

**ANNOUNCER:** The following content is provided by MIT OpenCourseWare under a Creative Commons license. Additional information about our license, and MIT OpenCourseWare in general, is available at [ocw.mit.edu](http://ocw.mit.edu).

**PROFESSOR:** Good afternoon.

**AUDIENCE:** Good afternoon.

**PROFESSOR:** So there I was in my car this morning as the pouring rain started, thinking if I make a dash for it -- I've got to take the computer. I need the coffee, because otherwise I'm going to fall asleep. I don't need anything in that bag, do I? So I took off without the bag. And it was true that I didn't need most of what was in the bag. But the lecture notes for today's lecture would have been a useful thing to take with me.

On the other hand, if there was ever going to be a day where I forgot the lecture notes, this is probably the one to do it. Because I'm going to talk about attention today. And attention research is what I do for a living. If there's anything that I should be able to just stand up and lecture about, this is it.

Now of course, that means that you should find this to be the most gripping topic in the entire course, and that you should decide that you want to do this for a living. Well, when you decide you want to do it for a sort of a living, like a \$10 an hour living, you can come and be a subject in attention research in my lab. I would again advocate that you sign up -- I saw there are still these notes around about signing up to be a subject generally in BCS. My lab is separate from the BCS business because I'm technically Brigham and Women's Hospital. But you can sign up with us, too, and we'll pay you \$10 an hour to do visual attention research. What could be better than that?

Is Kristen here? I don't see Kristen. Kristen, one of the TAs is also in my lab, and I was going to point her out. Anyway, send me an email, we'll sign you up. Talk to me. We'd love to have you. And you can do Where's Waldo experiments for \$10 an hour. You think I'm joking.

Let me try to explain why it is that I'm putting an attention lecture in between a sensation lecture and a perception lecture. It's not terribly typical. More typically, if people talk about attention, they go off and do it later, after doing sensation and perception. But why am I putting it in there? The core reason is that you simply cannot process all of the information that you

take in from the world. You're taking in a vast amount of sensory information. Your perceptual capabilities -- for instance, those that allow you to recognize specific objects -- are limited. You cannot recognize all of the objects in the world that you are looking at, all at the same time. It simply doesn't work.

And so, roughly speaking, there's the situation. You've got a lot of stuff coming in from the outside. And you've a box here that does, let's say, let's call this one a recognition box. And only one thing at a time gets to go in and come out of that box, basically.

So this is like the basic MIT metaphor about drinking from the firehose. Well, if you're really going to drink from the firehose, it's a very useful idea to restrict the flow in some fashion, and let some of that water just [SPLAT] and get you wet, or whatever. And so there's a severe constriction, sometimes called a bottleneck -- I think I've got some slides that call it a bottleneck later -- that takes all of this and only lets some of it through.

And that bottleneck is governed -- it's not just random what gets through -- it's governed by mechanisms of selective attention that allow some things to get through and leave other things on the on the floor. And so if you think of this as sort of sensation and perception -- which is a little bald, but -- then that's why you put attention in the middle there.

Now to motivate this a bit further, let me do a demonstration. Actually, this is the demonstration of why reading the *Tech* while listening to my lecture may not be a brilliant idea. Well, it may be a brilliant idea. It just depends on your particular goals in life. I need a couple of volunteer type people who wish to read here. All right, there's a volunteer person, and there's a pink volunteer person. Yes, you, MIT person. But you have to come up here. So kick a few people on the way by and stuff like that. You have to come up here and do a dramatic reading.

What I'm going to do is have these people both read to you at the same time. You're going to read from here, from where it says "Catherine." And you're going to read from here. And you're both going to read nice and loudly and steadily. At the same time, yes. That's the interesting part. And what you're going to do is you're going to listen to her -- for, actually, to "her" specifically. Listen for the third instance of the word "her." When you hear "her," say "her." For the third time, raise your hand. OK? Yeah, we got this? Yeah, all right. This her, not that her.

**NINA:** I could tell you my name.

**PROFESSOR:** That would help. You are?

**NINA:** Nina.

**PROFESSOR:** That's Nina. This is?

**ZAINA:** Zaina.

[LAUGHTER]

**PROFESSOR:** Right. OK. Zaina or Zena?

**ZAINA:** Zaina.

**PROFESSOR:** OK. At least it's not just one letter.

**NINA:** My surname's [? Navarre, ?] if that helps.

**PROFESSOR:** No, this is not going to help at all. Her. When Nina says "her" for the third time, raise your hand. OK? You got it? You got it? On your mark, get set, read. [OVERLAPPING VOICES]

**PROFESSOR:** OK, thank you. All right, that was excellent. All right, so what was she talking about? Yeah, something. No, no, somebody raise their hand. Hand, hand. What was --? A letter, thank you. That sounds good. She'd gotten a letter. Was it a nice letter? Who knows? It didn't sound too good. Her countenance wasn't doing good things. What was she talking about? Zaina? What? Oh, uh. She was talking about uh. OK. Was she talking?

**AUDIENCE:** Yes.

**PROFESSOR:** So what's your problem? What you were doing was -- so how many people -- well, obviously the hands suggested that everybody could manage to do the task. What could you pick up from -- Zena?

**ZAINA:** Zaina.

**PROFESSOR:** Zaina. I'm not going to -- otherwise it's going to turn into warrior queen and stuff like that. What could you pick up about Zaina's speech? Anything? No content. How many people knew she was talking? All right, so you can pick up something. What else did you know about her?

**AUDIENCE:** The tone she was speaking in.

**PROFESSOR:** The tone she was speaking in. If it'd been a male voice, if she had switched to a male voice, you would've noticed. Anything else, you think?

**AUDIENCE:** Was she reading from *Heart of Darkness*?

**PROFESSOR:** Was reading from *Heart of Darkness*? No, actually what she was reading from was Lucretius, *On the Nature Of the Universe*. A wonderful book. *De Rerum Natura* in Latin. He's a Roman author. This is sort of the first intro psych book. It's also the first intro physics book, intro everything. In those days, you could write a book called *On the Nature Of the Universe*, in verse. This is a prose translation.

She was actually reading Lucretius' theory of vision. And even she may not have noticed that, because it's all about thin films and cool stuff like that.

**AUDIENCE:** [INAUDIBLE] video.

**PROFESSOR:** A video. Yeah, well it's an ancient Roman video.

But only a very limited amount of stuff got in. So there was a certain amount of stuff that was getting in. But at some point your auditory system gave up on processing that stream. And in terms of extracting meaning, understanding the words, it went with Nina, because that was the job. You can't do both of them.

