

[SQUEAKING]

[RUSTLING]

[CLICKING]

NANCY

KANWISHER:

So this is the line-up for today. We're going to be talking about language today and on Wednesday. But I want to start with something that I gave very short shrift at the end of lecture last time. And I'm going to give it short shrift again, but in a slightly different way. You'll need this for the reading, which hopefully you've already tried, started.

Representational similarity analysis is subtle and rich and interesting. And it's taken me years of revisiting it to get its full force. So just keep going at it and hopefully every time you'll get it a little better.

So let me try another brief version of this. So representational similarity analysis is like a generalized case of multiple voxel pattern analysis that applies to other methods and characterizes a bigger conceptual space. So to remind you, multiple voxel pattern analysis with functional MRI is this business where you split your data in half. So you have one set of scans where people are looking at, say, dogs and another set where they're looking at cats, and a whole other separate replication where they're looking at dogs and cats.

You look at the pattern of response across voxels in each of those four conditions, dog 1, dog 2, cat 1, cat 2. And you ask if the pattern is more similar for the two different splits of the data in the same condition, dog 1, dog 2, and cat 1, cat 2, the diagonal here, than in the two cases where they're different, dogs to cats.

Everybody remember that? If you're having trouble with this, come see me or the TAs. That's not good.

So now, that's MVPA. And you can use that to ask of a given region of interest in the brain or the whole brain, if the pattern of response in that region can distinguish between class A and Class B. That's what it's good for.

So that's worth knowing, but it's impoverished. It's binary. I mean, cats versus dogs. It's a dopey example I choose. But whatever you choose, it's just going to be two things. It only takes you so far in characterizing what's represented in that region.

You can make it richer if you force it to generalize. So if these two are a smaller size and a different viewpoint from those, and it still works, then we show that there's generality. Train on one kind of condition, test on a slightly different version of them. That tests the invariants. That's richer and more interesting.

But even so, it's limited. So representational similarity analysis is a bigger, richer way of characterizing representations by looking at the pattern of response across multiple conditions, not just two and their variations.

So instead of something like this, we'd have something like this with a whole bunch of different stimuli or conditions that we scan people on. And then we look at all the pairwise combinations-- how similar is dog to cat, how similar is it to pig or horse or table or chair or whatever.

So then we have all of these pairwise similarities, which gives us a richer idea of what's going on there. And so now we don't have to choose a binary classification in there. We can look at that entire space. We can think of this whole space as our proxy for what is represented in that region of the brain.

So now that's cool. So everybody get the gist of how this set of pairwise similarities in a region of the brain is a richer idea of what's going on in that region and what it cares about? Everybody got that?

Now, chunk that matrix as one thing. That's a representation of what's represented in this part of the brain. But now we can take that unit and we can say, we can do the same thing on a totally different kind of data.

So here's what we just did. Here's like some region of the brain, voxels. We can do the same thing in behavior. Now we can say, OK, you rate for me how similar is a dog to a cat on a scale from one to 10. I don't know, six or something.

How similar is a cat to a pig? Four, I don't know. You can see, you imagine you get some similarity space. You could just get people to rate them and you could make a whole new matrix here.

Now you're characterizing your conceptual space over those same items behaviorally by asking people how similar each thing are. Here, we're comparing similarity of patterns of responses across voxels. Here, we're doing it by asking how similar it seems to people behaviorally. Everybody get how that's a similar kind of enterprise?

Or, we could record from neurons in monkey brains and show them the same pictures. And just look at the response across, say, 100 neurons in the monkey brain to a dog and a cat and a pig and so forth. And then we could, ask how similar is a response across neurons in the monkey to each pair of stimuli, just as we did that across each pair of stimuli across voxels. Everybody got that?

So in each case, we're getting a matrix like this. Now, we can do the totally cool-- oh, sorry, we're not quite yet. We can also do that not just on functional MRI voxels in the whole brain or in one region, but we can make separate matrices. These are obviously all fake data. I didn't take the trouble to make different matrices for each, right. But we can make different matrices for different regions of interest in the brain, one for each.

Voxels here, what's their pairwise set of similarities across those stimuli? Voxels over here, what's their pairwise set of similarities? Now, we can correlate these matrices to each other.

So we can say, for example, we had a bunch of people do ratings and give us their behavioral similarities based over these stimuli. And then we looked in some region of the brain and got the brain's similarity space and their responses across voxels. How similar are those to each other?

So it's like we've moved up a level. Each matrix is a set of correlations between each pair of stimuli. But then once we have that set of correlations, we can take the whole matrix and correlate it to another matrix.

This would be a way of asking in some region of the brain how well does the representation in this chunk of brain match people's subjective impression of that similarity space when you ask them about it. Everybody see how that's a way to ask that question?

We can also relate functional MRI voxels to neurophysiology responses across neurons. We can ask how similar is your FFAs-- let's not take the FFA-- your LO that likes object shape, how similar is its shape space in your brain measured with functional MRI to shape space in this part of the monkey's brain registered with neurophysiology. It's pretty cosmic, right.

We're asking if the monkey sees the world the same way you do, in a sense, for this method, by using these matrices and asking how similar they are across species and methods. Yeah?

AUDIENCE: So are the function for [INAUDIBLE] similarity. All of them are the same or?

NANCY
KANWISHER: You could do whatever you like. So you can do garden variety functional MRI like we've been talking about in here just like the Haxby thing from 2001. That's when it all started, right.

Just get a vector across voxels for one condition, a vector across voxels for the two condition, and correlate them. You can do that in responses across neurons.

But you can also do more exotic things. You can train a linear classifier on a bunch of voxels and say, how well can it decode the response to pig to the response to dog. And you can put that number in that cell.

So you can do it different ways, any measure of similarity. Or, very confusingly, there's an increasing trend to talk about dissimilarity, not similarity, by subtracting the r values from 1. I find that annoying, but it's all over the literature.

And who cares whether it's similarity or dissimilarity. Doesn't really matter. They're both ways of collecting a representational space. Yeah?

AUDIENCE: Are there any caveats into the [INAUDIBLE] that we should be available, since this is like a correlation of coordination.

NANCY
KANWISHER: Oh, a million. You're supposed to Fisher transform it and do all that garbage. And we're not discussing that in here. I'm just trying to give you the idea.

