[SQUEAKING]

[RUSTLING]

[CLICKING]

NANCY KANWISHER: Before we get on to the topic for today, I felt like last Monday's lecture was not my best. I don't know why. It's not that I didn't put time on it. I looked and I had the wrong lecture numbers on slides. There was all kind of chaos. I'm sorry about that. Sometimes you put in a lot of effort, and you still give a lecture that isn't all that clear.

So let me try to tell you what I thought were the main points. I started off by saying why it's really fundamentally important to be able to understand not just what people look like from the outside, but what we really care about people is what's going on the inside about their thoughts and their beliefs. And we are constantly making inferences about what people know, and believe, and want, and think.

And we do that all the time to understand why they're doing something and to predict what they'll do next. And so it's fundamentally important. It's the essence of being a human being in many ways. It's the essence of literature. And we do it all the time.

And a classic way that people have studied false beliefs is-- or thinking about other people's thoughts is with the false-belief task-- the Sally-Anne task that I described. And the reason people use the false-belief task rather than just a belief task is if the beliefs are true, you can answer what somebody will do based on the world, not based on their mind. And so to unconfound the two, we use a false belief that's different than the state of the world. So we can be sure we're asking people what will happen next based on what that person believes.

And through decades of use of the Sally-Anne false-belief task and variations of it, it's clear that there's a very distinctive developmental time course and ability to solve this problem. Five-year-olds solve it no problem. Three-year-olds systematically fail. And people with autism typically get-- pass this task late or not at all. High-functioning people with autism pass the task, just later, like 7, 8, 9 years, not five years old. OK?

So that's the kind of behavioral background evidence that there's something distinctive about thinking about other people's minds. Then we considered whether there's special brain mechanisms. And I argued that, yes, there are. There's a bunch of them. But the most impressively selective one is the TBJ shown up there.

And the evidence that it's specifically involved in thinking about other people's thoughts comes from the fact that that activation is a greater activation when you think about-- when you solve the false-belief type problems, when you think about another person's thoughts compared to when you think about a representation, a physical representation, like a photograph or a map, OK? So those are logically isomorphic problems. We have to always answer a question about a representation. It's just the representation is in somebody's head or it's a physical representation in the world. And in that difference, you get that region of the brain, OK?

So that's cool because it's a very nicely designed-- it's not quite a-- nothing's a minimal pair, but it's a minimal pair in some respects, right? It carves out whether the representation is mental or physical, but it doesn't solve everything. And a suite of other tasks have shown that that region is actually specific in a whole bunch of other respects.

It doesn't respond just whenever you think about a person. There are external properties. And most impressively, it does not even respond when you think about their visceral body sensations, like thirst, and hunger, and pain.

So the TBJ doesn't respond to just thinking about any mental states of another person. It's specifically thinking about their thoughts and beliefs. And that's pretty damn remarkable, right? It's, how abstract can you get? And yet, here's a specific brain region for that very abstract, very specific thing.

And the final little bit of evidence I showed you is that it also generalizes. It's not just about language because you can show people movies that have no words in them but that clearly show characters who must be thinking about each other's thoughts. And in those moments when the characters are thinking about each other's thoughts, that region turns on, more, for example, than when they're thinking about each other's pain, OK?

Yeah.

AUDIENCE: Did somebody look at what happens if other people think about me? Or if I think they thinking about me?

NANCY KANWISHER: You mean if you're thinking about other people thinking about you?

AUDIENCE: Yeah, exactly.

NANCY KANWISHER: Yeah, I would assume that that region would be engaged. I'm sure there are studies on that. Because you're thinking about their thoughts, right?

AUDIENCE: But more? Is it more interesting?

NANCY KANWISHER: Probably just because it's more salient, right? I actually, oddly, in this class, I talk about attention at the end, which is weird because attention is an issue with every study. But most of the brain regions we've talked about, surely, including this one, are modulable by how strongly you're attending to something. So if something's really salient, or important, or you're really paying a lot of attention to it, you're going to get more activation. And if it's more interesting to think about what other people think about me than to think about what other people think about each other, you'll find some modulation in here, I'm sure.

OK, and then finally, I talked about moral reasoning as a test case. It's not that the TPJ is selectively involved in moral reasoning. It's that many of the critical aspects of moral reasoning depend on what a person knew at the time. And so to use that information in moral reasoning, you need to pull that region in.

So I realized I wrote that question ambiguously on the quiz. I meant to ask, is the TPJ engaged specifically or only in moral reasoning, to which they answer-- the correct answer would be no. But I didn't put the "only" in there. And I decided it was ambiguous, so if you said "yes," you got the points.

Anyway, I gave several bits of evidence that, using the moral-reasoning case, that the TPJ is-- it's stronger evidence it's involved in thinking about other people's thoughts, first, that we showed that people with autism will have this difficulty in thinking about each other's thoughts. Even once they can pass the false-belief task, they put less emphasis, less weight on what the person knew at the time when evaluating the moral status of their actions, OK?

It's not that they make a mistake or that they're unable to morally reason. It's a pretty subtle thing. It's just a small difference in how much they weigh what another person knew at the time, OK?

That's also known as less forgiveness or less exoneration for accidental harm, right? You kick somebody by accident, and they go ouch, well, maybe you get a little bit of blame because you were a klutz and you should have thought or something. But it was an accident, so you should be exonerated. People with autism exonerate slightly less, right?

OK. So we then talked about the fact that if you zap the TBJ with TMS, stick a coil there, you do the same thing. You slightly reduce the weight people put on what the person knew at the time in their moral evaluation of the person's actions. And then I showed this bizarre fact that, as I think Gisella asked, well, shouldn't the TBJ be different in people with autism? Yes, absolutely, according to all of this, it should be. But the basic univariate measures-- how big it is, how selective it is, where it is-- do not find a difference in that region with that contrast in high-functioning people with autism.

That's surprising. But one possible answer to that is even though it's there and it's just as big and strong and all of that, it's-- that doesn't mean it doesn't represent different information. And I showed you an example that in typical people, you can decode from the TPJ whether the person is reading about another person's intentional harm or accidental harm. And you can't in people with ASD. OK? So that's my summary of last time.

And then all of that was focusing very particularly on the most fancy, quintessentially human aspect of social cognition, which is this business of representing each other's thoughts and beliefs. But I pointed out at the end that there are also lots and lots of other facets of social perception and social cognition, many of which have somewhat selective brain regions, lots of them other parts of the brain, and we just didn't have time to filtrate in to that, OK? Hopefully that was a little bit clearer than I was last time.

OK, so, so far, in this course, we've been focusing on all these bits of brain that seem to do very distinctive, often very selective things, OK? So the one we've just been talking about is that little guy right there. But we've talked about a lot of these things in here.

And the field of human-cognitive neuroscience has invested lots of effort to find these things and try to characterize what each of them does. And that's pretty cool, right? This is all stuff we didn't know 20 years ago, and it's nice, and it counts as real progress, I think. But it leaves lots of things woefully unanswered.

None of these regions can act alone, even though I've depicted them in a somewhat silly fashion, as nice little M&Ms on the brain. None of them act alone. None of them could act alone. They need information to process, so there has to be input to each region. They need to be able to tell other regions what they figured out or there's no point.

And probably, as they solve a problem, as they conduct their computations, they're probably interacting all the time with lots of other regions. So we desperately need to understand not just that this patch does faces. There's lots more we need to know. And one of the things we really need to know is, what is it connected to, and who is it interacting with? OK? OK. So that's what I just said.

And so that means looking at not just the cortex that we've been focusing on through this whole course, this dark matter that's on the surface of the brain up there. But today, we're going to do a figure-ground flip on the brain, and we're going to start paying attention to all that stuff that used to be background down there, all that white matter underneath, which is like a big heap of myelinated fibers that connect long-range regions of the brain to each other, OK?

So you might say, OK, just wires-- who cares about the wires? That was my attitude for a long time. I've gotten over it. We desperately need to know about the wires for all kinds of reasons.