We'd better let them go sit down. Thank you for being --

What would have made the task easier? What would make it easier to pay attention to one and not the other, do you think?

**AUDIENCE:** Amplifying one of them.

**PROFESSOR:** Amplifying one of them. Yes, if the warrior queen would have just been quiet, it would've been no problem at all.

**AUDIENCE:** If they read the same thing.

**PROFESSOR:** If they read the same thing. No, that's -- that's probably true, but not a deeply interesting true.

Well, for instance if she was male, it would be easier to segregate the two voices. If we moved them apart further it would be easier to segregate the two voices.

**AUDIENCE:** If one was singing.

**PROFESSOR:** If one of them was singing it would've actually probably been easier to segregate them. So if you change the sort of low level sensory information, it would be easier for you to decide which one to pay attention to. This is something that happens. Oh, so if you're sitting there reading the newspaper while you're trying to listen to this lecture, odds are you are missing one of the two messages. It's sort of dealer's choice there. But it's also not desperately polite, in case anybody was wondering. If you want to read the paper, you might as well go somewhere else.

But this happens in the real world all the time. There's a version of it known as the cocktail party effect. You go to a party and you're talking to someone, and you hear, typically, what? Like, your name, over there. So you do this selective attention thing. You listen to that conversation. You seem to be paying attention to this guy who's talking to you, but you're actually listening over there. The problem is eventually, this guy stops talking. And you realize, oh yeah, I'm supposed to say something now, right? I wonder what we're talking about. It can lead to a certain amount of embarrassment.

Now this happens ubiquitously in sensory systems and across sensory systems. So for example right now, until I mention it, you are not particularly aware of the pressure of your posterior on the seat. If I direct your attention to that, you say, oh, yeah, there it is. It was presumably there all along; I wasn't floating a moment ago. But until I direct your attention to it, it doesn't rise to the level of current conscious awareness.

And it shows up in vision, because the visual world is far too rich for you to process everywhere at once. And that's what makes these sort of Where's Waldo problems interesting and fun. If there was not a bottleneck like this, Waldo man would not have gotten rich. Right? Yeah, where's Waldo? There he is. Big deal. Have you found him?

**AUDIENCE:** No.

**PROFESSOR:** Oh look, I have a little laser today. Isn't that nice? So does it work? That's Waldo up there. So now you say, oh, that's really stupid, because I can't even see him now -- Oh, and we decided to exploit the technology by having it on three screens. There's no added information there, it's just it was too cute not to do it.

But if I say, where is the elephant spraying a car? You can find it. You might have noticed it before if you had been scrutinizing it. It was certainly visible all along, right? It wasn't that there

was a black hole here before. It's just that only when you had the desire to go in search for it did you manage to direct your attention to it in a way that allowed you to recognize these couple of objects. And it's that ability to constrict your processing that's really the focus, at least of the first part of today's lecture.

Let me show you the equivalent of the talking example, but now switch to reading. What you want to do here is to look at the little asterisks. And I'll put up two streams of text, columns, one on the left, one on the right. Nice and big so that you can read them. But what you should notice is -- keep your eyes moving down from asterisk to asterisk. What you should notice is you can read one or the other; you just can't read both at the same time, even though they're nice and big. Right? It just doesn't work.

It's not a visual restriction. It's a central -- it's a capacity limitation later on in the system. So this is by way of an answer to question one on the handout: what's the problem that attention is solving? Attention is solving this problem of having too much going on.

Oh, and attention is a grab bag term. I'm going to be talking about visual selective attention. Attention isn't one thing, like my laser pointer here. There are attentional mechanisms, selective mechanisms, all over the place in the nervous system. So when you are attending to the pressure of your posterior on the seat, you are selecting, probably using a different set of neural circuitry than when you're selecting one of these words. It's the same basic idea, but it's not like there's a single attention box in your brain somewhere.

OK. Some things, as we saw in that auditory demo, the reading demo, some things escape the bottleneck. Some things can be appreciated everywhere, all at the same time. Well, question two is, what is that set of things? And the answer is not babies. The answer is that there is a limited set of basic features that can be processed across the entire visual field at one time. Or, you could do it in auditory space. There'd be a set of basic features in auditory space, too, that could be processed at the same time. But I'm going to stick with vision.

So all these babies look alike. It doesn't take much to figure out that now there is -- da da da, where'd Mara go? Oh, there's Mara. If the baby turns green, you do something about it. Right? It's a highly salient stimulus. Or if the baby's head gets squashed, you know.

So they're a collection of simple, basic features, like color, size, orientation, that are not bottleneck limited in the same kind of way. You can find that if there's a single red thing in the field, you can find it anywhere without having to go hunting around.

Other things that you might think would be pretty obvious are not anywhere near so obvious. So as you look around here, you may notice that most of these baby heads are upside down and two of them are right way up. But it's not like the green baby head. You have to go hunting for upright versus upside down. Even though that's a very salient thing in the real world, whether or not you're upright, or whether your baby's upright or upside down.

So there are about, by last count -- last count was done by me, as it turns out -- 12 to 18 of these things that seem to escape the bottleneck. And that's probably about it. And they are a bunch of simple things -- well, seemingly simple things -- like color, orientation, and size. Things that you could imagine, for instance, the earliest stages of visual cortical processing doing. And then there are some other, more elaborate things that also escape this bottleneck. And they're things like -- well, if you believe my friend Chen from China, this would be an example of the importance of topology. He thinks that the distinction here is that this has a hole and this doesn't have a hole. The other possibility is that this has line terminations and that this doesn't. These are the sort of things you can fight about in this field.

But anyway, it's easy to find that among that. Curvy things among straight things are easy. Orientation in the third dimension works. So that cube is pointing up this direction; these cubes are pointing down over here. That turns out to be easy. Other examples would include motion.

Though actually, motion makes an interesting point. It's easy to detect the presence of something, but not so easy to detect its absence. So imagine the following. I didn't make a demo of this; I could have. Imagine you're looking at the ground and there's one little ant moving around. He's pretty easy to find, right? Because motion is one of these features that you don't have to go hunting for; it's just sort of there. On the other hand, imagine you're looking at an ant's nest, and there's one dead ant. How easy is it to find one ant who's not moving? Not easy. So the absence of a feature can be hard to detect. The presence of a feature, one of these 12 to 18 basic features, can be easy to detect.

