

The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high quality educational resources for free. To make a donation or view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

WINRICH

FREIWALD:

So my talk is going to be mostly about faces. And in many ways I'm going to connect to what Jim DiCarlo was talking about today and what Nancy talked about today. I just thought preparing for this that I should say a few things about primates and intelligence and how face recognition might be connected both to the species we are studying and to the overall question of intelligence. And I thought the most appropriate thing to do this at MBL would be to start with this kind of creature.

So you might have seen in the paper *Nature* that came out last week-- that the genome of the octopus was sequenced. It was really heralded in the public press that it's finally proving their intelligence. The argument is that there are 33,000 genes, 10,000 more than humans. That's of course not a very strong argument. There are plants with 45,000 genes. So that really doesn't tell you very much about intelligence. But amongst those genes that were found were lots of genes that are important for development of the brain. And there is very high heterogeneity in certain gene families that are controlling the development of the nervous system.

So it's one example of how a better understanding the intelligence of other creatures in neural terms, even in creatures like the octopus which is very difficult to study. The point I would like to stress about the octopus is there's really nothing social about the species. Actually almost simultaneously with this paper, there was one report about sexual reproduction in one particular species of octopus where it's not clear if it's violence or more affection. But this is really the one exception to otherwise a life that's pretty much non-social at all.

So in the octopus you have the egg stage and then this larva which are hatching very early on. There's no interaction with mom or anything like this. They go to the surface of the ocean and they start feeding and try to grow as fast as they can because only a few of them are going to survive. Then reproduction is a very scary enterprise. In many octopus, the male has to be careful not to be eaten by the female in the process. When it's successful, then actually the male stops eating. He's going to die even before most of the youngsters are going to hatch.

And all that mom is going to do is basically to make sure that fresh water is going to be delivered to the eggs. And that's it in terms of social life. So you can be very intelligent like an octopus and, who knows, Gabriel just mentioned extraterrestrial life. Maybe other species out there might also not be that social. And so there doesn't necessarily have to be a connection between sociality and intelligence.

Second thing is it can be very social, get a warm fuzzy feeling of being with others and you get lots of protection being in the group, but following your instincts in this way doesn't necessarily make you very smart. So I don't want to argue against the connection between sociality and intelligence in the form of social intelligence. But we have to be careful there's no necessary connection between those. However for primates there is this idea, this social intelligence hypothesis that really what made primates so intelligent is their sociality.

And so let's consider a little bit what the arguments are. It was most strongly put forward by Nick Humphrey in '76, but there are similar precursors to this idea. So who are these primates? The primates are a small group of mammals with about 400 plus species. They're very diverse. You can have very small primates, just 30 grams, all the way to 200 kilogram animals. They evolved, studying 65, 85 million years ago, there was this mass extinction. And so you have this high diversity within the mammals starting from this point on. So they have certain things in common with other mammals, but they're also very special in many ways.

All of the primate species are social. They're not [INAUDIBLE] social, but they are social. They develop very slowly. So it's really very different from the octopus. A lot of investment is made into the offspring. There are not very many offspring. And the lifespan of these animals is pretty long. The octopus life span is three to five years. They're very visual, rather than other mammals, which are very olfactory oriented. They have binocular vision, many have color vision, so vision is very important for primates. And they have on average larger brains than other mammals have.

So our understanding of the anatomy of primates and what makes this special is actually very rudimentary. But there are a few points that should be of interest. So if you look at the mammalian brains, obviously they're very, very complex mammalian brains. They're not primate [INAUDIBLE] ones. And the main factor really is body mass. So the bigger an animal is, this is elementary of body weight versus brain weight. And you can see there's a log-log relationship between the two. That's something that you find all over the animal kingdom. But if you compare primates to other non-primate mammals you can see that there is a larger

increase with body mass of brain mass. And if you count the number of neurons, the number of brain neurons, it is increasing more steeply with body mass than it is in other mammals.

This is obviously a very crude measurement. There are others. If you look for brains roughly the same weight, you can again see like in primates, you have many more neurons than in non-primates just by these two examples here. So there seems to be something different in the organization. There are other measures you can look at. So for example, the neuron size is increasing with brain size in rodents. So a larger brain in rodents is not necessarily one that has many more neurons, but these neurons are just getting bigger. And in primates this is really not so much the case. So the size of the neuron pretty much stays constant even if you decrease the size of the brain.

Or how much white matter do you actually need per neuron to connect other brain regions? In rodents, apparently the fiber caliber is increasing with brain size. So again, your brain might grow just because the anatomy of the basic element requires that it will need more space. But in primates that's not the case. If you get more white brain matter, that's likely because the connectivity is more complex. Primate brains also fold faster with increasing size than rodent neurons. So these are all just very coarse indications that maybe there's something special about the primate brain compared to other mammalian brains.

So along the anatomy, I mentioned that primates have forward-facing eyes. They can do binocular vision, they have color vision. They have skulls with a large cranium, that's something that makes them special from other mammals. They're also special in other ways that are important. If you think about embodied cognition, you don't have to buy into all these points, but obviously if you have a hand as complex as ours which we share with many primate species, there are lots of things you can do. And that requires you to be able to control it. And this gives you a power to interact with the environment that other animals might not have.

So the shoulder is more mobile, there's an opposable thumb in many species. And then in the face, there are changes that you might already have seen here. So if you go up to the more complex animals, the snout region is becoming increasingly reduced. And I will tell you later why this might be important.

So these are anatomical specialization in primates. Sociality is very important as well. And so there are four main organizational principles of sociality in primates. The dominating one is the

second one here. It's called the male transfer system. So it's a polygamous, multi-male organization. This is very important for the social life of primates because what it means is that the social behavior has to be complex. So there can be cooperation, like grooming, defense, and hunting which all the animals of a troop might engage in. But at the same time there's competition for food, mates dominance, hierarchies. And it's a function of the complexity of the social environment.

So primate social life was beautifully inscribed by Dorothy Cheney and Robert Seyfarth in this wonderful book *Baboon Metaphysics*. And I'm just going to quote from that. So they studied baboon monkeys in the wild and here's what they have to say about this. "The domain of expertise for baboons, and indeed for all monkeys and apes, is social life. Most baboons live in multi-male and multi-female groups that typically include eight or nine metrolineal families." Which means that the females stay in the group and they found families and they're going to stay constant over a long period of time.