So I'm going to go through a whole bunch of reasons. And there's a lot of little details, and I don't want you to panic. I just want to give you the gist of why this is worth paying attention to.

OK, so first of all, white matter makes up 45% of the human brain. So that alone tells you it's not like some trivial thing. It's a big part of your brain. This is all the more interesting because that's not true in other animals.

So I think white matter makes up a higher percent of the human brain than any other animal, or at least we're way up there. In mice, it's only 10%. And maybe that's a relatively uninteresting thing about scaling with brain size, but maybe it's something deeper about human brain-- what's special about the human brain and nobody knows. And here's a fun fact. If you took all the myelinated fibers in the human brain and you laid them out end to end, you could go around the world three times. So we've got lots of cableage sitting in here.

OK. So I briefly argued before that you simply cannot understand the cortex without understanding its connections of one region to another. It's just crazy to study one little patch of the brain and not know who it's talking to and where it gets its inputs from. And as I just said, the pressing need for that knowledge is heightened by the presence of this map, which we didn't use to have. Now that we have this map, it's all the more important and pressing to know what the connections of those regions are.

OK. So here's a nice quote making this point. This is Heidi Johansen-Berg and Matt Rushworth. They say, "Connectivity patterns define functional networks. The inputs to a brain region determine the information available to it, whereas its outputs dictate the influence that brain region can can have on other areas. Therefore, simply by knowing the pattern of inputs and outputs of a brain region, we can begin to make inferences about its likely functional specialization."

So I think that's a nice quote. It makes the point that it's not just that we need to know the connections, but the connections and the function are bound to be deeply enmeshed. One constrains the other. Yeah? OK.

Further, recall way back, which will quite possibly return on the final exam-- how do we define a cortical area? I gave you criteria for a cortical area. And one of the criteria was a distinctive pattern of connectivity, right? So it's part of the identifying properties of a cortical area is what it's connected to. And so that's another reason we should care.

A third reason is if we knew of a given cortical area what its long-range connections were to lots of other regions, that connectivity fingerprint-- remember we talked briefly about connectivity fingerprints a month or so ago? That fingerprint, the distinctive set of connections of that region-- you can think of it as a signature of not just how that region differs from other regions in that same individual brain but how we might find a homolog of that region in another species. And that would be a very interesting thing to do.

Wouldn't it be cool to know, is there a TPJ in macaques? Well, macaques can't solve an analog of the theory-of-mind task. Chimps-- we could debate a little bit. And narrow domains-- kind of sort of, a little bit, not really. Macaques-- no, OK?

So is there a homolog? Is there a corresponding region that-- maybe we took that region, and we adapted it and made it work better so we could do better things with it? And if so, what is it doing in macaques, right? I mean, I think that's just a totally cool question.

And in principle, one way to say what counts as "the same region" across species, which is kind of a weird question. They're different species, so what would "the same region" mean? One way to say what the same region means is to have a similar connectivity fingerprint, OK? So there are several studies that try to do that. I couldn't cram them into this lecture, but if you're interested in reading on it-- reading about it, shoot me an email. I'll send you some papers. OK.

I also mentioned that the specific set of connections of a cortical region, particularly its inputs, play an important role in development. Remember the rewired ferrets? If you redirect the input to what would have been primary auditory cortex in a ferret and you have that input come in from the eyes, you can get what would have been primary auditory cortex to become a lot like primary visual cortex.

So connectivity is important, not just in how a region functions and how we say what counts is the same across species but is probably also crucially important in the development of regions, OK? I also showed you evidence that the visual word-form area-- we can pick out exactly where it's going to land in an individual brain by the connectivity fingerprint of that region before kids learn to read, further evidence that connectivity determines later function.

OK. As if this is not enough, other reasons to care about white matter is that disruptions of white matter are at the root of many clinical disorders-- dyslexia, autism, developmental prosopagnosia, amusia, all of these things and others, for all of them, disruptions in long-range white-matter connections have been implicated as possibly playing an important role in the etiology of that disease.

Aging-- most definitely decline in white matter is prominent in aging. Sorry to say, there's a 10% decrease in white-matter fibers per decade starting at age 20. Use yours now while you have them. Let's change the topic. OK.

There's a lot of talk about how white-matter connections may change with experience, and learning, and plasticity. And that's a pretty patchy literature. And it's not a very impressive literature.

The classic thing you probably learned in 9.00, that when you juggle, you get changes in white-matter connections from juggling expertise. Maybe. Maybe not. There's some problems with a lot of that literature. So it's an interesting question, but it's not clear what the strong answers are.

And finally, I don't know about circuit design. Probably some of you do. But I gather that people who think a lot about circuit design-- one of the key features you need to take into account is wiring length. You want to keep wiring length short, right?

You have conduction delays. You have heating. You have space taken up in circuits. All of those things are bad in circuit design, and they're bad in brain design too.

So a lot of reason to think that a major factor in the design of brains, especially human brains, is minimizing wiring length because wiring length is very expensive metabolically. You've got to maintain ion gradients across cell membranes. It's expensive developmentally. These damn things need to figure out where to go, and if they don't go to the right place, you have a developmental disorder.

And so it's probably a real constraint on designs of brains. OK, so that was a whirlwind-- lots of reasons to care about white matter and connectivity. Oh, plus, at least in animal research, there's a whole suite of amazing new methods for looking at connectivity in animal brains. And [INAUDIBLE] can tell us more about that than I could because she's working in a lab that's right at the forefront of developing those methods. Someday, we're going to apply those methods to a human brain-- I can't wait-- and get the whole wiring diagram. OK.

Anyway, so what do we know about the connectivity of these regions? Well, you may be thinking, don't we already know all this stuff? After all, I showed you this diagram way back. And you've probably seen it every damn course you take in this department. It's in most textbooks in the field-- the whole wiring diagram of the visual system.

So don't we already know all this stuff? So what's the big deal? Well, here's the big deal. That's a macaque brain. And in macaque brain, you can get the actual answer to what is the actual structural connectivity of this patch of cortex to that patch of cortex.

There's a whole bunch of methods, but traditionally, you inject some dye here that's uptaken by neurons that travels along axons that goes here. You kill the animal, slice up the brain, and find that tracer over here. And then those two things are absolutely connected. That's the gold standard. And that's the basis of most of those studies, that method and variations thereof.

But we can't do that in human brains. And so we do not have anything like this information in human brains. Yes, David.

AUDIENCE: When was this done?

NANCY KANWISHER: Oh, that is a compilation. This was published, in, I think, 1991. But that was a compilation of heaps of studies that have been done before that. It was a big review article looking at all of this literature, where lots of classic neuroanatomy people would do these things where they would inject tracers, and slice up brains, and look in other places. And it was just a vast amount of literature that did that for many decades. It's sort of fallen out of favor, even though these things are, at least-- I don't know-- these are really crucial questions.

Now people use other methods to do that. You can use all kinds of optogenetic and other methods to map connectivity in animals. Yeah.

AUDIENCE: I have a question.

[INAUDIBLE].

NANCY KANWISHER: In here? Oh, that's a good question. Oh, it's probably dorsal and ventral pathways. Let me see here. Yeah.

Yeah, the red ones-- this is another thing I didn't even really mention, probably let alone give a short [INAUDIBLE]. That was lame. But anyway, the visual-- high levels of the visual system-- we focused on the ventral visual pathway coming down the bottom of the temporal lobe. But there's a whole other visual pathway that goes up into the parietal lobe.

Did I talk about that a little bit-- reaching and grasping? No, I didn't. Lame, lame, lame, lame. Anyway, a major part of the field I didn't get to.

Anyway, it's a whole other part of the visual system that seems to be more involved in visually guided action. And they're actually very interconnected, but they're trying to emphasize that the dorsal pathway is at least somewhat separable in monkeys. But my point is this is monkeys where they have the gold-standard methods, and they can actually discover the real connectivity.