I don't mean to be dismissive. I'm skipping over all of that stuff to just give you the gist of the idea. For purposes in this class, you could just eyeball that and that. And you'd say, oh, they're really ident-- no, they're not identical. I guess, I did switch it. I did switch a few of them, oh. OK, anyway, whatever.

For purposes in this class, you could just eyeball them. Mathematically, an r -value-- we're leaving out all the details. Yeah, OK.

And, of course, we can compare behavior in a person to physiology in a monkey, or behavior in a monkey to physiology in a monkey. And here's one thing you need for the reading. I hope it didn't already stump you. It's in a tiny part of one of the figures. We could make up a hypothesis of what's represented here.

We might say, hey, consider this patch of brain. Maybe it represents the animate/inanimate distinction. In the ideal case, that would mean all it knows is animals versus non animals. And so that would mean this should be the representational similarity space.

If these are all the animals, they're all exactly the same as each other. All the non-animals are the same as each other. But any animal and any non-animal are different. So this is a hypothesized similarity space of our guess of what's represented in a region, a model of what we think is represented in a region.

And we can correlate that to any of these matrices to ask whether our hypothesis of what's in there is right. Does that make sense? OK.

So why is that so-- oh, this whole thing so totally cool? It enables us to compare representational spaces across regions of interest in the brain-- the FFA to the PPA, do they have similar representational spaces-- across subject groups-- this batch of subjects and that batch of subjects-- without having to align voxels.

We're not aligning voxels. We've left voxels behind. We're only using these matrices. We can do it across species, across methods, and across hypothesized models of what we think is going on, like that.

So more generally, this probes representations in a richer way. We don't need to have just 10 or whatever I put there. We could have, if we keep subjects in the scanner long enough, or monkeys in the lab long enough, we can get hundreds of stimuli and really characterize a rich space.

And we're looking at not just two discriminations, but lots. The key requirement for representational similarity analysis, to be able to do all this cool stuff, is the axes need to be the same. So the stimuli that you're getting the similarity of need to be the same in the person doing behavior, the person doing MRI, the monkey doing physiology, the model.

If the axes are not the same, then there's no way to correlate the matrices. Make sense? We'll keep coming at this again and again. You'll see it in the paper for tomorrow night. And we'll come at it again in class on Wednesday.

So that was all catch-up. So today, we are going to talk about language. And let's start by reflecting on what an amazing thing language is. So right now, there's a miraculous thing going on.

I'm taking some weird, abstract, hard-to-grasp, even for me, ideas someplace in my head-- god knows where, somewhere in there-- and I'm trying to take those ideas and translate them into this bunch of noises coming out my mouth. That's already pretty astonishing.

Like, what? What does that idea look like? Who the hell knows? How do you take an abstract idea and turn it into a string of sounds? That's wild. Nobody really knows pretty much a damn thing about how that works, fascinating mystery.

But then that bunch of noises is going through the air and producing, let's hope, pretty similar ideas in your head. Wow. We do this all day every day. Big deal.

But it's astonishing. It's just astonishing that that works at all. So that's the essence of language. That's why it's so cool.

And let's think about how we're going to think about this. So the first thing to note is language is universally human. All neurologically intact humans have language. There are about 7,000 languages in the world. Sadly, this number is shrinking all the time.

They are all richly expressive, including sign languages. There are no kind of impoverished languages that don't capture the full richness of expressible human experience. They're all equally rich.

Language is uniquely human. Yes, chimps and parrots can accomplish all kinds of cool things, especially if you train them extensively. But what they have is not anything really like language.

And to give you a vivid sense of this, let's look at Chaser, the Border Collie. And what I want you to think about as you look at this little video of Chaser the Border Collie is what is the difference between your language abilities and Chaser's. Chaser is pretty damned impressive, but you are more impressive.

So watch it and enjoy and think about how it's different from what you do.

[VIDEO PLAYBACK]

- Some of us burst with pride if our dogs can respond to two or three commands. But what if we haven't begun to understand the possibilities of what the animal mind can really do? Our friend, astrophysicist Neil deGrasse Tyson, is host of *Nova Science Now*. And he brings us big news from the frontier.

- Walk up, walk up, walk up.

- Meet Chaser, beloved six-year-old Border Collie of Psychology Professor John Pilley.

- Good girl. She was born to live in the Scottish mountains-- Chase, toe, toe, toe-- and herd sheep. Go, go.

- John has taught chaser to tend an extremely large, if unconventional herd, of 1,000 toys.

And she knows the name of every single one of these?

- I hope.

- I find this hard to believe, so I test Chaser's memory with a random sampling.

Chaser, find Inky. Well, she got one right. Find Seal. Whoa, and that one too.

In fact, she got all nine right. But what about a new toy she's never seen or heard the name of?

Chaser's never seen Darwin, hasn't even ever heard the name Darwin. So we're going to see if she picks out Darwin by inference. Find Darwin.

I have to ask her again. OK, Chaser, Chaser, Chaser, Chaser, find Darwin.

(EXCITEDLY) Darwin! He's got Darwin!

She did it. Chaser's never seen that doll before, yet she settled on the one toy she didn't know by deduction. It's similar to the way children learn language.

But how does Chaser's ability compare with other species? Besides us, chimps and bonobos are the animal kingdom's top linguists, capable of learning sign language, but very slowly. They can solve some sophisticated problems, but they don't always pay close attention to humans.

- Is he coming?

- When I see my dog, my dog wants me to be around. Whereas a bonobo and chimpanzee, they need me. They're basically like, hey, you got any food. Can I get any food off of you? They're not interested in making me happy.

- Since dogs do like to please us, that humans need to find a way to tap the potential in all of our dogs.

OK, put it in the tub.

And dogs like Chaser are just waiting for us to discover all that they can do.

[GRUNTS] Smart dog.

- And Neil deGrasse Tyson is here with the astonishing Chaser here. Tell me what you learned about animal behavior and child behavior.

- Who would have thought that the animals are capable of this much display of intellect. I think we like thinking of humans as top of some ladder and don't even imagine that other animals could even approximate what we do.

- All right, I think we all want to see.

- You want the demo.

- Can we do it?

- A demo of this.

- Do you think we can it.

- Sure. We can try it.

- This is so astounding. Can we take away the stool.

