MITOCW | watch?v=Yo_GRNdkRV0

ABBY NOYCE: OK. So, the plan. The plan is that today, we are going to talk about-- we talked yesterday about models of working memory. We're going to talk today about some of the neuro stuff underlying working memory, what we know about how all of this is actually implemented.

And then, if we have time, we're going to segue into talking a little bit about longer-term representations of knowledge to get ready for next week, when we talk about longer-term memory, and we talk about decision making and working with the knowledge that you have, and things like that.

So just as a refresher, this is Baddeley's model that we talked about yesterday. Baddeley and Hitch said working memory is not a unified capability. Working memory has different pieces. In particular, they pointed out that you have a visual-spatial capability that seems to be separate from your auditory-phonological capability. You can do tasks that use both these pieces of working memory without one interfering with the other.

And we said that there is this central control executive process that coordinates both of these buffers, the phonological buffer and the visual-spatial buffer, decides what information goes into and out of those buffers, can work with it, do tasks on it.

So remember, yesterday we had the-- picture a D, rotate it to the right, picture a 4 on top of it. So that's, you're using your visual-spatial sketchpad to store that information, but what's actually manipulating it and moving it around is that central executive process. So that's an example of where one is working on the other.

Central executive can also post up from long-term memory. It also takes into account the sensory input that you're getting. Generally it's in charge of-- with some quotes around that-- what is actually in your mind at any given point in time.

AUDIENCE: [INAUDIBLE]

ABBY NOYCE: Yeah, well, OK. So the central executive doesn't necessarily get to control how things move in and out of long term memory the way that it controls what's going in and out of those two sensory buffers. So the relationship between them is different, which is kind of what I was trying to signify by drawing it slightly differently. It made sense. It did.

OK. So, we've got this lovely model, which you might notice says nothing at all about the brain. So what are we talking about when we say you've got these storage buffers? What's in them?

So, we generally believe that working memory-- memory when you're actively holding information, working with information-- is what's called activity based memory. So it depends on sustained activation in a population of neurons, a group of neurons that keep firing in a particular way in order to hold that representation in mind.

Having it kind of actively maintained in a group of neurons like that makes it right there. It's easy to work with this information. You don't have to pull up a new representation. But it's also less permanent. There's only so long that these neurons are going to keep firing in a particular pattern.

Longer term memory, like we'll talk to about next week, actually changes the waiting of synapses between neurons to put patterns into the way your brain is connected versus merely in the way it's acting right now. The reason it's less permanent is because as soon as these neurons stop being told to stay active in some way, then-- boom! They're going to stop firing in the pattern that is the representation of whatever you're holding in mind, and this short term working piece of information is going to be gone. This is a theory in the scientific sense. In the, this is how we believe it works and there is good evidence for it sense, like the theory of gravity.

All right.

OK. Brief project moment. So we've got this hypothesis. We believe that working memory depends on sustained activity of a population of neurons. Time to think about what we've been thinking about, about how we find things out in this field. Your mission is to take five or 10 minutes and think about how you would support or refute this hypothesis: Working memory depends upon sustained activity of a population of neurons.

So, write down answers to those four questions. You can talk to your classmates. This doesn't need to be a paragraph. You're not going to be handing it in, but we're going to talk about what we would do in order to make this work. Remember, an independent variable is the thing that you, the experimenter, are going to manipulate, are going to change. And then the dependent variable is the thing you're going to measure. You're going to see what happens. And then operational definitions are kind of like how you're going to change or how you're going to measure these things.

Let's look at some actual studies and we'll see ones that set this up both ways, both with working memory as an independent or a dependent variable. These are some of the classic single cell recordings. Monkeys, animal models. So single cell recordings is that electrodes into the brain measures the pattern of firing of individual neurons. Can't do it in humans.

So we've got monkeys that are trained in a delayed response task. And in this one, they were looking at a screen. And the screen had a fixation point in the center. And the monkeys looked at the fixation point. And they're trained that there will be-- a cue will appear and disappear. And they have to keep looking at the fixation point. And then shortly after a delay again, probably between five and 30 seconds is pretty typical for working memory tasks, then they get a go signal. And they move their eyes to where the cue was. So they have to hold the location of that cue in memory for the period of the delay.

