

NANCY So we won't get through the whole attention lecture. I saw this coming. I just felt like that last bit was important
KANWISHER: enough. We'll cut as needed. But I do want to talk about attention. And let me start by sharing with you some of the key ideas about what attention is.

OK, so to get into the mode of thinking about this, let's consider the following question. How do you feel about people driving while talking on their cell phones? Smart, not smart? And while you're thinking about that, also consider the case-- like let's suppose that you have had a hands-free setup, so you don't have to be looking down on your cell phone or typing away or holding your cell phone. Maybe that's OK.

What do you think? Is there no problem with driving while talking on your cell phone, provided you're not looking down and pecking away at it? Yes? What do you think? I want to know what you really think. I know what all the PR says. Is it fine? Yes?

AUDIENCE: I would say it depends on whether you're driving-- Like if the road is clear and it's like your usual path, like driving is pretty automatic at that point. But it it's like high traffic or a new area, then I think it would have a significant affects, so you shouldn't.

NANCY Totally. My rule is, if I am talking on my cell phone, certainly, if I'm turning left, I put the damn cell phone on my
KANWISHER: lap. And I tell whoever it is, I can't talk to you and turn left at the same time. All right, turning left is the hardest thing we do.

Self-driving cars can't turn left. Have I mentioned this? How does a self-driving car turn left? It turns right three times. Why is that? Because turning left is social cognition. It's like, do they see me? Do they know I'm here? Are they looking at my indicator? Do they realize I'm going to turn left? This is hard. This is social cognition. Nobody's mastered that yet. Anyway, it's hard for us too.

Anyway, the real question here is not for me to lecture you about driving but to think about, why is that a problem? Why can't you talk to a person and pay attention to driving? What is the big deal, right? And so the intuition is we have some kind of limited processing ability. OK, we're not like super supercomputers who can just do millions of things at once. We have some kind of limit in how much stuff we can handle, OK? And I think everybody will have this intuition.

You can't think about lots of different things at once. I like to think of this as the toaster model of cognition. That is, you plug in the toaster, and the lights dim, right? So if everything's on the same circuit, there's some kind of unitary thing, resource that's limited. And if some more of it goes here, less if it goes here. OK, so obviously we don't have simple circuits like that in our brain. And the relevant scarce commodity isn't just electricity. It's not electrons. It's some other kind of thing we don't totally understand. But there's still some sense in which of our mental processes are on the same circuit. OK?

I'm assuming everybody is sharing this intuition here. So think about this. How many people feel like you can listen to music and read difficult stuff at the same time? Raise your-- I'm curious. How many people feel like they can do that? Only a few. That's so interesting. Sort of. Yeah, I can't do it at all, like at all, like even background music that I don't even care that much about.

Someday-- I mean, I think these are really stunning individual differences. I don't know if there's any good work on it. I haven't seen it. But I think it's really interesting. Some people can. Some people can't. I don't know what's up with that.

How about recognizing faces and scenes at the same time? Can you do that? Yeah? How many people feel like you can recognize a face and a scene at the same time? All right. In fact, Michael Cohen, who gave that lecture on brain-computer interfaces a month ago or so, showed some beautiful work that, actually, you are better at recognizing a face and a scene presented simultaneously than two faces or two scenes. Can you think why that might be? Yeah, Isabelle?

AUDIENCE: Context.

NANCY Say more.

KANWISHER:

AUDIENCE: If you see somebody, even a familiar face, you can either-- it helps identify them as, oh, I know that person. Then, obviously, they're there.

NANCY Right. So they can go together into a thing, maybe chunk it as a thing or something. That's a good answer. It

KANWISHER: wasn't the one I was fishing for. Yeah, Jimmy?

AUDIENCE: If people and places are separate things that your brain calculates, if it's two faces, you have to calculate this thing with a face number one before going to face number two. Whereas if it's [INAUDIBLE]

NANCY Exactly. Exactly. And you do, right? You have an FFA and a PPA. And they're friends, right? So what Michael

KANWISHER: showed is not just that you can recognize a face and hold it in working memory, a face and a scene better than two faces or two scenes, but the degree to which you have that cost of doing two things in one category rather than spreading over two categories is linearly proportional to how different the activation patterns are for those categories in the ventral visual pathway.

So faces and scenes are totally different from each other. And so you get a big benefit for separating your two items across faces and scenes. Other things that are more similar in their pattern of response in the ventral pathway, like faces and bodies, you get a smaller benefit. OK? So it's consistent with this idea that, to the degree that we have separate processors, you can do, to some extent, processing in parallel, OK?

It doesn't explain why not everybody can read and listen to music at the same time. Because, as I argued a few weeks ago, these things don't overlap at all. So there are many mysteries. But there's some intuition there.

OK, to get a little more intuition about limited processing capacity, can you guys see? When I was sitting over there last time, I couldn't see the screen at all. Can you guys see it? It's OK? All right. Yeah?

AUDIENCE: Do you think the problem is [INAUDIBLE]

NANCY Not for me. Instrumental music, no lyrics still a problem. So I don't know. I mean, I'm sure there's a literature on

KANWISHER: this. I just don't happen to know it. I'm just trying to share the intuition. So I don't-- it's not the only thing for me.

