[DIGITAL EFFECTS]

NANCY KANWISHER: So to remind you, we've been talking last week about doing two things at once-- asking all sort of questions of what we might want to know about face perception in the brain-- there are some questions. But at the same time, the agenda has been to consider the different methods available in human cognitive neuroscience and what kinds of questions each one can answer. So last week, we talked about a bunch of them, and today, we're going to wrap this up talking about TMS and animal studies.

But first, I just want to remind you very briefly-- I won't go through in excruciating detail-- we talked about behavioral methods, which are great for characterizing internal representations, as you saw with face inversion effects and some of the other behavioral data, they have major disadvantages, which is that with behavior, you're just measuring the output. It's pretty sparse. And from that, you have to infer all the stuff that happened in between the retina-- or whatever your sensory modality is-- and the output. All that internal mental stuff you have to infer just from the output.

So it's amazing that that works at all, and you have to be really smart to do it. And lots of people have been doing that for a long time, but it's challenging. So why not look inside? And one of the best ways to do that, of course, is functional MRI. It has the best spatial resolution available for normal subjects. But as you guys all seem to pick up on, its temporal resolution is lousy, and its ability to tell you whether the neural activity you're looking at is causally involved in behavior is like nil.

Now, at least one of you-- I only read a few of the assignments-- but at least one of you was confused about which causal role we're talking about. And this is actually really important. So let's take a moment to talk about this. Causality, the idea of causality-- if x causes y, that means, essentially, y wouldn't have happened without x, or y happened more because x happened, when x happens than when x doesn't happen. So that's pretty basic. So that means if you want to test the causal role of x on y, you have to mess with x. That's the key challenge.

OK, so with that in mind, here's this whole causal chain. A stimulus lands on a retina. A bunch of neural activity happens, and some behavioral output happens. So there's a whole causal chain there.

Now, let's consider what kind of causality we're talking about. There's one kind which is that the stimulus causes neural activity in the brain. That's a kind of causality that we can absolutely test, even if we're measuring that with functional MRI, because we can mess with the stimulus. We can present different stimuli and produce different neural activity, OK?

So in that case, we can look at the causal effect of the stimulus on the neural activity. No problem. That's standard. That's what we do, pretty much, in every experiment with ERPs, or functional MRI, or so forth. Is that clear?

OK, on the other hand, if we want to know this kind of causality from some neural activity we measure in the brain to either a behavioral response, or a subjective feeling reported by a behavioral response, or something like that, that's the challenging part. That's the kind of causality that we can't infer from ERPs or functional MRI. Everyone got that? Yeah, OK, it's sort of obvious and not obvious.

OK, so let's talk a little bit more about that temporal resolution. I know I kept saying the temporal resolution of functional MRI is lousy, but I had run out of time and skipped through the key slide, So let me back up and do that here. This is the BOLD or MRI response as a function of time. This is an idealized version of it, but it looks kind of like that.

And sorry these things are tiny here, but these are seconds-- 5 seconds, 10 seconds. So let me show you what that means. If you're recording neural activity, back here, in the first stage of visual processing in the cortex coming up from the eyes-- which is where?

AUDIENCE: The occipital lobe.

NANCY KANWISHER: The occipital lobe, yes. What area?

AUDIENCE: Primary visual cortex.

NANCY KANWISHER: Primary visual cortex, exactly. OK, so suppose that we stuck an electrode in my primary visual cortex, and we flashed up a very brief visual display. OK, here's a visual stimulus, on for maybe a tenth of a second-- bright, flashing thing V1 loves-- V1, also primary visual cortex, right?

OK, the neural activity would happen in less than 1/10 of a second-- super fast. We know that from work in animals and even some work in humans, OK? So super fast after the stimulus. It just goes straight up from the retina, the LGN V1-- boom, there it is. So all the neural activity happens right there, and it ends right there.

But the MRI response is five, six seconds later-- this big, sloppy, slow thing as the blood slashes into V1 many seconds after the relevant neural activity. So what's relevant here is not just that it's delayed, but it's big and sloppy. And so both of those things are the reasons why functional MRI responses aren't good for distinguishing what happens on a fine temporal scale of, say, events less than a second.

All right, OK, in contrast, as I mentioned, when you glue electrodes on the scalp, or stick these fancy magnetic sensors in the big hairdryer device around your head, there you get beautiful temporal resolution, but it's like the Heisenberg principle of cognitive neuroscience. You want time, you don't get space.

OK, and similarly, here, we can measure the causal effect on scalp response neural responses, but not the causal role of those neural responses on behavior. Everyone clear with this? OK.

OK, then I talked about these rare cases where we can record directly from the surface of the human brain with electrical activity, where we now get both space and time at the same time. And the key disadvantage there, of course, is that it's extremely invasive. You have to take a big piece of skull off to get in there. And, of course, that would only happen in the case of people who are already in pretty serious medical circumstances.

OK, so now, when we have this incredible opportunity to record this amazing data from the center of the brain, does that enable us to make this kind of causal inference from neural activity to behavior? Yes? What do you think, yes? No? Isabelle? Is that Isabelle? Yes. Why are you shaking your head?

AUDIENCE: Because it just tells us which neurons are responsible for [INAUDIBLE].

NANCY KANWISHER: That's right. It's cooler, it's fancier, it's more impressive than functional MRI or ERPs, but it's still the same deal. We're just recording responses, OK? So we can do this causality, from the stimulus to those neural responses, but it doesn't tell us which of those responses are related to behavior yet. I showed you other methods that do, but this one alone doesn't. Everybody got that?