Now, how do you actually go about establishing that something is easy to find or hard to find? I've been doing this in very qualitative terms. But now let me explain how you actually go about studying this. What we would pay you \$10 an hour for if you show up in the lab. What we would do is show you a computer screen full of stuff, and ask you a question. A simple-minded question like, on the next one, is there a tilted line? And what you would be doing is sitting there with a couple of computer keys. Bang one key if the answer is no, bang another key if

the answer is yes. Do it as fast and accurately as you can, and we're going to measure your reaction time. The amount of time from the onset of the stimulus to the onset of your response. How fast can you do it? Well, I don't have keys for everybody here, so let's just do it verbally. Say yes or no as fast as you can in response to these guys. Tell me, is there a tilted line present? Ready?

**AUDIENCE:** Yes.

**PROFESSOR:** Ready?

**AUDIENCE:** No.

**PROFESSOR:** Ready?

**AUDIENCE:** Yes.

**PROFESSOR:** Ready?

**AUDIENCE:** No.

**PROFESSOR:** OK, that's pretty straightforward. What's the next thing? OK.

What you should have heard is that your answers were given crisply, in unison, and it didn't make any real difference whether there were lots of vertical lines on the screen or a few vertical lines on the screen. So if we were to collect real data and to plot the reaction time in milliseconds -- thousandths of a second -- as a function of the set size -- the number of items on the screen -- what you would get for any of these 12 to 18 items, if you did the experiment right, is an essentially flat line here. This would be the line for saying yes, it always turns out to take a little longer, or typically turns out to take a little longer to say no, but it's not dependent on the number of items on the screen. So is there an L, is there a green thing, is there an X among these pluses? All those things would produce similar looking results where the slope of this reaction time by set size function would be, essentially, 0.

Not all tasks behave that way. So let's do a different one. In this case you're looking for the letter T. It can be rotated by 90 degrees left or right, or -- maybe it can also be upside down; I don't remember what I put in. But it may not be an upright T. But it'll be a T. The distractor items are all L's. And I just want you to say as fast as you can, is there a T present? Ready?

**AUDIENCE:** No.

**PROFESSOR:** Ready?

**AUDIENCE:** Yes. Yes. [INTERPOSING VOICES]

**PROFESSOR:** OK, ready?

**AUDIENCE:** Yes.

**PROFESSOR:** Ready?

**AUDIENCE:** Yes.

**PROFESSOR:** You also heard the speed - accuracy tradeoff there. This is a known phenomenon in reaction time studies, which is, one can respond very quickly if you don't sweat the accuracy things. And people do that routinely. When people do that a lot in our studies, we call them bad subjects. And we don't invite them back. But what you should have heard there, and should have felt yourself, is that the responses were faster when there were fewer items present. And that the responses of the group, particularly for these larger set sizes, were spread out.

Why were they spread out? Well, some people got lucky. This thing came up and their attention happened to be around here. Oh look, there's a T. Some people were unlucky -- oh dee do dee dee, oh yeah, there's a T. And some people were trying to psych out the professor and said, there was a yes, there was a no, there was another yes. I know about this: there's going to be a no. And they said no without doing anything so boring as to actually look at the display.

So what you get for data in an experiment like this would look much more like this. As you increase the set size, now the reaction time increases in a fairly linear kind of a way. The slope on these is quite fast. I mean, this is 20 to 30 milliseconds, thousandths of a second, for each additional item to say yes, and about twice that amount to say no. Depending on how one exactly models this, this suggests that you're running through 20 to 40 of these letters a second. So you're going through it quickly, but you're having to search now. It's not simply obvious that there's a T there; you've got to go and hunt for it. Over here you can look for the 5, is another typical sort of task that would produce results like that.

I wanted to say one other thing about that, but now I don't remember what it was. Oh yes. What I wanted to say was that the speed of this tells you that you're not looking at the rate of

fixation on each letter. If you're doing this in the lab, you make sure that your stimuli are big enough that you don't have to move your eyes to look at each one. If you have to move your eyes, your eyes only move at a rate of about 4 per second. And so if you have to fixate each one of the items before you can tell if it's a T or an L -- so if you used little teeny letters -- this slope would be more like 250 milliseconds per item, not 40 or 50 or something like that.

Attention can move much more quickly than the eyes.

One of the things that tells you is that you can attend where you're not looking. Something that basketball players know very well. When you hear that a basketball player has great peripheral vision, what that really means is that he can be looking here and he can be paying attention to his teammate over there, and throw the ball and fake out the opposition. Because the usual assumption is that you're attending where you're looking.

Most of the time that's true. But OK, so now I'm looking at this guy wearing red up there. And he thinks that I'm actually paying attention to him. But I'm not, actually. Because of acuity limitations, I have no idea what I'm paying attention to here, but I think it's a woman person, and I think she just moved. Oh yeah, look, it is a woman person. I can move my attention away from the point of fixation. And I can move my attention much more rapidly than I can move my eyes.

Now, the find the red thing among green things is a case where the property of the target is one of these basic features and immediately gets your attention. The find the 2 among 5's, or the T among L's is a case where everything in the relevant display is essentially the same as far as the early visual system is concerned. T's among L's, it's a vertical and a horizontal line among other vertical and horizontal lines. There's nothing in this early processing that tells those apart, it turns out.

Most real world searches are not like that. In most real world searches, oh let's see, what do I feel like looking for? I'll look for glasses. If I'm looking for eyeglasses -- there are some right there, and there's some more. There's no process early in my visual system, you know, some huge chunk of cortex devoted to eyeglass detection. It just doesn't happen. At the same time, I don't search around randomly. No glasses there, no glasses there, no glasses there, no glasses there. I'm searching in an intelligent fashion.

Here's how you do that. Let's do one more basic search. What you're looking for here is a red horizontal line. Tell me as fast as you can whether it's present.

**AUDIENCE:** Yes.

**PROFESSOR:** Now how you do that is not by having a chunk of your brain devoted specifically to read horizontals. Oh, remind me later; I've got to check whether you still have a [? McCullough ?] effect, speaking of red horizontals. We'll check that out later.

The way you do that is, you use those 12 to 18 basic features to guide your attention around in an intelligent fashion. So if you're looking for red horizontals, you've got something that can do red. You know, give me all the red things. You've got something that can do vertical. Was I looking for red horizontals or red verticals? Well, anyway. This is a red vertical. You've got something that can do vertical. So I've got the red things, I've got the vertical things. I can do that early on in the system. All I need is something that will do something like an intersection operation. And if I were to guide to my attention to the intersection of the set of all red things and the set of all vertical things, that'd be a really good place to look for red vertical things. Oh look, there it is.