And what has been found is that for a task like this, the area that seems to be most involved is this dorsolateral prefrontal cortex. So, let's see, here's the central sulcus. So this is kind of, all of this is the frontal lobe, prefrontal can store everything from like here, forward. Dorsal on the top, lateral on the sides, dorsolateral prefrontal cortex.

And they found in there, both cells that were active while the cue was being presented, which isn't when working memory is really being engaged. But they also found cells that were most active during that delay period, that once the cue went away and they were trying to maintain that-- and when the monkey was trying to hang onto that information about where the cue had been, this group of cells fired sustainedly. Just, I'm not sure that's a word. So in particular in this area, there are neurons that fire during the delay only when the cue is to a particular location. So we've got neurons that are preferring one cue location or another, that fire then when that particular cue location is the information that's being held in memory.

All right.

I'll distribute those among you.

This is one set of the data from this particular study. Take those and pass them. So, you can see there's a diagram at the top of what the task was. But what I want to look at is along the bottom here, is a graph of the response of one particular neuron to cues in different locations. So, as arrayed around the fixation point. So, on the horizontal axis of each of these is time with the first section there being where the cue is given. And then, so this kind of area is where the cue is being presented. And then this is the delay period.

And you can see that this particular neuron doesn't change its firing rate much for most of these cases. But it strongly prefers a cue in one direction. Which direction? Towards the upper left here, right?

You can see that in this case, once this particular location is being held in working memory after the cue has been removed, this neuron goes nuts. And it just fires and fires and fires and fires. Its firing rate is much higher than these sort of baseline rates. You can see it gets activated a little bit for the straight up or straight to the left conditions. It's slightly higher than its baseline but it's definitely got a strong preference for one or the other. So this is a neuron that's involved in maintaining information about the location of the cue after the cue has disappeared. Does this make sense?

Cool.

AUDIENCE: Why is it top left?

ABBY NOYCE: That one? This neuron likes the top left location. There's probably other neurons that like all of the others. So this is just data from one individual cell, showing that the cells have directional preferences. It's like, think of it as being like the edge orientation neurons we talked about last week that like lines at particular angles. It's a matter of what's being, of how they're wired up to get input from cells earlier in the chain.

All right, so that's one kind of classic experiment they looked at. They took the same experiment, and they said, OK, but that's just correlational evidence, you know? We know that when working memory is happening, then these cells fire. But that correlation is not causation. What's some more evidence that would help make this look causal that would help draw these two things together?

And they looked at what happens when the monkey was wrong, when the monkey moved its eyes to the wrong location, not to the cued location. Was the activity of these cells that we believe to be working the same? Nope. So, in cases where the monkey did not do the task correctly, these working memory kind of maintenance activity cells were actually less active. Or they were active for a while and it dropped off.

And this might just be that the memory is decaying. It might be that the monkey just stops paying attention and all of the neurons in its brain kind slow down a little bit and stop firing so fast. We're still looking at correlational and not causal evidence.

So going one step further looking for a more direct connection, people looking at the same parts of the brain did some lesioning studies, so small lesions to particular chunks of that dorsolateral prefrontal cortex. That's the area where these cells that are active during the delay period seem to be located. And they found that if they did small lesions-- big lesions, working memory just dropped off. But if they did small lesions, you got monkeys who were OK at this task at most cue locations and then would have, like, one or two right next to each other where they were very poor. And the explanation that these researchers came up with was that we've managed to remove the neurons that respond to this particular location, while leaving intact neurons that respond to other ones.

Somebody did the same thing with cooling neurons. So instead of actually lesioning them, you slide a little piece of metal into their brain and then you put the other end of it in, like, liquid nitrogen. And it, therefore, brings the temperature of the neurons right next to it down enough that they don't fire. But it doesn't do long term damage.

