

[SQUEAKING]

[RUSTLING]

[CLICKING]

**NANCY**  
**KANWISHER:** Here's the agenda for today. As usual, a bunch of announcements in red. Assignment 4 was graded. There will be comments showing up online on Stellar soon on any of you who didn't get a near-perfect score on it. And I'll also be going over a little bit of it in a moment. And then once we do that, we're going to talk about navigation-- how we know where we are, and how to get from here to someplace else, which is much more awesome than it sounds at first, as you will see.

OK. So quick review. OK. So what was the key point? Why did I assign the Haxby 2001 article for you guys to read? It presents this important challenge to the functional specificity of the face area and the place area. What was that challenge? What was Haxby's key point? Yes, Isabel.

**AUDIENCE:** Well, he was just attacked-- it just has a preference for rectilinear and not seeing if it's actual scanning for. It's not truly detecting whether it's a face or not.

**NANCY**  
**KANWISHER:** Yeah. He wasn't worrying about rectilinearity so much back then. But his point was that we shouldn't care just about the overall magnitude of response of a region. Like, OK, it's nice if the face area responds like this to face, but isn't like that to objects. But even if it responds low and the same to cars and chairs, it might still have information to enable you to distinguish cars from chairs if the pattern of response across boxes in that region was stably different for cars and chairs. OK? That's really key. We'll go over it at a few more points.

But that's essential, right? A lot of the details that I'm going to teach you that go by in class don't matter, but I really want you guys to understand the PPA. And that's the nub of it. OK?

So the idea is that selective-- his claim is that selective regions, like the face area, contain information about non-preferred stimuli. That is, like, non-faces for the face area, or non-places for the place area. And because they contain information, those regions don't care only about their preferred category. So why does Kanwisher get off saying the FFA is only about faces and the PPA is only about places if we can see information about other things in those regions? OK. That's a really important critique. That's why we're spending time on it. OK? OK.

Next, what kind of empirical data might be an answer to Haxby's challenge? I presented at least three different kinds of data that can address this and say, hey, wait a minute. You know, you have a point, but what kind of data could speak to that and respond to Haxby? We didn't actually talk about this explicitly in class, but think about it.

Here's the claim he makes. What might we say, right? So that's empirically true. Like, you look in the FFA. Even in my own data, I can distinguish chairs from shoes a little teeny bit in the FFA. OK? So that empirical claim is true.

Why might it nonetheless be the case that the face area is really only about face recognition? What other data have you heard in here that might make you think that? Yes, Ben.

**AUDIENCE:** Because it's the presence--

**NANCY** Speak up.

**KANWISHER:**

**AUDIENCE:** Sorry. The presence of low-level stimuli that are generally in faces, but can also be sparse on chairs and cars in context.

**NANCY** Absolutely. So yeah, put another way, even if you had a perfect coder for faces-- like take your best deep net for face recognition, VGG face-- it can distinguish chairs and shoes too, right? The features that you use to represent faces will slightly discriminate between other non-face objects. So the fact that we can see that information in itself isn't strong evidence that that region isn't selective for face perception. Absolutely. What else? Yeah, OK.

**AUDIENCE:** Like transcranial magnetic stimulation? When you stimulate the epithelial and you look at it face it affects it, but when you are like looking at other objects, the effect is no longer.

**NANCY** Exactly. And so what does that tell you about-- OK, so there's pattern information in there about other things beyond faces. But? Apparently it's not used, right? Now with every bit of evidence, you can always argue back. People would say, well, TMS, those effects are tiny. Maybe there isn't and we didn't have power to detect it, blah, blah, blah, blah. But at least, absolutely you're right. TMS argues against them. What else?

Or at least is a way to argue against it-- and the Pitcher paper that I assigned and other papers that we've talked about in here provide some evidence that actually, at least the occipital face area really is only causally involved in face perception even if there's information in there about other things. What else? What other methods can address this? Yeah.

**AUDIENCE:** That is pretty new simulation, and even when you're pressing hand against your face, you can perceive faces in it.

**NANCY** Exactly. Exactly. So these are both causal tests, right? OK, there's information in there. But is it causally used in behavior? TMS suggests not. The little bit of direct intracranial stimulation data that I showed you also suggests the causal effects when you stimulate that region are specific to face perception, again suggesting that even if there's pattern information in there, it's not doing anything important because we can mess it up and nothing happens to the perception of things that aren't faces. Absolutely. What else?

We talked about it very briefly a few weeks ago. Yeah.

**AUDIENCE:** So if you remove the [INAUDIBLE], it just completely makes a person incapable of perceiving faces. That is causing--

**NANCY** Yes, but the crucial way-- yes, but the crucial way to address Haxby would be what further aspect of that? Yes.  
**KANWISHER:** And by the way, we don't remove the area in humans, but occasionally, we find a human who had a lesion there due to a stroke and then we study them.

**AUDIENCE:** So they're still able to do other categories.

**NANCY** Exactly. Exactly. So all three lines of evidence from studies of prosopagnosia, electrical stimulation directly on the brain, and TMS, all can provide evidence to various degrees. Again, one can quibble about each of these particular studies. But all of those suggests that even though there's information in the pattern, Haxby's right-- there's information in there about other things that aren't faces.

The only causal effects when you mess up with that region are on faces, not on other things. That suggests that pattern information is what they sometimes say in philosophical circles is "epiphenomenal." That is, it's just not related to behavior and perception. Make sense?

OK, moving along, how can we then use Haxby's method to not just engage in this little fight about the FFA and how specific it is, but to harness this method and ask other interesting questions from functional MRI data. How can we use it to find out, for example, does the place area discriminate, say, beach scenes from city scenes? We want to know what's represented in there. How could we use this method to find out?

Yes, Jimmy.

**AUDIENCE:** If I do what Haxby kind of did, and try the decoder, and see if the decoder could decide and differentiate between the city and or like an acre of shade.

**NANCY**  
**KANWISHER:** Exactly. Exactly. So we talked about decoding methods last time as a way to use machine learning to look at the pattern of response in a region of the brain, and train the decoder so it knows what the response looks like during viewing of beach scenes, train it so it knows what the response in that region looks like when you're looking at city scenes, and then take a new pattern, and say, is this more like the beach pattern or is it more like the city pattern? And that's how you could decode from that region. Yes.

**AUDIENCE:** That doesn't tell as much, in the sense that it's not telling you-- I mean, we know that there is residue of information nevertheless, and that this community can be varied on any region considered at any time, always.

**NANCY**  
**KANWISHER:** We have a true nihilist here. No, it's a good question. It's not the case that you can discriminate anything based on any region of the brain.

So there are some constraints. There are some things you can find in some places and other things you can find in other places. And they're not uniformly distributed over the brain.

However, the fact we just-- the point I just made about yes, there's discriminative information in the face area about non-faces but maybe it's not used, should raise a huge caveat about this whole method. How do we ever know? We see some discriminative information. How do we know whether it's actually used by the brain, part of the brain's own code for information, or just epiphenomenal garbage that's a byproduct of something else?

It's a really important question about all of pattern analysis. We do it anyway because we're beggars. We can't be choosers in terms of methods with human cognitive neuroscience. And we want to know desperately what's represented in each region.

So we do this. But whenever you see these lovely, "I can decode x from y," things, you should always be wondering. Who knows if that fact that you, the scientist, can decode it from that region means the brain itself is reading that information out of that region? Big, important question.