OK, so what I'm going to do-- this is a super low-tech demo. I'm going to show you an array of colored letters. Grab a piece of paper or something where you can write down a few letters. And I'm just going to flash up one array very quickly. And I want you to write down as many of the blue letters as you can. OK, ready?

There's only a few of them. So you can. OK, ready? Here we go. OK, write them down. Last year, everyone got all of them. So I tried to go a little faster this time. OK. OK, we're going to do it again. Ready?

There's going to be another display, and you're going to write down all the blue letters. Everyone ready? Here we go. OK, just wanted to share the intuition. The probability-- did anyone miss the N here? OK, how many people got the N? Nobody's going to admit it.

I wanted a few people to miss the N, right? I really did zip through. I probably cheated and made this one go faster. You missed it? OK, thank you. I'm sure a few others did. OK, anyway, the point is that, if you do this properly in a lab, the probability of detecting it here is much less than there, right?

So what does this show? It shows that we have limited capacity. We can't process all of those blue letters in parallel, right? When there's a single one, you get it just fine. But when there's a bunch, you don't necessarily get it, right? OK.

It also shows that we have ability to select the information, right? So you weren't bothered by the red letters. If I'd asked you right after about the red letters, you probably couldn't report any of them at all, right? It's a whole field that studies that. We won't get into that. OK.

So in fact, the probability of reporting the N in these two displays is independent of the number of red letters. They just don't matter. They are not entering into that limitation, whatever it is. The limitation is only for the ones you're paying attention to, the blue ones, OK?

OK, so why is our mental capacity limited? OK? Nobody really knows an answer. And I actually think this is a really unsatisfying part of attention research because it sort of seems obvious. But I think it's not. So sometime in the next 10 or 20 years, some smart computational person will analyze this well and get a more satisfying answer.

But right now, here's where we are. The standard story is, well, we have only so many neurons and so much capacity. And you can't do everything at once, right? So we can't process everything in the whole visual field at once. That's the standard story.

But the reason I find this unsatisfying is like, why the hell not, right? We have all this parallel stuff or the whole visual field, at least in the first few stages of processing. And we have quite a degree of parallelism after that. So I think that's the standard story. But I'm just marking that I don't find it very satisfying. Yeah?

A second answer, which I think is also not totally satisfying but also a little bit right at least, is that, typically, you're only going to do one action. Right? You're going around in the world. Let's forget colored letters on a display in a lab. Let's think about a person walking around in a busy city street doing something, right?

There's typically only one or a very small number of things you're going to do next, like walk down the sidewalk and not bump into people or pick up this object, right? You don't need to act towards all the things in your world. And in fact, if you had to, so given that you're only going to do one thing, why distract the whole rest of your motor-planning system by feeding it all this information that's just going to clutter it with garbage? Why not give it just the information it needs? OK, so that's the standard story. It's very squishy and vague but hopefully intuitive.

And there's some nice, compelling examples that illustrate this. So there's a fish called the pike fish that preys on smaller fish called sticklebacks. OK? And if you stick a pike fish in a tank with sticklebacks, it will catch the first stickleback faster if there's only one in the tank with it than if there are 10 in the tank with it. OK?

You might think, you've got 10 to choose from. You'll get a fish faster. But 10 is more distracting, right? And so maybe our action-planning systems also prefer the single stickleback. Again, there's no computational precision. I'm just sharing intuitions with you.

OK, so let me give you some more evidence that there really are significant capacity limits in perception and that, in fact, there's a lot of stuff right in front of us that comes right onto our retina that we don't see. OK? So I'm going to show you an example. I'm going to show you a picture for just a few seconds. And I want you to just look at it. And then I'm going to ask you some questions about it. OK? Here we go.

OK, so how rich a percept did you get? Did you see lots of stuff and get lots of detail? Or do you just feel like you just have the vaguest sense of a few buildings, a street, that's it? How many feel like they got a lot of pretty good detail? No one does. Or everyone's bored. I don't know which.

AUDIENCE: What do you mean good detail?

NANCY
KANWISHER: The colors of buildings, the presence of objects, details on the architectural styles, where an awning was, that kind of thing.

AUDIENCE: I feel like I got [INAUDIBLE]

NANCY
KANWISHER: Well, most people looking at this feel like-- I mean, who knows what it means? I mean, we're just sharing intuitions here. But most feel like, that was pretty rich. I have a sense of it. I have a sense of maybe what country that's in and what kind of stuff is there and what kind the style is. You get a feel for the place, all that kind of thing. Maybe not every damn detail, but lots, right? So the general intuition most people have is a fair amount of detail. OK?

Well, let's look at that some more. This is actually a very, very heated topic right now. Me and Michael Cohen published a paper a couple of years ago called The Bandwidth of Perception, which is an effort to grapple with this question of how much information is there in your current percept right now. And there are many different perspectives on this.

Let's find out a little more with a further demo. So what I'm going to do is I'm going to show you that picture again. And it's going to flash on a number of times. And each time it flashes on, there might be something that's different. So take notes on any differences you might detect. OK? OK. OK, what things changed? Yeah?

AUDIENCE: There was a woman walking down the sidewalk.

NANCY And?

KANWISHER:

AUDIENCE: That's all I got.

NANCY OK, but she changed. What changed about her?

KANWISHER:

AUDIENCE: Other than her clothes?