"They have a linear dominance hierarchy of males the changes often and the linear hierarchy of females and the offsprings that can be stable for generations. Daily life in a baboon group includes small scale alliances that may involve only three individuals and occasional large scale familiar battles that involve all of the members of three or four metrolines. Males and females can form short term bonds that lead to reproduction, or longer term friendships that lead to cooperative child rearing." "The result of all this social intrigue is a kind of Jane Austen melodrama in which each individual must predict the behavior of others and form those relationships that return the greatest benefits. These are the problems that the baboon mind must solve and this is the environment in which it has evolved."

Most of the problems facing baboons can be expressed in two words: other baboons. And so this is really important. So again, if you're social you don't necessarily have to be very smart. You can be very smart and not be social. But there's something special apparently about primates that links our intelligence to our sociality.

So again, the Social Intelligence Hypothesis, what works in its favor is that the primates have large brains. And primates, apparently group size is correlated with brain size across different species, and not the home-range size. This is an alternative hypothesis is that maybe you have to forage in more complex environments. And so a similar proxy for the complexity of your social life and your physical life predicts there's a better correlation of brain size with the social life than physical life.

The complexity of an individual's social relationships increases exponentially with group size. And groups are not small. And we're going to get back to this point in a little bit. The baboons and other primates know their peer's dominance, rank, and social relations. Everything that I mentioned in the previous slide is good evidence for behavioral work that actually primates know something about it. This social knowledge contrasts with surprising cases of ignorance outside of the social domain. Even when it's something as important as a predator.

So Cheney and Seyfarth, they also studied vervet monkeys. And what they observed is that vervet monkeys, one of the main predators to them is a python. And the python, if it crawls in the sand it leaves behind a trail. And there is absolutely no indication that the vervet monkeys make the connection between these trails and the presence of the python. Which is really striking, you would imagine this is the first thing that they would have to learn. So there are cases where they actually follow the trail into the bush and they're very surprised to find a python there.

Another example is, are leopards. So in the environment, leopards are of course predators who also feed on the vervet monkeys. Leopards have a way of putting the carcasses of animals they hunted down into the trees to protect them from larger predators like lions. So the presence of a carcass would actually indicate to you that likely there is a leopard around. And again, the vervet monkeys they don't make that connection. So if there is a carcass there, they're not particularly scared about it. It doesn't mean that they're ignorant. So they're even following alarm calls from other species as to whether an eagle is approaching, or if cattle is approaching. So it's not that they're generally dumb in this way, but there's really this very big contrast about all the details that they know about their social world and the obliviance that they can express for non-social factors.

Then we have specializations, we're going to talk about this in the brain for processing social stimuli. And then there's actually evidence that females who have better social abilities, that they are less stressed and they have better reproductive success. And so this all works to say that if you are socially smart in a baboon environment or in the environment of many primate species, you actually have better success of reproducing. Therefore there's a good argument to be made that your social intelligence will therefore go to the next generation. This is how social intelligence might improve.

I think there's one important point that's oftentimes not made. And that's if you become

smarter and smarter in interacting with your physical environment, your physical environment does not change very much. But if you are interacting with a social environment, and getting smarter and smarter interacting with the social environment, the elements in your social environment you're interacting with, they're also getting smarter and smarter. So you're actually setting forth an arms race where you're not only improving the situation by getting smarter, but we have to get smart in order to keep pace with the others who are outsmarting you.

And so you can see how there could be a connection between sociality and intelligence. That there's really this arms race does not occur for physical interactions but for social interactions. So that you actually have to be able to better predict the next move of someone else in your group and you have to know something about that individual for you to be successful.

So there are arguments against the Social Intelligence Hypothesis. So, in particular, we're ignorant about many other species. So there are other social species, hyenas for example have complex societies, but hyenas have not been studied as much as monkeys have. We also don't know about the complexity, mostly for this reason, where they're really compared to whales, whales for example, primate societies are more complex. And this of course would be a crucial conjecture of this Social Intelligence Hypothesis.

Then within the primate orders, there are actually some other correlations that do predict brain size very well. And so within the primate order, social learning, innovation, and tool use are strongly correlated with brain size and not with group size. So you could imagine a scenario where actually the evolution of basic, social intelligence is something that's very basic to primates. But then if you go to different species within the primate order and ask like why did they become so smart, like orangutans or chimps who can use tools, then it really might be the tool use that would be of more importance than the sociality.

So I have a movie here that actually illustrates these two different hypotheses. So you see social interactions here in a group of Tonkin macaque monkeys. You can see the facial displays, you can see that they are tending to each other. And well you might think that they are trying to figure out what's actually going on here. And here's the alternative hypothesis of this. This is tool use. So you can see this guy just invented a cool tool, a nose pick. And so it's anyone's guess what's more important to your intelligence, you been able to read the social significance of other individuals in your troop or your ability to invent a nice nose pick.

So the last point I wanted to make is a question. So are the primates' abilities in social knowledge really intelligent or is it just more like idiot savant-like abilities? So a unique specialization where they're good at. The argument that Cheney and Seyfarth made is the following. So the knowledge that they have should actually be true knowledge and just not learned associations. And the reason has to do with the complexity of the social environment. If you have 80 different individuals, which is the typical case for these baboon monkeys, you have 3,160 pairs of animals and 82,160 trails. It's going to be virtually impossible for you to learn all these different pairwise relationships and then behave intelligently based upon it.

Second, these relationships can change very fast. So it would not be very smart to try and make a list of all these pairwise interactions and then act upon that. No single behavioral metric seems to be necessary or sufficient to recognize associations like matrilinear kin. So human observers are apparently not very good to predict this if they don't know the animals very well personally. Then you might think, well, maybe it's not physically like a list but you don't really learn this much and you apply a simple rule to it. And that also doesn't seem to work very well because social relationships like friendships, they are intransitive. So if A and B are friends, B and C are friends, that doesn't mean that A and C necessarily have to be friends.