All right, so then I talked about studying patients with focal brain damage. And here, you really can make a strong causal link between a bit of brain and a behavioral ability. You lose that bit of brain, you can no longer do that task. That's a really direct kind of causal role.

I talked about double associations. I gave it short shrift, but it's actually really important. You should know it. A double dissociation is when you have one patient who can do A but not B-- say, recognize objects but not faces-- and another patient who can do B but not A-- say, recognize faces but not objects. And when you have in the literature two cases like that, now you're in a really strong position to infer that there's something fundamentally different about face recognition and object recognition in the brain.

OK, so that's really important-- the senses in which a double association is more inferentially powerful than a single association. OK, "more important" means I'm sure to test you on it. No, it's also important, whether I was going to test you on it or not. [LAUGHS]

OK, and so, of course, in focal brain damage, we can absolutely infer causal role from a bit of brain to a behavioral ability. Lose that bit of brain, lose the ability, yeah? OK.

And the case that I showed you with that amazing movie of the guy getting stimulated in his fusiform face area and seeing percepts of faces on top of whatever he looked at, that's a quintessential beautiful example of the causal role of neural activity there. We're basically directly manipulating neural activity. We're injecting neural activity there electrically and looking at the behavioral and cognitive result that occurs-- The guy. Sees a hallucinatory face.

OK, now, that is amazing data, but as I mentioned, they're very rare. We have no control over it. When we get those data, we celebrate and are all excited, but mostly, we don't get those data. Plus, those people have serious problems with their brains. That's why their brains are being opened up.

So is there any way to test a causal role of a particular part of the brain in a normal subject who doesn't have their skull open for neurosurgery and who has not had brain damage? Well, there's one way, and that's called transcranial magnetic stimulation, OK? So in transcranial magnetic stimulation, you take a coil of wire about yea big. That's a tight-wrapped coil of wire embedded in plastic, connected to a ginormous capacitor, and you hold it next to your head. Of course, that's what you would do if you were a neuroscientist.

And you discharge and make an enormous current through that coil that's very, very strong and very brief. The whole thing lasts less than one millisecond. And you guys know from 8.02, another case of the right-hand rule coming to our service. You have a hell of a current going in a coil. What's going to happen in brain tissue underneath?

AUDIENCE: Increase the magnetic [INAUDIBLE].

AUDIENCE: The electric field will [INAUDIBLE].

NANCY KANWISHER: Yeah, exactly. And so you'll get electric fields perpendicular to the coil sticking right into the brain like that. And what do you think happens if you stick a big, huge transient electric field-- boom!-- into your head like that.

AUDIENCE: Isn't that a magnetic field [INAUDIBLE]?

NANCY KANWISHER: Yeah, you're right. Right-hand rule is magnetic field. I was thinking I was misremembering, right? Electric current makes magnetic field, right? It was a long time ago I took 8.02. I did-- just a long time ago.

Anyway, for current purposes, doesn't matter. Either would do it. Actually, there's a variant of this where it's an electric field, but it's debated how well that works.

OK, anyway, what happens is you affect neural activity in tissue right underneath the skull, right? OK, so if you want to see a picture, a video of that happening, there's a video of me getting zapped with TMS on my website. You can check it out. It's kind of ludicrous. Yes, question?

AUDIENCE: What's the spatial resolution?

NANCY KANWISHER: Oh, we're getting there. We're getting there. OK, here's an early version of this. To generate these very strong and brief magnetic fields, they had these stacks of coils like this, and they rotated them around. It's a little crazy. Here's a more recent version. It looks like a big torture device, but it's actually no big deal. The guy's just holding his head on a chin rest to hold his head still, and there's a person holding the coil next to his head like that. And so that enables us to briefly and somewhat selectively disrupt a little patch of cortex there by sticking in this big random field.

Now, spatial resolution is not amazing-- maybe 1, 2 centimeters, something like that, OK? It's better than you might guess for such an incredibly crude device-- like something people would have done hundreds years ago, and yet we still do it today. You can also use a lovely method where you scan the subject with functional MRI first, find a particular functional region that you're interested in in that person's brain-- remember, these things can vary in their exact location across subjects-- and then find a way to register externally on the scalp, what is the closest spot to that region you found in their brain previously with functional MRI? And stick the coil right there, and exactly titrate its location with reference to that brain image. So that makes this whole enterprise more worthwhile.

So what can TMS tell us about face perception? Well, here's the problem. Here's my fusiform face area-- that guy right there. It's a few centimeters in from the scalp-- from the skull. So that's a drag. Unless we opened up my head, we can't reach it there with the TMS coil.

Believe me, the first time I had a chance to use a TMS coil, the very first thing I did was stick the coil there, crank it to the max, and try to see what would happen. Not a damn thing happened. It was very disappointing.

I knew lots of friends who tried the same thing. It was the most obvious thing. It just doesn't work; it's too medial. Yeah? Well, there was a question over here a moment ago?

AUDIENCE: If you use TMS near someone's brainstem--

NANCY KANWISHER: Yeah, that wouldn't be so smart. Luckily, the brainstem is kind of deep in there. So if you were really stupid and stuck it down, I don't know, way in here, you might be able to cause trouble. But mostly, people don't stick it back there. And actually, the subjects won't let you anyway, because there is a lot of neck muscles, and it really hurts when you do TMS over muscles.