So lesioning studies have a confound where animals that have le-- we know that when the brain is damaged, you'll see long term structural changes. As information and processes get moved to other areas, other areas take responsibility for what was originally lesioned. Whereas if you're just doing these kind of temporary knocking out one neural population, then you don't get that confound, because it's just a short term change.

Because with lesioning, you do surgery on your animals and then you've got to let them recover for a couple of days before you can drag them out and make them work for their food again. And so that's long enough for this sort of rewiring around the damaged part of the brain to start happening. So they wanted to correct for that and make sure that wasn't what they were seeing. Again, lots of data all kind of converging on the same point.

All right, so this is monkeys. That's cool and all, but really we want to know about working memory in humans. We can't sink electrodes into their brains. We certainly can't cut chunks out of their brains. But there's some evidence from humans nonetheless.

FMRI has been used for this. FMRI shows activity, as you might expect, in the dorsolateral prefrontal cortex. Also, for these kinds of spatial tasks where you are cued to a location and have to hold that location in working memory for a delay period, then you'll see activity in the dorsal parietal cortex, kind of up here on the sides. And again, just like in the single cell recordings, you'll see this activity increase sustained throughout the entire delay period.

So right now, it looks like we've got some pretty strong evidence saying that when you're working memory is working, when it's actively maintaining something, then you get an increase in firing among a set of neurons for as long as you're delaying it. So, as long as you're holding it throughout the delay period between you when you're given the item. And when you have to do something with it and can stop holding it.

All right, so far we've been talking about paradigms where your people or animals are holding a single item in working memory at one time. And to get down into the nitty-gritty of how exactly an item is in working memory, what that means, we've got to understand what happens when you're trying to hold more than one item at a time. Clearly, certainly, from a subjective perspective holding two things in memory is different than holding one thing in memory. You've got two things in there. So what sorts of neurological changes would we expect to see?

So there's two kind of possibilities when you're holding more items in working memory. You might see nothing change at all, but that seems unlikely. You might see that more parts of your brain get involved, get pulled in to help with this working memory task. Or you might see that pretty much the same parts of the brain are involved, but that the amount by which activity increases becomes more for each extra activity, each extra item that's being held in memory.

And the FMRI data seems to show that both of these things happen, both as you increase the set of items that you have to remember then more brain regions are recruited and pulled in to help keep that maintenance, keep that representation. And you also see activity in the regions that are involved increase. They actually become more active as they are representing more things.

One of the things that's interesting is we have this idea that different parts of the brain-- we have this idea that working memory has different storage spaces for different kinds of information that's kind of key to this Baddeley model that we're working with. And if the Baddeley model is correct, one of the things you'd expect to see is that for different kinds of information, you'd see different brain regions activated when you hold them in working memory. And that seems to be more, that seems to be at least somewhat true.

For example, like we said, for these spatial tasks you'll see the parietal areas in the parietal lobe getting recruited to help maintain that information. For, like, object recognition tasks, you'll see areas in the temporal lobe getting dragged in. So there's definitely at least some amount of activity in different brain areas for remembering different kinds of material in working memory.

All right, so what happens if you're holding onto more things? One of the kind of classic challenging working memory tasks that we've talked about before is this N-back task. There's a stimulus currently being presented to match the one that's one item back or two items back or five items back. This is, remember, the working memory task with that paper we read for the very first week of class used.

So, if you're doing a two-back task, then the first n gets presented. The answer is no, there's nothing two back. Get secondly for the first l. Actually, I think for two backs, are these all, no's? Yeah, it's all no for two-back.

But for three-backs, for example, that second k would be a yes, it matches the k that is three back from it. So, when you're just doing a one-back task, you've only got to ever remember the thing that was shown to you just before the thing you're now showing. And as n goes up as you do a two-back or a three-back task, this gets harder. It requires more working memory. You've got to hold more of the previous stimuli in mind.