NANCY Yeah. Over the successive presentations, was she there sometimes and not others? Did she change? Yeah, she

KANWISHER: appeared or disappeared. OK, good. What else, Isabel?

AUDIENCE: The awnings changed. The buildings changed colors.

NANCY Very good. Very good. How many people saw an awning change? Maybe half of you. OK. What else changed?

KANWISHER:

AUDIENCE: The car.

AUDIENCE: The model of the car.

NANCY The model of the car. Very good. Yeah?

KANWISHER:

AUDIENCE: I feel like the shops changed. But I don't know when or how.

NANCY OK. OK. You want to see how they changed? OK, here's how it ended. Here's how it started. Are you surprised

KANWISHER: how much? Raise your hand if you're surprised how much changed. Yeah, a lot, right?

OK, so what does that tell us? It tells us that either the sense we have that we really saw a lot of what's going on there was wrong. We didn't see as much as we thought. Because if we saw as much as we thought, we should notice massive changes like that. If you didn't see the awning, look over there. I mean, how could-- look at the color changes. Pretty major, right?

Or we perceive all that stuff in the instant, and it just goes poof by the time the next one comes along. That's one of the things the field is finding about those things are very hard to tell apart, not impossible, but hard. Yeah. OK, so most people feel like, wow, much more changed than I thought.

I can't resist one more hilarious demo. So have you already seen this like in 9.00? How many people have seen this before? Oh, I don't want to bore you. Do you mind seeing it again? We could skip it. It's kind of fun. Sorry? Do it? OK. OK. OK. So you're just going to watch this video and just track things because there may be changes happening here and there, and you want to notice them, OK? Here we go.

[MUSIC PLAYING]

[VIDEO PLAYBACK]

- Clearly, somebody in this room murdered Lord Smythe, who, at precisely 3:34 this afternoon, was brutally bludgeoned to death with a blunt instrument. I want each of you to tell me your whereabouts at precisely the time that this dastardly deed took place.

- I was polishing the brass in the master bedroom.

- I was buttering his lordship's scones below stairs, sir.

- I was planting my petunias in the potting shed.

- Constable, arrest Lady Smythe.

- How did you know?

- Madam, as any horticulturist will tell you, one does not plant petunias until May is out. Take her away. It's just a matter of observation. The real question is, how observant were you?

Clearly, somebody in this room murdered Lord Smythe, who at precisely 3:34 this afternoon, was brutally bludgeoned to death with a blunt instrument. I want each of you to tell me your whereabouts at precisely the time that this dastardly deed took place.

[END PLAYBACK]

NANCY Totally different guy.

KANWISHER:

[VIDEO PLAYBACK]

- I was polishing the brass in the master bedroom.

- I was buttering his lordship's scones below stairs, sir.

- I was planting my petunias in the potting shed.

- Constable, arrest Lady Smythe.

[END PLAYBACK]

NANCY It's actually an ad. But it doesn't hurt to get its message. All right, it's a British ad with an important message.

KANWISHER: But it's a pretty impressive demo too, isn't it? How many people feel like lots of stuff changed that they didn't notice? Yeah, that's intuition.

So the idea here is that all that stuff hits your retina. All of it kind of gets processed to some degree. But it's amazing how much of it goes unnoticed. OK? All right, oops. Sorry. OK, you might think, OK, is this something that only happens when it doesn't really matter? OK, you were looking for changes. But your life didn't depend on it.

What about commercial pilots? So here's a classic study where they brought in actual real commercial pilots with thousands of hours of flying experience. And they had them fly in a flight simulator. They had them land a plane on a runway in the flight simulator under foggy conditions.

And I forget what percent. I have to cheat and look at my notes. It doesn't say. It does. I'm just not seeing it. Some large percent of the pilots never saw this plane sitting right there on the runway. They landed the planes in the simulator right through that plane. And when they were shown it subsequently, they were shocked and couldn't believe they didn't see it, right?

They were using a heads-up display that tells them their orientation to where the runway is. And they're paying attention to this stuff, landing on the runway. And they just do not even see that very relevant thing, OK? So it's not just something that happens in like weird psych demos and funny movies. Even for very important things, people miss absolutely major, important information.

OK, so what all this tells us is attention is a filter that lets some information into our awareness but completely filters out of awareness a lot of other information that lands on your retina or your cochlea. OK? And there's two key properties of attention that go hand-in-hand. One, those capacity limits we've been talking about. You can't do everything at once. And two, selectivity, that is there's some way some subset of that information comes in some way to select it, OK?

All right, so let's say a little bit more about attention. There's lots of different ways and forms of attention. So let me give you a few key distinctions about attention. So far, it's just kind of a vague word that gets used lots of different ways.

So long ago, the great Helmholtz thought about attention and realized that there were two different kinds of attention. So he noted way back in 1860, our attention is quite independent of the position and accommodation of the eyes and of any known alteration in these organs and free to direct itself by a conscious and voluntary effort upon any selected portion of a dark and undifferentiated field of view. This is one of the most important observations for a future theory of attention. Indeed, it is, right?

So his point-- I think we did this at some earlier lecture. His point is you can fixate on my nose right now. OK, everybody fixate on my nose. Not that it's that fabulous a nose. It'll just serve for the demo right now.

And if I hold up different numbers of fingers out here-- keep fixating on my nose. No cheating. OK? How many fingers are off to the left side of my nose from your perspective? OK. How many fingers are off to the right side? OK.