All right, put another way-- so Jimmy mentioned just decoding in general and that's absolutely right. But to directly harness the Haxby version of this, what would we do? First, we would functionally localize the PPA by scanning them, looking at scenes and objects, find that region in each subject.

Then we would collect the pattern of response across voxels in the PPA while subjects were looking at, say, beach scenes. And so if this is the PPA, this is the pattern of response across voxels in that region when they're looking at beach scenes-- fake data, obviously, just to give you the idea. So we would split the data in half, even runs, odd runs.

That would be like even runs. Then we get another pattern for odd runs. And then we get another pattern for when they're looking at city scenes with even runs, and another pattern when they're looking at city scenes in odd runs.

So then, once we have those four patterns, what is the key prediction if using Haxby's correlation method? What is the key prediction if the PPA, if pattern of response in the PPA, can discriminate beach scenes from city scenes? What should we see from these patterns?

What's the key prediction? Claire. Key prediction-- you have these four patterns in the PPA, and now you want to know is there information in there that enables you to discriminate beach scenes from city scenes?

**AUDIENCE:** Is that like beach even and beach odd are more similar than beach even and city even?

**NANCY** Exactly. Exactly. Right. It sounds all complicated and it's easy to get confused. But the nub of the idea is really  
**KANWISHER:** simple.

It just says, look, the beach patterns are stable. We do beach a few times, we get the same pattern, more or less. We do city, we get a different pattern. And we keep doing city, we get the same pattern more or less. And the beach pattern and the city pattern are different.

So that's the nub of the idea. And so you can implement it with decoding methods or the Haxby versions, just to ask whether the correlation between two beach patterns-- beach even, beach odd-- is more similar than the pattern between one of the beaches and one of the cities. Just asking, are they stably similar within a category and stably different from another category? Does that make sense?

This is just a variant of this thing I showed you guys before. We just harnessed this to ask whether that region can discriminate. OK, and I just said all of this.

If you still feel shaky on this, there's a few things you can do. A version of my little lecture on this method is here at my website. You can look at that. It's just like six minutes and it's basically what I did before. But if you want to go over it again, there it is.

You can reread the Haxby paper, which I know is not super easy, but it's actually nicely written. And if you read it carefully, it explains the method pretty clearly. You can talk to me or a TA. And we'll get back to this question of whether we should do a whole MATLAB based problem set on this.

OK, let's move on and talk about navigation. This is a Monarch butterfly. It weighs about half a gram.

And yet, each fall the Monarch migrates over 2,000 miles from the USA and Canada down to Mexico. In fact, a single Monarch flies 50 miles in a single day. It's pretty amazing for this tiny, little, beautiful delicate thing. Even more amazing-- it flies to a very specific forest in Mexico that's just a few acres in size. And it arrives at that particular forest.

Now, that's already amazing, but here's the part that is just totally mind blowing and that is-- and it flies back north in the spring-- and that is that this whole cycle takes four generations to complete. And that means that the Monarch that starts up in Canada and flies down to that forest in Mexico-- one Monarch does that-- is the great-great-grandkid of his ancestor that last went on that route. Put that in your head and smoke it. That's pretty amazing.

Consider the female loggerhead turtle. She hatches at a beach, and goes out in the sea, and swims around in the sea for 20 years before she comes back 20 years later for the first time to the beach that she hatched at. Now, it's pretty amazing, but some mothers miss by 20 miles. They go to the wrong island or the wrong beach on the same island.

And so you might think, OK, it's pretty good. It's not amazing. But here's the thing-- the wrong beach that those mothers go to is the exactly right beach had the Earth's magnetic field not shifted slightly over those 20 years. They're exactly precise, but they just don't compensate for the shift in the Earth's magnetic field.

Here's a bat. This bat maintains its sense of direction even while it flies 30 to 50 miles in a single night in the dark catching food. And it maintains its sense of direction even though it's flying around in all different orientations in three dimensions, and even as it flips over and lands to perch on the surface of a cave. It doesn't get confused by being upside down.

This is *Cataglyphis*, the Tunisian desert ant. These guys are amazing. They crawl around on the surface of the Tunisian desert where it's 140 degrees in the daytime, and they have to crawl around up there to forage for food. And then because it's so damn hot, as soon as they find food, they zoom back to their nest and go down in the nest where it's cooler.

So here is a track of *Cataglyphis* starting at point A and foraging. He's meandering around looking for food going along this whole crazy path to point B. And then if he finds food at point B, boom-- straight line back exactly to the nest.

Now we might ask, how does *Cataglyphis* keep track as he's doing all this stuff of where his heading is back to his nest? The first thing you might think of is things like what it looks like. Maybe there are landmarks, maybe there are odors.

But no, he doesn't use any of those things. And we know that because when scientists who have set up this measurement device capture *Cataglyphis* after he goes out on this tortuous path and finds the feeding station, they capture him and move them across the desert-- on which they've drawn all these grid lines for the convenience of their experiment-- and they release them here. And what does *Cataglyphis* do? He goes on the exactly correct vector-- no landmarks, no relevant odors, and yet he's obviously encoded the exact vector of how to get home. Think about what that entails and what's involved.

**AUDIENCE:** The same vector with respect to north?

**NANCY** With respect to-- yes, well, with respect to absolute external direction, absolutely. So that's what I just said. So

**KANWISHER:** these feats of animal navigation are amazing.

And animals have evolved ways to solve all these problems unique to their environment. They've evolved these abilities because they really have to be able to find food, and mates, and shelter. And this is not just esoterica in the natural world. MIT students, too, need to be able to find food, and mates, and shelter.

So what is navigation, anyway? And what does it entail? Well, I'll argue over the next two lectures that there are two fundamental questions that organisms need to solve to be able to navigate. First one is, where am I? And the second one is, how do I get from here to there, A to B, wherever there is that you need to get?

So we'll unpack this. There are many different facets of each. But so for example, if you see this image, you immediately know where you are, and you also know where to go if, for example, it starts raining. You might rush into lobby 7, or if you're hungry, you might turn around and go back to the Student Center.

Same deal here-- if you see this, then you know where you are and where you would go to get to various things. Now, these judgments rely on the specific knowledge you guys have of those particular places. You recognize that exact place, and you have some kind of map in your head that we'll talk more about in a moment, that tells you where everything else is with respect to it.

But even if you're in a place you don't know at all you can still extract some information. So suppose you miraculously found yourself-- boom-- here. I wouldn't mind, actually, but that's not in the cards for a while. So you're here.

Even if you've just hiked around the corner, if you've never seen this place before, you have some kind of idea of what sort of place this is. Where would you pitch your tent? Where might you try to go to get out of this valley?

If it was me, I wouldn't. I have friends who would go straight up there and try to drag me along, complaining. If it was me, I'd rather look for some other route. But you can tell all of that just by looking at this image-- where you can go from there, not just what kind of a place it is, but what are the possible routes you might take.

So these fundamental problems that we solve in navigation, of knowing where am I and how do I get from here to there, include multiple components. In terms of where am I, the first piece is recognizing a specific place you know. So you might open your eyes and say, OK, this is my living room. I know this particular place.

But as I just pointed out, even if the place is unfamiliar, we can get a sense of what kind of place this is. Am I in an urban environment, a natural environment, a living room, a bathroom? Where am I?

A third aspect of where am I, a third way that we might answer that question, is something about the geometry of the environment we're in. So try this right now-- close your eyes. OK, now think about how far the wall is in front of you. Don't open your eyes, just think about how far away it is, how far away the left wall is and the right wall is.