- Sure. Let's try this.

- We'll give it a try. [INAUDIBLE]. Thanks. All right, so we get down?

- Let's get down on dog level. That's always better.

- All right, [INAUDIBLE].

- OK, Chaser, find Goose.

[STUFFED TOY SQUEALING]

- OK.

- Can I do this one?

- You can do this one.

- Chaser, Chaser, find ABC. ABC-- you did it!

We thank you. And we want everyone to know that it's a truly remarkable ~~NOVA~~ *NOVA* tonight. Four wheels reporting tonight on *NOVA Science Now* on PBS. And to you and your brilliant dogs at home, goodnight.

[END PLAYBACK]

NANCY OK. She's a very good girl. And she knows a lot of nouns, right, 1,000 nouns, apparently. But what can't she do
KANWISHER: that you guys can do? Is this language? Yeah?

AUDIENCE: It's word identification. It's not language. You modify actions [INAUDIBLE] language to be able to put verbs and nouns together.

NANCY That's good-- verbs and nouns together. What else? Yeah, [INAUDIBLE]?

KANWISHER:

AUDIENCE: It's fortification of things. If they were like a bigger ABC and a smaller ABC type of thing, that distinction wouldn't be possible.

NANCY Alex the Parrot can do that one. I don't have the video of Alex, and I don't want to get too hung up on this, but
KANWISHER: some animals can do that kind of stuff. What else? Yeah?

AUDIENCE: Yeah, it's probably closer to like sound identification. like, how I can identify the sound of a train or the sound of a car.

NANCY So just some rudimentary thing, like, visual form and sound. How about when she found Darwin?

KANWISHER:

AUDIENCE: [INAUDIBLE].

NANCY Sorry?

KANWISHER:

AUDIENCE: Wasn't that case just, like you said, deduction? It was just like, it wasn't any of the words.

NANCY That's right, that's right. But that's pretty impressive, isn't it?

KANWISHER:

Turns out, kids use that rule too in learning language. It's a whole set of studies of how kids use rules to try to figure out what people are referring to when they learn novel words. And that's one of the things that kids use. If there's a thing here that I don't know and somebody's saying a sound here I don't know, that thing probably goes with the sound. Yeah?

AUDIENCE: I was about to say, I took 9.85 last semester. We talked about like an exact experiment where kids were able to learn the words of toys that were like not English words, but like "dax" and stuff. But then when they were given like a new object, they would be able to identify it as different.

NANCY Exactly. It's called mutual exclusivity. And that's exactly what Chaser is showing here. OK, so pretty impressive,
KANWISHER: but not fully language, more like memorizing a bunch of nouns plus mutual exclusivity plus some other stuff, maybe.

She certainly can't understand who did what to who and why. This is not even in the ballpark. This is the essence of what we talk to each other about is this kind of stuff, all kinds of complicated relationships between different concepts that we communicate in language.

So animals in-- not just taught English, but animals in their natural environments communicate in rich and detailed ways with each other. But usually in each case, about a very restricted domain. What kind of danger is around? What kind of food source is around? Those basic kinds of narrow things that are of survival value, those are the things that animal communication systems usually deal with.

And in contrast, human languages are open-ended and compositional. Compositional means that we combine words to say new things, things no human being has ever said before. So that you don't see in animals.

So what is language cognitively? That is, what do you have to know to know a language? Bunch of basic things. One is phonology, the sounds of language. We've talked about this a bit in the case of speech perception.

Just hearing the difference between a ba and a pa, or seeing the equivalent gesture. American Sign Language is a fully expressive natural language. And there the phonemes are different pieces of hand movements rather than sounds, but function as phonemes all the same. And we talked about a region of the brain that responds very specifically to speech sounds in humans.

Moving up into the language system, that's just an input system-- and by the way, we also talked about the visual word form area, a very recent addition to the input system in language. But that's only a few thousand years old. It's really phonology that's the native form of language that's been around for tens, if not hundreds, of thousands of years in human evolution.

So semantics, we need to know what words mean. That's lexical semantics. But we also need to know how meaning arises when words go together. And related to how words go together, we need to know about the syntax of a language.

That is, the structure or grammar of a language. And so each language has a set of rules about how you string words together in that language. And usually central to that-- it's not the only thing, but a central part of that-- is word order. And that whole set of rules for how you string together words, following word order rules, determines the meaning of the string of words.

For example, shark bites man is different than man bites shark. And that just comes out of the syntax that we know that in English in this kind of construction the first word is going to be the agent, the one who's doing the thing. And the third word is going to be the patient, the one who's receiving the doing. And that's just built into your language system, that you know that implicitly.

There's also the pragmatics of language. That is, how we understand what somebody actually means when they say something to us, which isn't always just a function of the actual string of words coming out of their mouth. So if somebody says it will be awesome if you pass the salt, it's not all that awesome to have the salt. It really means, please pass the salt.

The pragmatics of the situation tells you the actual intent. And so to do pragmatics involves thinking about the other person's intent, what are they thinking, what do they want, what's going on in their head, and using all that background knowledge to constrain what do they mean by this particular utterance.

So let's just survey of the main pieces of what we mean by language. But for the next two lectures, we're going to focus on the core, which is syntax and semantics, this stuff in here. And I will sloppily use the word "language" to refer to this stuff, not all the other stuff. And we'll focus really on the sentence understanding.

So what do we want to know about sentence understanding? Well, the first thing we want to know is, is it even a thing. Is language a thing separate from the rest of thought?

Second thing we want to know is, if it is at least something of kind of a thing, does language itself have component structure within it? Are there different parts of the language system that maybe do different things? And if so, what is represented and computed in each of those parts? And third, how do we represent meaning in the brain?

So these are the things we'll address over the next two lectures. And let's start with this question that'll probably take up the bulk of this lecture. Is language distinct from the rest of thought?

Another way of putting this, a more familiar way, is to ask, what is the relationship between language and thought? Or even more pointedly, could you think without language? Probably, every one of you has wondered about that at some point.

So take like two or three minutes, talk to your neighbors about this, see if you can figure out whether you can think without language, and then let's pool your insights. Talk, think.

[SIDE CONVERSATIONS]

NANCY OK, if you guys all nailed it, I'm sure you solve the whole thing. People have been talking about this for probably
KANWISHER: millennia. So, what do you guys think? What were some of your reflections on this question? Come on, you guys. Yes, Carly?