You can do that. You can pull up this information or pull up that information without moving your eyes. OK? That's called covert attention because the input is identical. And you are just somehow adjusting the dials on your perceptual system to pull up this or pull up that. OK? So that's called covert attention.

To be unconfounded from overt attention, which is actually the most powerful kind of attention. Overt attention, you move your eyes from point to point to select different information. OK, so overt attention is a much more powerful filter.

And we make about two to four eye movements a second. And it's very powerful because the center of gaze has very high resolution information. Remember high density of photoreceptors at the fovea in the center of gaze, and the density tails off in the periphery. So we have much better information in the fovea than in the periphery.

OK, so it's both photoreceptor density, and it's also cortical area. So in retinotopic cortex, back here, primary visual cortex, V1, V2, those regions, you have about 20 square centimeters of cortex, like that area of cortex-- pretty big-- something like that processing the central two degrees of vision. OK? Huge cortical area devoted to a tiny little part of the visual world. And much less per degree for the periphery.

OK, so both at the photoreceptor stage and at higher-level processing stages, we vastly over represent the center of gaze. OK? That's why if you have loss of foveal vision and macular degeneration, it's really awful. You can't read or recognize faces. That's where you need this fine-grained foveal information.

OK, so that means that moving your fovea and parking it on different parts of the array is an extremely powerful way to select different kinds of information. So to give you a feel for how much people do this and how they select information, here's a video. This is taken from a head-mounted camera of a person who is watching these two people. And the yellow dot is where their fovea is.

So they're talking to my former postdoc Matt. And they're fixating on his face as he talks. And sometimes they take a peek at the other person. And you can see there's very fine-grained adjustments of where they look as this goes on.

And now they're walking down the hall in this building, and they're following her. Watch what happens when they turn the corner. Check out the signs. Read the little notices. Look down the hall.

And then watch what happens when they go around the corner. Oh, new person. Look at the new person. Right? Sorry, it's too dark to see what's going on here. Oh, this is a little fixation patch where they're recalibrating the eye tracker to figure out to make sure it's working.

Here comes a new person. They look at his face. Saccade back and forth between the two faces, right? Whoever's talking, you look at them. And you collect high-resolution information from that person's face. With both head turns, you can see their both head turns and eye movements. OK, you get the idea.

I didn't totally explain. So this is a head-mounted camera and eye tracker, OK, of a person who was walking around the building encountering people. And it was showing you both what was in their field of view and where they were fixating their eyes in that view. So you can see there's really fine-grained, moment-to-moment sampling of visual information all the time with overt attention with eye movements, OK? Make sense? OK.

OK, first, so that's the first distinction about attention, covert versus overt. OK? Now, you might say, why would we bother with this subtle, sophisticated covert attention when we can just move our eyes and do overt attention? And I think there's a bunch of reasons for that.

One is other people can see where you're looking. And there's sometimes you want to attend to stuff over there and not let people know, right? So there are many examples of this from elevator eyes, people who look you up and down. It's not politically correct. It's not nice. It's not considerate. People notice if you do that. So don't do it.

You can covertly attend below the face, if you like. But don't overtly move your eyes down there because they will know. OK? We primates are very, very attuned to where each other are looking. We are the only primate who has whites around the eyes. I just heard a talk last week by Michael Tomasello, who's one of the major primatologists. And he says that the fact that humans have whites of their eyes that make it so easy to tell where they're looking must mean that we mostly want to share information with each other about where we're looking. We're a very cooperative, communicative species.

But not always. Sometimes we want to sneak a peek at something where people don't know. The case that I notice all the time is I'm at a conference and some very familiar face comes along and says, hey, Nancy. How are you? And I'm thinking, who the hell is this?

And the whole time, I'm thinking I can see the name right there. I can't quite resolve it. I know that if saccade down for a half a second, they will see it. And I will be busted. And so I'm sitting there with a name right there racking my brain trying to think who the hell it is. Anyway, so there are many cases like this where we don't want to be caught looking. So those are all cases where you might want to use covert attention.

Also, sometimes you want to track lots of things in parallel, and you can't foveate all of them. You can only foveate one thing. So I don't know anything about sports. But think of your favorite complicated sports example where there are multiple players moving at once and you need to keep track of all of them. That would be a classic case of covert attention because you could covertly attend to several of them at once. And you have only one fovea. Well, you have two, but they go to the same place.

OK, so that's the second distinction overt versus covert-- or the first one. Right? So the next basic distinction of different kinds of attention are the kind where we decide where we want to look versus the kind where the stimulus draws our attention, OK? So there are lots of examples.

Like web pop ups are all designed to pull your attention, overt and covert. Even you don't want to look at that damn thing. But the weird, little dancing figure, they're all optimized to pull your attention over and make you read the ad, right? And they're pretty effective. OK, so that's called stimulus-driven attention or exogenous attention. It comes from the outside. And it pulls you in. OK? Another classic example is pop-out. If you look at this display, it's very hard not to notice and have your attention drawn to that red thing. Certain properties of stimuli that just automatically capture your attention.

OK, that is to be contrasted with voluntary controlled attention. That's like the case we did before where you were fixating on my nose. It was totally up to you whether attend-- I mean, I told you to. But you could have made up your own mind and attended to neither. And I wouldn't have known, right?

