels in that same patch of brain in that same monkey. And they got the MRI response to each of those 100 voxels and for each voxel to each of those five faces.

Everybody get that this is asking the same question? How well can you decode face identity from individual neurons or from functional MRI in the same animal?

And the answer is damn depressing. The answer is you can decode identity really well from neurophysiology. And you can't do it worth a damn with functional MRI-- big bummer. Yeah. So that's a drag.

It's just what it is. Presumably, remember, each MRI voxel has hundreds of thousands of neurons in it. So the real miracle is that we ever see anything at all. And when we can't see the neural code with the resolution that we need to to tell whether it's got information about face identity, that's just because we're averaging over so many neurons.

That was my lament at the end of the lecture on Monday, that there are so many limitations in human methods. And here's one of the key ones. What are the implications? It sucks.

Anyway, I want to get one more idea out. And that is-- yeah. Question?

AUDIENCE: Is that limited to fMRI? Or does it also translate to EEGs?

NANCY Oh, EEG's much worse.

KANWISHER:

AUDIENCE: Worse.

NANCY Much worse. Oh my god. Yeah.

KANWISHER:

The only thing that might be better someday is intracranial recording. But even there, you usually don't get enough electrodes. So you need these very rare cases where you have very-high-density grids of intracranial electrodes that some surgeon has decided by chance to put on a part of the brain where you happen to have hypotheses and you happen to be incredibly lucky to test your hypothesis. And that's very rare.

Did you have a question, Akwile? No. So I've been talking about neural decoding. And that's a way of asking what information is present in this batch of neurons or this bunch of voxels. And that's a really deep question to ask for cognitive science, because we're interested in information processing.

And we want to know what's represented in each region. It's really the crux of the matter in cognitive neuroscience. But we can also use it to ask in a richer way about the nature of that information in each region.

So suppose we want to know what exactly is represented there. We want to know not just that it can distinguish shoes from cats. That's OK. But suppose we want to know how is it doing shoes versus cats. Does it just know, for example, that shoes are elongated this way and cats are roundish?

And that's all it's using to do its classification. In other words, it's not really shoes and cats. It's this versus that or something. If we want to know how abstract those representations are or how invariant they are to variations in viewing conditions, then we can do the following cool thing. We can train the decoder on one set of stimuli and test on a different kind of stimuli.

So for example, we can ask, or are there representations of shoes that are invariant to, for example, color and viewpoint, chosen just because that was the nice shoe I could find when I was searching an hour ago? So if we train on these and test on that, is that going to work? Is it going to know that this is the same kind of thing as that? If it does, what have we learned about that shoe representation? Yeah.

AUDIENCE: It's kind of generalizable.

NANCY It's very generalizable. Yeah.

KANWISHER:

AUDIENCE: Different perspective.

NANCY Totally. It's not just this. It's something closer to shoeness. And we don't know exactly how far it is until we test more conditions.

KANWISHER:

But exactly-- we've shown that it's really abstract and generalizable. That makes it more useful. That makes it more cognitively interesting. We could even go off the deep end and say, OK, is it the concept of a shoe?

We could scan people reading the word "shoe" and ask, is that going to work? Anya's doing experiments like that. There are various people looking at this kind of thing. And so you can ask at any level how general or invariant is that representation.

So neural decoders are not just gimmicks to try to say, oh, I can see and I can read out what this person saw. They're powerful methods in science to characterize mental representations and to characterize how abstract they are.