AUDIENCE: I said I think that they could think without language because of like we talked previously about how [INAUDIBLE] babies are given very complex thought. But, like, he was arguing that the whale research, there's also the thing that babies kind of form their own language that we don't understand, but I don't think [INAUDIBLE].

NANCY Not really. If you take three-month-old babies-- not really. So perfect, absolutely, you can hear this. Babies can
KANWISHER: think. You take 9.85, you'll learn more. They can really think about all kinds of stuff. It's really amazing how much they understand.

And at three to six months, there's little or no language. So there's a beautiful case of thinking without language.

Yeah, David?

AUDIENCE: On the other side, if you don't give a name to something, if you don't give it word to something, then it's hard to really know it. Like, maybe there are 20 different types of the color green. And if you don't decide to call one of them olive and another one khaki green or something like that, then--

NANCY Then you can't see the difference?

KANWISHER:

AUDIENCE: Well, well, I don't know if you'd ever think of the difference.

NANCY OK, let's think about this. Do you think could see the difference? Suppose I held up an olive patch and a khaki patch to you. And for whatever reason, you had been raised with deprivation of the words olive and khaki.

KANWISHER:

AUDIENCE: But somehow it's not about just a perception question. It's about remembering.

NANCY Yeah, bingo, bingo. So that's roughly what the literature shows. Anya, help me out here. I forgot to look this up.

KANWISHER: The literature still show that perceptually you can discriminate them just fine. It doesn't make a damn bit of difference if you have words for it. But if you have to remember it-- sorry.

AUDIENCE: Faster.

NANCY Faster, faster. But accuracy in D prime, I don't think is different, maybe a little bit. Oops, caught. I meant to look

KANWISHER: this up. I knew this is going to come up. Write me an email to look this up and help me find the relevant stuff. Anyway, doesn't make a huge difference perceptually, but it does if you have to remember it for later. Yeah?

AUDIENCE: That's actually what I say because I'm actually reproducing the experiment that found that there was a difference in color [INAUDIBLE].

NANCY Aha, aha. What? In perception or memory?

KANWISHER:

AUDIENCE: So they found that-- I believe it was--

NANCY Because there's been a long history with this. They find one thing and they-- that's partly why I'm--

KANWISHER:

AUDIENCE: It's like a difference in the reaction time. Interesting enough, they found that if they introduce interference in their linguistic system, then that difference went away. So that's evidence that the language is causing the difference.

NANCY And that's in a perceptual discrimination. OK.

KANWISHER:

AUDIENCE: It's pretty small.

NANCY Yeah, yeah, well, behavioral, well, yeah, effects often are. Yeah, Isabel?

KANWISHER:

AUDIENCE: I remember one of the first neuroscience talks I went to in college was a woman who had been [INAUDIBLE]. She got in a terrible stroke and she's suffering from aphasia [INAUDIBLE] the speaking part and forgot all the language she learned. It took over her over a year to regain [INAUDIBLE].

And I remember the question that I asked was, you have this really terrible pain [INAUDIBLE]. But what did your inner voice sound like? And she said, well, I don't really have one, [INAUDIBLE]. And then she said, well, I must have thought in images and feelings.

And the interesting thing that I experienced when I was relearning to talk was that, the more English I learned, the more my thoughts was with grammar. So I still could have these thoughts, but they were formulated in a different way than they were when I had [INAUDIBLE] the structured language department.

NANCY

OK, that's great. So we're going to learn more about all of that, absolutely. OK, very good. So cool question, not obvious. Let's see what the data say.

KANWISHER:

So first of all, you guys talked about babies and how they can think. But animals can think too, maybe not fully as richly as we can, but they can think in all kinds of subtle, rich ways. And animals don't have language. And so that's another case, animals and infants.

And I'm mentioning numerosity because these are things we happen to have mentioned in here. Remember, the approximate number system. Animals are great at that. Very young infants are greater than when they don't have language at all. Also, by the way, people whose language do not have any number words whatsoever can do approximate numerosity.

So here's a cool study from Ted Gibson's lab a few years ago. They went down into remote parts of the Amazon to study this group of people, the Piraha. Here they are in their canoe. They are a hunter-gatherer tribe of just a few hundred people.

Their language is, as far as linguists can tell, unrelated to anyone else. And it has no number words. There's a whole dispute about that, but the current view is there are really no number words at all, not even for zero or one.

So how do they do at approximate magnitude? Well, let's see. So here is the testing session down in the Amazon. And this is the experimenter lining up a bunch of, I think, they're batteries.

And this guy is asked to match the number of balloons to the number of batteries. And he has to do it aligned this way so he can't just put them one next to the other. If you let him, he'll put them one next to the other. But this is designed to test it better.

And he puts down four balloons.

[VIDEO PLAYBACK]

[SIDE CONVERSATIONS]

Bingo, very good. OK, no number words in his language. What about this case?

- Hi, people.

NANCY

Oh, the plot is thickening.

KANWISHER:

- Six or five, [INAUDIBLE]

- [INAUDIBLE] lot of thread. [INAUDIBLE] of thread.

NANCY

He laughs, he thinks that's pretty funny. But watch.

KANWISHER:

- Five, five.

NANCY [? Valiant ?] goes ahead.

KANWISHER:

- [INAUDIBLE] I think it is [INAUDIBLE] five. Lots, lots. [INAUDIBLE] and intensifier, like lots and lots.

- You're doing well.

NANCY Right, right, right.

KANWISHER:

- There you go.

- Good.

- You can see which one.

- Nine-- nine, 10.

- 10? That was 10?

[END PLAYBACK]

NANCY So I think he gave nine for 10, or something like that. Anyway, if I had any of you guys do this task and I

KANWISHER: prevented you from counting by having you do verbal shadowing or something else to tie up your language system, you would do exactly the same as this guy does.

So the approximate number system doesn't require language, doesn't require number words in your language to get the concept. And it doesn't require use of language to do the task.

AUDIENCE: He saw him put [INAUDIBLE]?

NANCY Sorry?

KANWISHER:

AUDIENCE: He actually saw him put all of them? He saw?

NANCY Yeah, just like you--

KANWISHER:

AUDIENCE: [INAUDIBLE].

