Brain and Mind

May 14, 2004

Nancy Kanwisher, PhD
fMRI Investigations of Human Extrastriate Cortex: People, Places, and Things

Introduction by David Cohen

I'm David Cohen and it's my pleasure to welcome you to the third and final session of the "Brain [and] Mind" symposium. Having been a university administrator for the past 17 years, I'm not at all sure that I still have the neuroscience bona fides to deserve a spot on this illustrious program, but I appreciably accepted the assignment as a generously nostalgic gesture that recognized the time some years ago when I worked for an honest living as a neuroscientist.

Now this session is entitled "Biology of Mind," and in recognition of Richard Axel's interesting metaphor yesterday morning, I'd like to take a moderator's prerogative and re-title the session as "The Ghostbusters." Now in this session, complex cognitive questions in the phenomenon of consciousness are addressed as biological-search challenges. And as Eric Kandel and I were commenting to each other yesterday, a session of this nature is something that we could have only fantasized about when we began in the brain sciences some forty years ago, yet the questions that will be discussed today are the kinds of questions that initially attracted many of us to the brain sciences, even though the prospects of any kind of experimental answers at the time were exceedingly dim.

That we're here today in a session of this sort really is testimony to the explosive development of brain research over the past four decades. Perhaps a surrogate measure of this has been the development of the Society for Neuroscience. In its first meeting 33 years ago, there were two to three hundred attendees. At its meeting this last year there were 30,000, and that is a growth over just the last 33 years of the existence of the society. Neuroscience is clearly one of the most vital and exciting fields in contemporary science, and that's a fact that even the physicists would concede.

Now we have four distinguished speakers in this session who continue the impressive array of contributors to this rather remarkable event. And I might just make an idle observation, if you look at the program and look at the degrees that are held by the participants in each session, it appears that if you were trained as an MD, you work on the brain; if you were trained as a PhD, you work on the mind. And I'm wondering if there isn't some barrier in training of physicians that simply does not permit them to become ghostbusters.

Now the range in this session is extraordinary—it's from inferring the mechanisms of complex neurofunction from the activity of single neurons and functional neuroimaging, to biological, theoretical, and philosophical considerations of consciousness. It does indeed promise to be a provocative morning.

Our first speaker is Dr. Nancy Kanwisher, who's a professor in the Department of Brain and Cognitive Sciences at MIT and an investigator at the McGovern Institute for Brain Research. She received both her undergraduate and doctoral degrees from MIT, and after serving on the faculties of UCLA and at Harvard for some years returned to MIT to join its faculty in 1997. Now if you'll permit a brief apropos of nothing, whenever I hear MIT and Harvard concatenated as in Nancy's resume, it evokes a story I heard as an undergraduate at Harvard, and it's about a student who wields a cart full of groceries—this is in the Central Square Supermarket in Cambridge—he wields a fully loaded cart up to the register that says "Ten items or less," and the cashier looks at him in disgust and says, "Look, you're either from Harvard and you can't count or you're from MIT and you can't read." Since Dr. Kanwisher has been at both she can both count and read.

Nancy has been applying the noninvasive techniques of functional magnetic resonance imaging and magneto-electroencephalography, and if memory serves, I believe magneto-electroencephalography was developed at MIT some years ago, and I think it was by a namesake of mine, David Cohen, wasn't it? She's been contributing fascinating, compelling experimental results that identify and characterize regions of the human brain that mediate the visual perception of faces, bodies, places, and objects. She's been recognized by various awards including a MacArthur Foundation Fellowship, a Troland Research Award from the National Academy, and MacVicar Faculty Fellow Award from MIT. It's indeed a pleasure to welcome Dr. Kanwisher at Columbia, and I look forward to hearing about her fascinating findings on "fMRI Investigations of Human Extrastraite Cortex: People, Places and Things." Nancy.

MRI Studies of the Visual Cortex

Nancy Kanwisher: Thanks, David, and thanks so much to the organizers for inviting me. I've learned a lot and had a very good time so far.