And so if anybody had such a stupid ideas to try to zap the brainstem, the subject would probably object immediately before they got very far with it, because it would hurt. [LAUGHS] And you guys are all probably wondering, how safe is this? It's not totally clear. There have been lots of studies in animals--

[LAUGHTER]

--where they zap a rabbit 100,000 times or something like that and say, well, rabbit seems fine, hops around. And the best they can do in animal studies. When I first used TMS around 20 years ago, I read a few basic safety studies, and I thought, god, I don't know. But I also realized that if you look at the papers, the initials of the subjects were the same as the authors.

So I called them up, and I said, hey, tell me honestly, did you guys ever notice any ill effects from getting zapped? And the guy I talked to said, yeah, I've been zapped about 10,000 times, and I never noticed anything except for one thing. After a whole hour of getting zapped, it gave me a hell of a craving for ice cream. So I decided, OK, I can live with that.

We got it through the human subjects committee, and we do-- not a lot, but some TMS in my lab. And I'm probably now been zapped at least as many times as that kind, and I guess you guys can judge for yourself. So you don't have the before condition, so it's a little hard. Anyway, as far as anybody can tell, it's perfectly safe. Yes?

AUDIENCE: So there are some contraindications if you are prone to seizures or if you're on certain medications.

NANCY KANWISHER: Yes, yes, yes.

AUDIENCE: So if you ever sign up for a TMS study, read the fine print.

NANCY KANWISHER: Good point, yep. OK, so back to this. It would be lovely to zap that guy, but it's too hard to reach.

OK, so then, this guy David Pitcher came along, and he had a very good idea. And my 1970s synopsis of his idea-- paraphrasing still-- is if you can't zap the region you love, love the region you can. And so Pitcher said, hey, what about that other guy there? We haven't talked a lot about it. It's sometimes called the occipital face area.

I think of it as a kind of crappy version of the FFA. It's kind of face-selective. It's not as face-selective. It's more variable, so it's not as fun to study, but it's there in most people. I have a damn fine one, I have to say-- many people do-- and it is right out there next to the scalp, just asking for it. Right

OK, so here's what David Pitcher did. He gave subjects a-- you need a behavioral task, right? Because in this case, we're testing the causal role of a bit of brain on behavior. So we're going to measure behavior.

And so what is our task? OK, so here's his task. Sorry, it's a little tiny, but this is one trial.

Time is going this way. You present a face. There's a brief interval. You present another face. And the task is just, are those two faces same or different? It's your basic face perception task.

But then, what you can do is you can zap the occipital face area at different time-- during presentation of that second face, and you can do it at different time intervals. Remember, its effect is very brief. The actual magnetic change is less than a millisecond.

OK, so here's what David Pitcher found in that study. This is accuracy at the same different matching task when you stimulate the right occipital face area versus vertex-- that means you stick the coil up here, which is pretty far away from face regions. It's a control condition-- not a perfect one, but better than nothing.

By the way, TMS usually doesn't hurt unless you stick it over muscles. You stick it over the frontal lobes, and-- I don't know. Every time I try to disrupt my language abilities, it hurts too much, because there are muscles up there. But most places, like the top of the head, there aren't muscles, and it doesn't really hurt. But it still makes a loud cracking noise, and it's kind of like somebody went--

[TAPS SKULL]

--like that.

So you might imagine you need a control condition, right? Because if people are-- that also has a TMS pulse, right? If you bang somebody on the head when they're trying to do a task, you probably disrupt their performance. So you need to bang them somewhere else to see if it's specific to that location.

OK, so OK, so here's a little effective on the accuracy. It's not a huge effect size. So here, it's going from 85% correct to 78% correct when you zap occipital face area compared to vertex. Everybody gets what's going on here?

So that's good. That tells us something. Zapping here messes up face perception more than zapping here, OK?

OK, so that tells us something about causal role, but what else would you want to know? That's a beginning, but having just learned what I told you about TMS, what else could you do that would tell you more? Yeah.

AUDIENCE: You see the face [INAUDIBLE].

NANCY KANWISHER: Ah, well, that's a good question. It wasn't what I was fishing for, but it's a very good point. So this shows disruption, but I showed you with that video before that if you electrically stimulate the FFA, you see a face. Well, unfortunately, nobody has reported that when you zap a face area, you see a percept of a face. Boy, that would be fun if true, but it doesn't work.

And there's much debate about why. It probably has to do with the fact that your ability to target just that region is less good than it is with direct stimulation. There are many reports and many published studies where if you zap V1, you see a flash of light, OK? I don't see the damn flash of light. I've tried, and tried, and tried, and people in my lab who I trust promised me they actually see it. It isn't just BS. But I don't know; I don't see it.

Anyway, so probably, the question of when you get disruption and when you get a positive percept is a very interesting, complicated one. I think it will ultimately have to do with how those batches of neurons not only respond to faces, or light, or whatever, but how they code for that information, such that when you put a big artifactual, non-biological signal in there, will it have any meaning that the subject can interpret?

I don't know if that's helpful. I think nobody really understands that, when you get a positive percept. But I hope you can at least understand that at least if you mess with it and muck it up, you can disrupt. That logic is clear. When you will be able to actually stick in a signal and get a positive, coherent percept is a more subtle thing, OK?

OK, what else would you want to know? Yes.

AUDIENCE: Whether it messes up object perception, or not?