So what you've got is a front end that collects information that can be used to control this bottleneck to guide your attention around, to feed sensible things to the back end of the system. I think that's sort of pictured there.

And the result is that a search for something like a red vertical line, it's not as easy as finding a red thing among green things, but it's pretty easy. It's easier than finding a 2 among 5's or a T among L's, or anything like that.

Now this sort of guidance comes in two different forms. Or you can think of it as coming in two different forms. There's a bottom-up form that's sort of stimulus driven. And then there's a top-down form that's user driven by your desires. Let me illustrate that with a couple more searches for a T. Tell me as fast as you can whether or not there's a T in the next display. Ready?

**AUDIENCE:** Yes.

**PROFESSOR:** That was pretty crisp. How did you do it? Muhmuh. That's what I thought. Most people probably found their attention sort of automatically grabbed by this one oddball, which conveniently enough turned out to be the T. And so rather than having to search around, your attention was grabbed bottom-up to this item.

Top-down is based on what you know, or what you've been told, or instructions that you've somehow given to yourself. So I'm going to tell you, if there's a T in the next display, it's red. What happened out there? Oh, that was another -- that was also grabbing attention. It works in the auditory domain, too. If we set off an explosion, unsurprisingly, you would notice.

All right, you ready? Is there a T in this next display?

**AUDIENCE:** Yes.

**PROFESSOR:** Whoever said no was another speed-accuracy tradeoff, try to smoke out the professor who had a yes on the last one and therefore must have a no on this one. Look at the display!

Anyway, that's not as easy as the previous one. But if you searched around, you probably noticed, or you may have noticed, that you were searching through the red items. You're not going to bother searching through the black items if you know the T is going to be red. Right?

So let us suppose we did an experiment where the T could be either black or red. And I show you a bunch of displays like this, and I vary the set size, the number of items on the screen. Measure your reaction time. Let's suppose that the slope of that function was 30 milliseconds an item. If that's the case, and half the items are red in this display -- or on average half the items are red -- what's the slope going to look like if I tell you that the T is always red if it's present?

**AUDIENCE:** [INAUDIBLE]

**PROFESSOR:** Less steep. Yeah. Specifically how less steep?

**AUDIENCE:** [INAUDIBLE]

**PROFESSOR:** Very less steep. That's not specific. I want a number.

**AUDIENCE:** 15.

**PROFESSOR:** 15. Good number. Right? If you can eliminate half the items, the effective rate of search is going to be twice as great. So the slope will drop in half. And that's exactly what you get in experiments like this. They work very nicely. If you have only half the items on the screen relevant, subjects behave as though they are only looking through half of the items.

So by now I have answered question two: what escapes the bottleneck of attention? Well,

there are these 12 to 18 basic properties or features of the world that seem to escape the bottleneck. We can study this by measuring reaction time. There are other methods, too, of course, but I was telling you about the reaction time methods.

Oh, I see I put Anne Treisman and feature integration theory on there. Don't worry about the feature integration part. That's simply to allow me to give honor to Anne Treisman who really founded the modern study of visual attention, after having pioneered an awful lot of the auditory things. The auditory demo at the beginning was a classroom version of what's called dichotic listening. Typically what you do is put on a pair of headphones, and you would have one stream of speech in one ear and one stream of speech in the other ear. And you ask questions about, if you're attending through this ear, what can you still pick up through this ear? Anne was doing those things in the late '50s, early '60s. Went on in the '70s and '80s to really invent this field of the study of visual search, and is still doing great stuff, now at Princeton. She was not at Princeton when I was an undergraduate there, but she's there now.

All right, so I answered question three. And question four I answered by saying -- oh, conjunction search. That search for a red vertical thing is a conjunction of two basic features. It's not adequate to know that it's red; it's not adequate to know that it's vertical. The conjunction of those two sources of information is adequate, is what defines the target. And you can use this basic feature information, the basic attributes of the stimulus, to guide your attention around in an intelligent fashion. So that guidance comes in two forms. It can be bottom-up stimulus driven or top-down user driven.

All right, so what is that attention actually doing? Why is it that you need to have this -- what is attention making possible here that wasn't possible before? Oh look, it says that right there. Or what were those features doing before attention shows up?

Well, here is an answer to that. The answer is that you've got all those features. And in fact, early processes in the visual system seem to cut the scene up into what you might consider to be proto-objects. But those features are just sort of bundled together with an object. So before your attention arrives, something like this would be red and green and vertical and horizontal, and it's got points on it, or something.

What attention does is to bind those features together in a way that makes it possible for you to know that the greenness goes with the verticalness here, and the redness goes with the horizontalness. And those points are arranged, the whole thing's arranged into a plus. The

argument is that, OK, I need attention in order to recognize any given individual. Before attention arrives on that individual, that person isn't, you know, a black hole in space. That person is a loose bundle of features. That attention allows me to bind those features together in a way that allows me to understand how they interact, and what that recognizable feature might be.

So oh, there's Kristen. Hey, Kristen, stand up and wave. No really, I was plugging you before. So if you want to do this for \$10 an hour, go find Kristen. So all right up, now we can make fun of Kristen.

So before Kristen arrived -- no, before Kristen arrived, she was not visible. Before I attended to Kristen, there was presumably a proto-Kristen object out there that was a bundle of Kristen bits. Only when I got my attention to her -- even though she'd been visible all along, and I've looked over there a bunch of times -- even though she'd been visible all along, only when I got my attention to her could I bind those features together and make her into a recognizable Kristen.

Let me see if I can illustrate that to you with another demo here. And the way that's going to work is -- OK, so what you want to do in the next slide is to look for red verticals again. You ready? So tell me if you find a red vertical.

**AUDIENCE:** Yes.

**PROFESSOR:** Yeah. In fact, you might have noticed there are two of them. Very easy. What's the point? Well, this is a standard guided search kind of thing. Give me all the red things; give me all the vertical things; look at the intersection of those two sets, and oh lookie, there's two red verticals up there.

Now what I'm going to do is to simply take the horizontal here and jump it up to the middle of the vertical bit That's why this is in this sort of odd arrangement. I'm going to jump it up here. So I'm going to make a plus, like those pluses that we just saw. The reason for doing this is, I'm going to keep all the same pixels on the screen. Right? I'm just going to rearrange where the reds and greens are. And of course I'm going to change the location of the red vertical, because it's really boring if I keep it in the same place. But you're looking for red vertical again. Ready?