So the question that's motivated my research for a long time is this, within a fraction of a second of viewing a completely novel and unpredictable image of a complex scene we human perceivers have already extracted the gist of the scene, and if we know the specific individuals or places or objects in that scene we've also identified those individuals. If that doesn't strike you as remarkable, reflect on this:

Despite decades of effort, no current computer vision system is even close to what all of us can do in a fraction of a second. So the question is, How do we do it?

Now if I could answer that question I'd be as famous as Eric Kandel. I can't answer that question, I'm not that famous, but what I will do is dance around the edges of this question and show you the approach we've taken in my lab, which is this: Our general strategy is to try to get some clues into how vision works by looking at the functional organization of the parts of the brain that do it. So that's the stuff back here. Close to half the cerebral cortex is devoted to vision, construed broadly, including visually-guided action and thinking visually and things like that.

We are enormously visual animals. So how is all this expansive brain organized? Well we know a lot about the organization of visual cortex in macaques from decades of research where you can stick electrodes in and record from actual neurons, the gold standard in neuroscience. A great deal has been discovered, and so all of these visual areas up there have been characterized in great detail in monkeys, and even more impressively, the connectivity, the wiring diagram of the visual cortex has been characterized in great detail in macaques.

In contrast, until extremely recently, almost nothing was known about the organization of human visual cortex. So about ten years ago maybe two or three visual areas had been identified in the human brain. And then functional MRI came along. Functional MRI is just like regular MRI scans that you might get in a hospital for a knee or a kidney or a liver, except for two important things: One is that the images can be taken very quickly, upward of ten images per second, and the second is that the images are arranged in a special way such that they show, very indirectly but very well, a measure of neural activity. It works by way of blood flow. If a bunch of neurons in a focal region of cortex become active, that's metabolically expensive, more blood needs to flow to that region of the brain, and MRI looks at that change in blood flow. Also importantly, in contrast to other methods such a PET, which involve a radiation dose to the subject, as far as we know, functional MRI is completely safe. I certainly hope it's safe because I have very close to the world record of total amount of time spent inside these machines.

So the question that we've asked with functional MRI is very analogous to the question that Richard Axel talked about yesterday. He asked whether different parts of the olfactory bulb respond when animals perceive different odors. We ask whether different parts of the visual cortex become active when people look at different kinds of objects. So getting MRI machines to work is very fancy, I get zero credit for that, I had nothing to do with it. Using them to ask this question is really pretty simple. We stick subjects in the MRI scanner and we show them movies. So here's a diagram of the kind of movie that we show our subjects in one experiment. Time is going from left to right, each vertical stripe represents a period of 16 seconds during which the subject viewed a particular thing. So starting over here subjects viewed a fixation point for the first 16 seconds, then they saw a whole bunch of pictures of flowers flashing on the screen at a fairly rapid rate, then a whole bunch of animals, chairs, scenes, cars, faces, and so forth. We just tell the subjects to hold still, look at the pictures, and while they're doing that we're scanning their brain.

A Face-Selective Cortical Region

We then take all the data from the functional brain scans conducted during this kind of experiment and we ask statistically whether there are any voxels in the brain that respond significantly more strongly during the periods when subjects are looking at one of these kinds of images, compared to when they're looking at all the others. The kind of thing we see is shown here. Here's one slice through a subject's brain, this is a horizontal slice, like this. Left and right are flipped here and on all my images, that's the radiologists' fault, they started that convention. And what you see over here is a little teeny region in this subject's right fusiform gyrus where the statistics are telling us that the signal is much stronger; that is, neural activity is stronger in that part of the brain when this subject was viewing faces than when they were viewing anything else. Now you shouldn't believe these nice activation images if that's all you see, there are a million ways to cheat and get nice images like this that are entirely spurious. The reason to believe this image is if people show you the raw data. So here's the raw data coming from that part of the brain during this scan. And you can see just eyeballing it that this little teeny bit of this person's brain produced a much stronger signal during the two periods when the faces were presented than during any of the other periods. You can see the signal's higher when they were looking at all these different kinds of objects than when they were staring at a dot, so it's not absolutely silent for nonfaces, but it's a whole lot stronger for faces than for anything else. So this is what we call a face-selective region of cortex in the human brain, and what does that mean? Well, it could mean all kinds of things. From the data I've shown you so far this could be just due to, for example, greater attention that people may pay to faces. We're social primates, we care about faces, maybe we just pay more attention to faces than anything else. It could have to do with simple low-level properties of the stimulus that haven't yet been controlled. There are any number of accounts of what this could mean, and you can't make a strong claim about face selectivity until you spend a long time testing all those other accounts.