NANCY KANWISHER: Absolutely, absolutely. All we're showing here is it's messing up face perception. Maybe the guy can't see here. Maybe he's just globally blind. Maybe he'd have the same problem with object perception, absolutely. The assigned reading for Wednesday shows exactly that experiment, OK?

What else would you want to know? Remember, a TMS pulse lasts less than a millisecond. That enables us to ask a whole interesting kind of question. What else could we find out? Yeah.

AUDIENCE: Oh, just [INAUDIBLE].

NANCY KANWISHER: Yeah, yeah. I'm sorry. It's probably-- I don't mean to be insulting your intelligence. You're probably sitting there saying, this is too obvious. That's what I'm talking about. You can zap at different times and ask, when is the information going through there? When is that region playing a causal role in behavior?

And here is a very beautiful data that David got. And there's basically no effect at any point other than that interval between 60 and 100 milliseconds, OK? So that's cool. Tells you that's when that region is likely engaged in processing. Make sense?

OK, and is it Shardul? Yes, already made the point. I was going to ask you guys, does this tell us this region is specifically involved in face perception? Absolutely not. We'd have to test other things. It could affect every visual percept. OK, so you can read more about that.

All right, just so to collect all the advantages, it gives you strong causal evidence that a particular part of the brain is involved in perception or behavior. It has good temporal information, unlike studying patients with focal brain damage. And it is the only disruption method that can be used in normal humans, OK? So that's why, even though it's so crude and rudimentary, we use it, because it's the only thing that fills that niche.

A couple other unimportant things. Spatial resolution isn't as good as we'd like, but it's surprising how much you can learn nonetheless. And it doesn't reach very far below the scalp, although Ed Boyden-- the amazing Ed Boyden-- is working on a crazy new version of it that might. OK, so where have all this menagerie of methods gotten us to?

I won't go through all this in detail. We listed all these questions before. I gave you some of the answers from previous methods. The ones we've just talked about show, for example, that the fusiform face area-- or the occipital face area. In the case of TMS-- are causally involved in face perception, apparently not in object perception, pending the paper you're going to read. And so that's important, because it says when we try to come up with theories of how face recognition works, we might think about having a different theory for face recognition from our theory of object recognition.

OK, so this is all magnificent and wonderful, but I finessed this list of questions so that the methods would be able to address them-- at least a little bit-- and I sneakily left off a whole suite of other questions that are extremely important-- arguably more important-- that those methods don't address, OK? So we want to know not just that a region responds to faces. We want to know exactly what is represented in that region or other regions that respond to other things.

We want to know, what is the neural code for faces? We want to know, what are the actual computations that go on in a given region, how do they unfold over time, and how do those computations produce the representations and behavioral abilities that we measure? We want to know, what are the actual anatomical connections?

I showed you that little occipital face area right nearby but discontiguous from the fusiform face area. I've wanted to know for 20 years whether those damn things are connected anatomically. Shockingly, we still don't know that.

We want to know what is the causal role of each region in perception. And I showed you a few ways that we get little bits of data-- kind of, sort of-- but there's a lot of cases where we don't. And we want to know, how does all this stuff get wired up over development, right? What is the role of experience? Do you need to see faces to wire up the phase region, or is it there at birth before you ever see a face?

The sad truth is, for the most part, we don't have good methods to answer these questions in humans. So that's just a big bummer, but it's true. Most of these questions can only be answered by research in animals, or can be best answered by research in animals.

So I'm going to take a moment to talk about ethical issues in animal research, just to note that I think there is an issue. And I'll say that it's not unreasonable if you have qualms. I noticed in an earlier lecture, I started talking about recording from animal brains, and I didn't have time at that moment to mark this, but I do think it's important. If it makes you uneasy, that's totally legitimate. You should think about that, and respect that, and think hard about whether that's-- what you make of that.

Unambiguously, causing animals pointless suffering is just completely unacceptable, OK? So I think we can all agree on that. And I think there's a very difficult trade-off between avoiding suffering in animals and research that has saved countless lives. So people can legitimately come down on different sides of this, but many lives have been saved-- including mine-- based on animal research that enabled treatments that were life-saving.

And a few things to think about to help you inform how you handle that trade-off. First of all, know that animal research in the United States is very heavily regulated, OK? So animals receive excellent vet care-- shockingly, better than probably lots of citizens of this country. That's another topic. Also, there's a very major emphasis on avoiding pain.

So I think it's probably generally true that it's infrequent that lab animals suffer a lot of pain. Researchers and vets are very careful to avoid that. So the bigger issue is not so much are the animals physically suffering from pain, per se, but what kind of life is it to live in a lab and be a lab animal? And I think that's a legitimate question.

For monkeys, at least, where I, at least-- and maybe in my speciesist bias, being more sympathetic to similar species. I don't know if that's legitimate, but it's a natural. There are increasing efforts to improve the quality of life for monkeys in labs. Many monkeys are now housed in social groups where they can hang out with their families, and that certainly improves their quality of life.

Many monkeys basically play video games all day. In DiCarlo Lab, they're studying visual perception, and what do they do? They get the monkeys in there basically doing visual tasks in exchange for juice rewards. Not all that different from what, probably, lots of you guys do.

Now, maybe they'd be happier in nature. Probably, much of the time, they'd be happier in nature. But I think that's complicated, too. Nature can be pretty nasty. So it's not totally obvious that quality of life in a random lab is worse than quality of life in nature.