So you can decide what you want to pay attention to. And as I think I mentioned in I don't know what context a few weeks ago, thank God we can decide what to pay attention to. Because there's lots of stuff that goes on in the world that we don't want to have dominate our mental life. And controlled or voluntary attention is one way that we can have our mental life dominated more by things we want it to be dominated by and less by things we don't, right?

So that's that notion here. Like for example, you're sitting here. And it's like noon. And you're just thinking, OK, I'm really hungry. I'm really hungry. If you were my dog Charlie, your entire mental life for the whole next half an hour would be, I'm so hungry. I'm so hungry. I'm so hungry.

But you're a person. And so that thought may impinge now and then, but you can drive it away and think about something else. Because it's not going to get you anywhere to focus on how hungry you are. I know I just made it worse. I'm sorry about that. OK.

OK, the way that scientists have typically studied voluntary attention is they ask subjects to do this. And much like the case we just did, here's kind of very basic paradigm. You have subjects fixating on a cross.

And typically, in covert attention experiments, you want to make sure there aren't overt eye movements. And so you'll have an eye tracker to make sure the subject is keeping their eye on that cross. OK? And then you give them a little cue that says pay attention off to the right, but don't move your eyes, OK?

And then shortly thereafter, a little cue comes up. And you have to hit a button quickly that you detected the target. OK? It comes out at variable delays. So you can't just hit the button. OK, if you do that, you find that the reaction time to detect that target, if it's where you're attending, is really fast, not much more than 200 milliseconds, which is damned fast, right?

Now consider a case where you're instructed to attend over here. And the way you get these experiments to work is that these cues are valid 80% of the time. So there's a reason for subject to-- there's a motivation for them to follow your instruction. So they're trying to answer quickly. And there's like, OK, they're ready. It's almost certainly going to be over here. But oops. It's not. OK? Then reaction time is almost 100 milliseconds slower, which is a huge effect in psychology, 100 milliseconds. OK?

OK, so you're much slower to detect something if it's at an unattended location. OK? And if you have a neutral trial, you're somewhere in the middle. OK, so that is exogenous or voluntary attention. OK, and again, this is covert attention. No eye movements.

All right, so third distinction, the example I just gave you was spatial attention. You're attending to this location or that location in space. OK? That's probably the most powerful and common kind of attention is usually the things you're interested in are in a particular place. And you attend to that place. And there you go.

But it's not the only kind. You can have feature-based attention. So if your utensil drawer is a hell of a mess like mine is, like this, if your task is to find the black vegetable peeler, best of luck to you. It'll take quite a while. It's in there. It's right there. But it takes a while, right?

In contrast, if the task is find the purple rubber spatula, duh, right? OK, so that's feature-based attention. You don't know the location in advance. You know something about the feature you're looking for, that it's black or that it's purple or round or square or whatever the feature is. OK, so different kinds of filters you can set on your attention system.

OK, so that's just to give you some of the phenomenology of the different kinds of attention. How does all this work in the brain? I'm not going to tell you how it works. We'll just focus mostly on what the brain regions are.

OK, so let's go back to this case here, voluntary attention. So now the question is, what happens when you get this cue? You're not going to move your eyes. But you're kind of cranking up that part of space.

OK, what's going on in the brain right there before the stimulus comes on? Everybody get the question? OK, you can sort of feel it happening. But what does that mean? Well, this is very amenable to a simple functional MRI experiment.

And in the early days of functional MRI, a whole bunch of people at once realize, oh, right, we can answer this longstanding question of whether early retinotopic parts of the visual system are affected just by a cue like that, even before the stimulus comes on, right? And I think it was 1998. A whole bunch of people started doing this in 1997. I was one of them. But a whole bunch did it once in 1998. I think like six papers were published all doing versions of this next thing, which is one of those obvious ideas. We got scooped. But anyway, everybody got the same result.

So basically, if you just do a simple comparison-- this is the version they have. Subjects are fixating here. There's two stimuli. They're looking in V1 and V2. And if you just do a contrast of the case where you're looking here and you're attending, waiting for a target here versus waiting for a target there without moving your eyes, you get a big contralateral modulation in V1 and V2.

This is a slice near the back of the head here. So this is a piece of calcarine sulcus or primary visual cortex, right? And so you can see, even before the target stimulus comes on, just attending over there gives you a big baseline increase in neural activity. Everybody clear what's going on here? And of course, you can do the opposite one and see that it shifts to the opposite side of space. OK?

OK, so that's all before the stimulus comes on. You can think of it as priming the brain, kind of juicing up those neurons getting them ready to go, so that, when that stimulus comes through, boom. Your reaction time is faster. OK? OK, so that's modulating parts of visual cortex before a stimulus comes on to get it ready.

But we can also look at, what is the effect on a stimulus once it comes through, an attended stimulus or an unattended stimulus? To show you that, I'm going to show you a kind of feature-based attention and an ancient-- this is a paper we published in 1998, I think. This is a nonspatial version.

So we gave subjects-- guess what-- faces and houses and a little crosshair. And we had subjects in different blocks say whether the two arms of the cross were the same or different length, OK? So vertical one is different length. So that's a different case. Whether the two houses are same or different or whether the two faces are same or different. OK, so actually, it's spatial and feature, when you think of it, because they're in different places. OK?