Sadly, we can't do those things in humans. And in humans, we have only three methods, and none of them are very good. So we'll talk about them today anyway because this is such an important question, but the bottom line is-- this drives me out of my mind-- we basically don't know the connectivity of any of those regions for sure in human brains.

And somebody's got to solve that. Maybe one of you will go invent a method that works in humans that helps us solve that problem because it's actually, I think, really paralyzing to our field. So I'll tell you what we do know, which isn't much. But, you know, beggars can't be choosers.

OK, the first method has been around for a few years, and that's gross dissection. And I mean gross, like that kind of gross-- so only good for post-mortem brains, but it's really quite amazing. This is a bottom view of the brain back and front.

This is actually a physically dissected brain. Like, it takes a real serious neuroanatomist and lots of fancy methods-- I mean, not fancy methods but lots of careful, precise teasing apart of bits of brain. And you can actually see these big fibers coming up here.

So if this is the back of the brain and we're looking up like this, what do you think those fibers are connecting right there? Big fiber bundle coming from deep down in the brain up to right in there.

AUDIENCE: Is that thalamus to [INAUDIBLE]?

NANCY KANWISHER: Bingo. Exactly. OK, so that's the LGN right there. And this is called the optic radiation. It's this huge cable of fibers that come up. OK, first, here's the optic tract.

Actually, I forget. That's not-- I think this is the optic tract that's been snipped there. Then it comes up in here, makes a stop in the LGN, and then this big batch of fibers comes up right there to primary visual cortex, OK? Everybody got that? OK, so you can actually see it in dissecting a dead brain. OK, that's pretty cool.

But what if we don't want to wait for people to die? Often, we want to ask questions about a person right now in their brain. Do they have this disorder? Are they at risk of that disorder? What is their connectivity?

So for that, we have two methods, and I'll talk about these two methods in the rest of the lecture. The first one is diffusion imaging. OK, so I talked about this briefly before. But let me remind you of what the basic principles are.

So here is a picture of the optic nerve with a bunch of axons oriented like that. It's a big cable with a whole bunch of little fibers in there. And the basic kind of biophysics is that water wants to diffuse more along this length, following the orientation of the fibers, than it wants to diffuse this way, OK?

And diffusion imaging-- I'm not explaining any of the physics, but just take it for my word that what diffusion imaging does is give you a picture of the direction of water diffusion, OK? So you get a picture of a piece of brain, and it'll show you, for example, that right in there, all the fibers-- well, the water is diffusing this way. And over here, the water's diffusing that way, OK? And that's just what you see in a diffusion image, OK?

And so the inference people make is if you have all those parallel lines telling you there's lots of diffusion like this, there's probably a big fiber bundle going like that-- and there is. That would be the corpus callosum. OK?

All right. So this method works great for finding the big fiber bundles. OK? I'm going to dis diffusion imaging in a bunch of ways, but it is great for finding the big fiber bundles because in those big fiber bundles, axons are very parallel. There's a whole bunch of them, and you can really see it. OK.

And so people have been using this for over a decade to find some of the major fiber bundles in the brain. So you may have heard of the arcuate fasciculus that basically connects language regions in the temporal lobe up to Broca's area in the frontal lobe, OK? It's a big bunch of fibers that go-- I guess in me, they go like this, boom, right? [INAUDIBLE] And you can see those guys with-- this is a distant reconstruction, but you can see those with diffusion imaging.

Another one, the goes from the front of the temporal lobe up to the frontal lobe. You don't need to memorize these. I don't care about that. I just want you to get the idea of what you can see. Yeah, question?

AUDIENCE: So these are discoverable without the person having to do anything [INAUDIBLE]?

NANCY KANWISHER: Yes. Yes. These are anatomical images. So in diffusion imaging, you don't do anything. You can sleep, actually. That's ideal because it's long and boring.

Actually, it's not boring, but the scanner shakes like hell in a diffusion-imaging scan. It's pretty wild. We could charge admission for it. I don't know. I find it quite wild.

Anyway, this is the inferior longitudinal fasciculus. As it goes down the temporal lobe. So when we talk about the ventral-visual pathway-- face areas, place areas, all that stuff-- this is the big fiber highway that sits right on top of that whole chunk of gray matter that does all the processing. And it's a big pile of fibers that go straight down the temporal lobe, OK? OK.

And so here's more recent data. This is from Anastasia Yendiki over at MHG Charlestown over there. And she's developed this lovely piece of software that enables you to take diffusion images and identify, based on an atlas she's put together, 18 of the major fiber tracts in the human brain-- so nine per hemisphere.

And this is just showing you some of the big ones. This is the inferior longitudinal fasciculus I just showed you and so forth. Yeah.

AUDIENCE: So how does this compare to just a postmortem dissection? Could you see the--

NANCY KANWISHER: It's a good question. It's a good question. I don't exactly know. It wouldn't be easy. That thing I showed you works because you take the stuff off the top, and it's just kind of sitting there. But then you would have taken a lot of other stuff out, and you wouldn't be able to see that other stuff you had to take out. You know what I mean?

So here, you can surf through and pick out any of these. So it's definitely going to be better. But which of these you can see with postmortem dissection I'm not sure-- some of them, some of the bigger ones but probably not all of them.

OK, so here are some of the major tracts. OK. But you can do a little bit more than just find them with diffusion imaging. You can also characterize them a little bit. And this is a whole universe. There's people who spend their lives with all kinds of fancy measures, and I'm just going to tell you about the most common one.

So recall that the whole deal with diffusion imaging is it's looking for orientations of maximum water diffusion. So some parts of the brain have a systematic set of directions of water diffusion, and other ones don't, right? Inside a ventricle, the water can go any which way. There's nothing determining which way it goes.

So this is called isotropic because you have diffusion going equally-- in equal amounts in all directions and anisotropic, right? So it goes systematically more in one direction, one axis than others, OK? Diffusion can't see the axial direction, like left or right versus to left. It can just see that this axis is more prominent than this one, or this one, or this one, OK? So you're not actually seeing it move in a systematic direction.

OK, so that's the basic signal. So what you can do is in a little patch of brain, you can ask not just, what direction has maximum diffusion, which is what I've been talking about so far. You can say, how much more does it go in that maximum direction than any of the others? Is it more like this or more like that? And you can imagine a whole spectrum in-between.

OK, so you're just asking, how oriented is it? Is it totally oriented or just partially? And that measure is called Fractional Anisotropy, or FA. It's prominent enough in the field you should just learn this phrase. And you read any articles, especially any clinical articles, this is the first thing you'll see in any clinical papers that use diffusion imaging.

OK, and so fractional anisotropy-- there's a bunch of fancy definitions. And we don't care. We just want the idea. It's just, to what degree is that little patch of brain in this little part of a tract that you've identified more like this or more like that, OK? Is it anisotropic or isotropic? OK?

And so this has been used a lot to try to ask about the nature of fiber tracts in different groups-- young versus old, different clinical groups, autism versus typical, schizophrenia, you name it. Experience-- you train people up on a task-- do you change the FA, the Fractional Anisotropy, of some particular tract. OK?

So it's just a characteristic of tract is, how oriented is it along the way? Is it super clean, totally oriented? Or does it have some isotropy mixed in? OK?

All right, so this is all over the literature. Let me give you one cool example from Gabrielli lab that came out recently. So they identified the arcuate fasciculus that I showed you before, going from the-- basically, Wernicke's area curving around up to Broca's area up there so that you can identify it anatomically in each subject individually, OK? Now you've got it. You've identified which voxels are part of the arcuate fasciculus.

And then what they want to ask-- what they wanted to ask is, is the integrity, or the characteristics of the arcuate fasciculus-- is that important for language-- for dyslexia? So they measured the FA along this tract. How oriented is it all the way along here? And then you get some kind of average.