The third point, I'd say, is that the benefits of research are forever. You discover something major about how brains work, that's forever, right? So you've got to amortize whatever cost of animal suffering there is against the forever-ness of that insight. And so in my view-- not that you need to agree-- but in my view, animal research is vastly more justifiable than things like eating meat or buying leather, which is just transient entertainment or convenience, right? So anyway, you guys, I encourage you all to think hard about this and to come to different conclusions. I just wanted to note that these are issues that are worth thinking about.

That said, the methods in animal research are breathtaking, and they get more and more breathtaking every day. In this building, people are constantly inventing astonishing new ways to answer all kinds of questions. And I wanted to give you just a gist of some of the kind of stuff that you can do to answer that list of questions that I said we can't really answer in humans.

So just very briefly-- this used to be a whole lecture, but I've decided to cut it to one slide-- very briefly, about 10-plus years ago, these two amazing people, Doris Tsao and Winrich Freiwald-- who, mark my words, will get a Nobel Prize someday, or at least they should, and they might-- they popped a monkey in the scanner and did the very same experiment that we do on humans, OK?

So here's a monkey brain. Again, the cortex is unfolded so you can see the whole surface. The dark bits are the bits that used to be inside a fold. The little yellow patches are the patches that respond more to faces than objects-- just analogous to the FFA in humans, but there are six little patches in monkeys. OK, so that's so far, that's like, OK, fine, monkeys have them, too. That's cool.

But the thing is, because that's a monkey, you can then stick electrodes straight into that region right there, and you can record from hundreds of neurons in that region. And you can record the response of each of those hundreds of neurons to hundreds or thousands of stimuli. You can characterize the neural code for faces in monkeys in a way that you just can't for humans. In fact, Doris now published a paper last year called "The Neural Code for Faces," based on a decade of this research. It's quite breathtaking.

OK, second, you can watch those representations change, those neural population codes change over time. You can see, at one time point, what the code seems to be saying here, and then here, and then here, and you can watch that-- those codes-- change over time in each of those regions. And you can see different representations in each of those regions. It's quite breathtaking.

You can answer this question of, what are the anatomical connections between these regions, with a whole bunch of different methods that I won't go through here. But you can actually answer what's connected to what. And what these guys have found is that all of those yellow face patches are connected to each other by long-range connections that go through the white matter underneath the gray matter. Those regions are not connected at all to the intervening other patches of cortex.

So that set of six little regions is like a computational unit with different hubs that talk to each other. And you can see all that in monkeys in a way that we still don't know in humans. You can electrically stimulate, or disrupt with other methods, any one of those patches one at a time. You can disrupt them for 50 milliseconds here, 200 milliseconds there, whatever you like.

And you can study this whole system over development. How does it change from shortly after birth to monkey adolescence? And you can control experience during development. You can raise monkeys without ever letting them see faces, and ask whether seeing a face is necessary for the development of that region. We'll talk about that in a few lectures.

My point is just that with animal research you can answer vastly richer, more sophisticated questions than you could ever answer in humans, and that's just life. Yes, what's your name?

AUDIENCE: I'm Esther.

NANCY KANWISHER: Esther, hi.

AUDIENCE: So in these experiments, they showed them monkey faces, right? Not humans?

NANCY KANWISHER: Done all different ways. Remember, monkeys see other monkeys, but they see a lot of humans, too. And monkey face patches respond pretty similarly to human faces and monkey faces. Human faces respond pretty similarly to human faces and monkey faces, too-- even if you don't work in a monkey lab. OK, so just to say that there are loads of other methods, and we'll get these later in the course.

OK, so that snake assignment, I hope that seemed-- I thought you guys, for the most part, did very well and did exactly the kind of things that we had in mind. And I just want to go through a few bits of terminology, because I realized, some of you who messed up the wording, I hadn't really fully explained what the different words mean.

OK, so first of all, there's this incredibly boring words of independent variables and dependent variables. And frankly, I didn't know which was which until I started teaching this stuff a few years ago. But the concept is really important.

An independent variable, that's a factor that you, the experimentalist, manipulate and change, so that you can then measure what effect it has on a brain or behavior. The effect you measure is the dependent variable. The independent variable is called the independent variable because you, the experimentalist, get to mess with it, get to manipulate it, OK? The dependent one, you're measuring its dependence on the independent one. So just basically, in the experiment, you muck with something in the world, and you measure the consequences. The thing you muck with is the independent variable. The muckee, the thing you measure the effect on, is the dependent variable. Make sense? OK.

All right, so for example, the bold response, that's a dependent variable, and pretty much all the experiments we'll talk about here. All right, the hypothesis, most of you got that. The hypothesis is the statement about the world that you're trying to figure out if it's true in your experiment, OK?

A prediction-- most of you got this, but let me just say, a prediction is supposed to be extremely precise. It's the exact statement of what you will see when you measure your dependent variable if the hypothesis is true. What is the crucial thing you have to look for in the data you measure that tells you if the hypothesis is true or not? And the prediction is what you will find if the hypothesis is true, OK?

Confound-- we haven't talked about this yet. A confound is a difference between your conditions that you're manipulating other than the one you intend to manipulate. And hence, confounds give you alternative accounts.

Case in point, we compare the response in the brain when people look at faces versus when they look at a bunch of random objects. The fact that the faces have more curvy surfaces, or are animate, or are more interesting, those are all confounds with respect to the hypothesis that that region is responding specifically to faces. Everybody got that?