So the way we've done that is what I call a region-of-interest approach. What we do is first identify that region individually in each subject with a scan much like the one I just described. We find the bits in that subject's brain that respond selectively to faces and we say okay, there's the face area, then we run a new experiment and we ask how that region responds under some new conditions. And we can quantify the magnitude of the response in that area to the new conditions of interest. And this solves a whole bunch of technical problems with functional imaging. One is that those face areas, although they can be found in virtually every normal subject, there's just a certain amount of anatomical variability, you don't know exactly where it's going to land in the fusiform gyrus. So this way we don't worry about that. We find it in each subject and we look there. It also gives us a huge increase in statistical power, and has other advantages.

So having done that for quite a few years on the fusiform face area (FFA), here's a summary slide of some of our results. So on the left there are four different subjects' face areas, here and—I can't see from this angle—here, and that's me, I have this posterior annex to my face area, and here's another subject. So you see there's some variability here. So over here you see in each panel an example of a kind of stimulus, and a magnitude of response. A 1- to 2-percent signal increase is a reasonably strong response for functional MRI, and what you see is lots of different kinds of face images produce a strong response in this region.

Let me just point out a few things. I don't know if you guys can see the pointer. Can you guys see this image? Raise your hand if you can tell what that image is of. Very few hands. Okay. Raise your hand if you can tell what this image is of from the top row in the middle. Okay. So you guys behave like our subjects do. Most of the time they see the face in this image. If you didn't see it there's nothing wrong with you, it's a little tricky. It's a profile of a face like this with the nose sticking out on the left and the eyes and the mouth. The interesting thing is this is the identical image, it's just upside down. So what this enables us to do is respond to all the kind of serious vision geeks who want to tell us your face area, it just likes edges like this or it likes contrast or it likes this spatial frequency. We can say no, all those things are pretty much identical in these two images, and here where subjects usually see a face we get a much stronger response than here when they don't usually see a face.

You may be wondering, why does this region respond to cat faces? I'm looking in human brains and we humans care about human faces. Well we wondered about that too, and then we thought well you know, cat faces look a fair amount like human faces, in fact some animal faces look very much like particular people. Okay, those were totally irrelevant, I just like to show them.

The FFA and Facial Recognition

Okay, so to get more serious about it, all I've shown you so far is that this little bit of cortex is pretty selective for faces. Doesn't tell us what it does with faces, that's what I'd really like to know. So what I'm going to do is describe one experiment done with Kalanit Grill-Spector, who's now a professor at Stanford. She was then a postdoc in my lab, and we asked, well, our overall strategy to try to figure out what this region does is to give subjects a difficult perceptual task in the scanner, make it really difficult so that subjects make some mistakes, and then collect MRI data while they're doing the task and bin our MRI data by what the subject reports, trial by trial. Then we can look for correlations between the magnitude of the brain signal in different parts of the brain and the subject's response in order to get a closer causal connection between the MRI data and behavior.

So here's the experiment that we did: We chose famous faces like Harrison Ford and various other actors and politicians that our subjects knew the faces of. We presented them as a timeline from left to right, we presented them very briefly and followed them by visual junk, mask here, so that subjects would make some mistakes, and in this experiment we asked subjects for each face to give us one of three response options. Either it was the particular target face that they were looking for, in this case Harrison Ford, or it was some other guy, or it was nothing, just visual junk. So they have to make one of those three responses on each trial. We can then look at the magnitude of response in the fusiform face area on just the Harrison Ford trials as a function of what the subject tells us they see in those stimuli. So the stimuli are more or less the same, they're all pictures of Harrison Ford, and all that differs is what the subject reports.