So we then looked in a bunch of regions. I'll show you the data from the fusiform face area. And the key thing about this is, in all three tasks, we had the identical stimuli. So there's always faces and houses and crosses in the same place in your visual field, in all of those blocks. It's just a matter of which of those things you're paying attention to. OK?

So what do you think we'll see in the fusiform face area in this experiment? Is it going to care whether you're attending to faces or attending to houses or attending to the crosshair? Same stimulus all the time. Yeah, Jimmy?

AUDIENCE: [INAUDIBLE] when you're not attending. But when you're attending, [INAUDIBLE]

NANCY Exactly what happens. OK, everybody hear the prediction? You'll still get something. But you'll get more when
KANWISHER: you're paying attention to it. OK, so here's the FFA response over time. Here's the time course.

And these gray blocks are the blocks. These are the different attending to houses, faces, color, houses, faces, crosshair, houses, faces, crosshair. The gray bars are when the subjects are attending to faces.

Now, here, I actually can't quite tell. Yeah, I can sort of tell. You can't see it on the screen. But there are little rest periods between those blocks. So I'm trying to figure out if you're right that there's some response greater than baseline. I think there's actually very little selective response greater than baseline.

By design, we made this a whopping attentional manipulation, so that all of the tasks are really difficult. And when you're doing one of them, you just barely feel aware of anything else. In fact, I vividly remember the first time we ran this experiment, I was a subject in the scanner. And I did the first block. And I was doing just crosshairs.

And I went through the whole crosshair thing. And it wasn't till the end of the experiment I realized, oh, right, there were faces there. I was just completely unaware of them. So we designed it to really so tie up your mental capacity that you just didn't have processing resources left for anything else. And so if there is a response to everything else here, it's really low. But you can see, certainly, the attended thing is much more. It's much stronger in the attended case than the unattended case. OK?

So this is a general property. Like basically, all of the perceptual regions we've talked about are strongly modifiable by attention, OK, from retinotopic regions on up. OK, including V1 and, in fact, even including the lateral geniculate nucleus. So one synapse up from the retina, you're already modulating activity there by attention. OK? I know I said this, but it maybe went by briefly.

You have 10 times as many connections down from cortex down to the LGN as you have going forward. OK? One of the things they're doing is setting up selective filters, so that only the stuff you want to process makes it to higher stages. OK, Yes?

AUDIENCE: So does this sort of phenomenon generalize across the brain areas? I mean, it's like--

NANCY It's generally true of pretty much everywhere, yeah.

KANWISHER:

AUDIENCE: It might be the intuitive physics thing is justified. So it's better to see the same stimulus.

NANCY Exactly. It's an instance of this, exactly right. Exactly right. It gives us one of our paradigms. In functional MRI

KANWISHER: studies, have the same stimulus, vary the task. OK? Some things don't work as well, things that are just very dominant stimuli.

In this experiment, if we hadn't put the faces in the periphery and given subjects a damn difficult task, the modulation wouldn't have worked. Because faces are just so dominant, they're going to punch through even if you don't want to, OK? Like web pop ups. You also-- I'm going to skip this because we won't have time.

You could get very similar things in primary auditory cortex listening to high frequencies or low frequencies. These are all the same stimuli. This is voxel selected for low frequencies, voxel selected for high frequencies. You have a high frequency input in one ear, low frequency in the other. As you switch between, those responses toggle.

OK, so all of that tells you the effects of attention, how it modulates activity all over the brain. But what is the source of attention signals in the brain? And that source is a set of regions we have encountered before, sometimes called the frontal-parietal attention network, these blue and green bits up in here, back here in the parietal lobe, up here in the frontal lobe.

And they are active when you shift attention from one location to another, from one feature to another, or when you do any difficult attention-demanding task. We've also talked about them in the other context of the multiple demand system. It's pretty much the same set of cortical regions. That pretty much any time you do a difficult task, almost no matter what the task is, these regions turn on.

So they're not just about visual attention. They're kind of about basically controlling your mind, right, selecting information, making yourself do difficult things. Who knows what that is computationally. But these regions are very systematically engaged. And they're particularly systematically engaged when you're shifting attention over different locations.

OK, so just to remind you in contrast, everything else we've talked about, face areas, music areas, language areas, they're very specific for one mental process, domain specific. These are the opposite. They're ludicrously domain general. OK?

All right, I'm going to skip the video. When the system is damaged bilaterally, all kinds of weird stuff happens. OK? I'll post this. If you guys send me an email to remind me, I'll post this clip in. Balint's syndrome, where you have bilateral damage back there in these parietal regions, people can only see one object at a time. OK?

They're looking at a complicated thing like this. And they would see, I see Shosh. See anything else? No just Shosh. Nothing else there? No. Just Shosh, right? Totally weird. Anyway, so they get locked on one thing. They can't shift attention. OK.

OK, I want to talk at least briefly about awareness. So let's talk a little bit about neural correlates of awareness. And so first question is, if we want to study perceptual awareness and not just perception, how are we going to uncouple those things? We've sort of talked about a few examples with attention.

With attention, the thing you're attending to is at least much more dominant in your awareness than the stuff you're not attending to, even if the other stuff seeps in a little bit. So that's one way, where you have the identical stimulus. But you were varying the degree of awareness.