People like this task because they can change the memory load. They can show two subjects the exact same set of stimuli and have one of them do it as a one-back task and one of them do it as a three-back task. And the only thing that's changing is the working memory load. They're getting the same stimuli in the same order, so it takes out some confounds that can be based on just basic perceptual differences.

So, FMRI says that you get increased activity in the lateral prefrontal cortex and in the parietal cortex. And in both of these cases, it goes up linearly with n, so each time that you add another item to n. So if you go, the difference between a one-back and a two-back task is a certain amount of activity increase. And then the difference between a two-back and a three-back task is, again, the same amount of activity increase. It's a linear increase in the lateral prefrontal cortex and in the parietal cortex.

But one of the problems with this particular task model is that it's not just requiring you to hold items in memory. It's requiring you to actually hold a lot of information about them. And then keep changing that information. So, if you're getting this, you know, first of all, this n is the item you're looking at.

And then the n is one back. And then the l is one back and the n is now two back. And then the l is one back and the l is two back and the n is now three back. And you've got to kind of keep relabelling each of these thing that's in memory.

They're not static. They're changing. And that's really something that's an executive task. It's not just an item set. It's the actual working with the information. So there's a confound here between that executive task, where you're actively changing the information in working memory, and the maintenance process itself, where you're merely maintaining these representations.

So if you wanted to try to separate those, see what comes from just maintaining information in the brain and what comes from actively working on it, you'd probably want to find a task that doesn't have all that increased executive processes. Straight up ordinary item recognition tasks, so I show you one item or two items or three items. Look, it's a butterfly! Take it away again. Wait 10 seconds. Show you the item again, look, it's a cat. Is this the item I just showed you? Good. So this is a probe item. It's like a test item here.

AUDIENCE: The butterfly is inside the cat.

ABBY NOYCE: Possibly, but it is not the object I just showed you, nonetheless.

So this is a simpler task than the n-back task. It's less demanding. The executive requirements don't change with set size. You're merely maintaining a set of objects here. You're not working with them the way you are in the n-back task.

And you find pretty much the same FMRI data, which is that you get these linear increases in activity in prefrontal and parietal areas. So it looks like the change that we see in the n-back task is probably coming from the increased maintenance requirements of the larger n, looking at a three back versus the two back, and not just because you're doing all that extra work on these things to re-categorized them.

All right, so we know that holding items in working memory requires that this set of neurons keep firing for a sustained period of time. So how does that happen? Is there, like, some kind of command center that says, OK, I'm going to tell this set of neurons to fire in order to keep this piece of information in mind? Every time, as far as anyone can tell, as far as I can tell, every time in neuroscience you're tempted to believe that there is some kind of little command center that's making a conscious decision of some sort, this is a dangerous idea that you want to stay away from.

A lot more often what's happening is that a network of neurons works together in the way that they are wired up. They kind of all encourage each other or discourage each other and settle on a particular pattern. And that seems to be true of how this works.

So, Donald Hebb, who was one of the big guys in the '40s and '50s in neuroscience-- we'll talk about him a little bit more next week, another idea of his. He said, OK, so we know that the idea of having this neat and tidy command center that's directing everything the brain does is a bad idea. It doesn't seem to be how it works.

Can we just do this within a group of neurons? Can the neurons that are involved in maintaining a pattern be connected to each other, such that they maintain the pattern? That each one by firing excites the other neurons in the pattern and so, as they all fire together, it just keeps the firing rate up? Can this sustained activity be something that's just kind of self-reinforcing?

So each neuron fires. It stimulates other neurons and they stimulate it. And the whole thing keeps itself going simply by the connections between the neurons.

So this was his theory in the late '40s. And both our models of how neurons work and our knowledge about how neurons are connected to each other weren't as good then as they are now. So when he first came out with this, there was a lot of, we don't think this will work. Will neurons really do this?

Since then, people have started looking at it a bit more carefully. And it turns out that if you put together a model of a neural network, so a computer program where you've got simulated neurons that are connected to each other using information from how we know neurons interact with each other, how they work physiologically, what sorts of inputs they are likely to have, can you build a network that will self-sustain like this? Where once you get it going, it will just hold itself up there? And the answer seems to be, yes.