So what you see here is—this is a timeline here—this shows you the typical MRI response. It's delayed, this peak is about six seconds after the onset of the stimulus, that's because blood flow regulation lags well after neural activity. But what you see here in the different height of the peaks is in red where the subjects correctly said yes, that's Harrison, in blue are the trials where they said some other guy, and in black are the trials where they said I don't see anything. Remember it's always Harrison Ford. So what we want to say is that this difference between the correct identification and the simple detection of faces, that reflects a correlation between this MRI signal and face identification.

This additional effect down here shows that this region is also correlated with simply detecting the presence of a face, even when you don't correctly identify it. And this pattern of response was shown in all the subjects here. Of course I'm subject five who has generally the best data, but everybody else showed the same thing.

So we want to use these data to argue that fusiform face area's involved in both detecting the presence of a face and identifying individual faces. But you might be wondering when you get a response correct, when you detect what you're looking for, maybe there's like a yippee! effect, maybe there's a party in your brain and you just kind of . . . everything turns on more. Also we want to know how specific is this response to faces?

So in another set of experiments we did the analogous experiment on a bunch of different kinds of objects, in this case guitars. So here subjects had to say is that an electric guitar, another kind of guitar, or nothing. We analyzed the data in the same way, and here are the results for the guitars, broken down by individual subjects. You can see there's really no correlation between the response in the fusiform face area and successful identification of guitars. We found the same thing for tasks on cars and flowers and so forth. So this shows the specificity of the fusiform face area in detecting and recognizing faces, but it's also of interest to ask about the regional specificity in the brain. What's the rest of the brain doing? Maybe the whole brain is more active when you detect or recognize a face.

So here I'm showing you from one subject, as the physiologists say a typical subject, most of our subjects look like this. This is a bottom view of the brain, and what's shown in black are the outlines of the fusiform face area defined in our usual way, and what's shown in yellow and orange are the regions that showed a significant correlation with face identification in that subject. And you can see there's a pretty good match. And there isn't much else, you know, spills out a little bit over the back right there, but mostly it's that region that's correlated with face identification, not the whole brain. And when we do the same kind of analysis on this same subject looking at other kinds of objects, you see that most of the regions that are correlated with correct identification of other kinds of objects are outside the fusiform face area. So that says that other objects are primarily categorized and identified using other bits of cortex, nearby but distinct from the fusiform face area.

Okay, so all of that suggests that the face area, what it's doing for us is both detecting and identifying faces, and that it plays little role in detecting or identifying other kinds of stimuli. But as Nora Volkow said yesterday, these kind of data are just associations, they don't really show a causal connection. A correlation with behavior is kind of as good as we can get with functional MRI, but the real test is if disruption of this region disrupts performance on the task.

So these are some data—not mine—this is from a paper by Wada and Yamamoto and they report the case of a gentleman who had an unusually small stroke in his fusiform gyrus. And I've outlined it here in the right, fusiform gyrus right there. This gentleman had a remarkably specific case of prosopagnosia, that's a severe deficit in recognizing faces. This guy can't recognize himself in the mirror, he can't recognize family members, he's absolutely unable to recognize faces, and yet interestingly, like many other cases that have been reported previously, he's very good, in fact they claim absolutely normal, at recognizing objects. Okay, so his deficit is very specific, and although I've been trying to get an MRI scan on him to make sure that his fusiform face area was, in fact, affected by this lesion. For various technical reasons we can't scan him, but I'll just show you my fusiform face area on analogous slices. What you see is, here's this guy's lesion and there's my FFA. They're really similar, so I think it's a good bet that this guy's stroke took out just his fusiform face area and produced a very selective deficit in face recognition. So that suggests that the FFA is not only correlated with face recognition but necessary for face recognition.

Other Brain Areas in Facial Recognition

Okay, so enough on the FFA. I'll just briefly mention two other regions we've discovered since then. I should say the FFA has been fun to work on but it wasn't exactly surprising given the prior literature on prosopagnosia, and lots of behavioral and physiological work all of which suggested special neural mechanisms for face recognition. So it's fun to go in and find it and characterize it, but it wasn't a big surprise.