And how about the wall behind you? Don't open your eyes. How far back is the wall behind you from where you are right now?

OK, you can open your eyes. It's not rocket science. I just wanted you to intuit that even though you're presumably riveted by this lecture, and thinking only about navigation, you sort have a kind of situational awareness of the spatial layout of the space you're in.

So you might have a sense of I'm in a space like this and I'm over here in it. And we'll talk more about that exact kind of awareness of your position relative to the spatial layout of your immediate environment. It's something that's very important in navigation.

And another part of that is you might think, how would I get out of here? If I'm seriously bored by the lecture or for any other reason I urgently need to get out of here, you probably know exactly where the doors are in this space. It's just part one of those things that we keep track of.

So those are aspects of where am I in this place. What are the things we need to know to know how we would get from here to someplace else? Well, the simplest way to navigate to another location another goal is called "beaconing."

And this is a case where you can directly see or hear your target location. So you're sailing in the fog. You can't see a damn thing, but you hear the foghorn over there, and you know you're sailing to that point. So you just go toward the sound, nice and simple. You don't need any broader map of anything else. You just hear it and head toward it.

Or if you see this, and your goal is to get to the green building, well, there's a green building and you just head that way. Now, you're going to have to go around a little bit to get around those obstacles, but you know where to head because you can see your target directly. These are cases where you don't need a broader, long-term knowledge of the whole environment.

If you can see your target, you just go straight for it. So that's beaconing, simplest kind of A to B. And it requires no mental map, no kind of internal model of the whole world you're navigating in. But if you can't see the place you want to go, then you need some kind of mental map of the world.

So what do we mean by a "mental map of the world?" Well, this idea was first articulated in a classic experiment way back in the 1940s. So this was actually one of the original experiments that launched the cognitive revolution, when we emerged from the scourge of behaviorism to realize that it was actually OK, and indeed, of the essence, to talk about what's going on in the mind.

And a really influential study that launched the cognitive revolution by Tolman was done on rats. And it went like this-- he trained rats. He put them down in this area, and they had to learn that there would be food out there at the goal.

And so they just have to make the series of left and right turns to find the food. So you train them on that for a while till they're really good at it. And then he put the rats in this environment.

Now, the environment is similar, except there's multiple paths, one that seems analogous to the old route. So what are the rats do in this situation? They run down here, they run into a wall, and they realize, OK, that's not going to work. No surprises yet.

But then, the rats immediately come back out and they go straight out that way. What does that tell you? What did they learn? Did they learn a series of go straight, and then left, and then right, and then right, and then go for a long ways?

No. That wouldn't work over here. They learned something much more interesting. Even though they were only being trained on this task here, they learned some much more interesting thing about the kind of vector average of all of those turns. Everybody get this? It's really simple but really deep.

So from this, Tolman and others started talking about cognitive maps, whatever it is you have to have learned in a situation like this so you can abstract the general direction. We don't just learn specific routes as a series of stimulus and responses. So there must be some kind of map in your head to be able to do this, and rats have that, and so do you.

So let's consider this question right now. Where am I? Where are you? To answer that question to yourself, there's something like this in your head.

And it probably doesn't look exactly like that in your head, but there's some version of this information that's in your head that you're using when you answer the question of where you are. And you have some way to say in that map of the world, I know not just what the MIT campus looks like and how it's arranged, but I know where I am in it.

Now, if you want to know how to get somewhere else-- like suppose you're hungry and you want to go over to the Stata Cafeteria over there. What else do you need to know besides knowledge of the map of your environment and where you are in it? What else do you need to know?

You know you have this map, you know where you are, and you know where your goal is. Now you have to plan how to get over there. What else do you need to know? Yeah.

**AUDIENCE:** You have to know which parts are paths and which parts are buildings.

**NANCY**  
**KANWISHER:** Yes, exactly-- where can you go in there? Actually, where can you physically get through? Actually, our vector is right over there, but you can't go that way because you can't go through that glass, even though you can see through it. So knowledge of physical barriers, and what's an actual path and what isn't is crucial.

What else do you need to know? Suppose we had a robot in this room, sitting right here facing the front of a room like you guys, and we're programming the robot on how to get over there. What are other things we'd have to tell that robot to get it to plan how to get over to the Stata Cafeteria? Yeah.

**AUDIENCE:** Things to watch out for, like cars and traffic.

**NANCY**  
**KANWISHER:** Absolutely. We'd have to know about obstacles, like moving obstacles, not just fixed ones. Absolutely. What else? Yeah.

**AUDIENCE:** Initial orientation.

**NANCY**  
**KANWISHER:** Yes. He has to know which way he's headed. You're going to give this robot instructions on which way to go. It matters a whole lot if the robot is starting like this or starting like that. The instructions are different in the two cases, and likewise for you guys. To plan a route, you need to know which way you're heading.

If you guys have ever been in Manhattan, and you come up from the subway, and you see the street's going like this, and you know it's north/south, and you don't know if you're heading south or north-- really common thing. It's not enough to know I'm at the junction of Fifth and 22nd. You need to know, am I facing south or north?

Otherwise you can't figure out which way to go. That's called "heading direction." We just did all that. You need to know your current heading. You also need to know the direction of your goal in order to plan a route to it.

So in this kind of taxonomy of all the things you need to know to navigate, we've just added that if you're going to navigate in your own environment, you need to know not just where you are in it, but which way you are facing in that mental map. And we also talked about this business of what routes are possible from here, how do we move around obstacles, where are the doors, where are the hazards like cars, et cetera.

A final thing you need to know is that even if you have a good system for all of these other bits, it's still possible to get lost in all kinds of ways. You lose track, you get confused, you get lost. So we also need a way to reorient ourselves when we're lost. And we'll talk a lot about that in the next lecture.

So this is just common sense. We're doing a kind of low-tech version of Marr computational theory for navigation. What are the things that we would need to know or that a robot would need to know to be able to navigate? Just thinking about the nature of the problem. So that's what we need.

What's the neural basis of all of this? So I'm going to start right in with the parahippocampal place area, not to imply it is the total neural basis of this whole thing. It's just one little piece of a much bigger puzzle. But we'll start in there because it's nice and concrete.

All right, so this story starts about 20 years ago. I think I mentioned some of this in the first class when I talked about the story of Bob and I talked about Russell Epstein, who was then my post-doc. And he was doing nice behavioral experiments, and thought it was trashy and cheap to mess around with brain imaging.

And he was going to have none of it until I said, Russell, just do one experiment. Scan subjects looking at scenes. I know it's kind of stupid, but just do it. Then you'll have a slide for your job talk.

And he scanned subjects looking at scenes and looking at objects. And here is one of those early subjects, probably me-- I don't remember-- with a bunch of vertical slices through the brain, near the back of the brain down there, moving forward as we go up to here. Everybody oriented?

Sorry, it's not showing up very well in this lighting, but there's a little bilateral region right in the middle there that shows a stronger response when people look at pictures of scenes than when they look at pictures of objects. So we hadn't predicted this. Yeah.

**AUDIENCE:** Is the pink the eye color?

**NANCY** Yeah. Yeah, pink is-- all the colors are-- there's significance maps or P levels, right. So pink is higher than blue, **KANWISHER:** but blue is borderline significant. So this is kind of dopey.