NANCY I mean, that's the actual experiment being conducted right there. OK, so we've just argued that at least the

KANWISHER: approximate number system is present in animals who don't have number words, infants who don't, and adults who don't have number words.

What about other aspects of thought? And what can we learn from studying brain disorders, as Isabel mentioned a moment ago, a very rich source.

So here's the question we're considering. We're taking language and thought, or cognition, and we're asking whether they're totally separate in the mind and brain or whether they're totally the same thing or whether there's some relationship that they're somewhat different. So that's the question.

What do we learn from brain disorders? Well, let's start with developmental disorders. And there are unfortunately a large number of these.

For example, there are language savants, people with Down syndrome, Williams syndrome, Turner syndrome. These are all developmental disorders in which people have very low IQs, but, notably, in each of these cases, very good language.

Perhaps the most striking is Williams syndrome. These kids are remarkable. They have very low IQs. They can't do the most basic spatial reasoning tasks. They can't cross the street safely. They can't live independently at all.

And yet they're highly social. And their language is almost indistinguishable from any of yours. Not quite-- if you test them subtly, can find some differences, but it is rich and complex.

And it's bizarre. Because you'd think if your thoughts are so impoverished because your IQ is low, how could you have rich language. But that's the weird thing about Williams syndrome. Their language is extremely rich and, in fact, poetic and quite beautiful and expressive.

So that's really surprising and suggests that you can have quite severely impaired cognition and very good language. So that's the first crack that these things are more separate than you'd guess. Actually, I find this one more surprising than all the others.

But in cases of brain damage, which was the first mental function localized in the brain. So this is historically important. Way back in 1861, Paul Broca stood up in front of the Anthropology Society of Paris and he announced that the left frontal lobe was the seat of speech.

And this is on the basis of his patient Tan, who had a big nasty lesion right there in what became known as Broca's area. Tan was his name because, after that lesion, that was all he could say.

So this is back when the mainstream view was very much against localization of function in the brain. There were people like Franz Josef Gall who were going around saying that different parts of the brain did very different things, but Gall was kind of a nut and he was not taken seriously by the academic elite, whereas Broca was a fancy member of the French academic societies and a muckety muck. And when he announced that the left frontal lobe is the seat of speech, everybody had to pay attention. So it was big stuff.

Importantly, Broca noted that Tan wasn't globally impaired at thinking, that Tan could do all kinds of things, even though he could not speak. So he was already onto this critical idea way back in 1861.

And he's just the most famous in that group. There were a bunch of people before him in the decades before who were reporting similar kinds of associations.

So what would it be like to have intact thought despite impaired language? So Isabel mentioned, asking somebody who had a stroke. OK, Great. So here's another case. This is a case of this guy here, Tom Lubbock, who died a few years ago from a brain tumor in his temporal lobe that destroyed most of his language, but it destroyed it gradually.

And this guy was a writer. He was an art critic for a major English paper. And as he started to lose language, he wrote about it, and he wrote about it very beautifully.

And he said, "my language to describe things in the world is very small, limited. My thoughts when I look at the world are vast, limitless, and normal. Same as they ever were. My experience of the world is not made less by lack of language but is essentially unchanged."

So that's a very powerful and surprising piece of writing. It's a little bit mysterious, because here's this guy writing beautifully and telling us his language is impaired. So his idea of language impairment may not be mine. I wish I could write that well.

Nonetheless, he's clearly reflecting on what is a very big loss of his previous language ability. And I'm sure it was very painstaking to write these sentences. And he's still telling us that, even though he's lost a lot of language, it has not changed his experience. So that's just one subjective impression.

So that argues against this extreme view that they're the same thing, but it leaves a lot of slop. Yes?

AUDIENCE: [INAUDIBLE] because he had a [INAUDIBLE] of speaking and learning about the word before.

NANCY
KANWISHER: Yes. A very important point. Absolutely. So this is a case of somebody who had a lesion in mid-life 40, 50, something like that. He had a whole lifetime of using language to learn and bootstrap all of cognition.

So absolutely we have to separate two different questions. Do you need language to become a normal, intelligent, functional human being? Do you need it throughout development? Or, once you've developed, do you still need it to think? And those are two very different questions.

And, in fact, absolutely you need language to develop. If you reflect for a moment on all the things you know, take a quick mental inventory, survey all the things you know, it's a lot of things, almost all of those you learned because somebody told you.

Most of what we know we learn from language. Maybe you read about it. But that's somebody telling you in a different way. So language is crucial for development of cognition and for learning. Absolutely. But now we're asking a different question of whether you need it, whether it's the same thing in adulthood.

So this guy is a little bit complicated, because he obviously still has a lot of language left. Let's consider cases of people who have essentially no language due to brain damage. So this is known as global aphasia. And Rosemary Varley in England has been studying a group of three people-- I think she's got a few more, but here are her three main ones-- who have global aphasia. And she's been studying them for a few years.

And, sorry, it doesn't show here at all. Sorry about this lousy projector. Shows on my screen. They're big, nasty, lesions taking up a lot of the left hemisphere and basically knocking out essentially all the language regions in these three individuals.

And here's their performance on a bunch of different language tasks. They have to look at a picture and name it. They have to understand reversible sentences. That's like boy kiss girl versus girl kiss boy. They need to know who did the kissing and who got kissed, right, and a whole bunch of questions like that. And they are at chance at every one of these.

So these are people-- not just people who can't speak. They're people who can't speak or understand language pretty much at all. So it's as close as we can get to a case of a person who has no language ability.

So can these people think? So Rosemary Varley has done paper after paper in which she finds clever ways to communicate tasks to these people to find out what kind of thinking they're capable of.

Here's one. You have to order this series of pictures. So look at it for a second and you can figure out that it goes basically from top to bottom. So can people with global aphasia do this task? Yes, they're perfect at it, no problem whatsoever.

Now you might dispute, is that cause and effect. Is it knowledge of sequences? Are they different? I don't know. But anyway, it's a pretty rich task here.

Here's another task. Look at these pictures and tell which of them are things you know and which of them are things you have never seen before that I drew. Takes a moment, but you can figure it out. Top three things are real things, and those three things are things I drew.

So we could ask, does a person with global aphasia know the difference. Basically, do you have to be able to name things to know the difference of what's a real thing and what's not?

