

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

PROFESSOR: We had started to talk about learning, beginning with habituation, which led to a discussion of the effects of novel stimuli. And I had mentioned the autonomic as well as behavioral effects of novel stimuli. At the end of the session, I said that autonomic responses to stimulus novelty are the same as the responses measured by a lie detector. So what can explain that?

This is the way I explain it. That the autonomic nervous system and the orienting mechanisms of the midbrain respond to any conflict. Usually, when you're dealing with stimulus novelty, the conflict you're dealing with is between the expected stimulus and what is actually perceived. We expect things all the time, because we have an internal model of the world. We have a visual model of the world.

So every time you flick your eyes-- if I suddenly move my eyes over here, there's nothing novel, I expect. I have a pretty good model of this room. And so I always expect what I'm going to see. And if I don't see what I expect-- if suddenly there's somebody different over there I didn't expect-- now there's nobody there. But if there suddenly was somebody, I would have an autonomic response. Of course, my behavior would change too, in various ways.

So that's the kind of conflict we're talking about. Not the usual meaning, perhaps, of conflict. But it's a difference. We're always making this comparison between what we see and what we have expected to see.

But the conflict can also be between an internal representation of what is known to be the case-- what's true-- and a cognitive system. And it could be a system that generates output, like the speech system.

And they can differ. Why would the difference arise? Well, because these memories and expectancies are cognitive systems. And they're closely tied with motivational systems. So you can have two different systems active at the same time. And when they're in conflict, they're going to generate autonomic responses.

You should ask right away, is that always true? And I point out here that for some individuals, these systems can be pretty isolated from each other, and not cause any actual conflict. And we call such individuals psychopathic or sociopathic. That is, they don't feel any conflict. And it's possible that people can also learn to suppress autonomic responses-- although I can tell you, it's exceedingly difficult.

Let's look at some other learning topics. We know that scientists have often chosen to study learning in very simple animals like the sea slug, and snails, and the fruit fly. And there are other insects that have been used. They've been studied in some crustaceans as well.

What are the reasons for that? Now we know that the animal that dominates as a model is the mouse. But that hasn't always been, and it was exceedingly difficult to get at the cell on the molecular level in studying learning, before relatively recently. And for that reason, many people have studied other animals, like this animal. This is the sea slug-- the California sea slug. *Aplysia*.

So why do they choose such an animal? Why do they choose certain insects? There's several reasons. One is, the nervous system is simpler. There are fewer neurons, and it's especially important that many times, these nervous systems have individually identified cells. So that you can identify the very same cells from one animal to the next. And in many cases, they formed very similar synaptic connections when those are studied.

So you can look at the very same connections in different members of the same species. And that's very important in science, to be able to repeat. If you're studying the details of what happens at the synapse, you're looking at molecular changes-- that's very important. And most of these animals, in simple learning situations, are very predictable. Much more predictable than the higher vertebrates.

So those are the reasons people choose those animals. And people also have chosen to study some simpler forms of learning, like classical conditioning. And Scott discusses this. It's often called Pavlovian conditioning, or Type I conditioning-- as opposed to instrumental conditioning, or problem solving involved in, for example, trial and error learning.

Now, you should know the paradigm for classical conditioning. It's very simple. I just want to point out the historical point here that Ivan Pavlov and his followers in Russia interpreted all learning in terms of conditioned reflexes. And they wrote books about it, books explaining human psychology in these terms. Their teaching-- the Pavlovian teaching-- became state policy in the Soviet Union. And scientists who didn't go along were in danger of persecution.

I had the good fortune of winning a prize when I was young student, to go over to Moscow to a meeting. I'm also hesitant to tell you how long ago it was. It was actually in 1966. I remember I was pretty young at the time. And I felt there weren't other people my age going, except one other American, because we had both won the same prize. A guy from California.

And I still remember going to these meetings. And there would be some American and European speakers, and some Russian speakers. And the Russian speakers, invariably, would, say, introduce their talk with something about Pavlov. I asked them about it, and they explained to me that, yes, we have to do that. But then they would look around a little bit, and then say, in a lower voice, we don't actually believe it all anymore, but this is what was required.

Anyway, Scott, when he discusses it, makes a mistake. And I think it's good for you to be able to read your textbooks critically. No textbooks are written totally free of mistakes. And I actually quote the statement here. So look at that, and tell me what's wrong with it.