We didn't actually predict it for any deep reason. We hadn't been thinking about theories of navigation or anything like that. It was just one of those dumb experiments where we found something and we followed the data.

So we found this, and it's, like, OK, let's try some other subjects. So here are the first nine subjects we scanned. Every single subject had that kind of signature response in exactly the same place, in a part of the brain called "parahippocampal cortex."

So this is very systematic. And there's lots of ways to make progress in science. One way is to have a big theory, and use it to motivate brilliant, elegantly designed experiments. And another is you just see something salient and robust that you didn't predict, and you follow your nose, and try to figure it out.

So that's what we did in this case. It's like, OK, what the hell is that? So if you think about-- we eventually called it the "parahippocampal place area" after a little more work. If you think about what we have so far, we've scanned people looking at pictures like this and pictures like that.

And what we've shown is that little patch of brain responds a bunch more to these than those. So my first question is, is that a minimal pair? Tally, is that a minimal pair?

**AUDIENCE:** Sorry, I'm about my voice.

**NANCY** Sorry. Simple, simple. We're contrasting this with that.

**KANWISHER:**

**AUDIENCE:** Can you remind me what a minimal pair is?

**NANCY** OK, minimal pair is this thing we aspire towards an experimental design, where we have two conditions that are

**KANWISHER:** identical except for one little thing we're manipulating.

**AUDIENCE:** I don't really think it's a minimal pair, but I'm not really sure.

**NANCY** Well, I even told you what we were designing to manipulate, but--

**KANWISHER:**

**AUDIENCE:** There seems to be too many differences between a living room and--

**NANCY** It's ludicrous. I mean, it's a million differences here. So we don't know that we have anything yet. There's all

**KANWISHER:** kinds of uninteresting accounts of this systematic activation in that part of the brain.

So just to list a few that you've probably already noticed-- these things have rich, high-level meaning and complexity. So you can think about living rooms, or where you might sit, or somebody's aesthetic, home design, or there's all kinds of stuff to think about there, much more than just, OK, it's a blender. So there's just complexity in every possible way.

There are also lots of objects present here, and only a single object over there. So maybe that region just represents objects, and if you have more objects, you get a higher signal. There's another possibility, and that is that these images depict spatial layout, and that one does not.

So you have some sense of the walls, and the floor, and the layout of the local environment here that you don't have over there. And we could probably list a million other things It's a very, very sloppy contrast.

So how are we going to ask which of these things might be driving the response of that region? Well, a natural thing to do is just deconstruct the stimuli. So here's what we did-- this is actually way back 20 years ago.

There were better methods at the time, but I didn't know them, so I actually drove around Cambridge, photographed my friends' apartments, left the camera on the same tripod, moved all the furniture out of the way, and photographed the space again. Ha, ha. I know. And then these will be probably cut out with some horrific version of Adobe Photoshop that existed 20 years ago.

Anyway, we deconstructed the scenes into their component objects and the bare spatial layout. Everybody get the logic here? Just to try to make a big cut in this hypothesis space of what might be driving that region.

So what do we predict that the PPA will-- how strongly will it respond? Oops, how strongly will it respond if these two things are true? If it's the complexity or multiplicity of objects that's driving it, what do you predict we will see over there? We already know you get a high response here. What will we get over there? Yeah.

**AUDIENCE:** Probably get more biases to the furniture.

**NANCY**  
**KANWISHER:** Yeah, we'll respond more to this than that. Right. It's really simple-minded. If instead, it responds more to the spatial layout, what do we predict? Isabel.

**AUDIENCE:** It's going to respond to the empty rooms more.

**NANCY**  
**KANWISHER:** Yeah. And that seems like a weird hypothesis because these are really boring, this kind of nothing going on here. And there's just lots of stuff going on here. I mean, it's not riveting, but it's a whole bunch, whole lot more interesting to look at these than those.

Believe me, I got scanned for hours and hours looking at these things. And whenever the empty rooms came on, I was, like, oh, my god, I'm just so bored. There's just nothing here, whereas here at least there's stuff.

But that's not what the PPA thinks. What the PPA does-- oops, we just did the localizer-- it responds like this. This is percent signal change, a measure of magnitude of response, to the full scenes, way down, less than half the response to all those objects, and almost the same response as the original scene when all you have is a bare spatial layout.

Pretty surprising, isn't it? We were blown away. We were, like, what? What?

But can you see how even this really simple-minded experiment enables us to just pretty much rule out that whole space of hypotheses? It's not about the richness, or interest, or multiplicity of objects. It's something much more like spatial layout because that's kind of all there is in those empty rooms.

I mean, it could be something like the texture of wood floors or something weird like that. But one's first guess is it's something about spatial layout. Does this make sense? It's just a way to take a big, sloppy contrast, and try to formulate initial hypotheses, and knock out a whole big space of hypotheses. Yes. Is it Alana?

**AUDIENCE:** Yeah, I'm sorry, I might have missed the design. So people who are looking at the empty room would not have the furniture?

**NANCY**  
**KANWISHER:** Good question. I skipped over all of that. We did-- yes, that's true. We did mush them all together and one could worry about that, that when you see this, you remember that that's a version of this. Absolutely. Absolutely.

And so maybe-- yes, nonetheless, if what you were doing-- that's absolutely true, but if what you were doing here is kind of mentally recalling this, then why couldn't you also do that here? Maybe you could. You might argue that this is more evocative of that than this is, but it's also got lots of relevant information. Yeah, Jimmy.

**AUDIENCE:** For the furniture, did you guys try placing them in the exact position as the scene and seeing if that--

**NANCY** We did both versions for exactly the reasons you guys are pointing out. And it didn't make a difference. Yeah.

**KANWISHER:** Sorry, Cooley.

**AUDIENCE:** It'd be-- you would transfer if they were just responding to the things, like more stuff? Like in the empty room, there's more background, but there's still more background.

**NANCY** Totally. You're absolutely right. This is taking us pretty far, but it's still pretty sloppy. This stuff goes all the way

**KANWISHER:** up to the edge of the frame, and here there's lots of empty space.

Is that what you're getting at? Absolutely. I took out those slides because I felt I didn't want to spend the entire lecture doing millions of controlled conditions on the PPA. I thought you'd get bored.

But actually, another version that we did was we then took all of these conditions, and we chopped them into little bits and rearranged the bits, so that you have much more coverage of stuff in the chopped-up scenes than the chopped-up objects. And in the chopped-up versions, it doesn't respond differently at all. So it's not the amount of total spatial coverage. It's the actual-- something more like the depiction of space. Was there a question over there? Yeah.

**AUDIENCE:** I was wondering if there would be any difference between looking at images as 2D or 3D scene, and actually being there to see the 3D inside of the scene.

**NANCY** Totally. Totally. It's a real challenge. With navigation, navigation is very much about being there and moving

**KANWISHER:** around in the space.

And this is just a pretty rudimentary thing where you're lying in the scanner, and these images are just flashing, flashing on, and you're doing some simple task, like pressing a button when consecutive images are identical. It's not moving around in the real world. You don't think you're actually there.

But here's where video games and VR come in because actually, they produce a pretty powerful simulation of knowing your environment, feeling you're in a place in it. And so lots of studies have used those methods to give something closer to the actual experience of navigation.

So where are we so far? We've said the PPA seems to be involved in recognizing a particular scene. So this just says it responds to scenes and something about spatial layout, maybe. Does it care about that particular scene or do you have to recognize that particular scene to be able to use the information?