Here's another task. Which of these is the plausible event? That's more complicated, because here we just need to know, is that a real thing that I know. Here we need to know, who's doing what to who, and does it make sense?

So it taps world knowledge, figuring out who's doing what to whom, which many people think is at the core of language. So how do people with global aphasia do? Perfectly at both of these things. Well, not perfectly, but the same as control subjects. Yeah, Carly?

AUDIENCE: I'm just confused. Like, how do you get the question across what they need to do?

NANCY I don't know exactly, but you do something like, for example. Do you ever play charades? Like that.

KANWISHER:

AUDIENCE: So, like, it's not exactly-- couldn't someone argue that there's actions that you're doing or some kind of form of language?

NANCY They're communication. They're not language. So when we say language, we really mean language. Not

KANWISHER: necessarily noises coming out of the mouth, because American Sign Language counts.

And I didn't have time to put that in this lecture, which is a damn shame, because it really does count in every way and is very interesting and uses similar neural structures and all that stuff. But language is different than communication. There's all kinds of ways of communicating. Yeah?

AUDIENCE: And how old are these patients again?

NANCY I don't know exactly, but it's almost always strokes. They're probably 40 to 60.

KANWISHER:

AUDIENCE: So it's definitely an adult. I mean, it's not a infant thing.

NANCY No, no, no. These are all people who had brain damage in midlife or later in life.

KANWISHER:

So that's pretty impressive, OK. So, basically, these people with global aphasia are able to do every single task that Rosemary Varley has tested them on. So I just showed you causality, nonverbal meaning.

Here's a cool one. Remember reorientation-- you should. May well be on the final exam. To remind you, I did a whole most of a lecture on this thing about reorientation.

Remember, rats and infants, if you hide food there and put them in this box, they later go 50-50 to the two corners, even though that wall should disambiguate which is the exactly correct corner. They should always go here. They have the knowledge that it's there, but they go 50-50.

And, remember, I said that Liz Spelke has this interesting argument that the key thing you need to be able to solve that task is language. Because, in fact, if you test adults and you tie up their language system, they behave like infants and rats.

But if you don't tie up their language system, they can do the task, which is pretty suggestive that language is a crux of the matter. However, the global aphasic do this task just fine.

So now, we have to go to [? Min ?] Young's hypothesis, which is that maybe the role of language in reorientation is learning about that whole spatial system during childhood, which the global aphasics could do, not maintaining the ability once you've gained it.

All right, they can do-- I won't give you all the data on this-- but they can do arithmetic tasks, logic tasks, algebra tasks. They appreciate music. They can think about what other people are thinking.

So everything that all that kind of high level abstract quintessentially human abilities that we are impressed with ourselves for being able to do, these people can do without language. So language and thought are not the same thing. You can still think in lots of different ways, even after you lose language.

On the other hand, as has already been brought up, global aphasics had language during development. So saying that you don't need it as an adult is not the same as saying you don't need it during development. You absolutely do need it during development, because it's a key way we learn about the world.

And for example, there are studies from Rebecca Sachs' lab showing that deaf kids who learn language later-- for example, if they're born not to deaf parents but to hearing parents who don't cotton on to the fact that it's important for them to learn ASL early, and hence they don't get language until later, those kids are not as good at understanding what other people are thinking, something that we usually learn about through language.

Further, even though I'm making a big deal about how you can think without language, I'm not saying that language is irrelevant to thinking. Every time I write a grant proposal, I think, oh god, I have all these ideas in my head and now I have to waste weeks and weeks and weeks, blah, blah, blah, putting them all down on paper to try to get money to fund my habit.

And then I get into like sentence 3 and I suddenly realize, oh, oh no, I haven't been thinking about this clearly at all. So this is my very informal introspection on the role of language in my own thinking. Like, even when I think there's a clear thought, the same thing happens when I go to prepare a lecture.

It's, like, oh yeah, I know this stuff. But I put together some slides, I'm like slide 2, no, I don't really know this stuff. So there is some role for language and thinking. And I'll give you one example here.

One of the many things that language can do is to make information more salient. So right now, close your eyes, everyone close your eyes. I mean it, I see if they're open.

While keeping your eyes closed, point south. You may not exactly know where south is, but make a good guess. Use your whole arm so everyone can see when they open their eyes.

Keep pointing, but now you can open your eyes and you can look around and see where everyone else is pointing. You guys are not bad, not bad, not bad. But we got some over here, we're a little turned around here.

Anyway, it's roughly over there. So, yeah, hang on, wait a second. Yes, right. Well, hang on. Yeah, right, it's over there.

So your vector average was closer to the true thing than a random vector, but not so hot. If your language forced you to keep track of this, you'd be better at it. And we know that from the case of the Pormpuraaw, these guys here, who live in Australia. It's Aboriginal people.

And they spend a lot of time going around in the remote outback of Australia, where they need to know where they are. And who is going where and when is really of the essence in their lives and in their social interactions.

So when they run into each other, they don't say, hi, how are you. Instead, they say, which way are you going. And a typical answer might be, "North northwest in the middle distance, how about you?"

They don't talk about things being left or right or behind them, reference frames that have to do with the person's own body, which are frankly really stupid reference frames. Because I can say this thing is to the left. And then I turn, and now it's not to the left anymore. Like, how stupid is that, right?

These guys have a much better system. They would rather say, oh, "You have a bug on your southeast leg." Right, OK.

So these guys, people who speak this language, they have to be aware of absolute compass directions all the time just to speak. And so they're oriented all the time, unlike us. And in that sense, their language makes salient certain kinds of information.

It's not that we can't think about direction. It's just that most of the time we're not aware because our language doesn't force us to think about it.

So interim summary. We've been asking this question of whether thought is separate from and possible without language. Before you guys take off, you wrote it on the board, this board right here? Awesome.

You guys need to tell me when there's time to take the quiz. So you're going to have seven minutes because there's seven questions. And so at 12:18 let me know and I will turn the board around. OK, 12:17 because it'll take me a minute to turn around.

All right, thank you. Take notes, tell me about that time.

So here's a question we've been engaging in, is thought separate from and possible without language? And the literature from neuroscience psych patients says, yes, absolutely, they're totally separate. Global aphasics have many forms of thought without language.