But another way-- oh, I need you guys to pass out the glasses. You want to both do that at different sides? OK, so I'm going to show you a stimulus. These guys are going to pass around glasses. And let me say in advance, I want them back. I use these every year. So drop them off here when you leave.

OK, so what you're going to do, you put them on either way. Doesn't matter. OK, so first, we're going to do optics. Nothing interesting. So look at this stimulus and close one eye. OK, now look at it through the glasses. Look at it through the glasses, and close one eye. And you should see just a face or just a house. And if you close the other eye, you should see the opposite.

Does everybody get that? OK, that's not psychology. That's optics. The glasses are just filtering one image into one eye and another image into the other eye, OK? This is all this is is a way to get different information to your two eyes.

OK, now just look with both eyes. And don't do anything. Just kind of look. And if anything cool happens, you can kind of go, ooh, aah. Or you can tell me what's happening. Just watch. Doesn't always happen immediately. Yeah, Kwylie, what's happening?

AUDIENCE: [INAUDIBLE]

NANCY Sorry. What's that?

KANWISHER:

AUDIENCE: Every time I blink, it changes from the house to the face.

NANCY Aha. OK, it changes from the house to the face only when you blink. Yeah, Kwylie?

KANWISHER:

AUDIENCE: So at first when I first put it on, I was like I saw the house and the face superimposed. And then you told us put one eye. And then I opened both eyes, and it went from the house to the face.

NANCY Uh huh. Uh huh. And then did it stay there? Or does it keep flipping? OK, if this is not working for you, don't

KANWISHER: panic. There's nothing wrong with you. It's hard to get everybody's red-green balance the same. Yeah. Sorry. Question? Shardul, yeah?

AUDIENCE: [INAUDIBLE]

NANCY So I share your intuition. I think I can choose too. But there's actually a whole literature on that. And the guy I

KANWISHER: was collaborating with on this project, 20 years ago, insists that, actually, that's a wrong intuition. You can't choose. And I said, oh, the hell. The hell I can't.

I'm in the scanner being your best subject ever because I'm switching exactly when you want me to switch. Don't tell me I can't switch. He said, no, the literature says you don't. Anyway, I don't know the latest on that. Yeah?

AUDIENCE: Kind of related, it feels a lot harder to see the house than the face.

NANCY Yeah.

KANWISHER:

AUDIENCE: I can see the face a lot easier than switching.

NANCY Yes. It's not optimally set up, so that it's perfect for each. Really the reason I do this is to amuse myself looking at

KANWISHER: all of you with your-- thank you. OK, so you can keep looking or not, but I'm going to tell you what's going on here.

The cool thing about this is when-- if it didn't switch for you, the person-- is there anybody for whom it didn't switch? One. OK, so maybe your color vision or who knows what. But anyway. OK, if it doesn't switch, there's nothing wrong with you. The percept everyone else has is it just flips every few seconds from one to the other. OK.

So what's cool about this is, when it switches, nothing changes on your retina. Only your state of awareness changes. And that gives us a wonderful lever to study what goes on in the brain when awareness switches, unconfounded from the stimulus. OK? It's like varying attention, but it's more powerful.

So of course, we put people in the scanner. We put mostly me and a few others. But anyway. We popped ourselves in the scanner. We just taped these things on our forehead, nice low-tech experiment. And we looked at that exact stimulus. OK? And we scanned the brain while people sat there watching the stimulus flip back and forth. OK?

So this is a stimulus for the whole experiment. This is the percept. It switches back and forth. The subject has a little button box, so they can say, now I see the face. Now I see the house. Now, of course, the choice of a face and a house was not random.

Why did we choose a face and a house? So that we could look at the response in the FFA and the PPA. FFA loves faces, hates houses. PPA loves houses, hates faces. Perfect. Right?

So the question is, are they going to switch when your percept switches, even though nothing's changing on your retina? What do you guys think? Switch? Switch? How many people think it's going to switch? How many people think it's not going to switch? A few. OK, all right, so that's what we did.

Here is now the raw MRI time course averaged over the FFA and averaged over the PPA for one five-minute experiment or three-minute experiment in one subject. Probably me. I forget. Again, stimulus is the same the whole time. The letters are the times when the subject pressed the button.

And you can sort of get a sense that there's a little bit of maybe like here's a phase the FFA response goes up. Here again the house goes up. But it's kind of hard to tell. So really, the way to analyze these data is to take all the face to house flips, right, and align them and signal average to get rid of noise. OK?

So we can look at, what happens before and after a switch from face to house? And what happens in the opposite direction switch? Everybody get the idea? It's just a way to clean up the noise here. And here's what happens.

This is time zero. The subject presses a button saying that their percept has flipped between a house to a face. The response in the FFA shoots up. And the response in the PPA shoots down. Isn't that cool?

OK, now why are these peaks and valleys after the button press? Yeah. Evan, right. There's a delayed signal, right? OK? This is harder. Why do they come back together out there? Yeah, because it flips back. It flips back.

And all we've done is signal average to one button press. But there are different durations of percepts. Sometimes it lasts three seconds. Sometimes it lasts 10 seconds, everything in between. By the time you're 12 seconds out, most people have flipped again. That's why it goes back. OK, make sense? Yeah?