We can do that. You get these patterns of reverberating activity where each neuron excites the other neurons in this cluster that are maintaining whatever information. And so if you look at individual neurons in one of these models, you get firing rates that are very similar to what you'd see in a real life electrophysiology measure of the behavior of a particular neuron that's involved in working memory.

All right, want to see one? So this is a model made by a guy named Wang at Yale. So you've got, that's time moving along the top here. This is a set of neurons, simulated neurons, that are firing. Ready? So there's about to be a cue. So the cue is to the left. And you can see this population of neurons that responds to that area. And the cue has gone, but the firing among the population of neurons there is holding itself up. You can see that the spike just maintains over time. It just keeps going, and going, and going.

AUDIENCE: This guy's name's Wang?

ABBY NOYCE: So there's. Yep, wanglab.yale.edu, I believe. So he's working with one of these models of how you can build a neural network that will self-sustain like that. You want to see it again? You want to see it again.

So, again you'll notice there's a baseline firing rate. None of the neurons are really firing a whole lot more than the other. And then when that cue is presented, these are populations of neurons that are tuned to different locations. So all the ones that are tuned to that location there, 180 degrees from kind of the default point, are activated when the cue appears and then they continue to be activated after the cue is gone.

So they're not only excited by the actual perceptual input from that cue. They're also responding to each other, they're exciting each other and keeping the firing rate up among the circuit.

AUDIENCE: What does d mean?

ABBY NOYCE: D is for delay. So this is a delay period. It's like in your classic working memory task, you'd have a cue and then a delay and then a test of some kind. And this one, we don't care so much about what the test is per se. We're interested in how the neurons behave during these different time periods.

And it's a cool video. Questions? Cool.

So that's an example of the kind of model, of this kind of model showing that you can get a group of neurons that just keep a circuit like that maintained.

AUDIENCE: So, does this mean that we hold things in our memory because our cells do it for us? Or not because--

ABBY NOYCE: Everything you do in your brain, you do because your cells do it for you. You want to do things because your cells do it for you. This is one of those really dangerous, this is one those things that's really hard to think about.

AUDIENCE: OK, in this one study, like, my dad showed me in the Wall Street Journal, where they tested people and went about decision-- it was about decision making. And they had to, like press a button with their left hand or right hand. And that was, like, basically, the basics [INAUDIBLE]. And, by the end of the study, the testers could actually, by looking at the neural patterns could predict which hand they would use to press the button, like, two seconds before it actually happened.

ABBY NOYCE: Yeah, there's some classic work on this and I don't remember what the guy's name is who kind of did the pioneering stuff on this. But he set up a paradigm where he had people just sit there with a button. And I don't remember what measuring, what imaging they were using, but they were using an imaging technique. And he, and they were looking at a big clock, and he said, OK, at some point decide to press the button and notice what time the clock says, like with just the second hand going around it, when you decide, when you make that decision.

And what he found out is that consistently, you'd see changes like in people's premotor cortex, which, like, coordinates your motor activities and stuff, before they consciously experienced having decided to move. So your brain starts doing, preparing for the action you're going to take before you feel like you've made that decision, which is weird and a little bit scary.

But once you start getting down into the nitty gritty of neuroscience talking about free will becomes a little bit depressing. So, I mean--

AUDIENCE: What about, like, when I think there's a brand new study, did a study with epileptic patients. He'd stimulate one part of the brain and make one, he'd say OK, don't move your right arm, then he'd study what part of the brain that [INAUDIBLE].

ABBY NOYCE: I believe it. That's not as, I don't find, personally I don't find that as disconcerting as this idea that when I make a decision, my conscious experience of making a decision is preceded by a bunch of unconscious already made decision stuff. That's, that's a little bit weird. I mean I know that my actions are caused by firing of neurons in my brain, that doesn't throw me as much personally.

AUDIENCE: I think that [INAUDIBLE] that, like [INAUDIBLE], other times, patients would say, I didn't didn't move my arm, you did.