In contrast, the next area that we discovered in collaboration with Russell Epstein, who's now a faculty member at U Penn, we call the parahippocampal place area (PPA). It's about a centimeter anterior in the brain and a little medial, toward the middle of the brain from the face area, and it's bilateral. You can see it here in four subjects, right there, these two little blobs. It is just behind the hippocampus but not overlapping with the hippocampus in parahippocampal cortex, and here's what it does: It responds very strongly when you look at pictures of places. It can be outdoor scenes or indoor scenes, it doesn't matter if there's any stuff in those scenes or if it's just a bare spatial layout. These pictures are really boring to look at when you're in the scanner, but the PPA is very excited about these stimuli. And when you take them apart the visual properties are very . . . here you can still see the spatial layout and the response is the same. And here the spatial layout information is disrupted even though the visual properties are very similar. Signal drops. Can't reach it with the pointer, but the abstract scenes made out of LEGOs on the far right there produce a very strong response in this region there to objects made out of LEGOs. So this leaves completely open the question of what this region is contributing to place perception, but it does provide preliminary evidence that this region's very selective for place perception.

A little sidebar. It's of interest both to identify and characterize the function of these regions, but it's also a lot of fun to use them to address other questions in visual cognition. So just to mention a few of the kinds of fun and games we've had with the FFA and the PPA, we've shown that the activity in these regions is very strongly modulated by visual attention. So if you have exactly the same retinal information hitting your eyes from a stimulus, if you choose to pay more attention to faces than places you can essentially crank up and down the dials on your own visual system. You can control the operation of your own visual system by visual attention. We can see this reflected in the magnitude of response in those areas, which is modulated by visual attention. Second, we can show using a bunch of perceptual tricks. In this case we used binocular rivalry that when exactly the same perceptual information hits your retina, if the stimulus is ambiguous such that sometimes you see a face and sometimes you see a place, we can see the activity toggling in lock step with the subject's reported percept in these two areas. These regions reflect not just what is striking the retina but what the subject perceives from that information.

And finally, we can remove ourselves from the stimulus altogether and just ask subjects to lie in the scanner, close their eyes and imagine a face or a place. And when we do this we see selective activation in the very same regions, in the face area for imagining faces, and in the place area for imagining places.

Okay, so back to these selective regions. The third one, and the final one, I'll describe only briefly that we've discovered. This is work—with Paul Downing, who's here, now a faculty member at the University of Bangor, Wales—is what we call the extrastriate body area. This is pretty weird, but what can I say, the data dragged us here kicking and screaming. And what the data show is a strongly selective response in every subject we scanned to bodies and body parts. So here it is in four subjects, these are now vertical slices like this. This activation on the lateral surface is right next to visual area MT that's interested in visual motion, for those of you who know it, but not overlapping with it. This region here responds much more strongly to all different kind of bodies and body parts, compared to lots of different kinds of control conditions. To get control of the low-level visual properties of the stimulus here we compared stick figures of bodies versus scrambled stick figures. Much higher response when you see a body in the stick figure than when you don't, and silhouettes showing bodies versus scrambled silhouettes. Also very similar visual properties, much stronger response when you see bodies than when you don't. All of this, of course, leaves wide open the question of what this region is doing with information about bodies. It may be used in recognizing individuals, it may be used in knowing where your own body parts are for better visually guided action, it may be used in social cognition, in understanding other people's actions. So all of this is wide open and under active exploration right now.

Experience and the Visual Cortex

Okay, so I've described these three category-selective regions of the brain, shown diagrammatically here. The body area is not out in space, it's just around on the upper surface of the brain where I can't depict it here. Now, these regions are found in pretty much the same rough location in essentially all normal subjects, so this is showing us something very basic about the functional organization of the human brain. Second of all, the selectivity of these regions is very strong. This is not like lots of brain imaging studies where you scan a large number of subjects and you just barely reach statistical significance with a tiny effect size. These are all enormous effect sizes. And my intuition is big effect sizes are telling you more important things. But of course it raises a huge number of questions, of which I'll just briefly describe and discuss two.