And they measured that in a bunch of kids with dyslexia and in a bunch of kids with no reading disability who are matched in other dimensions of non-verbal cognitive ability, OK? And what they found is the fractional anisotropy was higher in the typical kids than the kids with reading disability, with dyslexia, OK? And from that, they implicated that this connection may play some role in dyslexia. OK?

It's not totally obvious because one would think this region is connected to that region-- those are languagey regions, right? It's not visual regions. You think of dyslexia as a problem seeing the letters and which ones are oriented which way. And this suggests that, at least, higher-level connectivity between language regions may be implicated, OK? Yeah.

AUDIENCE: So this just talks about the architecture and not the per unit information being [INAUDIBLE]?

NANCY KANWISHER: That's right.

AUDIENCE: There's no information [INAUDIBLE].

NANCY KANWISHER: That's right. It's just saying, where are the wires, and how organized are the wires? Period. Yeah.

AUDIENCE: So they won't flow through some technique as-- if I can mess with it, like I'm measuring it, if there's some means to miss the [INAUDIBLE] diffusion, will that have any effect?

NANCY KANWISHER: I didn't quite get that. Say it again.

AUDIENCE: So there's this water diffusion that's happening that I'm going to measure. But what if I have some way to intervene and change the rate of diffusion some such?

NANCY KANWISHER: I don't know how you'd do that. I mean, it's a pretty basic physical property-- diffusion of water and how it's constrained by lipids, right?

AUDIENCE: But that shouldn't affect, like the function of information flow [INAUDIBLE]?

NANCY KANWISHER: Ooh. I have no idea. I mean, that's a biophysical question I don't know about. But notice, this is a pretty distant proxy. You're mostly looking at water between axons, not even within them. And so it's just a proxy for, how well can we see those fibers and how they're oriented, OK?

Yeah, it's pretty removed from actual signals going along the wires. Anyway, so everybody get the sense that-- you know, it's one little finding. But it implicates something about that tract in dyslexia.

OK, so that's interesting. But first, it's just correlational. Lots of things are just correlational. Most of the stuff in this class is just correlational. Same is true here.

But a little more seriously, it's not totally clear what fractional anisotropy means, OK? So there's a real tradition of treating high-fractional anisotropy as if that's good. After all, we're in a fiber bundle. Shouldn't all the axons be oriented nicely in there and not all scrambled?

Surely, oriented is good, and scrambled is bad. Well, maybe, but sometimes, fibers cross a fiber bundle. So you can have a fiber bundle like this with other fibers crossing it.

And so when that happens, maybe that's good. And so people use FA as a proxy for good fiber. People say "fiber integrity" even. But there's a whole question about what exactly it means. People will spend their lives looking at the biophysics of fibers and all the different things that fractional anisotropy and the other measures might mean. And it's actually pretty complicated and unresolved.

Another challenge with fractional anisotropy is it's extremely vulnerable to artifacts. All of diffusion imaging is extremely vulnerable to artifacts, OK? And I'm going to give you an example of a study we did a few years ago.

So I was, for a while, trying to work on autism. I've, more or less, given up because it's, as far as I can tell, impossible. But back while I was still trying, we scanned a whole bunch of kids with and without autism with diffusion imaging, OK?

And at the time, there were about 50 published papers, almost all of which said one of the things you find with autism is that there's an underdevelopment of long-range connectivity and an overdevelopment of short-range connectivity. And so then people would free associate with all kinds of speculations about, OK, this explains aspects of the autism phenotype. They can't put different ideas together because their connections across the brain aren't as good. And they're obsessed with little details because they have too many local connections and all kinds of-- suggestive, but very, very fuzzy ideas like that. So 50 papers pretty much all found underdeveloped lower fractional anisotropy and long-range connections, long-range tracts in autism, a very established finding

So we went in not to raise hell but just to kind of replicate some of those basic findings while studying some other things. And, in fact, when we did what everyone else does-- that is standard analysis-- you collect your diffusion-imaging data, and you eyeball it loosely, and if it really looks terribly tainted with artifact, you throw that subject out. And, otherwise, you keep it, and you analyze your data. And you look at, here are the 18 fiber tracts that I showed you before. And you ask, which of those have higher or lower fractional anisotropy in kids with autism compared to typical kids?

And the basic finding-- we replicated the usual finding. And that is, overall-- this is column A-- most of those tracts showed lower fractional anisotropy in the kids with autism than the typical kids, OK? Many of those differences were individually significant in individual tracts. Those are the ones with the asterisks.

OK. So that's the standard finding in the literature, and we replicated it. However, we noticed that a lot of the data really seemed suspect. And we started measuring the amount of head motion between the kids with autism and the kids without. And guess what. Kids with autism move in the scanner more than kids without. And guess what. Diffusion imaging and fractional anisotropy in particular are highly influenced by head motion.

So then we said, OK, let's get a little more careful. And so we did a more stringent analysis. And we looked at the kids. We had quite a few of them in each group. And we took the subset of kids who we could match for head motion, OK?

So now we've got the kids with autism and the typical kids, but we've now got the subset we have to choose to match for head motion, OK? It usually means the typical kids who move a little more and the autistic kids who had slightly less head motion. That's what you need to do to match them.

And when we do that, the usual pattern disappears. Now there's only a single tract that shows lower fractional anisotropy in the kids with autism than the typical kids, OK? The inferior longitudinal fasciculus. So that's worrying.

But then further, we thought, OK since many of those kids we had, especially the typical kids-- many of them we had scanned twice. So we thought, OK, let's really make this case. This is clearly a problem in the field. In fact, it's a broader problem. Pretty much any comparison across age groups or clinical groups-- one group moves more than the other group. Uh-oh.

What about the entire literature? Hundreds, probably thousands of published papers, essentially, none of which pay attention to this-- this is 2014. People have cleaned up their act since, but up until 2014, almost none of them paid any attention to this whopping problem.

So we figured we better make this point salient because there's a lot of money and time being wasted publishing garbage. And we want to make the point saliently. So we took the typical kids who we had scanned twice, OK? And sometimes a kid will-- the same kid will move more in one session than another session.

So we said, OK, let's compare the very same kids on the session where they move more than the session where they moved less. And you know what? We replicated the autism phenotype. Those were typical kids, not autistic kids.

The point is head motion alone will reduce fractional anisotropy and will look a whole lot like a clinical disorder. And so every time you see that a clinical disorder is marked by some anatomical difference, your first thought should be, how carefully did they deal with head motion and other artifacts that are going to differ between groups? OK?

So I say that not to dis the entire literature but just to alert you that these things can really matter. The paper from Gabrielli Lab that I just described-- I looked, of course, before I presented it in here, and they cited us, and they used our methods for matching head motion-- good for them. So these things are changing, and I think the field will start cleaning up its act. But it's amazing it took this long.

OK. All right, so finding fiber tracts and characterizing them with fractional anisotropy are nice, but, really, what we want to know is what's connected to what, OK? Which of these things are connected to each other? Which other brain regions are they connected to?

And so to find that out, we need to not just study white matter itself and the tracts, but we have to get out of the tracts and into the gray matter. So we need to start in a patch of gray matter and figure out where we can go by following those axons, OK? And so the method for doing that is called tractography. So there are many versions of this. I showed briefly these pictures before. I'm sure you've seen these.

The simplest idea of what you do, leaving out all the details, is you start in some gray-matter region, some voxel in the gray matter. You want to know what it's connected to. You just follow those little orientations, and you can-- you see where you can go. OK?

And so that's basically how you make this diagram. OK? You're just following those. You start over here you follow the orientation-- go do, do, do, do, do, do, do, right? OK? I mean, you do that in a computer, right? An algorithm does that, follows along-- ch, ch, OK?

OK, so that's called tractography. And the idea's awesome-- how great to be able to see what's connected to what. And there are many, many thousands of papers that do this for good reason. We need to know what's connected to what.