ABBY NOYCE: Yeah, so your perception of, because you don't have that decision to move coupled with the actual motion. Anyway, back on topic.

One of the things that was consistent about all of those different studies that we were talking about, about different parts of the brain involved in working memory and maintaining a representation, seemed to have, hey look, this region of the prefrontal cortex is heavily involved. Prefrontal cortex does all the cool stuff.

Now prefrontal cortex seems to be heavily involved in what's called executive functions. Think about Baddeley's central executive, here. And it certainly seems to be really important in maintaining representations. It's not the only area of the brain that does it. You'll see this happening in temporal lobes and in parietal lobes. But the prefrontal cortex seems to be behaving differently than those other areas.

So, this is a study. They had that same image recognition task. Look, a butterfly. Take it away. Look, another object. Is it the same one as I just showed you or different? But during the delay period, they actually showed other images.

So I'd be like, look, a butterfly. Remember the butterfly. Look, a beach ball. Look, a crescent wrench. Look, a chair. OK, test time. Look, a butterfly. Is this what you just saw?

And so you'd have all of these images presented. Butterfly is the one you just saw. But what's holding onto that image? So if you don't have the distractors, if it's just a butterfly, delay, butterfly again, then you'll see sustained activity in both the temporal areas and in the prefrontal cortex. But if you get these distractors in between, when the first distractors are shown, that representation in the temporal cortex goes away. It just stops being there. Whereas prefrontal cortex's activity remains high throughout the entire delay period.

Prefrontal cortex seems to be special, not so much in terms of what kind of material it's storing. It's not, you know, just visual spatial or just phonological or anything like that. But it's better, almost, at storing it. It seems to be specialized to hold on to this stuff when you want to hold onto it, when you're consciously trying to hold onto a particular piece of information and not something else. It seems to be better at resisting distractions, better at maintaining that representation. And it's, of course, also running all of that executive stuff.

So prefrontal cortex definitely seems to be involved in the storage stages of working memory. And it seems to be better at it than some other parts of your brain. But Baddeley's model says that, hey look, executive and executive control and storage are different functions, but we're seeing them both happen in more or less the same portions of your brain. Is this a problem?

So there's two kind of possibilities that don't conflict with the model. There's a third possibility, which is that the model is wrong, but the model, thus far, has been really, really useful in trying to break down what working memory is. Nobody is quite willing to throw it out. It could just be that there are different subregions or different subpopulations of neurons that are involved in storage and maintenance of working memory items versus in this kind of executive control of working items.

And there's another possibility, which is this kind of tweaking of our model. Which says that the prefrontal cortex is specialized to maintain information that is related to a particular goal. So if you have a goal, you want to do something. Prefrontal cortex hangs on to that goal, modifies, works with information relative to that goal versus merely just holding items in memory in the sense where-- like, walking through the world requires that you have some amount of visual information and working memory. You've got to remember where things were when you just looked at them and all of that.

So this is called the Goal Maintenance Model. This is the prefrontal cortex maintains the information about what your goal is and it directs your attention and behavior to attain that goal. So this is classic executive function there. Directing attention, coordinating behaviors, prefrontal cortex is involved.

For example, if you are, so for example, I used to work about half an hour from my house. And that was my route every morning. I went down the end of the road, I turn right. I went a couple miles, I turned left. This was my driving route.

If I tried to go somewhere that was going to be about 2/3 of the way to work and then make a different turn, if I'm not paying attention when I get to that turn, I'm going to go straight past it and stay on my usual route to work. Prefrontal cortex not doing its job. If I, in order to remember to make that left hand turn, I've got to remember that, hey, I need to do this. Keeping a goal in mind.

Prefrontal cortex is maintaining that I need to make this left hand turn that's not on my usual route. And it's got to direct my behavior. It's got to say, OK, the prefrontal cortex says, OK, get in the left hand lane. Move over. All of that is being modulated by this goal-related ability of the prefrontal cortex.

AUDIENCE: If it doesn't do its job, and you forget a lot of things that you want to do, like, five minutes after?