**AUDIENCE:** Yes.

**PROFESSOR:** Who said no?

**AUDIENCE:** I said woah.

**PROFESSOR:** Oh, woah. OK. Woah's good, woah is good. Particularly by the time it says find the two red vertical lines. Anyway, you should have found both of them.

Let's let's check intuition here. How many people vote that it was easier to find the red verticals when they were in pluses? How many vote that it was easier when they were ripped apart? That is the correct intuition.

Actually, I think I put the data -- I think I realized earlier than I put half the data on a slide. This is the data for looking for the pluses. Quite steep slopes of about 50 milliseconds an item here and about 140 here. Just looking for the red verticals when they were in the disassociated pluses would have been down here, with a slope of about 10. But I somehow left it off the slide.

Why is this? Why are the pluses so much more difficult? The answer is that before attention arrives on the object, these two pluses are essentially the same thing. They are red and green and vertical and horizontal. And without attention, you just don't know the difference between them. This thing, this square has red and green and vertical and horizontal in it, but it's in two objects. And so since you direct your attention to objects -- to things that are objects; I've got too many "to"s in there -- this is not a problem in the way that these guys are a problem.

In fact anything that you do -- I don't think I brought the demo, but anything that you do to make this less like a single object makes the task easier. So if I was to put a little shadow on here, so that it would look like this thing, the vertical piece was sticking out in front of the horizontal piece, it would get easier. Because now you could direct your attention separately to different planes in depth.

So attention is directed to objects, and objects are available ahead of time as just sort of these loose configurations, constellations of features. Once attention gets there, they get glued together into recognizable objects.

All right. So what happens when you move away from an attended object? That's not a unreasonable question in this framework. So let's see. I need -- Rachel. There's Rachel. I thought I recognized her. All right, I have now recognized Rachel. Limited number of people

who I actually recognize by name in here. And they come to regret it. But anyway, all right.

So she was here all along. I happen to have attended to her and bound Rachel into a recognizable Rachel object. I now, without moving my eyes in fact, I'm attending elsewhere. And somebody's up there, again, my peripheral vision's lousy, but I can see that somebody was moving up there. They waved a piece of white paper a moment ago. The question is, when I moved my attention elsewhere, what happened to Rachel? Did she remain bound, or did she collapse into Rachel bits again?

**AUDIENCE:** [INAUDIBLE]

**PROFESSOR:** What? She collapsed into Rachel bits. How could you tell?

**AUDIENCE:** [INAUDIBLE]

**PROFESSOR:** That's why I was deliberately still looking at her, to avoid the issues of blur.

But the way to do this is not to continue picking on Rachel, but the switch to dancing chickens here. There we have -- you can tell we're back in the realm of my artwork. Oh, I like this, with the chickens on three screens. This is so good. Anyway, I like those a lot.

Now, so you know you know what you're looking at here. You're looking at a bunch of chickens, right? And they're doing this little leggy thing. You would think that, having recognized that there's a bunch of chickens there who are doing this little dance, that if one of those chickens fell apart into chicken bits, that you would notice, right? Seems reasonable. How many of you noticed?

Ooh, ooh, very slow group here. It should be -- how many chickens are there here, about 20? It should be about one in 20 of you happen to be -- you have all seen that already. So one of these chickens fell apart. Well, if you think, quite apart from the fact that the artwork is a little lame, the implications are non-lame. The implication is, all right, I'm looking at you guys. I think I'm looking at a bunch of humanoid life forms. They're moving a little bit, stuff like that. And you would think that if one of you just went to pieces here, that I would notice. The data strongly suggests that that's not the case. That I would eventually notice, as my attention roves around the room, if it turned out that, oh my god, not only has that person not dozed off, but her head fell off, I would notice that and react with according shock and amusement.

The way this experiment is actually done is not with the cute little dancing bits. You'd be looking at a screen like this, and you'd hear, beep. And the question would be, is there a destroyed chicken? Yeah, it's there, right? Beep. Beep.

**AUDIENCE:** Yes.

**PROFESSOR:** Beep.

AUDIENCE: No.

PROFESSOR: Beep.

**AUDIENCE:** Yes.

**PROFESSOR:** And

so on. You can do it. It's not a difficult task at all, particularly with a few big chickens. But you have to search. You have to search through the chickens each time. And you're no better with a display that's got the same fixed number of chickens up there all the time, compared to a display which has *de novo* chickens popping up out of nothingness each time.

Oh, the feet are moving around. for the demo, why are the feet doing this little chicken dance? Remember I said that motion is one of these things you can pick up automatically? If you don't have something like the little moving feet, then when you have a chicken fall apart - boink - the movement of the contour, compared to all the ones that aren't moving at all tips you off that there's something there. And that tells you that motion's important, but it doesn't tell you the interesting fact that you're not aware when an otherwise coherent object falls to bits.

By the way, it turns out you're also not aware when previously incoherent material coheres into a chicken. We did the classic chicken soup experiment. We had a screen full of chicken bits like this, and you heard beep, and you had to figure out whether or not there was now a chicken present. And you had to search for that, too. So chickens emerging from the chicken soup, which you might think would be striking, don't turn out to be striking either.

All right. Well the chickens are kind of ugly and complicated. How bad is this problem?

So let's get basic here. No more trying to fool you. Well, of course I'm trying to fool you. No more dancing around chickens, and then oh, did you see -- after the fact I ask you whether you saw something that fell apart. These are what? Red and green dots. If you weren't sure

about that, it says so at the top. All I'm going to do is, I'm going to cue one dot -- I don't care about any of the other dots. All I want to know is, did that one dot change color? Say yes or no. Whoops. Where'd it go?

**AUDIENCE:** [INAUDIBLE]

**PROFESSOR:** Well, the answer turns out to be no. This is such a great exercise in applied statistics, right? How many -- he can't really be -- he said the last one, so no.

**AUDIENCE:** Yes.

**PROFESSOR:** Oh yeah, but he can't possibly be doing three in a row, right?

**AUDIENCE:** No.

**PROFESSOR:** That does turn out to be a no. Look, you can hear people going both ways. People are terrible at this. They're just barely above chance. And the barely above chance is consistent with them sort of sitting on two or three dots. Because you're not just doing a couple of these, you're doing hundreds of these, for \$10 an hour. So you can sit on a couple of them and say, if I get really lucky and he cues the one I'm looking at, I'm going to get this right. And if he doesn't, I'm clueless. I mean, it's red and green. It doesn't get more basic than that. Yup?

**AUDIENCE:** I have a question. Do people's reaction times change? Because red and green, they have the same after color, or afterimage.