OK, it's very common, amid undergraduates, to use confound to mean anything bad about an experiment. That's not right. A confound is a very particular thing. It's another dimension that co-varies with the thing that you care about. It's like a nuisance variable that's correlated with the thing you're manipulating, and hence is giving you a difficulty inferring a clean inference from your data.

All right, a contrast. We talked about activations in the brain, like those little yellow patches I showed in monkey brains a moment ago. That shows you the bits that responded more in functional MRI when that monkey was looking at faces than objects. The contrast is faces versus objects, right? It's looking for a higher response in one condition than another.

Make sense? OK, these should all be fairly clear. I just know that not everybody got this.

OK, now, the point of a contrast is to isolate a mental process, OK? So let's talk about that for a second. So how do we decide what contrasts to use?

OK, well, first thing you have to do is get clear about your hypothesis. State it explicitly. Most of you guys did that really well.

Often, your hypothesis-- with functional MRI, at least-- will concern a particular mental process that you're studying-- like face recognition. Now, remember, importantly-- I said this briefly way back-- functional MRI can only tell you about differences between two conditions. The absolute number, you're going to measure the MR signal intensity in one condition-- say, when people are looking at faces-- and it's going to be something like 726. And it's totally meaningless. That's just how strong the MRI signal is from that point. It doesn't mean a damn thing on its own.

But then, if we also measure, in that same part of the brain, the MR signal intensity when the subject is looking at objects, and it's 720, then now we're in business, OK? All right, so everything is a difference. So that means that in any imaging experiment, you'll need to compare two or more conditions. One condition will never get you anything.

And if you want to isolate a particular mental process, you need to turn that mental process on or off, or you need to vary how strongly it's turned on. So this is all in the service of, how are we going to decide what contrast to use? That's our goal, is to turn on or off one little thing.

OK, and here's the problem. If I told you, OK, look at my face, and don't process low-level visual information, and don't think about what I'm saying. Just see my face. It's like, what? You can't do that, right? There's a whole processing chain. You can't just do one little mental process at a time.

And so that means we can't just have a task where you do only mental process x, and a task where you don't do mental process x, If you're not doing other stuff. So what that means is we need to choose two tasks, each of which has lots of mental processes, but that differ in only one. And then, we can compare those two.

So this is called subtraction logic, and it comes from work over 100 years ago in cognitive psychology and people who were just measuring behavior. This dude, Donders, he's a Dutch physiologist, and he invented the subtraction method to measure reaction times in humans, way back. And so with functional MRI, we're doing the same thing.

So we're going to come up with two different tasks which involve the whole suite, from input, to mental processing, to output. And yet, we're going to try to make them differ in just one particular mental process. Everybody with the program here?

OK, all right, so what you aspire toward in the contrasts that you choose is something called a minimal pair, right? So the idea is we're going to have these two tasks that are identical in every respect, except for that one thing we care about, OK? So here's a task, and here's a task. This one involves snake perception, and this one is identical to this one, except for snake perception. That's what we want.

OK, and if you get those two things, that's called a minimal pair. And this is the single most important thing in experimental design. All the other stuff-- like how you arrange your stimuli over time and all that kind of stuff-- OK, it matters a little bit, but this is the crux of the matter. What are those conditions, and are they the right kind of minimal pair?

And you guys got the gist, but I felt like most of you didn't really engage. OK, what exactly were those non-snake conditions? So that's really the crux of the matter.

So the most common problem with imaging experiments is not that the scanner wasn't as fancy as it could have been, or they didn't use the latest cutting-edge analysis method. The most common problem is that people's contrasts-- their conditions-- were not designed beautifully enough to isolate a single mental process, OK? That is that the conditions were not minimal pairs. Any other difference between the two conditions other than the one you intend is a confound.

All right, so let's engage on this. Now, if we ran a whole experiment only on male subjects, is that a confound? No. Why not, Isabelle?

AUDIENCE: Because it's not a difference between the two experimental conditions.

NANCY KANWISHER: Yeah, it's just a bad design feature, or something that limits your ability to draw inferences. Again, sub-optimal design it's not the same as a confound. A confound is this very particular thing.

OK, if all the snake pictures have grassy backgrounds, and all the non-snake conditions do not, is that a confound? Yeah, exactly a confound, right. OK, so I just said all this, so I'll stop boring you.

OK, and the reason that the grassy background thing is a confound is it gives you an alternative account of that contrast. Maybe it's grassiness, not snake-ness, that's the key difference. You don't know.

OK, all right. OK, so all of that said, minimum pairs are like a platonic ideal of experimental design. What you aspire toward, but you can never really do it. If the two conditions were identical except for this one little thing, they'd be identical. You can never totally pull it off, but you can track the little ways in which you fail, and you can test them one at a time in later experiments, OK?

All right, good. All right, so here's what we're going to do. We're going to break into groups, and you guys are going to think how to take the kind of designs that you already put together and turn them into actual experiments-- which is going to require deciding on a whole bunch of other things, and then we're going to discuss the things you come up with.

OK, what are the exact conditions you'll run in your experiment? So we could spend a whole class talking about this. So I'd love to hear your best-ofs, but I don't want to engage on that for a whole class.

A lot of the keys, some of you guys had very clever non-snake conditions to test to get close to minimal pairs. I want to hear about those. But then, beyond that, here's something that probably none of you mentioned. It's understandable; I don't think I said much about it. What are subjects doing in the scanner?

Are they just lying there, and the stimuli are just flashing up, and they're going dumdy-dumdy-dum? Are they doing something with the stimuli? Go think about what you would want to have happen, OK? So what is the task?