Now, our subjects mostly didn't know those particular scenes. But we wanted to do a tighter contrast asking if knowledge of the particular scene matters. So what we did was we took a bunch of pictures around the MIT campus, and we took a bunch of pictures around the Tufts campus. And we scanned MIT students looking at MIT pictures versus Tufts pictures. And then what else do we do?

**AUDIENCE:** Get the Tufts students.

**NANCY** Yeah, why?

**KANWISHER:**

**AUDIENCE:** Oh, just to make sure that it's not all about that weird architecture of the set.

**NANCY** Exactly. Exactly. So this is called-- yes, whose weird architecture? I think ours is weirder. So it's not just about the

**KANWISHER:** particular scenes or the particular subjects.

So everybody get how with that counterbalanced design, you can really pull out the essence of familiarity itself, unconfounded from the particular images? So when we did that, we found a very similar response magnitude in the PPA for the Tufts students, for the familiar and unfamiliar scenes. Really didn't make much difference. Yeah.

**AUDIENCE:** Taking a step back, so we started off with the one question of navigation and it involving all these different components. I just want to place this--

**NANCY** We're getting there. We're getting there. There won't be like a perfect answer. We're not going to end up with

**KANWISHER:** that slide, with the exact brain region of each of those things. We'll get some gisty, vague senses of what this is.

OK, so this tells us it's not about-- whatever the PPA is responding to in a scene, it's not something that hinges on knowing that exact scene. So it can't be something like, if I was here and I wanted to get coffee, what would my route from this location be, given my knowledge of the environment. Because otherwise, we wouldn't get this result.

So whatever it is, it's something more immediate and perceptual to do with just seeing this place. So where are we? We've said that there's this region that responds more to scenes than objects, that when all the objects are removed from the scenes, the response barely drops. And its response is pretty much the same for familiar and unfamiliar scenes.

So all of that suggests that it's involved in something like perceiving the shape of space around you. Doesn't nail it yet, but it kind of pushes you towards that hypothesis. Yeah, was there a question here a second ago? No? OK.

**AUDIENCE:** I was talking about experiment, but is it accurate when you look at a map?

**NANCY** Oh, great question. Not very much. Yeah, if you take pictures of places from above versus this kind of view, you

**KANWISHER:** get a response in this kind of view, but not above. Yeah, very telling.

OK, so I'm going to skip. We're not going to do the 30 other experiments. We're going to skip to the general picture, that here's the PPA in four subjects in this very stereotyped location.

And here are some of the many conditions we've tested. It's not just abstract maps like this. They don't produce a strong response.

Oh, this is an answer to Cooley's question way back. Here's the scrambled-up scene-- much lower response. So it's not just coverage of visual junk. And it responds pretty strongly to scenes made out of LEGOs compared to objects made out of LEGOs, and various other silly things.

So all of that seems to suggest that it's processing something like the shape or geometry of space around you-- visible space in your immediate environment. Nonetheless, there's always pushback. And there's pushback on multiple fronts, and there should be. That's proper science.

So one of the lines of pushback was this paper by Nasr, et al. that I didn't assign. I assigned you the response to it. Anyway, what Nasr et al. Did was scan people looking at rectilinear things like cubes and pyramids versus curvilinear, round-y things like cones, and spheres. And what they showed is the PPA responds more to the rectilinear than the curvilinear shapes.

OK, that's the first thing. And so then, they argue that in general, scenes have more rectilinear structure than curvilinear structure. And they did a bunch of math to make that case. And so they argue that maybe the apparent scene selectivity of the PPA is due to a what of scenes with rectilinearity? Yeah.

**AUDIENCE:** Confound.

**NANCY** Yes, exactly, a confound. This is exactly what a confound is-- something else that covaries with the manipulation  
**KANWISHER:** you care about that gives you an alternative account, namely it's not scene selectivity. It's just rectilinearity. I mean, that might be interesting to other people, but it would make it not very relevant to navigation and much less interesting to me, at least.

So that's an important criticism. And so then the Bryan et al. Paper that you guys read starts from there and says, let's take that seriously. Let's find out.

And so you guys should have read all of this, but just to remind you, they have a nice, little 2 by 2 design-- remember we talked about 2 by 2 designs-- where they manipulate whether the image has a lot of rectilinear structure or less rectilinear structure, and whether the image is a place or a face. And what they find in the PPA is the same response to these. And it's higher to the scenes than the faces, and rectilinearity didn't matter for the scenes.

So evidently, even though it does matter with these abstract shapes, in actual scenes and faces, it doesn't seem to be doing much. It's not accounting for this difference. Everybody get that?

OK, let's talk about this graph. Are there main effects or interactions here? And what are those main effects or interactions? Yes, Cooley.

**AUDIENCE:** There's many different scenes.

**NANCY** Yeah, of category, scene versus face. Anything else?

**KANWISHER:**

**AUDIENCE:** What's the first one? What was the first thing? In PPA category, what's the subtype?

**NANCY** Oh, wait, this here-- these are scenes and those are faces. I'm sorry, and this is the code here. These are

**KANWISHER:** rectilinear versus curvilinear. Just one main effect, or is there an interaction, or another main effect?

Just one main effect. These guys are higher than those guys. That's it. So that just tells you there's nothing else going on in these data other than scene selectivity.

Rectilinearity doesn't interact with or modify scene selectivity, and it doesn't have a separate effect. Nonetheless, as we've been arguing with all the whole Haxby rigmarole, does the fact that there's no main effect of rectal linearity in here mean that the PPA doesn't have information about rectilinearity? No, Josh, why?

**AUDIENCE:** This little, tiny moment that could be-- you know, this is not the right experiment to--

**NANCY**  
**KANWISHER:** That's right. This is a big-- well, it's the right experiment, but not the right analysis. It's the big, average responses are the same, but maybe the patterns are different. That wouldn't directly engage with this, but we wanted to know, was there information in there about rectilinearity.

So how would we find out? So this was your assignment, and I think most people got it right. But in case anybody missed it, we were zooming in on this Figure 4 here. So again, this is just the same basic design of experiment two.

And now, let's consider what's going on here. So you guys read the paper and you understood what was going on here. What's represented in that cell right there?

What is the point of this diagram? What are they doing here? And what does that cell mean in that matrix? You can't understand the paper without knowing that. Is it Ali? No, sorry. What's your name?

**AUDIENCE:** Sheldon.

**NANCY**  
**KANWISHER:** Sheldon, I've only asked you six times. Yeah, go ahead.

**AUDIENCE:** So they want to see whether the activation patterns can better discriminate between rectilinearity of the same category of things or between categories of things with the same rectilinearity. So the first thing I said is to the left and the second one is to the right. And they--

**NANCY**  
**KANWISHER:** Sorry, wait, here and here? No.

**AUDIENCE:** Right side. Yeah, so that part is discriminating between rectilinearity, and that side is discriminating between categories. And they take the differences of-- well, not the differences, they take how well it can distinguish between each of those categories and plot them down there.

**NANCY**  
**KANWISHER:** Right, OK. That's exactly right. So this is how well it can discriminate plotted down here, but based on an analysis that follows this scheme. So what does that cell in there represent, that dark green cell? What is the number that's going to be calculated from the data corresponding to that cell?