**PROFESSOR:** They'd better not have the same after -- they have the opposite. Yes.

**AUDIENCE:** No, no. But the opposite of red is green, and the opposite of green is red. So if you do yellow and blue or something else --?

**PROFESSOR:** Well, yellow and blue are also opposite in the same sense. But it doesn't matter. The color does not matter. In fact, we can do another one with different colors. Look at this new. More cool colors.

But maybe I was just being nasty to you. Because there were a lot of dots up there for you to choose among. So I'll tell you the relevant dots. What I'm going to do here is I'll ask you about the color of specific dots. I won't change them. I'll just put them up there and ask you about particular dots. And what I want you to do is tell me the color. So if I say, what color is that dot, the answer is --

**AUDIENCE:** Purple.

**PROFESSOR:** Good. If I happen to cover it up with a black blob, tell me what color it was before I covered it up. OK? Ready? All right, here we go. You'll see how this works. Where'd it go? There we go.

**AUDIENCE:** Red. Yellow. Blue Green. Green.

**PROFESSOR:** Good. See, you're not -- I put this in because at this point you might be sitting there saying, I'm so hopeless! And I wanted to prove to you that you're not. Well, you are, but not that hopeless. All right, ready?

**AUDIENCE:** Purple. Red. Blue. Yellow. Red. Green. Yellow. PROFESSOR: Ooh,

a few people actually got it. A bunch of people did the, urp. But yes indeed, that was yellow. It was cued before, so we know you paid attention to it. But it was cued about five items back. And so you'd paid attention to it. It didn't take much binding to say, that's yellow. You'd already done all the work on it. Five blobs later, by the time your attention is somewhere else -- it wasn't invisible during that time, right? You don't really know what it is.

All right, try this.

**AUDIENCE:** Red. Green. Red. [MURMURING]

**PROFESSOR:** A couple of people caught on. He changed it. This is what happened here. Whoops, not that way. Go back. OK.

So this makes a useful and important point. So, red. While your attention was diverted, I changed the color. Why is that important? What that tells you, with a very basic sort of stimulus, is that the following ought to be true: I attend to Rachel, I attend away. While I've attended away, Rachel is replaced by a kangaroo. I am now asked, what was there? I say, you know, it was Rachel. The fact that, even though, you know, still visible in the visual field and everything, until I attend back, I would simply not know that something had changed there.

So in fact, if you're worried that -- the trick here, obviously, since there are 300 of you or so, you want to convince me that you're paying attention in this class, you draw my attention early in the class, and then you subtly sneak out. And presumably I think you're here attending the whole time. Because how often do I get back to each individual person? Well, actually, it's not that good. Because at 30 to 40 people per second, I can get back to you pretty quickly. So

forget it.

But don't forget the basic point here, which is that you're only aware, you're only updating your knowledge about the world, through this narrow bottleneck of attention, for the current object of attention. Everything else, you're basically working on your hypothesis based on the last time you checked up on it.

So here is actually what the data for an experiment like this look like. So if you didn't pay attention to the colored dot, right? If I never asked about it at all. Here's chance. 50% in this particular experiment. Because this is a two color version of it. Is it red or is it green? You've got about a 50-50 chance of getting it. You do a little bit better than that. If it was recently cued -- if I just asked you whether it was red or green -- you do pretty well. But as soon as it's four items ago, or eight or 12 ago, you're back to being pretty pathetic. So you don't keep a good record of this. You're only updating in the current object of attention.

This suggests that your memory is pretty small here. We'll talk about memory more extensively later. But let me illustrate that your memory is actually fairly small. Here what we're going to do is, I want you to remember these guys. Got them? OK, take them away. Are these the same?

**AUDIENCE:** No.

**PROFESSOR:** OK, well your memory isn't that small. That's good. How about these guys?

**AUDIENCE:** No.

**PROFESSOR:** No, no, no. This is a new set.

[LAUGHTER]

Ready? Boink.

**AUDIENCE:** Yes.

**PROFESSOR:** Whoops. Sadly, I can't remember. Remember these.

**AUDIENCE:** [INAUDIBLE] They look the same, don't they?

**AUDIENCE:** Yes.

**PROFESSOR:** OK. So, well. How about these?

**AUDIENCE:** No. No, this is a new set.

**AUDIENCE:** Yes.

**PROFESSOR:** Yes, something changed. So this time I transposed the red and the yellow. That's a little more difficult, because I didn't introduce a new color. How about this?

**AUDIENCE:** Yes. Yes.

**PROFESSOR:** People aren't quite sure. The answer is that the capacity of this sort of memory is about four. And some of you will have gotten the fact that there was another transposition, right? Of the yellow and the green? Whoops. The yellows and the greens. Yeah. The yellow and green guys are -- whoops! -- switching there. Some of you will have gotten and some of you will have not gotten it, because some of you were sitting on the right four and some of you were sitting on the wrong four. But it's only about four.

Four what? It turns out to be four objects. Look at this. Tell me if anything changes. So here we have at least color, shape, and orientation going on.

**AUDIENCE:** Yes.

**PROFESSOR:** Yeah, most people will know here that the red thing flipped from pointing up to pointing down. That would seem to suggest that you can keep track of 12 things, because there are four colors, four shapes, and four orientations. But if I spread those out across 12 objects, you'd be very bad. It's that you can keep track of about four objects. You can keep track of multiple features of each of those objects, but it's only about four objects that you can keep track of.

Now let's see. How are we doing in question land? OK, so the answer to question six, at least to the first part about it, is that the objects don't seem to stay bound. That you need to continuously update the visual world in order to have some idea of what its current state is, and that you're only updating the current object of attention.

After a brief break, we will establish what the Sistine Chapel has to tell us about that fact. But those of you who wish may study this image for the next couple of minutes or so. And everybody else can just sort of stretch. And then we'll come back. While I apologize to Rachel for picking on her. You're not traumatized for life or anything? OK, good.

[? [CROWD NOISES] ?] AUDIENCE:

Have you seen this video they have where it's a bunch of people bouncing balls to each other?

**PROFESSOR:** Yeah. That's now gotten to be so common that I'm not using it.

[PRIVATE CONVERSATION]

**AUDIENCE:** Do you know who did that?

**PROFESSOR:** Yes, Dan Simons. Then at Harvard, now at University of Illinois. I will describe a different Dan Simons experiment in a minute. OK, let's get back together here.