So given that, what would you predict from functional MRI? So if I told you, which is true, that these are the brain regions that are active during language tasks, for example, when you understand the meaning of a sentence, what would you predict? Should they be activated only by language, not by non-linguistic tasks? What do you think? Take a moment to think about it.

These are the regions that are engaged when you understand the meaning of a sentence. Would you expect them to be engaged based on what I've just told you when you do mental arithmetic, when you think about spatial orientations, when you appreciate music?

No, right. If they're separate, they're separate. They should go on in different brain regions. Everybody have that intuition? No, you don't have that intuition?

AUDIENCE: No. I mean, do you think about things in terms of words, even as a mental crutch, even if you didn't have to?

NANCY
KANWISHER: OK, fair enough, fair enough. So it doesn't nail this case. It could well be that you have separate systems for all those other things, but you still lean on the system. Not necessarily, but you use it sometimes.

In fact, there's evidence for that that we won't get to today. But the initial thought is, you don't need to activate it.

Well, here's a surprise. Up until recently, pretty much the whole brain imaging literature says that language overlaps with all of these things in the brain, that the activations overlap in the brain. They're all the same thing. That's been the received story for 20 years or so of brain imaging.

And that just does not fit with the patient literature. So we have a conundrum. Here are just a few examples. Stan Dehaene says, "arithmetic recruits networks involved in word association processes." People who study music say regions such as Broca's area and Wernicke's area, which have been considered specific to language, are also activated by certain aspects of music. Thus, the idea of language specificity has been called into question, and on and on. There's a million of these. I just put a few of them up there.

So what's going on. How are we going to resolve this contradiction? On the one hand, the patient literature suggests that language is separate from the rest of thought. And on the other hand, most of the neuroimaging literature says that if you look at those language regions, you find them activated in all these other kinds of things.

One hypothesis is David's, that they're activated but not essentially so. But there's another hypothesis. And that is that there's a methodological flaw with most of the prior research.

What is that methodological flaw? It's an inappropriate use of something called a group analysis. I've alluded to this a few times briefly, but let me do it for real now.

Let me first say, it's not that a group analysis with functional MRI is an evil thing that should never be done. They have uses. But particularly for the question of asking whether common regions of the brain are engaged in two different tasks, it is not a good method, for the following reason.

So, first, let's say, what is a group analysis. With functional MRI, it just means-- and, again, I'm going to be very sketchy with this because this is not an actual hands-on methods class. I'm just trying to get you to understand the gist of the methods. You take a bunch of scanned brains and you align them in a common space as best you can.

You can't do it perfectly because brains are anatomically different from one person to the next. But you do your best to align them as best you can. Then you do an analysis across those aligned brains. And you ask, what is consistent across this group of subjects.

That's a very useful question to ask. If we want to know overall what are the brain regions that are consistently activated when you understand language across this whole group of subjects, that's a good use of a group analysis. You'll find that picture I just showed you before with stuff going down the left temporal lobe, a bunch of left frontal lobe stuff. And that will be a very blurry picture of the regions that are most consistent across subjects.

Yes, [INAUDIBLE].

AUDIENCE: Do you line them anatomically as you new [INAUDIBLE] side to each other? Or do you line them functionally, so you can look at the scans in the functional [? region? ?]

NANCY
KANWISHER: So therein lies a universe of options. What I'm talking about now is a group analysis is aligning them anatomically. And that's where the problem comes in. And where we're going to go from that is you need to align them functionally.

If you just align them anatomically, then the following can happen. So you do a standard group analysis and you say, for example, let's do a language task and arithmetic task and a music task. And let's suppose you find this-- basically, Broca's area vicinity is activated in an overlapping fashion in all three. Each of those is based on an analysis of 12 or 20 subjects aligned as best we can.

So that's basically what the literature shows is lots of stuff like that. But here's the problem. You can get that result in a group analysis, even if the actual data looks like this in each individual subject. No overlap at all in any subject, but those regions are in slightly different locations. And so if you average across this, you get that.

Everybody see the problem? So it's not that it's a bad idea to do a group analysis. It's a nice, initial, blurry picture of the approximate consistent locations in the brain for a given task. The problem is when you say, oh, there's overlap, therefore they're the same thing, because you can get this result even if there's no overlap in any subject at all.

So the whole literature did this for 20 years and made all this talk about how language is on top of everything else in the brain. And for a long time I was sitting by the sidelines going, oh my god. And then, eventually, Ev Fedorenko came along and she knew about language. And I said, let's figure out, maybe they're right, maybe that's true, or maybe it's like this. Let's find out.

So how do we do that? What you do is exactly what [INAUDIBLE] mentioned a moment ago. You align them not anatomically but functionally. That's a whole reason to use functional reasons of interest. We've encountered this before when I was carrying on about why we do functional localizers with the fusiform face area. This is the same deal.

It's just that that insight started in the back of the head and hasn't reached the front of the head yet, or it's about here. So some people get it here. And the farther forward you go, the less people realize this is an issue, which is really ridiculous, because it gets more and more important as you go this way. Some stuff is actually aligned in the back and nothing is aligned in the front.

Anyway, so what do you do? You do just what we did with the FFA and all the other regions. One, in each subject individually, you identify those language regions.

You run some localizer. It's like, OK, I got this and that and that. And then once you identified them, you can ask, OK, does that region in that subject show activation for arithmetic. No, that's next door, right, et cetera.

Everybody got this? This is really important. I guess just because I'm obsessed with it. I honestly don't know if it's globally important or if it's just my personal obsession, but you need to know it for this course. We'll leave it at that. So this is standard in people who study vision and it's less standard in people who work in other domains. But they're slowly cottoning on.

So how do we identify language regions in each subject individually? There are lots of possible ways to do this. But here's the way I'm going to show you that's been used a bunch by Fedorenko and others.

So we start by saying, OK, let's find candidate brain regions that respond to language, which I told you, by language, I mean sentence understanding for present purposes. So if we want to look at sentence understanding, we've got to start with sentence understanding.

So if you look at the screen, you'll see some of the stimuli we use. So subject is lying in the scanner and they see that. And then we can either give them a task or not. And we'll talk about that in a second.

What are we going to compare it to? Well, there are lots and lots of different things we could compare it to that control for different things. But we started off with this, if you read this here.