**AUDIENCE:** Similar piece of same rectilinearity and same pattern.

**NANCY**  
**KANWISHER:** Exactly. Exactly. So just as if you want to distinguish chairs from cars or something else, if you want to know is there information about rectilinearity in there, you take these two cases which are the same in rectilinearity-- both high rectilinear, both low rectilinear for run one and run two-- and that's the correlation between run one and run two for those cells.

That's the within rectilinearity case. And if there's information about rectilinearity, the prediction is those within correlations are higher than the between correlations, just as we argued a bit back with beaches and cities and everything else-- same argument. This is just presenting the data in terms of run one and run two, and which cells do we grab to do this computation.

So each of the cells in there-- for each of the cells, we're going to calculate an r value of how similar those patterns are. A pattern for rectilinear scenes in run two, a pattern for rectilinear scenes in run one-- this cell is a correlation between those two patterns. How stable is that pattern across repeated measures?

All right, so that's what that r value is. The two darker blue squares here are the r values for stimuli that differ in rectilinearity. And remember that the essence of the Haxby-style pattern analysis is to see if the within correlations are higher than the between correlations. In this case, the within correlations are within rectilinearity versus between rectilinearity.

And so then they calculate all those correlation differences and they plot them as discrimination abilities. And so what this is showing us here is that actually, the PPA doesn't have any information in its pattern of response about the rectilinearity of the scene. However, if we take the same data, and now choose within category versus between category, ignoring rectilinearity, and we get the same kind of selectivity correlation difference within versus between for category, there's heaps of information about category.

Does that make sense? Again, if you're fuzzy about this, look back on that slide. I have lots of suggestions for how to unfuzzy yourself on it.

So interim summary-- PPA responds more to scenes than objects. It seems to like spatial layout in particular. It does respond more to boxes than circles, but that rectilinearity bias can't account for scene selectivity.

That's all very nice, but what is a whole other kind of fundamental question we haven't yet asked about the PPA? So we've been messing around with functional MRI, measuring magnitudes of response, trying to test these kind of vague or general hypotheses about what it might be responding to. Yes.

**AUDIENCE:** Causation.

**NANCY** Yes, what particular causation?

**KANWISHER:**

**AUDIENCE:** I guess like how the scenes, like with how the PPA with what role it plays in the person being seen.

**NANCY** Exactly. Exactly. Again, we can test the causal role of a stimulus on the PPA, all of the stuff I talked about did

**KANWISHER:** that. Manipulate the stimulus, find different PPA responses. But what we haven't done yet is ask, what is the causal relationship, if any, between activity and the PPA and perception of scenes or navigation?

So far, this is all just suggestive. We have no causal evidence for its role in navigation or perception. All right, so let's get some. I'll show you a few examples.

So one, as, you guys have learned by now is these rare cases where there's direct electrical stimulation of a region, and there's one patient in whom this is reported. This patient again, is being mapped out before neurosurgery.

They did functional MRI in the patient first. This is his functional MRI response to, I think, houses versus objects. Houses are not as strong an activator as scenes for the PPA, but they're pretty good. PPA responds much more to houses than other objects.

And so that's a nice activation map showing the PPA. And those little circles are where the electrodes are, little, black circles. So they know they're in the PPA because they did functional MRI first to localize that region. Now those electrodes are sitting there.

And so first thing we do is record-- or first thing they did-- is record responses. They flash up a bunch of different kind of images, and they measure the response in those electrodes. And so what you see is in those electrodes right over there, 1, 2, 3, that correspond to the PPA, you see a higher response to house images than to any of the other images.

And you see the time course here over a few seconds. Everybody clear? This is not causal evidence yet. It's just amazing, direct intracranial recordings from the PPA-- I think the only time this was ever done, because it's pretty rare to have the electrodes right there in a patient who's willing to look at your silly pictures, and all of that.

But now, what happens when they stimulate there? So let's look at what happens when they stimulate on these sites 4 and 3 that are off to the side of the scene selectivity. And this is just a dialogue. We don't have a video, unfortunately. The videos are more fun, but this is just a dialogue between the neurologist and the patient.

And the neurologist electrically stimulates that region and says, did you see anything there? Patient says, I don't know. I started feeling something. I don't know, it's probably just me. No, it's not you.

And then they stimulate again. Anything there? No. Anything here? No. So that's right next to the side of the scene selective electrodes, right next door, a few millimeters away.

Then, they move their stimulator over here. They don't move anything, they just control where they're going to stimulate. Patient, of course, has no idea.

Neurologist says, "Anything here? Do you see anything, feel anything?" Patient says, "Yeah, I feel like--" he looks perplexed, puts hand to forehead-- "I feel like I saw some other site. We were at the train station."

Neurologist cleverly says, "So it feels like you're at a train station?" Patient says, "Yeah, outside the train station." Neurologist-- "Let me know if you get any sensation like that again."

Stimulates. "Do you feel anything here?" "No." And then he does it again. Did you see the train station or did it feel like you were at the train station? Patient, "I saw it."

These are very sparse, precious data, but that's so telling. It's not that he knew he was at the train station abstractly. He saw it.

So then, they stimulate again, right on those scene-selective regions. Patient says again, "I saw almost like, I don't know, like I saw-- it was very brief." Neurologist says, "I'm going to show it to you one more time." Really what he means is, I'm going to stimulate you in the same place one more time.

"See if you can describe it any further. And to give you one last time, what do you think?" "I don't really know what to make of it, but I saw, like, another staircase."

The rest I couldn't make out, but I saw a closet space, but not this one." He points to a closet door in the room. "That one was stuffed and it was blue."

"Have you seen it before," neurologist, "Have you seen it before at some point in your life, you think?" "Yeah, I mean when I saw the train station." "Train station you've been at?" "Yeah." Et cetera, et cetera.

So it's not a lot of data. But it's very compelling. What is the patient describing? Places he's in that he sees, and then he describes this closet space and its colors. Interestingly, colored regions are right next to scene regions, so that's kind of cool, too.

So it's causal evidence. It's sparse. Ideally, we'd like more in science, but it's pretty cool. Yeah.

**AUDIENCE:** At this point, the patient is just staring at a blank wall?

**NANCY** I actually forget in the paper. I've got to go look that up. I forget exactly what the patient was doing, whether-- I  
**KANWISHER:** think he's just in the room looking out. Usually, they don't control it that much because it's done for clinical reasons, and the patient is in their hospital bed, and they're just stimulating.

So he's probably just looking out at the space he's in. In fact, he must have been because at one point, he says, "The closet, not like that one over there." So if he was staring at a blank thing, he was also looking out at his room. So yeah.

**AUDIENCE:** This may be a little bit off topic. You said that the region for color perception is very close to this, it seems like. Is there any relationship between functional proximity and--

**NANCY** That's a great question. Nobody in the field has an answer to this. People often make hay about the proximity of  
**KANWISHER:** two regions, like there's some deep link because this thing is next to that thing.

The body selective region is right next to, and in fact slightly overlapping with, area MT that responds to motion. It's like, bodies move. Well, faces move and cars move, too. So I don't know. It's tantalizing.

It feels like it ought to mean something. And people often talk as if it does. And maybe it does, but nobody's really put their finger on what exactly it would mean. But it's useful.

So when Rosa Lafer-Sousa who you met in the color demo, and I showed that in humans, you get face, color, and place regions right next to each other in that order, that was really cool because Rosa had previously shown that in monkeys, the monkey brain it goes face, color, place in exactly the same order. And so we thought that's really interesting. That suggests common inheritance because that's so weird and arbitrary. Why would it be the same? So it can be useful in ways like that, at least.