One is, where does all this functional organization come from in the brain? That's a very hard question to answer because—let's go back for a second—there are two very different kind of caricatured simple stories you could tell about the origins of these regions, and they're both quite plausible. One says look throughout primate evolution. If you had to choose three categories of objects that were probably really important for us to be able to perceive clearly, and that was probably faces, places, and bodies. That's one story, so maybe natural selection has produced special-purpose machinery that's hardwired into the brain, and maybe that's why those things land in systematic locations in the temporal lobe. But here's another equally plausible story: Each of us individuals looks at these stimuli very frequently in daily life throughout our lives, and we know that the cortex is very shaped by experience, as several of the preceding lecturers have described. And so maybe the visual cortex just kind of does statistics on the input, and says, "I'm seeing a lot of these and those and those, and let's allocate big chunks of cortex to the things we see a lot." So those are both plausible stories. I would love to figure out which of those or what kind of compromise position is correct for the face, place, and body area, but because they're both plausible here, and because I'm working with humans, I don't know how to answer that question.

So I'll address a different question, which is closely related, and that is whether experience alone is ever sufficient for the construction of a category-selective region of the cortex. This is work with Chris Baker on the right and Jia Liu. So here's how we tested this: We used the test case of visually presented words. Why words? Well, our experience with words is very high, especially for geeks like my students and me. Our experience with words is pretty close to our experience with faces, although importantly, it does start later in development. In contrast, unlike faces, people have only been reading for a few thousand years, and this is not usually thought to be enough for natural selection to produce special purpose machinery. So what that means is that we can use words as a kind of test case. If we find regions of cortex that respond selectively to visually presented words, there'll be a kind of existence proof that extensive experience can produce a selective region of cortex.

So lots of people have reported the existence of a visual word form area, and there's a lot of work on this, but the critical questions haven't really been adequately answered. One is, How selective is that region? And the other is, Is it really resulting from experience? So I'm going to briefly describe an experiment that Chris and Jia and I have done. We scanned subjects while they looked at words and pictures and lots of other things, and our first question was, Can we find a region in their brains that responds selectively, or that is significantly much more strongly to words than pictures? So here is such a region in one subject. It's small but we see this in most subjects in the left fusiform gyrus, and it's producing a much stronger response to words than pictures. I told you before not to believe activation maps, so what we do here is replicate that with a further study, where we make sure we get the same result in the same subjects in the same scanning session, and we throw in some other conditions to see how selective this region is. So here first is the replication of the higher response to words than line drawings. Well what about other stimuli? This is the response to Chinese characters. Our subjects don't speak or read Chinese. Here's a response to strings of digits. So it was quite selective. Strings of digits are a lot like letter strings. But here is this response to letter strings. So already we've learned one important thing: This is not, as many have claimed, a visual word form area. It's maybe a letter area or a letter string area. But for the purposes of testing the effects of experience that's okay, I'll take letters, same deal. So this region is quite selective.

However, I haven't really shown you yet that this is due to experience. You could say well, maybe there's a whole kind of space of different feature selectivity in the cortex, and whatever features happen to be present in the words but not the digits and character strings maybe they would be there even if you never learned to read. It's hard for me to find subjects in the Cambridge area who don't know how to read English words; however, we can move to a different language. So what we did was we looked at the response to people looking at both English words and Hebrew words as a function of whether they were readers of Hebrew or nonreaders of Hebrew. So first these are the data from these subjects showing you the response in that region to English words versus pictures in our non-Hebrew readers, and English words versus pictures in our Hebrew readers, who also read English. Okay, so that just shows you the kind of thing I already showed you. Now the critical question is what is the response to Hebrew words in these subjects? Okay, so here are the non-Hebrew readers. You can see that the response to Hebrew words is quite high, quite a bit higher than to pictures; however, it's significantly lower than the response to English words. What about the Hebrew readers? It's a whole lot higher. So I should say we only have three subjects here so far, so we need to run a few more subjects. But it's so clear in these three subjects that I'm already reasonably confident that this result will stick. And what it seems to be telling us is that the difference in the response to Hebrew in this case, in this case shows us a very strong effect of experience on the selectivity of this region.