Third, some of you mentioned baseline conditions but didn't really say what they are. What would a baseline condition be? And do you want them, or is it a waste of scan time? Think about that.

OK, next, suppose you get to scan 10 subjects for one hour each. Now, think about how that design is actually going to go. Are you going to assign different conditions to different subjects-- so these five people will see all the snake images, and these five people will see all the non-snake images? Or, are you going to have snakes and non-snakes within each subject?

Next, it's nice to not make the subject do their task non-stop for an hour. We usually give subjects breaks. So we break an experiment into pieces of 3 to 10 minutes-- or whatever I wrote, yeah. And so those are called runs.

So think about how you want to allocate those conditions to runs. And how many runs will you include? And then, think about what's going to happen within each run.

So if you're going to have multiple conditions within a run, are you going to stick all of the snake conditions in the first half and all the non-snake conditions in the second half? If not, why not? And if there are multiple conditions within a run, yeah, are you going to clump them all together or interleave them randomly, and what are the trade-offs there? And, what is the order of conditions within a run? And we won't get to number 10 for the moment.

OK, so we're going to break you guys into four groups, and you're going to talk amongst yourselves and try to come up with your best answers to these in five, 10, minutes, something like that. And then, we're going to pull your thoughts on this when we get back, OK?

OK, so part of my agenda in doing this is just to break up the monotony of me going blah, blah, blah, because experimental design is like, it's important, but it's not the most riveting thing. The other thing is, experimental design is basically just organized common sense. And so most of this stuff, you guys just answered all these questions just by thinking about them.

You need to know a few things about the methods, but really, in experimental design, the biggest, the best guideline, the best way to think about design is think about, OK you're the subject. You're lying in the scanner. You're doing that. Does that work? Are you actually going to be doing what you're supposed to be doing? Are you going to be selectively turning on and off this one little mental process you care about, or are you doing a million other things, like falling asleep, and getting bored, and all of that, and predicting what's going to happen next, and all that kind of stuff?

OK, all right, so let's just take a few examples. What were some good kinds of control conditions-- that is, non-snake stimuli that are good to compare to snakes that maybe aren't perfect minimal pairs, but that get partway there? I saw a few, just in the few papers that I looked at.

Yeah, I've got a-- I'm sorry, I've asked your name like six times. But I'm going to-- on my trusty sheet, tell me again how you say it?

AUDIENCE: Achay.

NANCY KANWISHER: Achay. OK.

AUDIENCE: So for ours, we compared snakes to worms.

NANCY KANWISHER: To worms, yeah.

AUDIENCE: Because they have really similar shapes.

NANCY KANWISHER: Awesome, and they're both animate. That's great, love it. What else? Who else had a good control condition? Or who had an interest in control condition? Yes, sorry, your name is--

AUDIENCE: Lauren.

NANCY KANWISHER: Yes, OK.

AUDIENCE: Yep, our group had pretty much the same baseline background, and we would just superimpose images of different objects on it so that remained consistent throughout.

NANCY KANWISHER: Uh-huh, and the background was like what?

AUDIENCE: Forest floor.

NANCY KANWISHER: Uh-huh, OK. So you stick a toaster on the forest floor or something like that, versus a snake or something, yeah?

AUDIENCE: The idea was more like other animals, or stuff that would make more sense.

NANCY KANWISHER: That's good. That deals with the grass confound problem, right? Absolutely, very good. What else? David, you had interesting ideas.

AUDIENCE: Well, we were talking a lot about animate versus inanimate things, so like comparing to a garden hose.

NANCY KANWISHER: Yes, garden hose! Love it! I actually ran this experiment a bunch of years ago, and we used a garden hose-- or a bunch of garden hoses, coiled up in the grass. We tried to make them slither and all that.

Anyway, but garden hose is great. Say more. You had other good ideas in your--

AUDIENCE: Yeah, we also-- we talked about some of my ideas were looking at videos, with motion.

NANCY KANWISHER: Why?

AUDIENCE: What'd you say?

NANCY KANWISHER: Why?

AUDIENCE: Oh, because when you get a snake, it kind of slithers and has this very distinctive thing, that it feels like the motion is what creeps me out when I see a snake.

NANCY KANWISHER: Totally. Me, too.

AUDIENCE: And if you have a rigid thing that looked like a snake, but it was just sliding rigidly, then it wouldn't really creep me out.

NANCY KANWISHER: Exactly. This is a key insight, right? So think about if we're interested in how you perceive snakes, we want to know not just how you do it in some weird lab environment. We want to know how you'd actually do that.

The whole reason to choose snakes is it seems like something that could be biologically relevant. There might be special hardware. When I'm out hiking and I see even a curved stick, I, like, jump and shriek before I can censor myself.

It's horrible. I find it very embarrassing. It's not consistent with my self-image. But I have no control over it; it just happens.

And so I've thought for a long time, there's some damn bit of my brain that's making me do that, and it pisses me off, and I'm going to find it. Well, we looked and didn't find it. But anyway, you go from those intuitions. Often, your own introspections are very informative, and I think your intuition is exactly right. There's a very characteristic motion that snakes have, and it could be that that's the cue.

So then, the trick is you have slithery motion versus what? I don't know. That's hard. What other kinds of motions could you have?

AUDIENCE: Right, so you could have just even rigid motion, where it's not slithering, or it's not changing shape. It's just sliding or rotating.