This is our currently best method for looking at the structural connectivity of different gray-matter regions to each other. And so you can ask, for example, OK, let's put a seed in the fusiform face area and see where it goes. Wouldn't that be cool? Right. Wouldn't it be cool?

Unfortunately, it doesn't work. So I have to tell you that I don't know if I'm the best person to report to this because I'm not-- I've only been trying to do this for a few years. But, I've been collaborating with the best people in the world over there at MGH Charlestown who are working closely with us. And we can't get this thing to work worth a damn.

And so now I'm actually confused whether the entire literature is garbage. I don't think it's entirely garbage. But I think it's full of overoptimistic evaluations of what you can tell from tractography because in our hands, we started with reality checks, put a seed in the lateral geniculate nucleus. Let's make damn sure we can get up to V1. Well, you can get up to V1, but you can get up to V2, and V3, and V4, as well, which are all wrong, right? LGN only goes to V1.

Worse, you stick a seed next door in the medial geniculate nucleus, which is the part of the thalamus that goes up to auditory cortex, you also end up in V1. Wrong. Wrong. Wrong. Wrong. There's not very many anatomical connections in the human brain where we actually know the right answer where we can do these reality checks, but of the ones we know that we've tried, it doesn't work.

And we're using the best diffusion-imaging scanner in the world. It's right over there. So maybe I'm doing everything wrong. But at the very least, I think there are a lot of problems with this method.

This is not just me worrying about this. And many people have been worrying about this for the whole 15, 20-year life of diffusion tractography. And some of the challenges are, like, famous.

So to follow those little orientations, you need to-- you can see, like there'd be lots of places where, OK, there's a bunch of different ways you could go. It's ill posed, right? So people use heuristics to constrain those solutions. And those heuristics are based on assumptions about how fibers bend in the brain, namely that they don't make really sharp angles, right? That's reasonable. Most of the time they don't, but sometimes they do.

And in particular, when you're going from white matter to cortex, often, you make a very sharp turn. And so it's very, very difficult to figure out how to get from a given gray-matter patch into the underlying white matter exactly what the connectivity is. So that's one problem.

Another famous problem with tractography is called the crossing-fiber problem. So imagine a bunch of axons somewhere in the brain that cross like this versus imagine a bunch of fibers in the brain that come up to each other and then go apart, OK? Everybody get this? The connectivity's totally different here-- no way you're ever going to distinguish those with diffusion tractography. So people try to get higher and higher resolution to see down those individual things, but they're not there. Yeah.

AUDIENCE: Why would something like that happen?

NANCY KANWISHER: Yeah, why would they do that? Weird stuff happens in the brain. So it's not incredibly common, but it's not unheard of.

Yeah, remember, the brain wasn't dissolved-- designed now optimally to solve all the problems it needs to solve with the optimal solution from scratch. It evolved gradually over time. And so there are all kinds of weird things that are workarounds for pre-existing decisions that evolution made earlier. And so both brain and body have lots of bizarre attributes that aren't how you would design it from scratch. They're just the fix that evolution made at that point given what had already been fixed. And so there's weird stuff like that.

Anyway, I mention this to say I'm more negative about diffusion tractography than probably anyone else because I've spent a lot of the last two years trying to do it, and it's big bust, and I'm cranky. So it's probably not as bad as I'm laying out. Plenty of people do it. They get some kind of answers out of it, but it's problematic at least.

My best guess is that it is OK for fingerprints. If you're asking, OK, here's some patch of brain. How much does it connect to, say, these 85 other regions? And is that different than the fingerprint for this region?

That's probably OK because a lot of those individual solutions might be wrong, and there's still enough left over to see a kind of difference. So I feel like you can-- conductivity fingerprints are probably worth doing. But, actually, just answering the question of, is there a structural connection from A to B-- I don't know. I can't get it to work. OK?

OK, so I think I did all this before just to-- conductivity fingerprints-- do you remember this? You start with one place, and you measure how well you can get from each location to each of these other ones. And I showed you before that the work of Zeynep Saygin and a bunch of other people has shown that you can-- actually, in an adult, you can predict where that adult's fusiform face area is just from their diffusion tractography data alone because it has a distinctive connectivity fingerprint, OK?

I don't want to go through all that again. Do you guys remember that, more or less, the gist? OK, so that just tells you that there is this systematic mapping between the connectivity of a region and its function. And connectivity fingerprints, despite all these problems I've been carrying on about, have enough signal left in there to predict the function of a region and maybe to say something about homologies across species. OK, blah, blah, blah. Right.

OK. So where do we get? You can find the major fiber bundles with diffusion imaging. That's worthwhile. You can characterize fractional anisotropy. I don't really know what it means, but it means something. And you can find very approximate connectivity fingerprints good enough the predict function.

OK, so that's worthwhile. But actual structural connections of one particular cortical area-- not very good. At best, it's a weak signal. So that's a drag.

So let's consider the other method people have used to try to work these things out. And that's resting functional correlations. So let me describe where this story starts. This story starts with a paper in 1995 by Biswal, and this is the figure from his paper.

So, first he had people move-- they had-- he had people in the scanner doing finger-tapping. So they're lying in the scanner. He's scanning their brains while they tap both fingers or not, or tap both fingers or not. And you get these-- it's hard to see-- these two little bits of motor cortex corresponding to the finger-motor region. OK, no surprise there. We're just mapping a little bit of motor cortex.

But then he does something cool. He looks at the time course over that experiment in one of those motor regions-- in one hemisphere, and he looks at the time course in the other hemisphere when the subject is at rest, not doing anything, OK? Sorry, I left this out. You scan them doing this, and then you scan them just lying there going, dum-de dum-de dum, or whatever you do when you don't have a task, OK?

And he finds that these very far-apart regions at rest, when the subject is not tapping their fingers, are extremely correlated. So that's very not obvious. These things are centimeters apart.

The subject isn't doing anything in particular. You're not telling them what to do. And they certainly aren't tapping their fingers. So why are these two bits of finger-motor cortex going up and down in lockstep like that?

Well, nobody knows, actually. I mean, this was however-- 20-some years ago, right? Still, nobody really knows why those damn things are going up in lockstep like that. But it's systematic, it's tantalizing, it makes you want to play more, and many people have.

OK, so I would say we still don't know exactly why those things are going up and down together. But the pattern of brain regions that go up and down together has proven to be a whole fascinating window into the brain. OK, so that's our next topic here.

OK, so here's another depiction of more exactly what you do. OK, so step 1-- you find a seed region in here, in left somatosensory motor cortex, OK? So that's that region there. You get its time course, OK? Sorry, at rest.

You find that region, and then you scan the person while they're just told to do nothing in particular. You get the time course averaged over all those voxels at rest. There it is, OK? Now you take that time course. And you correlate it with the time course of every other voxel in the brain. And you say, show me all the voxels that are correlated with this region at rest. And you get this-- lots of systematic brain regions that are highly correlated at rest with that region you started with. Everybody get what we just did?

OK, totally non obvious. Well, you might say, OK, fine. This is finger-motor cortex. This is the other one. That's what I showed you from Biswal before.

But why this thing? Why this thing way down deep in the brain? Why that thing down in the cerebellum miles away in the brain? Why are they all in cahoots with each other. I like using "in cahoots" when talking about correlations because nobody knows what the correlations mean. So "in cahoots" is as technical as I think we should get. Yeah?

AUDIENCE: --correlations without a time shift. Just [INAUDIBLE].

NANCY KANWISHER: Good question. Good question. Wouldn't we love to know about the time shift? But here's the problem. There should be a time shift because it takes a while to conduct down axons from here to here, probably a few milliseconds. But a few milliseconds we are never going to see with functional MRI. So, surely, there is a time shift, but this method can't exploit it, OK? Yeah, I'll just leave it at that.

OK, but does everybody get what this map is? We've just chosen a seed region, a starting point just for the hell of it. And we've asked, what other bits of the brain are correlated with that region at rest? It's a pretty weird thing to do. And you wouldn't do it if you didn't find systematically replicable answers that are repeatable across subjects.