So we just went through all of this. So how does this go beyond what we knew from functional MRI? I'm insulting your intelligence. You know the answer to this. It goes beyond it because it tells you-- it implies that there's a causal role of that region in place perception, some aspect of seeing a place.

Now, all of this about the PPA I just started in there because it's nice, and concrete, and easy to think about. But no complex mental process happens in just one brain region. Nothing is ever like that.

And likewise, scene perception and navigation is part of a much broader set of regions. So if you do a contrast, scan people looking at scenes versus objects, you see not just the PPA in here. Again, this is a folded-up brain, and this is the mathematically unfolded version so you can see the whole cortex. Dark bits are the bits that used to be inside a sulcus until it was mathematically unfolded.

So there's the PPA kind of hiding up in that sulcus. And when you unfold it, you see this nice, big, huge region. But you also see all these other regions.

Now there's a bunch of terminology and don't panic. I don't think you should memorize everything about each region. You should know that there's multiple scene regions.

You should know some of the kinds of ways you tease apart the functions, and some of the functions that have been tested, and how they're tested. But you don't need to memorize every last detail. Because it's going to get a little hairy.

So here's a second scene region right there called retrosplenial cortex or RSC. And actually, Russell Epstein and I saw that activation in the very first experiments we did in the 1990s, but we really didn't know what we were doing back then. And we knew that this is right near the calcarine sulcus.

Remind me, what happens in the calcarine sulcus? What functional region lives in the calcarine sulcus? It's just a weird, little fact, but it's kind of an important one that we mentioned weeks ago.

V1, primary visual cortex-- that's where primary visual cortex lives. And remember, primary visual cortex has a map of retinotopic space, with next door bits of primary visual cortex responding to next door bits of space. And in fact, that map has the center of gaze out here and the periphery out there.

So when Russell and I first saw that activation, we had the same worry that Cooley mentioned a while back. And that is the scenes are sticking out. There's stuff everywhere. The objects, there isn't that much sticking out.

And we thought, oh, that's just peripheral retinotopic cortex. But it's not. It's right next to there and it's a totally different thing.

And it turns out to be extremely interesting. You don't need to know all that. It's just silly, little history.

There's a third region up there that's on the outer surface out there that used to be called TOS and is now called OPA. I'm sorry about that. You don't need to remember this. Know that there are at least three regions.

But TOS slash OPA is interesting because there's a method we can apply to it that we can't apply to the others. What would that method be?

**AUDIENCE:** TMS.

**NANCY** Yeah, TMS-- it's right out on the surface. You just stick the coil there and go "zap." So of course, we've done a lot of that. Can't get the coil into the PPA or RSC. It's too medial.

**KANWISHER:**

And there's another region that we'll talk about more next time called the hippocampus. You saw the hippocampus when Ann Graybiel spent all that time digging in the temporal lobe to find that bumpy, little, dentate gyrus, approximately right in there. And so all of these-- and probably other regions, but these are the core elements of the scene selective regions that are implicated in different aspects of navigation.

So when you have multiple regions that seem to be part of a system, that's an opportunity. Because now we have the possibility that maybe we could figure out different functions for different regions. And then maybe that would really tell us more than just scenes and navigation, end of story. It's kind of rudimentary.

It would be nice if different aspects of the navigation story engage different parts of the system. So really what we want to know is, how does each of these regions help us navigate and see scenes. And I'm not going to answer that fully. The field is still trying to understand all of this, but I'll give you a few tantalizing little snippets.

So let's take retrosplenial cortex right here. So this is first the response of the PPA right there, and retrosplenial cortex, which is just behind it. This is just its mean response to a bunch of different kinds of stimuli, showing you that it likes landscapes and cityscapes, scenes, more than a bunch of other categories of objects. And that's true of both the PPA and RSC. No surprises here-- they're both somewhat scene selective.

But then in a whole bunch of other studies summarized in this graph here, Russell Epstein and his colleagues had subjects engage in different tasks while they were looking at scenes. In some tasks, they had to say where they were. He's at UPenn, and he showed his subjects pictures of the UPenn campus. And they had to answer all kinds of questions about what part of campus they were, where they were on campus, and also about which way they were facing given the view of the campus they were looking at.

Then he also showed people familiar scenes and unfamiliar scenes, much like we did with our Tufts study. And he had object controls. And you can see the PPA doesn't care about any of that, doesn't care, really, if they're familiar or unfamiliar, doesn't care what task you're doing on the scene. You're looking at a scene, it's just going.

So we didn't really tease apart functions there. But RSC responds differently in these conditions. It's engaged in both the location task and the orientation task. It responds substantially more when you look at images of a familiar place than an unfamiliar place. So this is the first time we've seen that in the same network.

And so now, think about all the things you can do when you're looking at a picture of a scene and you know that place. You have memories of having been there. You can think about what you might do if you were there, how you would get from there to someplace else. All of those things are possible things that might be driving RSC.

Another thing that might be driving RSC is that if you're looking at a picture of a familiar place, you orient yourself with respect to the broader environment that that view is part of. So what I showed you that picture of the front of Stata, you immediately imagine, I'm out on Vassar Street facing that way, roughly northwest, I think. If you look at a picture of a scene and you don't know that scene, it doesn't tell you anything about your broader heading in the broader world. So all of those are things that the RSC, its function seems to depend on knowing that place.

Perhaps the most telling case comes from a patient who had damage in retrosplenial cortex. And the description in the paper of this says that this patient could recognize buildings and the landmarks, and therefore, understand where he was. So lots is intact-- can recognize scenes and know where he is.

But the landmarks he recognized did not provoke directional information about any other places with respect to those landmarks. So this person can look at a picture and say, yeah, I know that place. That's the front of my house. But then if you say, in which direction is a coffee shop two blocks away, he doesn't know which way it is from there.

So this should sound familiar. This is my guess of the bit that my friend Bob got messed up. Yeah. This is exactly his description-- he could recognize places, but it wouldn't tell him how to get from there to somewhere else.

And so the best current guess about retrosplenial cortex is that it's involved in anchoring where you are. You have this mental map of the world, and you have a scene, and you're trying to put them together. Given that I see this, where am I on the map, and which way am I heading in that map?

Again, think about the problem you face when you emerge from the subway in Manhattan. You look around. Where am I, and which way am I heading? That's what you need retrosplenial cortex for.

How about this TOS thing? There's lots of studies of it. I'll give you just one little offering.

So this is a causal investigation because as we discussed, the TOS is out on the lateral surface. So we can zap it. And so of course, we do.

And so in this study, we were asking whether TOS is involved in perceiving the structure of space around you. So we took scenes like this from CAD programs, and we just varied them slightly. So for example, the position of this wall moves around, the aspect ratio, the height of the ceiling moves around, and we make this subtle morph space of different versions of this image.

And then for control condition, we do the same with faces. We morph between this guy and that guy, and make a whole spectrum in between. And then in the task, what we do is here's one trial. One of the scenes or faces comes on briefly, and then shortly thereafter, you get a choice of two, and you have to say which of these matches that one.

And then what we do is we zap people right after we present this stimulus. And so the idea is this is as close as we can get to a pretty pure perceptual task. How well can you see the shape of that environment or the shape of that face?

You don't have to remember it for more than a few hundred milliseconds. So it's really more of a perception task than a memory task. And what we measure is, we actually muck with how different these two images are in each trial, and measure how far apart they have to be in morph space for you to be about 75% correct.