Others like family relationships are complex, they're non-associative. So if A is the mother of B, that actually means that B is not the mother of A. So there's a more complex interaction there as well. And then finally, there can be simultaneous membership in multiple classes. And again, for you to be able to keep track of this, you better have a cognitive model of what's going on rather than just a list of associations that you learned. It's very difficult in experimentation to prove it's not association, but I think these are very good arguments to consider that actually these primates have active knowledge of their social environment.

So there's one example for this where you can actually make the point very nicely. And this is the story of Ahla. Ahla is a baboon monkey and she was actually living with farmers in South West Africa. So there was a habit at the time to actually replace dogs who were herding goats with baboons. And so you can see Ahla sitting here. You can see her here adopting some of the behavior of the goats. So she's licking salt here, which is something that baboons would naturally not do. She would continue to engage in social behaviors that are typical for baboons. So she would groom the goats, for example.

But the most amazing thing about her, which is a little hard to see as you see her here, she's

carrying one of the little yearlings here, and brings it to its mother. And the description of what this animal was doing was when the goats were brought home and then sometimes they were separating the mother goats from the offspring, she would actually go manic and then try to pair them up and she would not stop until she was finished. And this would even happen as multiple goats were calling for the yearlings and then vice versa. And so she would have to put order into the social world that she was living in and she wouldn't stop until this order was restored.

And the farmer said that they were not able to tell any of the adult goats or the yearlings or know any of the pairs. But this was like her world. This was like a social environment that she was in and she would structure it according to her cognitive demands. So the point is that primates have intricate social knowledge. So they know about the status of individuals, like their age or their gender. They know about the interactions of individuals. They recognize them very simply like grooming and mothering. And then based on these observations of the different individuals in the social world, they built these cognitive structures like friendship, kinship, and hierarchy, have an interesting, complicated structure to them.

All of this is rooted in the concept of the person. And this is very important and as I'm going to be talking about face recognition, I have to emphasize these are two different things. You can recognize a person from their face, but if you can recognize a face it doesn't mean that you know who it is, who is behind the face. So the person concept would include something like this. It's a juvenile, female monkey. It's the daughter of x and so on and so forth. So this is knowledge that we have. It's actually been shown in rhesus monkeys, the monkeys that we work with, that they actually have this person knowledge.

So why do we study faces? So for us faces really are the ideal intersection between object recognition, the study of which Jim DiCarlo talked about, and social cognition. So as Jim was alluding to yesterday, vision is really important in primates. About a third of the primate brain is thought to be involved in visual function. So this is a lot. And this is testament to the fact that there's a lot of information to be gathered from the outside visual world, but also that it's difficult computationally to gather this. And so Jim was explaining some of the computational challenges that object recognition has to solve yesterday and then we'll come back to some of his points a little later.

So what is an object? So Jim was saying that it's the basic unit of cognition. And just very quickly, what it actually is it's more than a collection of features. So the Gestalt Rules of

Perception actually emphasize this. So if you have proximity of elements, you group them together. If these elements share similarity, you group them together to larger entities. If there's good continuation, you group these local elements together into lines. If that's common fate, you group them again together to larger entities.

And something similar is true for faces. If you have different face parts, then the wrong organization but now you put them together correctly, you can suddenly recognize a face. So there's a larger scale organization to not that goes beyond just being a collection of features. And maybe something similar is going on for a higher order condition that I'm sure other people are going to talk about. And it's in the physical interactions, we infer causality from just the sequence of events. Or social interactions, like in Heider-Simmel movies where you are telling yourself a complex social story unfold even when there's just simple, geometric shapes moving around.

So this creation of higher order representations I think is essential for object recognition. It's a constructive process that the brain imposes on the piece of information it gets from the eyes. It's not just a collection of features. It's kind of the basis of symbolic representations. It can create meaning, especially if you think about the face, if it's the face of someone you know that's very meaningful. And it makes information actionable. And these are really, I think, the important links between object recognition and social cognition, and faces are smack in the middle of this.

So I already showed this movie. To use again I emphasize the social communication that's taking place here. You can see the facial displays here. So the older male who's chasing the younger animal is making these facial displays. By the way, most of you will never have seen a Tonkin macaque and still you cannot understand what's going on there. This is something very special, again that you don't have in all animals, these facial displays. And Charles Darwin was actually again one of the first people to notice it in 1872 is that you use your face to express your emotional state. Otherwise your emotions are private to you but you use body language and then facial language to suppress your emotions. And oftentimes you do it even if you don't want to. It just happens automatically.

And that's not possible in all animals. So if you are a fish or a frog, there are lots of really cool things you can do. You can sit on the front porch and enjoy the day. So lots of things that you can have in common with primates. But facial communication really requires something more that's very mammalian specific. In mammals you actually have in the face musculature that it's

not attaching from bone to bone, but it's now attaching to the skin. And so if you look at these two rats here where their whiskers, I hope you can see it here, in the end are labeled. You can see they're actively exploring each other's faces in a somewhat sensory fashion. That's possible because they can move their whiskers because of this musculature.

And that's a specialization that's becoming more and more refined in primates. So in rhesus monkeys and chimps and humans, we have 23 different facial muscles. They're becoming more and more flexible. I mentioned before that the snout region is increasingly reduced in primates. So you have some simpler primates where there's still a strong snout, which is limiting the ability of the face to move. But the more complex the primate is getting, the more flexible these muscles become and the more expressive the faces become. So the face now becomes richer and richer with social signals that can read out. And in rhesus monkeys which are shown here, you have a fixed set of facial expressions that, again, for a system that can analyze these, it's very important information about the emotional state of another animal.

Primates are also very interested in faces. I'd very much like to show this movie which is showing a three-day-old macaque monkey. And I'll tell you what the point of this study is. So you can see that he's attending very closely to the face of the experimenter. Of course, if there were bananas you might think he would also be, this isn't proof that there is this specialization for faces that Nancy was alluding to before, but it's at least intuitive. The second thing why I like to show this movie is exactly what happened right now. You're all getting really excited about this absolutely adorable, little critter, right? And I've seen this movie now hundreds of times and it's still like this. It's still very emotionally charged.