So in answer to the first question—How do these regions arise?—well, what we could say is at least some of the selectivity in this general part of the brain is based strongly on experience. It doesn't tell us that that's the origin of the face area, the place area, and the body area, but it's a kind of existence proof that says that might be the origin.

Selectivity in High-Level Cognition

Okay so the final question I'll address only very briefly, and that is whether this kind of high degree of functional specificity that we see in the face, place, and body areas is really something about the visual system. That it's very kind of modular, it parcelates the problem of vision into all these separate little functions allocated to different regions of cortex. And maybe that's special about vision.

Or, do we ever find that kind of selectivity for very high-level cognitive functions? And what I'll say is I think the answer to this question, this is very much an empirical question, I think you have to test each case. So, for example, there are a lot of very interesting claims about a region in the parietal lobe up here in the left hemisphere that's been argued to be selectively involved in understanding number, that's a very high-level function, and there's quite a bit of evidence for that. However, in my lab, what we show is we can find that region, it is involved in understanding number, we can replicate that. But it's also involved in lots of other things, like visual attention, like selecting responses that are appropriate for the visual stimulus, like lots of other functions. So I think the answer in many cases will be no, there aren't highly selected bits of brain allocated to high-level cognitive functions, there are general purpose bits that do a lot of cognition for us. And by the way, that's also true in vision. I focused on the category-selective areas but there are also what appear to be very general-purpose areas that represent the shape of more or less any old object, as far as we can tell.

So some things are very general like this, but I'll show you just very briefly one region that is remarkably specific, beyond my wildest expectations. So this remarkable and brilliant student arrived in my lab four years ago, Rebecca Saxe, and she said, "I want to study theory of mind." This is the topic that Mike Rutter mentioned yesterday. And she said she wanted to do this with functional MRI. I said, "Well, that's lovely, but one, I don't know anything about theory of mind; and two, I don't believe for a moment that it's localized anywhere in the brain. And if it's not, then functional MRI is the wrong tool to use." But luckily she was undaunted and she spent three years in a very focused fashion approaching this question. So she ran a version of a Sally-Ann task that Mike Rutter described yesterday, in which subjects read a little verbal description of a person who's presented with some information about another person and what they know, and they are then asked a question that requires them to distinguish between what that person knows about the world and what is true about the world. The way you pull those apart is to make the belief false. Otherwise, you can answer the question by reporting on what's true of the world. So she used the Sally-Ann task, and she showed that this region in the temporal parietal junction is very strongly activated when subjects perform those false-belief tasks. Several other people had shown related things before; however, what they hadn't done is get obsessive, like we like to in my lab, and really address each of the alternative accounts we can think of. So Rebecca did this and she ruled out every plausible alternative explanation I can think of, a few which are shown here. One is this region does not simply get engaged whenever you do identical problems about nonmental representations. So instead of doing Sally-Ann task, instead of having Sally watch Ann and having the subject report on what Sally knows, instead of Sally we have a Polaroid photograph. So now we have instead of a mental representation, a physical representation. So you have the logically identical question, you ask the subject in the scanner about what's in the photograph when it differs from reality. No dice, this part of the brain is not interested. So it's not just responding to any representations that differ from reality. It's not simply false representations. So if you ask about true beliefs you get the same high response in that region. If you ask people to reason about hidden causes, you describe physical events for which they have to infer some hidden cause, you don't get much of a response in this region. And if you ask people questions about people that don't refer to their beliefs, like about their physical appearance or other properties, or their actions, you don't get the strong response here. So despite my extreme skepticism before I went into this, I'm now convinced that this region is very selectively involved in representing the contents of other people's beliefs. And that's a very high-level cognitive function. So I'd say *yes,* at least some very high-level functions do get small regions of cortex allocated to them.

And I'll stop there. Thank you.

Questions and Answers

I guess I can take maybe one question. Yes, back there.

[Question inaudible.]