So the idea is, it's visually similar. You can hear the sounds in your head. You can pronounce those things to yourself. But there's really no syntax and no meaning-- not perfect, but a first pass.

So when you do that, you get activations that look like this. Here are four different subjects. And you can see they're very systematic things. See these three blobs-- boom, boom, boom, boom, boom, boom-- in each subject, and a bunch of stuff in the temporal lobe like that in each subject. They're quite systematic but absolutely not identical.

All right, so, yeah, it's just what I did. So now what do you do next? Well, we just made this up, sentences versus non-word strings. Well, who says that's a good thing to do?

So the next thing you do is you've got to validate your localizer task to make sure it isn't just like trivial in some sense. So the first question is, is it reliable? So, here's session 1, three different subjects' activations. Well, just scan them again. There's a lot of talk about fancy statistics, blah, blah blah. Just scan them again.

Wow, look how similar these two little hot spots, this elongated one. I mean, it's remarkable, extremely reliable within a subject, and yet somewhat different across subjects, so check one, reliable.

More interestingly, does it generalize across task and presentation modality? So before we just had people reading sentences. And I keep saying, reading is not the native form of language.

So let's replicate that reading. And now we're adding a memory task. So at the end of each string, a little probe comes up and you have to say, was this word or a non-word in the previous thing-- sequence.

And let's compare that to just listening to the sentences. Wow, look how similar. So that tells us that we're not studying reading or speech. We're studying language after those things converge.

Those regions don't care if you saw a word or heard the word. They just care if you're representing the meaning of a sentence. Everybody with me why that's important? All right, check, check.

Does it generalize across languages? Suppose you're bilingual and speak two different languages. Here's two subjects who speak both English and Spanish. Wow, look how similar. So it's really language in general, not English or Spanish or a particular language.

Does it generalize across materials? So we could have reading sentences versus non-words that we've been talking about here with two different runs in one subject.

Are we going have subjects listening to speech versus degraded speech, like this? Here's the speech case.

[VIDEO PLAYBACK]

- During my days of house arrest, it felt as though I were no longer part of the real world.

NANCY

OK, versus this.

KANWISHER:

- [INAUDIBLE]

[END PLAYBACK]

NANCY

OK, so very degraded. You can't understand what's being said, but it has similar prosody and some similar

KANWISHER:

structure. And the point is, you get very similar activations with those very different kinds of contrasts.

So now we have really validated this thing. It checks out in all the ways it should. It doesn't care about modality. It does care about meaning. And it's highly reliable.

So now we can put it to use. Now we can ask, what does each of those regions do? All right, so to do that, in each participant then we find those regions with this localizer.

Now let me just step back a second. There's nothing magic about this localizer per se. When you want to study something, you use common sense. You try something, you validate it.

It may turn out later that of the thing that we thought we were identifying language with this localizer, it's got this other stuff. And then maybe you refine your localizer into something different. So it's not that this is the only possible way. It was just a sensible approach.

So you use this to find those regions. Here they are in these four subjects. And now, you can say, let's find. So you have to figure out some way to say that thing corresponds to that to that to that. And there's a whole bunch of math that was invented to do that.

You can basically see it with your eyeballs that those guys roughly correspond and those guys roughly correspond. The math is just a way to do that. And then once you've found that region, you can measure its response in a whole bunch of new conditions and ask what it does.

And in particular, so this is different from a group analysis where you don't identify those regions. You just choose regions anatomically

So if we just align them and said, OK, that's a region, well, we don't have much of the language stuff there, not much there, a lot there, not much there. OK, that's not great.

Then we take another one and we define this. This is a problem. No language stuff here, lots of language stuff there, none and lots. Not good.

Everybody see how that's a problem? OK, I guess I'm flogging this. We can move on now.

But the main problems with the group analysis are you might fail to detect neural activity that's actually there, because it doesn't align well enough across subjects and so it doesn't reach threshold. It's not consistent. But for present purposes, the more relevant problem is, you might fail to distinguish between two different functions, because they invariably coexist within that region or not.

So we're not doing that for present purposes. Instead, we're going to now go back to the conundrum of why do the patient studies suggest that language is distinct from the rest of thought, but the past functional MRI studies suggest that language overlaps with other functions in the brain. And we're going to consider the hypothesis that if you study individual brains and localize those regions individually in each subject, then the story might be different. And it is.

So here's the task that Fedorenko and I did a few years ago. We came up with seven different tasks. I won't bore you with all the details. It doesn't really matter. We just had lots of stuff, arithmetic, spatial working memory, various cognitive control tasks, working memory tests, all kinds of stuff, focusing on things that-- music, focusing on stuff that other people had said overlaps with language in the brain.

And so first thing is you've got to make sure those other tasks actually produce activations, because it's easy to make up a task and have it not do much, and then that's not very interesting. So, yes, each one of those tasks produce lots of activation. Look at all that red stuff. Looks like a bunch of pizzas.

So they produce good activations. Now the question is, do those activations overlap with the language regions. So let's consider two of them. This is basically Wernicke's area and Broca's area, two well-known language regions, identified individually in each subject and now averaging the response over all the conditions.

Here's a response when subjects read sentences and non-word strings, sentences and non-word strings. That's how we define those regions, but this is in data that wasn't actually used to define those regions. We held out some data and just cross-validated it.

Now the question is, how do those regions respond to all of these other things? They don't, pretty much at all. So notice what's happened here. The prior literature shows massive overlap between language and all these other things. In our data, when you identify those language regions in each subject individually and measure the magnitude of response in those other things, they don't respond.

So this shows stunning specificity of the language regions consistent with the picture that comes from the patient literature, from studies of brain damage. Language really is separate in the brain from all of these things. Everybody get that picture? And the reason the literature had it wrong is they were mushing brains together and blurring the hell out of their data and drawing wrong conclusions.

I'm speeding up because I don't want to run out of time. So we started with these questions here. Is language distinct from the rest of thought? I'm saying, yes, language may be necessary to learn to think. And it is indeed.

But the evidence from the neurological patients is pretty powerful. Global aphasics with pretty much no language can think in myriad, sophisticated ways. And when you do your functional MRI studies right, you find that the language regions in the brain, in fact, are not active during non-linguistic thinking.

Make sense? Questions? Wow, I finished on time.