ABBY NOYCE: That would be bad.

AUDIENCE: Is that Alzheimer's or is that different?

ABBY NOYCE: Alzheimer's is different. Alzheimer's, prefrontal cortex is one of the areas that degrades in Alzheimer's. It's certainly not the only one. And there's a particular pattern of behavioral deficits that you see in Alzheimer's. And the short term memory loss is really only one of them.

But, yeah, definitely prefrontal cortex not doing all of this executive control that, like, inhibits irrational behavior and keeps you, and does attention, and helps you maintain goals, all of those are abilities that you see decreased in Alzheimer's.

All right, so, think about this in context of that paper on motivated seeing that we did last week. Where subjects were given an ambiguous stimulus. They had, consciously or not, a particular preference to perceive that stimulus in one way or another, right? Remember the horse-seal?

What do you think? Is prefrontal cortex involved in how you perceive that ambiguous stimulus? The question here, the question that you would ask, I think, is does that preference to not drink the viscous chunky foul smelling vegetarian, vegan, smoothie, is a goal? And it's probably not, well, it's probably actually something that's pretty close to the front of your mind as you watch your score get lower and lower, but--

AUDIENCE: Was that a pun, or?

ABBY NOYCE: No. I tend to be the one who makes puns by accident and then other people laugh at them.

AUDIENCE: Front of your mind, prefrontal--

ABBY NOYCE: Prefrontal cortex is the stuff in front of the frontal cortex.

AUDIENCE: So it could actually be a pun.

ABBY NOYCE: It would be.

AUDIENCE: [INAUDIBLE]

ABBY NOYCE: So one of the things that, one of the questions here, and it's one to which I don't have an answer to, is whether kind of, for something to be considered a goal in this goal maintenance model, depends on it being a goal that you are consciously aware of or not. Just something to think about. Anyway, all of this stuff is interrelated. Cognition requires lots of different capabilities. I like looking at connections.

OK, so we've been talking about, moving right along, we've been talking about how different parts of your brain are involved in maintaining representations of information. We haven't talked much about what exactly these representations are. We know that there are patterns of firing in particular groups of neurons.

So, we're going to switch gears. We're going to go back to talking kind of information theory a little bit about what kinds of representations the brain uses. Remember, a representation is a physical state that stands for something, for very broad events of something. And it carries information-- wow, I can't type today. I'm sorry guys-- an object, event, or concept, and carries information. Representations must be intentional and information carrying.

And this is intentionality in the philosophy sense, not in the, like, you must intend to make a representation. That use of intentional is from its original Latin root. And it means that it has to refer to something. It's from a Latin root that means to point.

So the representation must be, must refer to something, it can't just-- and a representation has to carry information about whatever the something is that it's about. So, if, for example, your representation of this chair might hold the fact that it's usually in this classroom, you've seen it for three weeks. It's red. It can be sat in. It's kind of funky looking.

All of these are pieces of information about this chair that would be contained within your representation. Your representation of this chair might include the fact that it is a member of this broader category, "chairs." You'd have connections like that.

So the information the guys, who are like, hey look, cognition is about information. What kinds of information might we be dealing with? Have four main types of representations that people talk about. And the most basic is this kind of modality specific image.

So you take a photograph of something. That's a representation. It causes changes in the CCD on your digital camera. It contains information about the amount of light out there in the world. It points to a picture of whatever you took a picture of. And so, an image type representation in your brain is, again, very simple, information about the light that is reflected off of stuff and onto your retina at each point out there in the world.

We talked about, last week, that the brain stores this basic image information in early stages of visual processing. Remember the monkey brain with the bullseye pattern overlaid right on the brain, there? That retinotopic organization where information about each point just lines right up on the back of your brain. So the brain doesn't do a whole lot with this information, but it definitely has this particular type of representation.

There's an early visual processing format. At later stages of processing, then you're using a more sophisticated representation, again, of a visual stimulus of some sort. And this is a feature record.