That's the standard psychophysical measure. The details don't matter. But our dependent measure is, how different do the stimuli have to be for you to discriminate them as a function of whether you're getting zapped in TOS or not.

And so here are the data. So let's take the case where you're doing the scene task here. What this threshold is, is again, how different the stimuli need to be for you to discriminate them. So the higher the bar, the worse performance. They have to be really different or you can't tell them apart.

And so what you see is when you zap OPA, that lateral scene selective region, discrimination threshold goes up a bit. That means you get worse at the discrimination. The stimuli need to be more different.

Compared to zapping the top of your head-- remember, you always want a control condition, and there's no perfect control condition because it feels differently to be zapped in different places. But getting zapped up here is a better than nothing control.

And then here's the occipital face area. That's the lateral face region we talked about before when I showed you another TMS study. Basically, whenever there's anything lateral, we zap it because we can. And see, it's not affected here. Zapping the occipital face area does not mess up your ability to discriminate the scenes.

However, in the face task, we see the opposite pattern. For the face task, zapping the occipital place area doesn't do anything compared to zapping the top of your head, but zapping the face area does. This is a double dissociation.

If we just had the scene task, it would be like, yeah, maybe. Who knows. Maybe, who knows why. But it's not very strong.

But when you have these opposite things, then we really have much more strong evidence that these two regions have different functions from each other. Everybody get that this is a double dissociation, in the same sense of when you have one patient with damage in one location and another patient with damage in another location and they have opposite patterns of deficit, then we're really in business. Then we can draw strong inferences.

So we just said all of that. So that's just a little snippet. These and other data suggest that that region is strongly active when you look at scenes, and it seems to be involved in something like perceiving-- just directly online perceiving the structure of the space in front of you.

So we already did retrosplenial cortex. And next time, we'll talk about the hippocampus in there, and its role in the whole navigation thing. Now, since I have ended early-- a rare event-- I actually put together a whole other piece of this lecture, and I thought, no, don't always have a part you don't get to.

But then it turns out we do get to it. We're going to go over this more later, but we're going to start with this business right here. So anybody have questions about this stuff so far?

OK, so I've spent a lot of time talking about multiple voxel pattern analysis, because it's the only method I've mentioned so far that enables us to go beyond the business of saying how strongly do the neurons fire in this region to the more interesting question of what information is contained in this region. But I also ended the last lecture with this kind of depressive note-- that you can't see much with MVPA applied to face patches, even when we know there's information in there with electrophysiology data.

Remember, I showed you that monkey study where they tried MVPA in the face patches in monkeys and they couldn't kind of read out a damn thing. And then they try MVPA on individual neural responses of the same region, and they can read out all kinds of information. And that tells you the information is there and we just can't always see it with MVPA.

Now today, you've seen cases where can see stuff with MVPA in the scene region. So sometimes it works, sometimes it doesn't. And when it doesn't work, we're left in this unsatisfying situation that we don't know if the information isn't there or if the neurons are just so scrambled together that we can't see the different patterns.

So bottom line, we need another method. MVPA is a whole lot better than nothing, but we want to be able to ask, is there information present in this region even when we think the relevant neurons are all spatially intermingled? So let me just do a little bit of this and then we'll continue later.

So goal-- this new method is called "event-related functional MRI adaptation." And we use it when we want to know if neural populations in a particular region can discriminate between two stimuli, two stimulus classes. So for example, do neurons in the FFA distinguish between this image and that image?

So if we want to know that, we could measure the functional MRI response in the FFA and find this would be an event-related response, similar responses to the two. And as I just mentioned, that wouldn't mean that there isn't information in the FFA that discriminates that. It just says they have the same mean response. Everybody get that?

Now, if we zoom in, and think about what might neurons be doing, it's still possible-- even with the same mean response-- that neurons could be organized like this, with some of them responding only to this image and some of them responding only to that image. But it's also possible that all of the neurons respond equally to both. And we kind of desperately need to know-- I mean, not in this case. This is a toy example, obviously.

But we often, when we're trying to understand a region of the brain, we need to know which situation we're in. So that neural population can discriminate these two and that one can't. How are we going to tell which is true?

Well, we talked before about multiple voxel pattern analysis, but as I just said, it only works when the neurons are spatially clustered on the scale of voxels. So imagine you have these situations here. This is getting more and more of a toy example, but just to give you the idea.

Suppose where those neural populations land with respect to voxels is like this. So if each of these is a voxel in the brain, a little, say, 2 by 2 by 3 millimeter chunk of brain that we're getting an MRI signal from, if you have the different neural populations spatially segregated enough that they mostly land in different voxels, then MVPA might work here. Is that intuitive? Do you guys all see that?

Then we'd get a different pattern in these voxels if we're looking at those two different images. But even if we have the situation here, which is kind of informationally the same, if they're spatially scrambled so that they're in roughly equal proportion in each voxel, MVPA won't work. Does that make sense?

And so that's when we need this other method called "functional MRI adaptation." Make sense? I'm going to go one minute over probably. So the point of functional MRI adaptation is it can work even when there's no spatial clustering of the relevant neural populations on the scale of voxels.

So let me go through it quickly and we'll come back to it later. So here's how it goes-- the basic idea is, any measure that's sensitive to the sameness versus difference between two stimuli can reveal what that system takes to be same or different. So for example, if a brain region discriminates between two similar stimuli like these, then if we measure the functional MRI response in that region to same versus different trials-- so this would be a different trial.

You present Trump and then the chimp back to back. That's one trial, compared to a same trial, chimp and then chimp. And of course, we counterbalance everything, so we also do chimp and then Trump in another different case and then Trump and then Trump in another same case.

If we find that the neural response is higher when the two stimuli are different than when they're same, then we know that that region has neurons that respond differentially to the two. So remember, we started with a case where the mean response is the same to this image and this image if you just measure them alone. But now we want to know, do we really have neurons that respond differentially?

So we're using the fact that neurons are like people and muscles. If you keep doing the same thing to them, they get bored. Been there, done that.

So you present this back to back. You get a lower response than if you present this and then this. That's called "functional MRI adaptation." It's like that waterfall MT adaptation we talked about before, but just crammed into a fine time scale.

And so then if you do that, you can ask what a region thinks is the same. So then, we could ask, what about these two images? Does it think those are the same?

And if we find a response like that, what have we learned? So if these two respond like that, what have we learned about a region that shows? This is all fake data, obviously, but if we saw that, what have we learned? And then I'll let you go, as soon as I get a nice answer to this. Yeah.

**AUDIENCE:** So if it's the same between two pictures of the same stimuli, that means that it's activated. It can discriminate. But if the yellow is at the same degree as the red, it would just be the brain reacting to different pictures.

**NANCY**  
**KANWISHER:** You totally get that. It's probably right, and you totally get it. Key point-- just because I don't want to torture you guys and go way over-- but key point is, it's the same response is the lower response.

We tell that with this case, and we actually give it a same one. So same is lower than different. That's just how this method works.

Then we're basically asking, does that count as the same to this brain region? And we're finding, yes, it does. That tells us that those neurons are invariant to all kinds of things-- viewpoint, facial expression, when he last dyed his hair, who the hell knows, all these other things.

So we'll talk more about this. But the idea is, now we have another method in addition to MVPA that can start to tell us what neurons are actually discriminating. OK, sorry to go over.