And when that happens, you go, OK, I don't know what this means, but it's pretty systematic. Let's keep following the thread. OK? Question?

AUDIENCE: What do you mean by correlated [INAUDIBLE]?

NANCY KANWISHER: OK, so let's do this again. You scan people moving their fingers. You find little-finger region here. Now you scan the same person just in the scanner. You say, I'm going to scan you for five minutes. Just close your eyes and don't do anything in particular. You lie there. I scan your brain.

Now I take that region, which I found before. And I take the time course of that region while you were just lying in the scanner doing nothing, and I get some randomish-looking thing like this. Now I take that time course, and I say, let's see if there are any other voxels in your brain that were correlated when you were lying there with that time course. And I color them in, and there are lots of them, even regions that are far away. OK? It's really not obvious. You wouldn't have predicted this would happen, yeah? Yeah?

AUDIENCE: So if you look at the brain just as some underlying resting rhythm and like just all regions of the brain just have some resting rhythm, wouldn't it be just always be [INAUDIBLE]?

NANCY KANWISHER: Yeah. OK, so there's been a whole suite of speculations of exactly that kind. Are there endogenous rhythms that are characteristic of particular brain regions and so those things go together? Maybe, but so far, that doesn't seem to be the main answer. For a long time, people thought, OK, is it just blood-flow supply? Maybe the blood-flow supply to the brain branches and feeds those regions, and that somehow regulates the bold response in those regions.

There have been many accounts like this, and none of those seem to really capture it. It really seems like, probably, those neurons are firing in sync with each other, right? Yeah, question, Nava?

AUDIENCE: Yes, before a question in direction-- if you have a time delay, I guess the question was because if you would have a time delay, you could see what's further right. Can you, instead of the time delay, see if you measure it in one of the regions that seems to be-- seems to be correlated-- if you measure from one of those, if you could estimate the distance based on how strongly they correlate [INAUDIBLE]?

NANCY KANWISHER: Oh, with the thought of maybe there's not a time delay, but maybe you lose some of your correlation with distance. You could. But just looking at this, these guys are pretty far apart. So it's certainly not that there's just things are-- that nearby things have a similar.

AUDIENCE: No, but I mean, if you would say those are five different regions, did you measure from region 1?

NANCY KANWISHER: Yeah, yeah, I got you. Yeah. Yeah. Yeah, you could. I think that's not going to work because there are big, big correlations that people find between very distant regions. I'll show you more. Yeah. I mean, you could try that, and I'm sure people have done that. I can't tell you exactly where.

Actually, they do it as part of their-- one of the common ways you normalize your data is to take this and normalize it for distance from the seed region, which would be a way to build that factor in. And once you do that, you still get lots of stuff.

AUDIENCE: They take the distance and the image space for that, right?

NANCY KANWISHER: You can do it different ways. There's an algorithm that somebody at MGH wrote that is distanced by, most likely, a white-matter path or as the crow flies, not that the crow can fly straight through the brain, but you see what I mean. Yeah. Sorry, go ahead.

AUDIENCE: Is the result different in if you measure the correlation when they're doing the finger-tapping action versus [INAUDIBLE]?

NANCY KANWISHER: Yeah. OK, so this is a really important question, and it's a whole part of this field that I'm leaving out of this lecture because I'm sort of suspicious of it. But your question is a good one. So you're saying, would they be correlated while you're finger-tapping?

Well, certainly, if we did the paradigm while they're tapping both fingers, they're going to be correlated because we built the correlation into the task. We said, while you're doing this, do that. And so they will surely be correlated. And so there's a whole enterprise where people try to factor out those things and ask, even after you account for the activation of the task, are there changes in these patterns of correlation with the task you're doing?

And that's called PPI for physiological interactions. And lots and lots of people do it-- hundreds, thousands of papers. It's probably pretty respectable, but it drives me nuts because I don't feel like there's any way you could know that you're fully accounting for the task. And so I think those correlations may be largely reflecting regions commonly activated by the task, and that's why I didn't put that in this lecture.

But surely, task will also produce correlations, right? Let me just put it another way. If I flash up a bunch of faces versus-- it's like faces versus nothing-- and then we look at the correlations during that period, well, you'll find correlations between V1 and the FFA because when their face is on, both V1 and the FFA turn on. And when there aren't, they both turn off. That's just a task response, right?

So to be able to look at how these endogenous correlations are affected by task, we would have to be certain we could siphon off the entire task effect so that we could look at just the residual. And I don't think any of our analysis are good enough to siphon off an entire task effect. And that's why I just don't go there with PPI, even though everyone else does. If you didn't follow that, it doesn't matter. I'm just trying to give you an answer.

I'm going to take just questions of clarification now because there's a couple of things I really want to get to and I'm running out of time. OK. But everybody should understand this-- an activation map that's made by asking, which brain regions are correlated at rest with a given region I choose a seed region I choose? OK.

OK, important caveat-- even though people call this "resting functional connectivity," we will not be using that phrase in this class because we do not know that it's connectivity in the structural sense. It's just a correlation, OK? And I'll say more about that later. But if you read about resting functional connectivity, it's the same thing. It's just, I think, people are making a mistake using that word.

OK, so let me get this idea across here. You may have heard of the default-mode network. There's heaps of papers on this. It's a thing. There's a lot of discussion of it. And it's bizarre. It has arisen from two independent findings,

OK? So let's do these findings one at a time. The first one is people started noticing around 15 years ago that across lots of different kinds of tasks, if you looked at not the intended direction, like, say, reading sentences versus staring at a dot, or doing a demanding working-memory task versus a really passive-viewing task, anything where there's a really engaging task versus an easy task, you would find a bunch of regions that were activated in the reverse contrast, regions that are more activated when you're doing less mental activity, typically, regions that are active when you're just lying there at rest compared to doing something difficult.

And so originally, people were like, what's up with that? How can that be? It seemed paradoxical, impossible. But, in fact, it's not impossible, right?

Suppose I had you do a bunch of mental-arithmetic tasks, and they're pretty demanding. And I compared that to just having you lie there in the scanner doing nothing. It's like, OK, do mental arithmetic for 20 seconds, rest 20 seconds, mental arithmetic for 20 seconds, rest for 20 seconds. Now imagine we find parts of your brain, systematic ones, that are more engaged at rest. What might that mean? David.

AUDIENCE: That part of the brain could be, I think, like daydreaming.

NANCY KANWISHER: Daydreaming, exactly. Yeah. You can't turn your brain off. You don't turn your brain off at rest. You daydream. Absolutely. What else? What are the things-- some of the things you do when you daydream? What are the typical contents of daydreaming?

I guess it depends who it is and what you're daydreaming about, but there are very systematic things people do. They recall episodic memories. It's like, oh, yeah, before I got in here-- you replay things that were happening. And what else do you do? You think about people? Why? Because we're social primates, and that's what we care a lot about.

You don't think only about people. Some of you guys might be trying to solve a math problem that you couldn't solve before. But most people in the scan are asked to do nothing, are recalling events, which usually involve people, or thinking about people, OK?

So this whole suite of brain regions that was called the "default-mode network" is just the regions that are more engaged when nobody tells you what to do than for a whole bunch of things when they tell you what to do. And so it's some weird mix of daydreaming and other stuff. And the interesting thing about it is they're reasonably systematic.

So those are the-- I keep getting confused here. Hang on. Let me get this right. Did I label this backwards? All right, they're the green guys. Yeah, deactivated during demanding tasks. Yeah. There's too many negatives for me here.

The green guys here-- does that look familiar, that patch? What does that look like kind of, not exactly but kind of? Sorry?

AUDIENCE: Visual, the visual system.

NANCY KANWISHER: It's sort of near the visual system, yeah. It is, but it's also like something else we've been talking about recently. Our TPT. It's a little further back, but it's right in the same region, OK? And here are these medial regions. There's a medial view of the left hemisphere, like, "take my right hemisphere out and look at the inside" view. That's this-- [INAUDIBLE] and sulcus, right? All these medial regions.