So here's the third reason. The experiment is based on facial movements here. You can see it's getting really excited about it, it's getting very active. And now he's reproducing these facial movements as best as he can. There's a specific facial interaction that's happening in human babies, I think for three months. It's happening in these rhesus monkeys for two weeks. And you can see that there is an intricate connection between what they're perceiving and what they're acting on in an automatic fashion. But this emotional part I think is really important. It's just at a certain point that you can't control these things. Faces really get very deep into your emotional and social brain automatically. And so one of the lines of research in my lab is to try to figure out the circuits that make that possible and then to use this to get an inroad into the social brain of function beyond face perception.

So amongst the signals the faces are sending, you recognize Charles Darwin, so it's identity,

they're getting their social communication, there's emotional responses, and there's also face following. So the direction that the eyes are looking to, we are following automatically. We can control this later, but this initial automatic response. Here's one very nice illustration from the British TV show. You have people wearing these glasses, actually it's a large background to people wearing these glasses where their eyes are drawn on these glasses that are going to one direction. And so you know that these are not real eyes, but what's happening is your attention is drawn constantly to this upper right region. And it's getting annoying over time because you know that there's nothing there. You know that they're not really paying attention there. But automatically your attention is drawn there and then you're going back again and your attention is going out there again. And so this is another thing that comes from the face that gets deep into your attentional control system.

So social perception can start with faces, but faces are the most important visual sources of information. We get gender and age, of course identity, and things like perceived trustworthiness or attractiveness from just a very brief look at the face. And then there are these dynamical signals, like mood and overt direction of attention that we also get from the face.

So how does this all work? So Jim was already explaining some of the challenges of object recognition to you. And so here are some of the challenges. So first of all, the social scene like this one here, lighting conditions can sometimes be non-optimal. And so the first thing for you to analyze the facial signals which are in this scene, is to localize where the faces are. And I'm going to tell you a little bit about what we understand about the mechanisms of that.

Then once you know where the faces are, you want to analyze them further, you want to know who these individuals are. And I just realized that the images that I had from this, which are of course also taken from *The Godfather*, might not be the best. Where is the other picture of this individual here in this display of these five faces? Upper right. And then there's another individual, there's Don Corleone and there's another person down here with two different directions. And then if the lights were down a little bit more, you could see this better. The cool thing is that we have a way of relating these two pictures to each other knowing that they are from the same person, even though physically on a pixel by pixel basis these two actually are much more similar to each other. And so we'd like to figure out how the brain is doing that, achieving object recognition, in this case face recognition, in a manner that's invariant to transformations that are not intrinsic to the object.

This is just a reminder that face recognition actually is very difficult. So this is of course just made up from *Curb Your Enthusiasm*, but there's a condition that many of you will have heard about, prosopagnosia. And to a prosopagnostic person who is face blind, the social world might look like this. So a prosopagnostic has great difficulty telling one individual from another. This is at least the most typical condition. And you can imagine that your social life would be really difficult and your enthusiasm about socially interacting would really be curbed if all the individuals looked at this and looked all the same. So there must be something about the new mechanisms that's very precise.

So what's the neural basis of face recognition? So the story really starts with Charles Gross many years back in the late '60s, early '70s. He was recording from the inferotemporal cortex. He was showing pictures of monkey faces, other social stimuli like the monkey hand, he would scramble the face, and then look at the responses of cells. And he was the first to find a face selective neuron. Here's one. So these vertical lines have the action potentials the cell is firing. This is the period of time his face was shown. And this is the period of time the control object, the hand, was shown. And you can see the cells responding selectively to the face and not to the head.

So this was a very nice finding. It actually took him some time to convince himself that he could publish it because he thought people would not believe it. It's recording an anaesthetised animal. But luckily, he did publish it many, many years later. And so this is the first evidence that there is a specialization in the brain for faces. He found many other cells that like other things and then faces and I think people thought that they were intermingled with other objects.

So this was the view. That this is the side view of the macaque brain. This is the superior temporal circuits, the one big circuit in the monkey brain. And all these symbols here indicating positions where people found face selective neurons. The thought was they intermingled with object recognition hardware. It's basically the view that there is a big IT cortex where everything in object recognition can happen. And yes, of course you would have some cells that are face selective, you would have other cells that are non-face selective. And the mixture and the complex pattern of activity really is what gives you the identity of the object.

Then Nancy used fMRI to discover face selective areas. So first, these are views from face areas. We now know multiple face areas that she was talking about before. So here are different slices to this. And the thought from these images really was that, no, that maybe

within this large expanse of object recognition hardware there might be very specialized regions that are really there selectively to process faces. And so you give the FFA, and so the question you would ask is, is this really a region that's devoted to face processing and face processing only?

Are these regions really face processing? Modules devoted to face processing or is it just the tip of the iceberg based on your statistical analysis that this region just looks a little bit more face selective than the neighboring regions? And second, do monkeys also have these localized face areas like humans? And you've got the answer already. Yes, they do. And then what is the distribution of cells within these regions versus outside? So this is really the research Doris Tsao and I engaged on many years ago. We used fMRI on macaque monkeys, same technology as in humans, slightly different coils. And this is the picture that we got. Very consistently across different animals.

Here in the temporal lobe, you have six, face selective regions that you find that anatomically specific regions there's some variation from one individual to the other. But with the exception of the most posterior area, you actually find all these areas in all individuals on both hemispheres. There are also three areas in the prefrontal cortex which are a little harder to find. But the one in orbitofrontal cortex is actually as reproducible as the one in the temporal lobe. So, yes there are. Monkeys have localized face areas like humans. And as Nancy was alluding to, we actually have quite a bit of evidence by now that these systems might be homologous. Very, very difficult to prove that they are homologous. But all the evidence we have so far is really pointing in this direction.

So how selective are these face patches? And what Doris and I did was to lower recording electrodes into these face areas and record from cells inside this fMRI identified areas. And I'm going to show you a movie of one of the first cells we recorded from one of these regions. So it's actually a video we took from a control monitor. So it shows the same thing the monkey shows. The quality is not great because it's an actual video camera we took to take this image. In addition to what the animal saw, you will also see this black square which is indicating where the animal looked on the screen but the animal did not see this. And you're going to hear clicks if everything works fine when the actual potential is fired.