We talked about this when we talked about object recognition and vision last week. So a feature is a meaningful sensory aspect of a perceived stimulus. So, edges, for example, are definitely a feature of the visual world. And we know the brain is very good at finding edges. It's got cells that are specialized to detect edges. And so there's probably at least, and so a neuron that detects edges is a representation. It's a representation of a particular sensory aspect, of a meaningful sensory aspect of a stimulus.

Frogs have cells in their, cells in their retinas that respond to dark dots on light backgrounds that move around. Their bug detectors. They're responding to a meaningful thing out there in the world. To particular aspects of it. So that, those cells are a feature record representation. A representation of a bug, of something I want to eat. If you are a frog. Probably not if you are a human.

AUDIENCE: My brother ate a fish food once.

ABBY NOYCE: I think I have sampled fish food once. I think I will admit. I was young, I knew no better.

AUDIENCE: Was it yummy?

ABBY NOYCE: No. It was, like, the little flakes. And I tried one flake. And I said, I'm not doing that again.

AUDIENCE: He said at the time, it tasted good.

ABBY NOYCE: That's intense. Older brother? Younger brother?

AUDIENCE: He's my older brother.

ABBY NOYCE: He was probably trying to get you to do it. Speaking as an oldest child.

OK, so now let's talk-- there's, the other two are these kind of more abstract types of representation. So for a long time people had this idea of amodal symbols. So the feature records and the images that we were talking about before are tied to a particular sensory modality. They allow you to represent things in their particular sensory terms, but they won't help you to, for example, connect your image of this chair with the word "chair" with, like, your broader concept about chairs-- that they exist, that they are for sitting in, that they generally have a seat and back and four legs. All the things that are about chairs, which aren't necessarily visual sorts of things.

So most people use amodal representations to talk about the relationships of different objects in the world to one another. So you'll see people talking about frames and semantic networks, which are both ways of specifying. And just in case it wasn't clear, we're back into information processing theory and not so much into what the brain does here. But they specify relationships between objects. So there's, you know, water bottle on the table, and that kind of "on the" relationship would be an amodal symbol in your brain. It doesn't have anything to do with the visual representation, per se, but it's information about how objects are arranged in the world.

Basically, frames and semantic networks are storing roughly the same pieces of information. And these are, there are subtle differences that I don't entirely understand exactly how they're wired up. And scientists write nasty letters to journals about each other, about which one is right.

And then there's also what's called a property list amodal symbol. So for example, for this water bottle, you might have the property that it is blue, which is a little bit too concrete, really for this. But you might have "contains water" as a property. You might have "made of plastic" as a property. You might have "unbreakable" as a property. These aren't strict perceptual ideas, but they are things that are connected to one particular object.

So amodal symbols are kind of necessary for any good theoretical model of how the brain stores information. There's no real evidence that there are cells in your brain that do this, per se. There's no good neurological evidence for these types of representations. So they're kind of up in the air.

More recently, people have been trying to wrap their heads around these not directly perceptual representations with a slightly different model, which is as a pattern of activation in a neural network. So you have a group of neurons, or group of populations of neurons. And for any given concept or idea that you're trying to represent, then, some of them will fire and some of them won't fire. So you might have a particular pattern of neurons that fires for this "on top of" concept. Whether for the laptop being on top of the table or the binder being on top of the chair, you'd have a pattern of neurons that represents this.

And one of the things that's nice about this model is that you can have, for example, I have two chairs here. Chair. Chair. And they're different. You probably have representations in your brains of these chairs that are not entirely the same. But also, those chairs are more like each other than they are like the water bottle. So from a purely intuitive point of view, it would make sense that your representations of each of these chairs are more alike than either one of them is like your representation of the water bottle.

And a model like this, where you're representing concepts as a pattern in a group of neurons firing rather than in a particular neuron responding to one thing or the other, lets you take advantage of these differences among entities within a category while still giving you a way to have category representations. There's not good empirical evidence for this yet. But it seems to fit what we understand about how neurons work better than the straight up amodal symbols older school model. OK, questions.