It looks a whole lot like the social-cognition network that I talked about last time, that you identify with the contrast of belief task versus-- the false-belief test versus a false-photo test. So that's weird finding number one, that there's a systematic set of regions that are engaged at rest. They're called the default-mode network because they're what you do by default when nobody's controlling you externally. So that's finding number one, and there's finding number two. But first, Jack, did you have a question?

AUDIENCE: Yeah. I was just wondering, does deactivated during demanding tasks necessarily imply that it is activated during not-demanding tasks?

NANCY KANWISHER: We're not distinguishing between those. We're just taking those two conditions. You'd have to have some third baseline to figure out whether those two were different. And that's very problematic because we're are we having a problem saying, what counts as a baseline, right? So we'll just compare those two, OK?

So that's weird finding number one. And it's not that weird when you think about it further because, of course, you're doing stuff when you're lying there, right? But the further finding that really put the default mode network on the map was when people started putting seeds in parts of that default-mode network up here and finding that they got the whole rest of the network at rest, OK?

So all of those things are correlated at rest. It's not just that they're all activated at rest. Their time courses are correlated at rest.

So, actually, what people mean by default mode network now is not, I took the reverse contrast and all the stuff that activated more for rest than task, I call that default mode. Actually, what they mean is I stuck a seed in there during my rest scans, and I took all the stuff that was correlated with that position because those pick out, more or less, the same thing. OK?

So that's led to a whole lot of discussion about what the default mode network is, and what it means, and what we can learn from it. That's what all this says. I'm just trying to figure out how I'm going to do this because I'm going to run out of time. Maybe I won't run out of time. We'll just go for it. OK.

So people started messing around with these correlations at rest. And they found that you could find other systematic sets of regions if you stuck seeds elsewhere. And so another systematic region, a set of regions, is all the hot-color ones, the yellow and red ones here. And so if you look in there, you see various things-- the interparietal sulcus, a bunch of frontal regions, Visual-Motion area, MT, or other visual regions down there. And they found that set of regions was strongly correlated with each other at rest. You stick a seed in here, and you get all that yellow stuff, OK?

And then they looked at it, and they said, yeah, right, we've seen those regions engage. Whenever people do demanding-- potentially demanding tasks-- they've seen that before in other task contrasts. So all those things that turn on when you really have to pay a lot of attention and you're doing a really hard task-- all those regions do that, and they're also correlated with each other at rest. OK? And so there's this convergence of these two different lines of work-- task contrasts that just say, what makes a given set of regions turn on or off, and correlations-- which things are correlated at rest? OK? And so they're both converging here with these two different networks.

OK, so I need to do a little sidebar on this other hot-color network, not the default-mode network but this other one. It was originally called "task-positive" because it turns on more when you do tasks than rest. I mean, that's a really vague statement, OK? But it also has lots of other names, and the name that we're going to refer to here is the multiple-demand regions, OK?

Multiple demand comes out from another line of work that just converged with us. They're picking out pretty much the same set of brain regions. But "multiple demand" means lots of different kinds of cognitive demand activate those same regions. OK?

So I can give you a difficult spatial working-memory task. I can give you a difficult perceptual orientation-judgment task, a difficult arithmetic task. In each of those cases, I can compare them to an easy version of the same task, and I'll get, more or less, those regions there, which are-- it's getting a little vague here, but they're pretty similar to the task-positive ones, OK?

So this is both interesting and scandalous. It's scandalous because-- to me only, not to anyone else-- because unlike all the regions we've been talking about so far that have these very specific functions-- they just do face recognition or just theory of mind-- these ones will do anything, almost, anything difficult, OK? So whenever you engage in a difficult task-- I'm skipping over a whole literature-- it's a big literature on this-- but lots and lots of totally different kinds of tasks that have nothing in common other than they're very demanding-- you engage those regions.

And in some ways, that's an even more fascinating puzzle. Like, what the hell would those operations be? What is in common between spatial working memory and arithmetic and line-orientation judgment and all the other things that have been shown to activate these regions when they're demanding? Nobody knows. I think it's a big, fascinating puzzle. Someday we'll have a computational story about what's actually computed in those regions, but we don't yet, OK?

There's a lot of stuff on the multiple-demand regions, and they're interesting. But I can't resist one little thing. All right, I can't. I have no self-control because I'm not-- I don't have enough multiple-demand activity right now, so I'm going to have to just tell you these things because they're cool.

OK, so a guy named John Duncan has spent the last 15 years arguing that the multiple-demand regions first are really truly multiple demand. He's tested lots and lots of different tasks, and they're very, very domain general. But second, he thinks they're implicated in fluid intelligence. Fluid intelligence differs from crystallized intelligence.

Crystallized intelligence is stuff like your vocabulary, just stuff you've learned and cached away, and facts you've stored, and abilities you've-- specific abilities you've stored. Fluid intelligence you measure with stuff like Raven's matrices, where nothing you know is going to help you do it. You just have to be smart and see some abstract pattern or something like that.

And so Duncan thinks that these regions are related to fluid intelligence. And one of his measures of that is if you find people with brain damage-- he had a big set of around 80 people who had brain damage who he'd been studying in all different parts of the brain.

And what he found is if you have brain damage in those regions, your IQ goes down as a result of the damage in proportion to the amount of cortical volume destroyed by the damage. If you have damage anywhere else in the brain, your IQ is unaffected. You may become paralyzed, or aphasic, or prosopagnosic, or akinetopsic. You may have any of these very specific deficits according to where it lands, but it won't affect your IQ.

And so the picture here is that in addition to all these special-purpose processors that this course has been focusing on, we have this thing that's kind of like the brain's CPU or something like that. And it seems to live in approximately those regions, and it seems to under-- it seems to be essential for fluid intelligence, OK? So I'm skipping over lots of literature just to heighten that this is a particularly interesting set of regions here, OK? Yeah.

AUDIENCE: Do people study novelty?

NANCY KANWISHER: Yeah.

AUDIENCE: The regions that specialize in novelty.

NANCY KANWISHER: These guys will be interested in novelty. These guys will be interested in novelty but not only. You can do the same boring but difficult task on, and on, and on, and they'll keep going.

OK. The reason I went on that sidebar is that you can identify these regions, not just by scanning people while they're doing difficult tasks but by sticking a seed in any of those regions and getting the others, OK? That's the task-positive network pretty much, OK? So we're getting this convergence between sets of brain regions that we find with a task contrast and sets of brain regions we find by finding what's correlated with what.

And the bigger picture of this whole thing is that I've been focusing on individual regions and what they do, and the gist of this whole resting functional-correlation literature is a very relevant level of organization of the brain is not just an individual cortical region but a set of cortical regions that seem to be in cahoots. And, again, we don't know what that means exactly, but they're correlated at rest, and they have something to do with each other, OK? So we're finding this higher-level organization, and the multiple-demand system is part of it.

OK, so how are we going to look at this? So there's a bunch of these. I've talked only about the default mode network and-- it's another name for the same thing-- executive control. Don't worry about it. You think of that as multiple demand. It doesn't matter. But there's a bunch of them that you can find by sticking seeds in different places.

And so yes, I just said the big idea here is that networks are an interesting level-- an interesting kind of unit in thinking about brain organization-- bigger than an individual region. It's a set of regions that have something in common, OK? But I've sort of backed into in this awkward way of saying, OK, here are things that are correlated with each other, and here's what we know about the same regions from previous task analysis. Most of the literature on resting functional correlation just looks at correlations and doesn't try to put it together with what we know about those regions, and that just seems deeply weird to me.

So for years, I ignored this whole thing because it's like, I don't know what these resting correlations are. And if I don't know what they are, I'm not going to work on them. And then Idan Blank came along. And when he was a first-year grad student at Fedorenko, said, hey, we have all these resting-functional data. Let's have Idan, for his rotation, for a month, analyze the resting-functional data.