Anyway, here's the quantification. So with 96 different stimuli in this image set, 16 faces and 18 non-face stimuli, this is the average response which is normalized between minus 1 and 1.

And you can see that the biggest response of this particular cell actually of course to the 16 faces and not to any of the control objects. You can see though that there are some stimuli here in the gadget category, for example, that are eliciting responses that are quite respectable relative to the faces. But really the biggest responses recorded were to the faces.

So then we color coded this so you have a response vector of the sale, where red is now symbolizing response enhancement and blue is symbolizing response suppression below baseline. And the advantage of using this format is you can now stick all the responses you get from all the cells that you're recording day after day after day from this one face area. And you get a population response matrix.

And the way that this works is cell numbers are organized from top to bottom, page number from left to right. And you can see very quickly that most of the cells here are either selectively enhanced or selectively suppressed by faces. There's a small group here between, something like 10% of the cells, where it's not so clear what they are doing. But if you do the population average, you can see much bigger responses to all the faces rather than on face objects. If you look more closely, what these pictures are eliciting in these intermediate responses, these are like clock faces, apples, pears, there are things that have physical properties in common with faces. So you can kind of fool the system to give a partial response. And this is one clue to what this area might be doing. It should be doing a visual analysis of the incoming stimuli to try and figure out if these are faces are not.

So these are cells in the middle face patches. I was actually going over this pretty fast but I'm going to use this later quite a bit. And so let's wind back. We have one posterior area here, to middle face areas, see the middle face patches, and then three anterior ones. I'm mostly going to talk about this one here, AL, and this one here, M, in addition to the middle face patches.

So we think that actually this is another automatic face recognition feat. We can't stop feeling sorry for these peppers. They've been just cut in half. And so they seem to be screaming, and then you know they are OK but still you feel like something really bad just happened. And so we can't stop having these inferences about peppers where they look like faces. And one reason could be that we have this specialized circuitry that's just getting active with right features, even if you know these are not faces.

OK so when the faces were discovered by Charles Gross, this really fell on very fertile

grounds. And I should just discuss some of the implications. So David Hubel and Torsten Weisel just discovered a few years before orientation selectivity. So it was a big jump from early processing where I could see how selectivity of cells was getting more complex. More concentric representations to elongated ones, from simple cells to complex cells, but complex cells are as selective as simple cells but don't really have a special location. All the way up to the opposite end of the visual system and now you find a face selective neuron.

Jerome Lettvin just had coined the term grandmother neuron which some of you brought up yesterday. The idea is that there should be one neuron in your brain, or this is the hypothetical situation you came up with, one neuron in your the brain that's firing if and only if you see your grandmother, no matter what she's wearing, which direction you see from. That's the neural correlate of you perceiving your grandmother is the activity of this one neuron.

And there were other concepts like Jerzy Konorski gnostic unit that made the same point. Then Horace Barlow came up with this idea maybe it's not one cell, but multiple cells. But gave us a sparse representation of pontifical cells, a few of them at the top of a processing hierarchy. And that's actually how we recognize faces. And then of course there's the opposite view of Donald Hebb. He talked about cell assembly since there's no-- things are there like large assemblies] of cells. Or Karl Lashley who talked about mass action and actually were completely against functional specialization.

If you look at a plot like this, I think one of the things you want to emphasize is that these cells really don't fall into any of these categories. You can have cells that are very, very face selective but they don't have to be very sparse. They will appear sparse if you poke them over and over and over with non-face stimuli because they're not going to respond to those. But within the domain of faces, they're going to respond to pretty much all faces. There are differences between these different cells. I'm going to come back to that as well. But it's one example where we can actually ask them these deep questions about what is the neural code and quantitative matter by focusing on the right stimulus and the right place to look at it.

So we have some evidence that monkeys, like humans, have face regions, and the monkey face patches appear to be dedicated domain specific modules. The practical implications of this is that now we have unprecedented access to function homogeneous populations of cells coding for one high level object category. And we know this category, we can make stimuli. And we can modify the stimuli sometimes in parametric fashions. And so we can have very deep insights into how these cells actually are processing interfaces, how they are restricting

properties from these faces. And we can do causal tests and actually show whether these cells are involved in face recognition behavior.

And we're just going to go over this very quickly. This is work of Srivatsun Sadagopan, he actually gave me this picture of himself. This combines the front view and a profile view. The logic is very simple. So we wanted to inactivate one particular region in a male. I'm going to tell you in a second why a male. While the monkey would be engaged in a task like this where it is to find a face in a visual scene. The visual scene that we constructed looks a bit like this. It's displayed on a touch screen monitor so the animal's free to move around. It has to find the face in the scene and the scene is composed of a pink noise background embedded in which there are 24 different objects. And the target object, in this case the face, is going to be varying in visibility across 10 different levels. We would have other tasks where the monkey body was included, which you will not be able to see here but there is a monkey body here or the monkey was looking for a shoe.

Then we would infuse muscimol which is a pharmacological agent that is inactivating cells along with a contrast agent gadolinium which you can measure an MRI. And then this yellow region here is the face area now and this white region is the actual injection site that we used. And so this gives it a way for every experiment to control, are we inside the face area or are we outside. And we can use the outside injections as controls. What's been found is shown here. So we have a psychometric curve. So in normal behavior you're getting better and better at finding the face in the scene as you increase its visibility. If you inactivate, you get a reduced face detection behavior.

I should emphasize that we only inactivating one face area out of 12 in the temporal lobe. We are only activating one hemisphere. And what Jim DiCarlo was emphasizing yesterday this retinopathy at this level of processing. So the animal can actually use a scanning strategy to go one direction and overcome this deficit. And so likely this effect would be much stronger if we had inactivated on both hemispheres or would have controlled precisely for the eye movement. But we are here for natural behavior.

And the controls of bodies and shoes are, there is no effect there. We put lots of controls and we went to the next state of the behavior. We did the injections outside as I mentioned, it's very specific for inactivation inside the face area that the most basic of face recognition abilities, face detection is impaired. And so there's a way you might visualize it. So one way to explain this behavior would be that the visibility of a face like this would actually, with an

activation, look something like this where it's going to be harder to detect.