All right, to briefly review. the story I have been developing thus far is that even though you are looking at this scene from the Sistine Chapel, and this is the expulsion from Eden, there's Adam and Eve, and this very cool snake. And there's Adam and Eve getting chucked out, with the angel poking them in the head and stuff like that. Even though you are looking at this, you know what you're looking at, that at any given moment the only thing that's really coming through from the world to recognition is whatever is currently being fed through the bottleneck, the current object of attention. And that maybe three or four objects, the recent status of three or four objects is currently held in this visual short term memory.

The implication here is that I could change this scene and you wouldn't notice. So let's find out. What did I change?

**AUDIENCE:** [INAUDIBLE]

**PROFESSOR:** I need a hand or two here. Yeah, sure, what?

**AUDIENCE:** [INAUDIBLE]

**PROFESSOR:** Oh, the fig leaf. The fig leaf, yes. The originator of change blindness, which is what this phenomenon is known as, is Ron Rensink, now at the University of British Columbia. And he refers to what he calls "areas of interest." If you change something that people are paying attention to, they notice that. But of course I knew that. And so how many people picked up the other three changes?

**AUDIENCE:** [INAUDIBLE]

**PROFESSOR:** Oh, some. We have a few people picked -- what did you get? I can't hear you.

**AUDIENCE:** [INAUDIBLE]

**PROFESSOR:** The stick thing. And what? Sorry?

**AUDIENCE:** [INAUDIBLE] Something showed up at the top that's funny. The stick thing moved, and something showed up at the top that's funny. So now with that information, we can go -- whoops.

**AUDIENCE:** Right there.

**PROFESSOR:** You got the stick. See, the reason for the blank is the same as the moving chicken legs, which is that you don't want to have motion transience giving stuff away. But if you have motion transience -- do do do do do -- you would think that if you were in the Garden of Eden and the branches were moving from tree to tree, or for that matter Eve's foot was moving to Adam's body, you would notice. But if you're not attending to it, you don't notice.

So this is part of a large set of phenomena that come under the general heading of change blindness. At the break, somebody was reminding me of one that you may have seen because it's made it onto Nova and things like that. Done by Dan Simons, where you're watching people apparently play a weird game of basketball in front of the elevators, it turns out in the psych department at Harvard. And while you're doing that, a guy in a gorilla suit -- actually, Stan reminded me, a woman in a gorilla suit. It's hard to tell; she's in a gorilla suit -- walks in, walks into the middle of the game, waves, walks out. And then afterwards you ask -- oh, and you're doing a demanding task. You're supposed to count how passes there are, or something like that.

And you're asked, did you notice the person in the gorilla suit? Well, first you're asked, did you notice anything weird? Eh, no, very boring. Notice the person in the gorilla suit? Yeah, right. What person in a gorilla suit? Show them the video again. Oh my --

Another great Dan Simons experiment was done when he was at Cornell, actually. You're on the street in Ithaca, New York, and some guy walks up to you and asks you for directions. Actually it's Dan Simons walks up to you and asks you for directions. And so, since you are a nice person, you start giving Dan directions. Now you're standing there on the street and, who knows why, but these two guys with a door are carry a door down the street. And they walk between you and Dan. Which is kind of rude. And then they're off down the street somewhere.

And the question is, do you continue to give directions once you see Dan again? Of course, the real question is, did you notice that when the door went by, Dan Simons ducked down and left with the door, and his then-student Dan Levin popped up in his place? And it's a different guy.

50% of the subjects in this study kept talking. A surprisingly large number of these, on being debriefed later, claimed to have noticed a change. Which is a little strange, right? I'm talking to this guy and the door, and now I'm talking -- there's another guy here, but what the heck? He probably wants the answer to the same question. I don't know what that's about.

But the important finding there is that 50% of the people behaved as though they hadn't noticed the change from one person to another, who they were talking to. What's going on here? Now people aren't completely stupid. The experiment has not been done, but we kind of absolutely know that if I'm talking to Dan Simons, short white guy, and now the door goes through, and a tall black woman is standing there -- hm, you know? Probably that's, again, the sort of front-end stuff that people tend to pick up on.

But if what you're doing is, I don't know this guy, but I've got a sort of a model of this guy. I'm talking to kind of a short, white guy person. And da da da, I'm still talking to a short, white guy person. It's not the same one, apparently, but that turns out not to be a problem.

This has given rise to a notion that perception is what Kevin O'Regan has called a grand illusion. That the only thing that you actually see is the current object of attention. That I think I'm seeing all of you, but all I'm really doing at the moment is paying attention to the guy with the grey stripe on up there. Yeah, there he is. And now that he's riveted my attention by waving at me, the rest of you are just not there. You are just some sort of grand illusion floating around in my head.

Now in some sense, that's correct. That what you are seeing is a creation -- the burden of the lecture next time will be to say that you're always seeing a theory about the world. You're not seeing the world directly. You're always making an interpretation, your best guess about what the stimulus means. And all the evidence I've been showing you for the past hour or so suggests that you're only updating that theory through this very narrow bottleneck. So in some sense, you are only seeing this creation of your mind, and the only object that you are currently updating is the one that you are currently attending to.

But to call the whole thing an illusion, it seems to me, misses an important aspect of the

experience. Let's take a very old example. The French philosopher of the, I'm thinking early 18th century, whose name I will now proceed to misspell. Does that look -- any good philosopher sorts? That about right? Condillac, I believe is how you pronounce it properly. But anyway, Condillac wrote a number of very interesting things about sensation and perception. He's most famous for his statue. His statue that he proposed as an entity with no senses at all. And he asked what would the mental life of this statue be? And argued that, in the absence of any sensory input, there would be no mental life.

And now, he said, let's imagine opening up, I think he opens up the statue's nostrils. And argues that the entire mental life of this statue is now the smell. Whatever, I think he waves a rose under it or something like that.

But a little further on he has a different example where he says imagine, you're in a dark -- a dark chateau, I believe. And it's completely pitch black, because of these heavy curtains. And it's morning, and you throw open the curtains. If it were the case -- this is not what he's saying, but if it were the case that all of vision was nothing but a grand illusion, you only saw the spotlight of attention, this one thing that you're attending to at any one moment, your experience of this brand new scene ought to be like sort of a weird paint brush. Initially, I don't see nothin'. Because I haven't attended to anything. Now I attend to an object. And now this person, object, is the only thing in the scene. And boom, boom, boom. And I slowly fill you in.

That's not the impression you get ever when you see a new scene. You may not know what you're looking at, but you see something everywhere instantly. And the grand illusion thing misses the fact that you're somehow sensing something about the entire visual field all at once.