I said, resting functional-- we don't know what the hell it means. Let's not bother. She's like, get over yourself. Let's let him play with it.

Well, thank God she's not as stuck in her ways as I am because Idan spent just a month playing with some of our data, and what he found blew me away. So here's what he did. He said, OK, let's start with actually identified regions of brain where we know something about what they do, like the language system and the multiple-demand system. OK? Let's identify those regions in each subject individually, and then let's scan subjects at rest, OK?

First, you scan subjects with sentences versus nonwords. You find the language regions. Then you scan them with a difficult-versus-easy spatial working-memory task. You find the multiple-demand regions.

Then you scan them at rest, and you get the average timecourse from each of those regions at rest. These are fake data, just to give you the gist. Makes sense? Are you with me now? Now you can ask, OK, which of these things are correlated with each other at rest? And this is now a more interesting thing to do because we're asking this principled question of regions we know something about rather than random seeds in some random location, right?

OK, so now what we do is we examine those correlations. And we just ask, for example, how correlated are those timecourses of two different language regions or those two different parts of the multiple-demand system? And how strong are the correlations between the systems, some little piece of the language system and some little piece of the multiple-demand system? Makes sense? So that's a cool question to ask, and that's what Idan did.

And here's what he found. Let me first orient you. So here are lots and lots of regions of interest that were identified functionally-- a whole bunch of language regions up here, a whole bunch of multiple-demand regions down here. The details don't matter, OK?

But what I'm going to show you is in each cell, we're going to have a correlation between a given-- so this would be a given part of the cell over here-- given part of the multiple demand system. Sorry, a given part of the multiple-demand system and some other part of the multiple-demand system, or over here, a cell would be some part of the language system and some part of the multiple-demand system, OK?

So when you do that, here's what you see. Here are the correlations at rest between all of these pairs of conditions. And if you squint a little-- you don't even have to squint much-- but the black ones are the ones that are not significantly correlated at all. Blue means a negative correlation, and hot colors mean a positive correlation.

And so what you see is here are all the language regions. These are the right-hemisphere language regions, which barely even count. They're just there for the hell of it.

This is really the core language regions, and you can see they're all correlated with each other, even ones that are really far apart-- Broca's area and Wernicke's area-- 10, 12 centimeters apart-- strongly correlated at rest, OK? And if you look at different parts of the multiple-demand system, they're all strongly correlated at rest, even regions that are far apart-- something way up in the frontal lobe, something way back in the parietal lobe-- strongly correlated at rest.

And so yeah. If you zoom in here-- so what this is really-- does everybody see that this is revealing a lot of structure? Yeah, question?

AUDIENCE: So the diagonals would be, like, self correlation, right?

NANCY KANWISHER: Yeah, that's why it's black. Actually, what is it? Oh, you know what? It's split in half. It's split in half. It's actually a better way to do it. You take your data, and you have two different halves of the data. And so it gives you a baseline for the-- no, that doesn't make any sense, does it? Never mind.

AUDIENCE: Yeah, [INAUDIBLE].

AUDIENCE: [INAUDIBLE].

AUDIENCE: [INAUDIBLE].

NANCY KANWISHER: It is maroon.

AUDIENCE: If you look at the spectrum on the bottom, [INAUDIBLE].

AUDIENCE: Yeah.

NANCY KANWISHER: All right, I'm going to have to solve this offline because I'm now confused, unless Anya can figure it out right now and bail us out.

AUDIENCE: [INAUDIBLE]

NANCY KANWISHER: Yeah, but as they're pointing out, it's not black. OK.

AUDIENCE: [INAUDIBLE]

[INTERPOSING VOICES]

AUDIENCE: I'm saying it should be high. I was just wondering [INAUDIBLE].

NANCY KANWISHER: Oh, it's a correlation of 1. Is that what it is?

AUDIENCE: Yeah.

NANCY KANWISHER: OK, because it's correlated with itself. All right. OK. Thank you. Anyway, does everybody get the gist that all of these different pieces of the multiple-demand system that we identified individually-- they're all correlated with each other at rest?

All these different pieces of the language system are correlated with each other at rest. And there's no correlation at all between any part of the language system and any part of the multiple-demand system at rest. All of these-- OK, maybe a couple. These cells are all either black for not significant or inversely correlated. That's the cool colors.

So do you see how this gives us a totally cool way aside from just the functional localizers we ran to identify these regions to show us that these things are functioning as a system, right? It's not just that Broca's area is a cool little thing that does a piece of language and some bit of the temporal lobe is a cool thing that does some piece of language. But those guys are part of a broader system, and these resting functional correlations are revealing this broader system and the integrity of the parts within it, as well as the distinction between those parts and parts of another system, or network.

Does everybody get that idea? Good. That's a big idea for this lecture. I think what I'll do-- so you can-- I'll skip that.

So this is basically the correlation between all of the cells within the language system, all of the cells within the multiple-demand system, and any pair of cells between systems here. So that's just averaging over the matrix I showed you before. And so this is a cool way to ask about broader systems in the brain. And I was going to show you some data published just a month ago that asked this question not just about the language and multiple-demand systems but also, about the theory-of-mind network, which is basically really similar to the default mode network.

But the theory-of-mind network we can identify that we talked about last time. And you can go think offline. Actually, I'll take a suggestion.

What do you think? Should the theory-of-mind network be correlated with the language system, with the multiple-demand system, both, or neither?

AUDIENCE: Neither.

NANCY KANWISHER: Neither? Why?

AUDIENCE: This is something different.

NANCY KANWISHER: OK, that's a totally reasonable answer, and that's largely true but not 100% true. What else might you think? It is a different system, and it does function quite independently but not perfectly. Yeah.

AUDIENCE: Because it's [INAUDIBLE] analyzing [INAUDIBLE]

NANCY KANWISHER: All right. That's a lovely speculation and an intelligent one, and it's half true. I'll just skip to the data. So in this very recently published paper, Alex Paunov, and Idan Blank, and Ev Fedorenko looked at the language system, the theory-of-mind network, which is not just the TPJ but these other regions that I mentioned briefly that you also get in the contrast of the false-belief test versus the false-photo task and the multiple-demand network. Same deal-- identify each of those regions in each subject individually, then scan the subject at rest and see what's correlated with what.

And here's the answer. You see, again, replicate the language system, especially in the left hemisphere. The theory of mind system is all a system. And the multiple-demand system is a system, separate system. But if you look in at the cell where you have theory of mind and language, it's slightly above chance, not theory of mind in multiple demand but theory of mind in language, probably for just the reason you said, the whole essence of language, even though it's a different thing, than thinking about the contents of someone else's thoughts. They are so enmeshed in each other.

The reason we have language is to take our thoughts and put them in your head and take your thoughts and put them in our head. And so it makes sense that those things are a little bit correlated.

AUDIENCE: [INAUDIBLE] language?

NANCY KANWISHER: Yeah. Yeah. But neither of them is correlated with multiple-demand. It's 1:26, so-- or-- wait. Am I reading the wrong-- yeah-- oh, 12:28. Sorry.

There's a number here. That's how long I've been talking. So 12:28-- so if you need to go, that's fine, but I'm happy to answer questions. So go ahead.

AUDIENCE: I was just going to say, there was a review paper about [INAUDIBLE] talked about how there's a behavioral connection between theory of mind and, like, when children started to learn [INAUDIBLE].

NANCY KANWISHER: Lots of links, especially developmentally, yeah. Yeah. Did you have a--

AUDIENCE: By an extension of that, [INAUDIBLE] should be correlated to the theory of mind [INAUDIBLE]?

NANCY KANWISHER: Yes, wouldn't you think? Not really.

AUDIENCE: [INAUDIBLE]

NANCY KANWISHER: The FFA is irritating. It's not strongly correlated with the things it ought to be correlated with.