The second way we can take advantage of this is that we have now access to individual cells. We actually ask more precise questions about how they're processing faces. And actually there's an activation study that was motivated by earlier work we had done on selectivity of these cells for features that should be relevant for face detection. This is what Shay Ohayon did when he was a grad student with Doris. He's actually now a post-doc with Jim. And again, the question is how can you detect faces even when the lighting conditions are very difficult.

There's beautiful work from Pavan Sinha that's emphasizing that coarse, contrast relationships in the face of very good heuristics to do that. The reason is that the 3D structure of our face stays the same even when the lighting conditions are changing. And so no matter whether light is shining from, typically the eye regions because they're recessed relative to nose and forehead, are darker than nose and forehead. And so he found that in the human psychophysics you have 12 heuristics like this, forehead brighter than left eye, forehead brighter than right eye, nose brighter than mouth, and so on and so forth. Twelve of these characteristics that actually together can allow you to detect the face. And in fact, your face detector on your cell phone is using a very similar strategy as well. Trying to find the coarse, contrast relationships in the scene.

So what Shay did was he started with a real face, he would parse it into 11 parts, and then randomly assigned 11 different luminance values to these 11 different parts and change this rapidly. And then the analysis would look like this. So no matter what the overall pattern looks like, he's going to look for a particular contrast relationship with the forehead versus the left eye. And he's going to ask is the neuron responding differently in these conditions where the forehead is brighter than the left eye versus in these conditions where the left eye is brighter than the left forehead.

And you can do this for all pairs of combinations of 11 different parts. And in these 55 different combinations, we can mark by arrows the prediction from human psychophysics. So human psychophysics told us 12 of these contrast pairs are going to be important. It also told us what the polarity was that was going to be important for detecting a face. Again, forehead brighter than eye, or eye brighter than forehead.

OK and this is what we actually found. So what this shows is a population diagram of all the cells that Shay found and it was half of the cells that he recorded from that showed some

selectivity to some of these contrast features and to some of these contrast polarities. What's plotted here upwards is for one contrast polarity and one particular contrast pair, the number of cells he found that were selective. And as you go through the entire diagram, you see that there are only a very few examples of contrast pairs where actually different cells like different polarities.

Like here for example you have more than 60, 70 cells that all like one polarity and not a single one that likes the opposite polarity. And this is true for all these polarities. So it's a very consistent pattern. Second, we can explain all the human psychophysics preferences here. Not only that, these are important dimensions for these cells. But in all these cases the cells care for the exact same polarity that they would have predicted from human psychophysics. In addition, there are other contrast pairs that apparently don't matter this much in human psychophysics, but these cells also care about.

So they seem to be using these coarse contrast features. And again, they're very useful for face detection. Now we got the behavior that we know that the areas involved in the face detection with stimuli where it's actually hard to make out the detail. Then Shay did a control and I thought this was really the coolest thing ever. He's a computer scientist and so of course he knew about databases and how to use them. And he would say, OK, can we actually fool these cells into responding to non-face stimuli that comply by the rules of this coarse contrast of faces?

So here are some examples. This is just a pattern where there are some dark regions where the human face eyes might be and so on and so forth. This is a pattern that only has one of these 12 contrasts correct. But also in human faces, you can find some-- when a person is smiling, wearing glasses, most of this contrasts are actually not in the face that should be there. So there are some very contrast correct faces and there are some faces that are not very very contrast correct. And now you can ask how does the response of the cells in the middle face patch change as you are increasing the number of correct contrasts, either on the face or non-face stimuli.

And the answer is this. If you increase the number of contrasts in the face, the cells respond more and more. If you change the contrast of non-face objects the cells don't care. So there is something else that caused contrast that the cells care about, they're not easily fooled to respond to things that clearly aren't faces when the coarse contrasts are correct. And we could actually have predicted something like this should happen from an earlier study. So the first

study we did where we took advantage of the fact that in a face selective area we can record from faces over and over again and that they have similar properties was a study where we looked at the effect of part and whole.

This is one of the central features in psychophysical of human faces is that you can get information from the face without any detail just from the gist of the face. An example is again from Pavan Sinha. If you have a blurred face of a familiar individual, like some of you might recognize Woody Allen here, of course with the glasses it's a little bit cheating, but you can recognize him. And the other examples in the study were people who don't wear glasses, so you can recognize a famous face just from the gist of it. You don't need the details.

On the other hand, we can process details. We can focus on details. And so how do these two things relate to each other? What we did was we constructed a face space, a cartoon face space, based on very, very simple geometric shapes. So these faces are just made out of ovals, and triangles, and lines. So very simple geometric shapes. But if they are put together they actually look like faces. Now we can parameterize this face space, we can vary certain parameters. So we had faces that change in aspect ratio. They go from Sesame character Ernie here to Bert. We have pupil size, like no pupil here, very big pupils. We have inter-eye distance here. So these eyes are close together almost like cyclopedian fashion or they can be very far apart from each other, stretching the outside of the face and so on and so forth.

And we would now randomly change these features, all these different features dimensions, randomly choosing a new value every time we show this face. It looks a bit like a cartoon character who is trying to talk to you. But the way we analyze this is very simple. So we just asked no matter what these other features are, for the first feature dimension mention does the firing of the cell change as we're changing this feature dimension. Then we asked this for the second dimension, the third dimension, and so on and so forth for all these slightly different dimensions.

What we found is shown here for one example cell. So we had 19 different tuning curves. And of these 19 different tuning curves, for four are significantly tuned. For this particular cell it was face aspect ratio so didn't like Ernie, it liked Bert. It liked the eyes very close together, not far apart. It likes the eyes a little bit narrow, not wide. And it liked big irises, not small ones. What's very typical for how the cells are processing these features are these ramp shaped tuning curves. So more than $2/3$ of the tuning curves have this ramp shape. Which means that these cells are relaying the information that they are measuring almost in one to one fashion.

This is not what the cells are actually doing. It's just a metaphor. But it's almost like they're taking a ruler, measuring eye distance, and they're relaying this feature in almost one to one fashion in their output. Another implication is that most of your coding capacity is actually at the extremes because many cells have big responses there, many cells have small responses. So most of the capacity is there. That's the range where caricatures live. And oftentimes we are better able to recognize individuals based on caricatures than the individuals themselves.