We've tested that. We had what we call—okay, the question was, Is there a snake area? And I have to apologize, I tried to get on the Internet last night to show you a snake picture but the Internet connection in the hotel wasn't working, so sorry. We did test that condition. In fact I skipped over this because I didn't have time. But we tested pretty much every other visual category we could think of for which we could make an even remotely plausible story that there might be a special-purpose bit of cortex. And that included one condition where we had both spiders and snakes. Paul Downing, who did the work with me, is like, not phobic about snakes but he is about spiders, and I'm the opposite, so we threw them together in one category. We couldn't find any selective responses. Of course, importantly, the fact that we don't see them doesn't mean they're not there. It's the easiest thing in the world with functional MRI to not detect a selective response that's actually in there for any number of reasons, including for example, the neurons that are selective may be interleaved with other neurons below the resolution of your MRI scan. But the main story is we tested lots of other conditions and we didn't see anything with the kind of size and robustness of selectivity that we see in each subject for faces, places, and bodies.

[Question inaudible.]

I haven't, but it's . . . okay, so the question is, What happens in autism subjects who generally don't like to look at faces, what are their responses? I haven't studied this, but Mike Rutter mentioned yesterday a couple of studies of people who have looked at it. The reports in the literature, there are two or three studies that I know of, argue that you do not get a fusiform face area-selective response to faces, and instead nearby regions of the cortex that respond more strongly to objects seem to be engaged when autistic subjects look at faces. That's suggestive, but I'm very cautious about this for two reasons: One is I know of at least one other lab, I think two, that is getting different results, and I want to see what happens with that. And the other is there are many ways in which that experiment could fail to find the face selective response. For example, the very fact that the autistic subjects don't like to look at faces may mean that in the scanner they're not even [inaudible] the faces. I don't think the published studies tracked eye movements to make sure that the subjects were looking at the faces. This is critical. Even if they're looking at the faces, you need to make sure they're attending to the faces, so they need to be given a task that requires them to perceive the faces. And until all that is done and replicated I'm a little bit skeptical.

Man: Since science is a progression of errors, does the work, in terms of organization, relate back to the work of Lashly and [inaudible] and others who have done work in the area to vindicate some of their concepts of fifty years ago? And the second thing, have you done any work in the area of auditory perception—music, for example? There's an area of the brain with perfect pitch and so on, which you can perhaps do some specific work. I wonder if you can elucidate that. Thank you.

Nancy Kanwisher: So the questions were, first, How does this relate to the historical work of Lashly and others? The second question was, What about music and auditory perception?

On the first question, yes absolutely, I'm addressing the oldest question in neuroscience, which is the pendulum on this topic. The fashions on this question have swung back and forth reliably every few decades for a very long time, and that question is, To what degree is function localized in the brain? And I showed you some functions that are very localized in the brain, but I also noted that many others are not. I gave the example of understanding number, but there are many other examples of things that are not localized in the way that I described. So yes I'm still pursuing the same old question. I'm sure that neuroscience will still be at this for a long time. And I should say that my position on this is a little bit idiosyncratic. Many of my colleagues who run very similar experiments would stand up here and tell you that what the data show is massive overlap in those functions, so there's not any agreement on this right now.

Second question on music and auditory perception—I haven't studied this but there are a number of interesting studies. The closest one to my work is a study by Pascal Belin that was published in *Science* in 2001. It actually came out of a conversation that Pascal and I had where he was studying tones and auditory perception. I said, "Well, what would be the auditory equivalent of the face area?" Well easy, it would be a voice area, right? Not speech because that's different, but voices, because you can recognize individuals from voices, you can . . . like faces you can pick up emotions and all kinds of other information, age, about the speaker. And sure enough he went and did a number of studies and found in the superior temporal sulcus, up closer to auditory cortex, a region that's quite selectively responsive to human voices, and this included stuff like laughing and crying and sighing, it did not include stuff like speech, because, that is, it did not require speech, the nonspeech sounds also activated it. But it wasn't just any human sound, so when they had people clapping or walking or making other sounds that were clearly human did not activate the area, it was specifically voices. And there's lots of other work on music, but I don't know that literature that well. I'm probably supposed to stop.