NANCY KANWISHER: Right. Right, exactly. Anyway, all those are all good ideas. Good, so what should the subject do in the scanner? Should they lie there and go dumdy-dumdy-dum? Should they do a task? If so, what task?

Oh, if you guys don't volunteer, I'm going to start calling on people-- even though, as the Jenkins study, showed it's nearly impossible to look at these damn photographs and figure out who's who. David, in the back.

AUDIENCE: So--

NANCY KANWISHER: Task. Task, or no task? What task?

AUDIENCE: Yes, we talked about having the subjects find a way to indicate that they're paying attention and not just dozing off.

NANCY KANWISHER: Right.

AUDIENCE: So one idea was to have them essentially indicate the source, make [INAUDIBLE] think about a bunch of problems like, well, we don't want them thinking about snakes for the entire our experiment. So--

NANCY KANWISHER: Also, if they're going to tell you they're seeing a snake, maybe by pushing a button, then on the snake trails, they're pushing a button, and on the non-snake trials, they're not.

AUDIENCE: Well, on the non-snake trials, they would push another button.

NANCY KANWISHER: Ah, but then they have two different motor responses.

AUDIENCE: Yeah, so then we would have them run the experiment again, but switch buttons.

NANCY KANWISHER: Good, good. Smart, very nice.

AUDIENCE: Ideally, we might just have them perform a task that's completely unrelated to looking at snakes or thinking about snakes, just so that they're not affecting--

NANCY KANWISHER: Absolutely, and the point you made before is a good one. If you're looking for snakes all the time, maybe even if it's apples and dogs, you're thinking, is it a snake? Is it a snake? And maybe you're using that region, and it's a mess, right? Absolutely, yeah.

All right, so this is a common challenge in experimental design, and these things don't have clear right answers. What I want you to do is just see the trade-offs. On the one hand, just passive viewing, lying there, is good in a way. The things are just impinging on your sensoria, and it's doing whatever it will do. But the downside is subjects fall asleep and get bored, and you don't know if they're awake.

So that's a problem. OK, but the key thing is whatever the task is, you don't want the task to engage asymmetrically with the stimulus condition, because then you're building in a confound, right? So in the group I was in, we were talking about, well, you could have people-- well, we were talking, actually, about faces and objects in that case. So you could have people name the things, but if they're naming snakes versus non-snakes, it's not very good if they're going snake, snake, snake, snake, snake, dog, toaster, apple. One is easier than the other and more repetitive. There are all kinds of problems there.

All right, baseline conditions. I didn't really say what a baseline condition was. Sorry about that. What I meant by a baseline is different from a control condition.

The control condition would be like non-snakes contrasted with the snakes. Baseline tends to be like a minimalist condition that's supposed to turn the brain off. Can we turn the brain off? No, of course not, but we can aspire toward it. We can go partway out there.

We can say, OK, if we're studying vision, let's minimize activity in the visual system as best we can, OK? So you could just have a blank screen that feels like a pretty minimal thing. You can have people close their eyes. The reason that, in vision experiments, people tend to have fixation, where there's a tiny dot and subjects are supposed to hold their eyes on it, is that in natural- left of their own devices, people move their eyes a lot-- several times a second. And moving your eyes produces all kinds of activity and lots of neurons. And so it's a very active visual thing, even if there's nothing on the screen. And so staring at dot is closer to shutting off your visual system, even though it's not shutting it off.

OK, so given that most of the contrasts we've talked about are like faces versus objects or snakes versus non-snakes, and all the activations that I've shown you guys are contrast between an experimental condition and a control condition, why are we bothering with baseline? It doesn't even figure in that contrast. Yes, Jimmy?

AUDIENCE: Well, if the region is truly selective for only snakes, you could use the baseline as, in this sense, like a control, because you can compare the other control to it. If it's really selective for snakes, then the non-snake object should respond in the same as the [? minimal. ?]

NANCY KANWISHER: Awesome, everybody get that? So that was exactly right. And this is, I think, a very interesting point.

So suppose we have-- remember, with MRI, you just have two numbers. So here's the snake response, and here's the non-snake response. If we don't have a baseline, that's all we have-- two numbers, OK? And that's fine. If we run enough subjects, that could be significant.

But now, let's think what else we know if we have a baseline. Suppose we have a baseline of staring at dot. And that's down here. We'll call that fixation.

Are you impressed? And you've run enough subjects, so that's significantly different. Are you impressed? Yeah.

AUDIENCE: Less so than if the fixation were higher up.

NANCY KANWISHER: Exactly! Why?

AUDIENCE: Because then, if the results are higher up-- or if the fixation is the second one that you just drew-- then the response to a snake is twice as much as non-snake.

NANCY KANWISHER: Exactly. Does everybody see how-- yeah, that might be significant, but who cares, right? Some tiny little ratty-ass effect, versus if it's like here, or even-- this is the case Jimmy was talking about, like that. No response at all more than staring at a dot to the non-snakes, and yet this response to the snakes.

That would even more impressive. So there are different degrees of selectivity, right? Not just does it respond differentially, but how selective is it? Oh, boy, I'm going way over time. I'm sorry.

So you guys did great thinking through these things. And, of course, I didn't get halfway through my lecture. That's OK, we'll roll over the best parts for later, and the ones that aren't that fun will just go by the wayside. I will put notes on the rest of some of these things, but I think all of you guys pulled out-- just thinking hard about it and using common sense, you can see that a lot of experimental design is common sense. All right, see you guys on Wednesday.