So in the middle the face patches, they're causally and selectively relevant for face detection. The cells are virtually all face selective. Based on these two findings we actually suggest and it's a little like Nancy said as a strong signal you're putting out that you get backlash if there is. So we do think that there are modules that are there for face processing and face processing only.

The gain of the tuning curve is modulated by the presence of the entire face. There's this ramp shaped tuning, which is very useful. The cells are sensitive to contrast relations which is very useful for face detection. So we can really get mechanistic about understanding face recognition. It's not just that we can say, OK, these cells are responding more to faces but we can say why they're responding more to some faces than to others. In fact, you can predict from the cartoon results how the cells are responding to pictures of actual people with the very fine details physically.

And so at the level of the middle face patches we already have some of the requirements for face recognition system. So we have mechanisms for face detection, we have some encoding of facial features, and we have encoding of configurations.

Nancy said I should talk about this. I'm going to talk about this. Sebastian Miller was a wonderful grad student with Doris and me. He asked the following questions. So whether the face pictures are related to each other or not. If you look at the overall organization these face errors are very far apart from each other. So the most posterior to the most anterior is one inch apart, it's a third of the entire extent of the primate brain. They live in different cytoarchitectonic environments. So you could also imagine that maybe the connectivity is mostly local.

On the other hand, they are interested in faces and so you would imagine that maybe there are specialized connections between them. The way we addressed this was with micro

stimulation inside the scanner. So we would first image the face areas. We would then lower an electrode to one of the face areas here, we record from cells to make sure it's face selective, but then we would use this electrode to pass a current through inside the scanner. Passing a current through an electrode is going to activate cells. That in turn is going to change blood flow and oxygenation. So things we can pick up on the scanner.

And so yes, if this worked you should get a swath of activity around your stimulation site. But if these cells at your stimulation site have predictions that are strong and focal enough to drive down some neurons, you might also find activation that spatially remote locations and they can then see where these locations are related to the face areas. So here is a computer flattened map of the brain of one macaque monkeys. The green areas indicate by outline the extent of the face areas. We placed our stimulation electrode in one of the face areas and this is the map we got from micro stimulation, versus no micro stimulation.

So there's no visual stimulus there, actually that works during sleep during complete darkness. So yes, we get a swath of activity around the stimulation site, and we get multiple spatially disjunct regions that are activated. So they're strongly driven by the cells in this region and these overlap with the face areas. And this was either be found very consistently across different phase areas if you stimulated outside you also got this patchy pattern of connectivity, but now it's outside of the face system. And so this is the picture that we got. Is that yes these face areas are actually part of a network of face areas that are strongly interconnected with each other. There's now data from retrograde tracer studies, we find that 90% of the cell bodies that are labeled after an injection inside a face area are inside other face areas or in the same face areas. So it's a surprisingly anatomically specialized and closed network.

So what's happening in these areas, and again, my movie isn't going to work. So in this area AL was more anterior, also virtually all of the cells are face selective. But you have a property emerge here which you didn't have before. And that is mere symmetric confusion. And it's something that we did not expect. We're still puzzled by it. We have no explanation why it's happening. But in this area you have cells that like a profile view. And if they like one profile view, they also like the opposite profile view. And this region here as I mentioned initially that some of the cells did not really seem to be face selective. It's a small percentage. But actually these cells are selective for facial profile. But if they like one profile right, they don't like left. If they like left, they don't like right.

In AL, this is being confused. And then if you go to AM, you have cells that respond to all

faces. It doesn't matter where they are, doesn't matter who they are, doesn't matter how big they are. And the other cells that also don't care where they are and how big they are, but they care exquisitely about identity. So they can be very, very finely tuned to identity in particular to people that the animals never see in real life. So there seems to be a computation going on from here to here where in Jim DiCarlo's conceptual framework you could imagine that there's a manifold that's now becoming sort of flatter and more like a more explicit representation. And for some reason creating this has to go through this mirror symmetric confusion.

I just want to highlight. So we meant to touch upon the question whether a face area should do different computations from non-face areas. Actually, my intuition about this was actually quite the opposite. So I thought they would likely do the same computation or hopefully the same computations as outside the areas just in different material. So why in other non-face areas? Why don't you want to mix these cells? We have one study that's a little too complicated to explain here that gives some clues. But some computation work from Joel Leibo was a grad student with Tommy actually gave some clues to that.

So Joel and Tommy were thinking about invariance. And Tommy told me he's going to talk about this at some later point in the course. The different kinds of transformations, easy ones and difficult ones. The easy ones are affine transformations. We're just shifting something in space or in size, or we rotate it in plane. And so if you learn how to correct for this transformation for just three dots of light, you can look for any image that you can ever see. So this is relatively easy.

But then they're non-affine transformations that are actually changing the picture and they are very difficult. So if you change your facial expression for example, or if lighting conditions are changing, or if you're turning your head in depth, this is a non-affine transformation. So it's not predictable from just three dots. And you can learn something there that could tell you how this picture would look like under this non-affine transformation. And one of the insights from Joel was that if you learn this non-affine transformation on a particular object category-- let's say faces-- you actually have learned nothing about another object category like cars. That's actually quite surprising to me but there could be a reason why you might want to have all the cells that have to learn representations across one transformation put them all in one location.

The second insight they had, and I think it's still very surprising to me that it actually works so easily, is to give the computation a count for the system that I just described to you in qualitative terms. So we have three levels of processing. So we have a front end where cells

are useful for face detection, they're all very face selective. So you could think of this like a three level processing hierarchy where the level one is like a face filter that's just going to tell you it's their face or not. The top level you want an identification. And I didn't show the examples, again I hope with a connection to the monitor I can show you the actual movies. You have some cells that are very, very finely selective of facial identity. With pattern readout techniques you can read out identity extremely reliably.

If you now have like Hebbian Learning Rule, maybe Tommy is going to explain this to you more, you actually get something pretty magic. So you do get invariance at level number three, which is kind of what you wanted and might not be surprised by. But as a byproduct, you are getting mirror symmetry at level two. And that's something you didn't stick into the system and it just happens, not like magic, but there's an explanation for why this happens, out of very general assumptions about the system.