AUDIENCE: [INAUDIBLE]

NANCY I don't know. And I'm not sure that that was generally true. I hadn't noticed that before. But we should scrutinize
KANWISHER: the images. I can't actually remember if this was averaged over subjects or if this is one subject. I suspect that's a fluke. But we could look at it and see.

OK, so the point of all of this is that these regions, the face area and the place area, care about what you are experiencing, unconfounded from what's hitting your retina, right? They are reflecting the contents of your awareness, not just what's coming into your eyes, which is cool.

OK, so that shows that these regions kind of track awareness. OK? But can we find any evidence for perception without awareness using neuroimaging? OK? So to make this point, I'm going to accelerate slightly, but I think there's just barely time.

There are lots and lots of studies of this genre. And I'll tell you about one. So I'm going to show you. I'm going to flash up a very rapid series of digits. You have to get ready to write stuff down. Your task is to see if there are any letters in the digit. There might be zero. There might be one. There might be a few. Write down any letters you see. It's going to go by really fast. OK, everyone ready?

OK. Here we go. OK, write down any letters you may have seen. OK? All right, let's do another one. Everyone ready? OK, write down any letters you may have seen. OK?

OK, raise your hand if you saw both the A and the P in the first sequence. A few of you. Less than half. Raise your hand if you saw both the X and the H in the second. Almost everyone. All right, well, I cheated slightly. But never mind how I cheated.

The way you cheat in demos-- you guys will be teaching someday. You have the one you want people to get wrong first. It's a total cheat. So it's slightly cheated. But if you do it right-- it didn't work last year. That's why I cheated.

Anyway, even if you do it right in the lab, you get a really strong effect. There was only one digit between the A and the P in the first sequence. And it's like your brain is still dealing with the A when the P comes along. You just don't even see the P. You're looking for letters. You don't see it. Whereas there were three digits between the X and the H in the second sequence.

And just to show you some real data. So this has been done lots of different ways. Here's an example of this experiment. And this is the time that people get the second letter when they don't have to report the first. Well, the first one's colored, right? So you either have to report the colored letter and detect the letter after it or just detect the letter after the colored thing.

So if you don't have to report the first one, there's no dip in accuracy. Oh, sorry. This is a function of the distance between the two items in the sequence. But if you do, there's a big dip. That's maximal with one intervening item. Yeah? OK, well, here it says two or three. But anyway.

OK, so that's called the attentional blink. And the idea is there isn't a physical blink of the eyes. But your attention system is tied up processing the first target. And it doesn't get the second one, OK? So there are dozens, probably hundreds of papers on this phenomenon. It's pretty cool.

But for present purposes, we're going to use it to say, does that unseen second target get into the brain? Presumably, it lands on the retina because you didn't blink. So it could probably get there. Probably got to V1. How far up the system did it get?

Can you think of how we might test this using functional MRI. How would we see how far those stimuli go in an experiment like this? What would we use? What would we measure the MRI response to? How could we design a version of this that you could do with functional MRI?

We need some MRI response, where, if we get that response, we know that it's a response to a given stimulus. What would we measure? And what stimuli would we use?

AUDIENCE: Well, we could use faces or something.

NANCY
KANWISHER: Yeah. It's always the same. Yeah. Faces and houses. Lots of ways to do this. But it's not even my experiment, and they used faces and houses. OK, so they did a version of that very same thing. This is a bunch of garbage flashing on.

Early on, there's a face. The subjects have to say which of those three faces they've studied it is. And then after that, a scene comes on, or it doesn't come on. OK, they have to just say, was there a scene? First of all, which face was it? And was there a scene afterwards? OK, so it's a variant of the thing you just did.

So here's the behavioral data. If you don't have to report the face, then you are very accurate reporting the scene. If you do have to report the face, then you're very bad at it if the scene comes right after. But if there's a big interval in between, you do OK with the scene. Same thing we saw before but now with faces and scenes. OK, that's the behavioral data.

What do you see in the PPA, right? The second target is always a scene. So what you see in the PPA, here's a response in the PPA to that second scene if it's a hit. That is if you detect it, right? You correctly say, yes, there was a scene in that trial. Here's a response with a correct reject. There was no scene in that trial. And you correctly said there was not.

OK, so that difference shows you the response in the PPA to a scene that you've seen, that you detected consciously. But here's a critical case. I don't know why there's a little pop up there. Anyway, this is a case where there was a scene, but you've said there wasn't. OK? That's called a miss. You missed it. It was there, and you didn't see it.

And what it shows you is two cool things. Well, first of all, it tells you that the more aware you are of the stimulus, the stronger response. But the crucial thing it shows you is this difference right here, which is significant. This is a kind of perception without awareness.

You did not detect the scene. You said there was no scene. But your PPA detected the difference. Everybody get that? So we have evidence for perception by the PPA at least of a scene without awareness. That make sense?

So there are many, many studies like this that kind of take it every which way. This is just one little example of how you can use neuroimaging to ask, how far up the system does an unseen stimulus get? Make sense?

OK, so we are about out of time. And I just gave you a little taste of some of the work on neural correlates of awareness, showed you, with binocular rivalry, a neural response that's correlated with awareness unconfounded from the stimulus, uncoupling, again, perceptual awareness from what's hitting your retina, and, in the case of the attentional blink, some evidence for perceptual representation without awareness. OK?