Let me offer a way of understanding that that will then tie back to the visual physiology that I was talking about in the last lecture. Here's the idea. Early in your visual system, you've got the processes that, sort of a big river of information that tells you about those 12 to 18 features or attributes that you can get out -- these are eyes. This is my drawing again. So from your eyes, you've got this big flow of information up into your brain. And at some point, it hits this bottleneck that's taken care of by attention. Object recognition, the ability to tell that that's a branch, that that's a snake, and so on, only one object at a time can go in and come out and rise to the level of some sort of perceptual awareness, populating your visual experience.

And that bottleneck is guided by these collection of basic features that you've got. If you know

you're looking for red stuff, you set these settings for red. And maybe vertical, and big and moving and so on. And so you can regulate what gets through here. And only the one thing at any one time is getting up into there. And so the current object of attention gets to rise to awareness, and you know what you're looking at.

That's the story that I've told you to this point. That's the story that gives rise to the notion that everything else in the visual field is some sort of an illusion. But look, when I was doing that red and green dot thing, it wasn't that you didn't see the other red and green dots. They were there. You just somehow had a very impoverished ability to tell me anything about them. And a way to think about that is to propose that there's another pathway, another big fat river of information about, say, these 12 to 18 attributes, that isn't limited by the bottleneck. But that it doesn't let you -- it's not a cheat. This doesn't now let you go and recognize objects everywhere all at once. It can only do a few things. It can sort of give you the statistics of the world.

You know, I'm looking out at you guys and I'm seeing a sort of texture of people amongst purple. And that sort of impression of purpleness, of a tilted plane, is the sort of thing that you might get out of this big, broad, unrestricted, nonselective, as it's labeled on there, pathway.

There's evidence that you can get a little bit of semantic information. Semantic means the meaning, when you're talking about language, it's the meaning of the utterance, let's say. When you're talking about vision, it's the meaning of the stimulus. So I might get the notion that I'm in an enclosed space. This pathway by itself is not going to tell me what enclosed space I'm in. But I'm in a space. There's a tilted surface there. And so on.

But this is going to give me, that broad pathway is going to give me the feeling that there's something happening everywhere. And this pathway is going to tell me what's happening specifically here, now. And between the two of them, I can build up an idea in my head of, oh, I'm in 10-250. I'm talking to this bunch of people, some of whom I know by name, some of whom I recognize because they've been here before, and so on. And I can keep updating that 20, 30 times a second through this pathway. And I can keep experiencing something, that sort of wallpaper of the world effect, through this other pathway.

Now that ties back, it might tie back to things that we talked about before. If you remember the idea that you can broadly cut visual processing, visual cortical processing, into two big pathways, a what and a where pathway. A what pathway going down into the temporal lobe,

and a where pathway going up into the parietal lobe. This selective pathway, this thing that only does one object at a time, would then be mapped onto the what pathway. What am I looking at, what am I attending to right now?

If you were to lesion that, if you were to lesion it, or you were to have damage to the temporal lobe of your brain, you might well end up with an agnosia. That's not a term that ended up on the handout, so you want to write that one down. An agnosia is a failure to know, if you like. To know what something is.

So an agnosic, if you have a person with a fairly global agnosia, visual agnosia, they would be able to say, yeah, I'm looking at a bunch of objects here, but I don't know what they are. Here's this object. It's sort of orange. It's got orange and brown and white blobs on it. And it's got this very long part, and there are these four pointy things coming off the bottom of it. I've got no idea what that is; maybe it's furniture of some sort. You'd look at it and say, that's a giraffe. An agnosic would be able to tell you about it, but not know that it was a giraffe.

Smaller lesions produce rather specific agnosias. There are reports in the literature of agnosias specific to, say, fruits and vegetables. More common is a form of agnosia called prosopagnosia, which is a specific inability to recognize faces. You know that it's a face, it's got two eyes, it's got a nose and mouth. You don't know who it is. Small lesions down in that pathway can produce that sort of damage.

That would suggest, then, that the other pathway ought to be mapped onto the where pathway. And if you get bilateral damage, for instance, to the parietal lobe, you can end up with a disorder known as Balint's syndrome -- might as well write the word down here. Named after Balint -- that has as one of its properties what's called a simultagnosia. This is a situation where you can recognize an object if you can get your attention on it. But that's the only thing you can respond to, in some sense. It is as if the grand illusion theory was really right, that you can only see the current object of attention.

So you do something like this with a simultagnosic, say, what's that? Draw his attention to it. That's a book. OK, what else have we got here? OK, what's that? That's a cell phone. What's that? That's a cell phone. Anything else? No. What's that? That's a book. What's that? That's a book. Anything else? No. So one object at a time. As if the where of the world had disappeared.

If you get damage -- we'll talk about this more later in the term -- but if you get damage to the

parietal lobe on one side, particularly on the right side, what you can end up with is a disorder known as neglect. It comes in a variety of flavors, again depending on the particular lesion. But the characteristic is, you ignore the contralateral, the other side. Now that can be the other side of space, so that if I'm a patient with a right hemisphere parietal lesion and I'm looking at MIT volleyball here, everything in the left visual field, I would simply ignore. I would behave as though it did not exist. If I took away everything else and put a stimulus in my left visual field, I could show that the patient could still see it. But with a full visual field, he behaves as though there's nothing there at all.

Patients with neglect will do weird things, like -- they're in the hospital, typically, because they've had a stroke. You give them their dinner. They eat everything on the right side of the plate and leave everything on the left side of the plate. Why? Because they didn't like the mashed potatoes? No. If you rotate the plate, they'll eat the stuff on the other side of the plate. It's as if it just didn't exist, in some fashion.

Now, you'll remember the parietal lobe is also where you get the representation of the body surface, and stuff like that. So neglect patients can also be patients who neglect one half of their body, and deny that part of their body is theirs. This is a little easier to understand if you figure that the stroke might well have also knocked out the ability to control that side of your body. So a stroke on the right might leave you paralyzed on the left.

But you can end up with situations like one described, I think, by Oliver Sacks in one of his books, where a patient is saying, "This is a cheap hospital. This is a really cheap, lousy hospital." How do you know it's a cheap, lousy hospital? "Because they're doubling up on beds." What you mean they're doubling up on beds? He says, "Look at that leg. That's not my leg." So you can get, this is somebody looking at their own leg and denying that that leg belongs to them. That's another aspect of neglect.

OK, what I'm going to do next time is to talk about the way in which you make hypotheses about the world.