So the point I want to make here, this particular model could be wrong, but it's something about how knowing something about the overall organization of the system might actually reveal underlying relationships that you might not think about. So the fact that there are three levels of processing and not four or five or six might actually impact whether you find mirror symmetry or not. Or whether you find mirror symmetry at one level or another level. And you don't necessarily get this automatically. Just put any processing system together.

Because I was mentioning a facial motion, I would like to give a brief vignette of that. And so I was emphasizing here transformations along this direction. But you can see that at least two levels of processing there are actually two face areas here, one lateral to the STS, and the other one deep inside. And so one of the questions we had was what's going on here? How are they different? And so one way you might think about this is again, connecting faces to social perception. So there are some faces out there that are not really faces. And so the faces of dolls are just one example. So physically they are faces but you can actually tell that they are not really real, they're agents.

And people like Thalia Wheatley are wondering about questions like why dolls are creepy. So there is an expectation that the face should belong to a real agent. And there are different clues that can give it away. So if the face is on top of a body, more likely it's an agent just like an artificial stimulus. If a face is moving that's another clue it's an agent. And this will change fundamentally how you interacting with it. Your interaction with the doll is likely going to be very different than with a baby. And so again, objects and their meaning are making them

actionable in different ways. And we have to understand what the circuits are that actually make this possible.

So one way to look at this again is to think about these facial displays. I showed the Tonkin macaques and Clark Fisher was an M.D. PhD student in the lab was actually addressing this question. And so he made movies like this one here. Luckily there's no sound so we can actually play them. These are movies of macaque monkeys making facial movements of all different kinds of facial expressions. And we then also have stills that are just changing from time to time. And we have controls of toys that the animals also know that are either moving or that are jumping every second from one state to the other.

And we would ask as we had in an earlier study, are these areas responding to this motion differently than to the static images. So here's what he found. So he had six different face areas he was looking at. If you're looking at static form selectivity we're just reproducing the way that the area was found just with a different stimulus. So all the six areas respond more to faces than to objects. If we now compare moving faces to static faces, all the arrows are responding more to the moving faces than the static ones. Some quantitative differences, but overall the same pattern, more responsive to moving than static stimuli.

If you now compare on the right inside the modulation by moving objects through static objects. You can see that also all the areas or almost all have a slight advantage of moving objects over static objects. There seems to be a general motion sensitivity there. But if you do the interaction of shape and motion, you can see that all the areas are selectively more enhanced by face motion than by non-face motion. But they all look pretty similar.

So now you can actually wonder if you have a contrast like this, like moving versus still, there a couple of things that are different. So is it really about motion or is it just about the content? If you just show a picture every one second you can say, well, there's less content there, therefore you might have more adaptation, therefore less response across the board. Is it about update frequency? You know, a fast update versus a slow update. So what Clark did, he was creating another stimulus. If you think about creepiness, it's actually like a little bit creepy. So like a scrambled version of a motion that's shown here. Shows us the same frames of the movie but just randomly associated with each other. So if anything, now the motion energy in this thing is higher than in this one here. And we can look now for the contrast of those two and also the contrast through here.

And now what he finds is actually something where the face areas are qualitatively different. He finds two face areas here which are responding more to the natural motion than unscrambled motion, and three face areas that are responding more to the scrambled motion than the natural motion. So they shift opposite preferences. What we think is going in the areas here that remember also the benefit for facial motion over static faces, that they just like a fast update of content. Ideally in a way that's not predictable. If you show me something new I'm going to respond. And if it's something that I can't predict even better.

So this is what these guys are doing. But these ones here, they seem to be really sensitive to facial movement and naturalness of facial movement. And that was not smart. And these areas are located more deep inside the STS, more dorsally, and these areas are located more ventrally. So there's an organization and to discover the new face area they didn't know before. It's a seventh face area which he called MD which is really like a face motion area.

There are lots of reasons why we're excited about this. I mentioned the link between face perception and agency interpretation. This is one possible link, there are more. It might be a second phase processing system, I'm going to go through all the evidence. It's only indirect, though that this area might not be connected to the other ones. I told you before that the six face areas are intricately integrated to be one network. This area never showed up in stimulation, therefore it might be separate. And this is kind of nice because it fits very nicely to the human situation. So in the human brain you have the posterior STS face area which is exquisitely sensitive to facial motion. Actually you often don't even get it for static faces. But it's very sensitive to facial motion and actually Nancy has a beautiful study on that.

And this area by several accounts is not like the other face area. So it seems to be like a specialization. That's another thing, another reason why we think these systems might be connected to each other. Just a cool thing in the end, who can recognize this actor here? Show of hands. OK, who can recognize him now? So facial motion gives away a lot of things like identity. Jack Nicholson has very typical facial movements. So it's not just agency, it's not just facial expressions, it's also identity and lots of things that can go away. So we actually don't know yet what these areas are doing.

So my summary, and I'm sorry I'm going into lunch. So we can do fMRI on macaque monkeys just as in humans. What we find we can apply to lots of domains within attention studies, found new attention areas. Here we applied to face processing, we find face selected areas for face processing. Recording the micro stimulate-- so inactivating these regions is supporting the

notion that these are likely modules that are selective for processing faces and faces only.

These are interconnected into a face processing network. It looks like all these areas have different functions and specializations. So fMRI experiments are notoriously under powered in the way of number of different dimensions. So if we call them face areas it is not to say that they're all doing the same, but they likely all have different functions and likely sub regions of these different functions. So again, there's no contradiction here to the view of a fine organization. Then there's a seventh face area which doesn't seem to be connected. Could be separate and separate face processing system. We have evidence for processing that we can understand now in computational terms. And this is one way that it can link, sometimes causally sometimes correlationally, activity of single cells. Two different levels of organization to a very complex social behavior. And that's, again, I think a very cool opportunity to have in the domain of social cognition that you can actually control stimuli very well because faces are so powerful to get into your social brain, you can likely take this approach deeper and get insight into actual social intelligence beyond face perception this way.