If you know the paradigm for classical conditioning, you would be able to see this. So what's the basis for classical conditioning? You have a response that occurs automatically. It's a reflex response. Take, for example, leg withdrawal. What causes leg withdrawal because of an inherited, innate reflex? Well, something painful. You step on something painful. You pull your foot back. And that's automatic.

You become aware that you did that, and you become aware of the pain after you're already pulling your leg back. That's the withdrawal reflex. And you can elicit in your friend-- if you hold the hand and pinch him at an unexpected time, he will pull his hand back. And then he will yell out, what did you do that for? But he's already pulled it. So this is an unconditional, or unconditioned, reflex. The original meaning-- the better translation-- is actually unconditional.

And the unconditioned stimulus is that painful stimulus. It can be heat. It can be other kinds of sensory induced pain. Basically, an intense stimulus.

So then, how do you condition that in the classical manner? You cause a stimulus to occur just before the unconditional stimulus. So the unconditional stimulus, or unconditioned stimulus, is the UCS. The conditioned stimulus is the CS. The conditioned stimulus has to occur first. Because then, the animal or the person learns that that conditioned stimulus is always going to be followed by-- it's not a cognitive thing. It happens pretty automatically.

So what will happen? Well, after a number of pairings of the conditioned stimulus and unconditioned stimulus, they'll withdraw the leg just in response to the conditioned stimulus. Could be a tone.

All stimuli don't work equally well, although the early writers-- the early students of classical conditioning-- thought just any stimulus would work. It's been the ecologists who've shown that's not actually true. There are specific differences in different animals.

Anyway, he mixes up the direction. He says the application of the US, and then the CS immediately after it, every five minutes. He's got it reversed. That wouldn't make any sense. You see, the conditioned stimulus-- it's something that normally predicts the stimulus.

What happens if you always present the stimulus just afterwards? I won't lead to any conditioning.

And he says it leads to a much stronger response to the unconditioned stimulus. No. Not at all. It leads to a response to the conditioned stimulus.

And this is the paradigm, then. You can say, US. In Scott, it's the unconditioned stimulus, but UCS is more common. So originally, before any learning, the UCS leads to the unconditioned response, the UCR. And after learning, the conditional stimulus, or conditioned stimulus, leads to the conditioned response. Which is usually almost identical to the unconditioned response. If you measure the details, you might find out it's not exactly, in every way, identical.

And I'm just giving some examples here of laboratory studies that conditioned leg withdrawal that I was talking about. The innate flexion reflex in response to painful stimulus. They've done condition eye closure. That's very common. They usually use a little puff of air as the unconditional stimulus. And you can get conditioned eye blink.

And you can train someone, or an animal, to blink their eyes every time a little tone comes on, or some other stimulus. And of course, we are conditioned just by our experience. Anything that predicts anything coming at our eye, we will blink.

Now, there's other kinds of learning. This is the topic of social learning. When I say, in your own development, how have you shown learning by imitation or by mimicry? And what's the difference between learning by imitation and learning by mimicry? It's actually pretty hard, but there is a difference that's spelled out in Scott. Does anybody know the difference?

How do birds learn to sing? Male songbirds will pick up the songs of adult males that are singing. And they learn those songs before they actually start signing. Now, when they're learning, there's no reward. So we call it mimicry. They learn by mimicry, but they don't even have to start mimicking initially, while it's happening. In many cases, they do.

But with imitation-- like, a young cat will imitate its mother in hunting. It's got the fixed action patterns to hunt, but without the imitation of its mother, it doesn't acquire normal motivation to hunt, as we've learned. So with imitation, there's immediate reward; with mimicry, there's no tangible immediate reward. That's the way we discriminate them.

And they've been observed in many non-human primates. Not just monkeys. We think monkeys are the ones that imitate and mimic the most. But in fact it occurs in many other animals. How have we brought that up in the class already? Well I mentioned kittens imitating their mother. That's certainly a non-human animal.

But it occurred then in jackdaws. Remember the Konrad Lorenz story of how they learn to identify enemies? How do they do it? They imitate-- or mimic. It's a little hard to know; there is some social reward here, so I suppose we can call it imitation. But the adults, when they detect a dangerous animal that has been seen dangling something black that could be a jackdaw, the young learn from them.

And you know, from the adult encounters with that person or animal-- like, it could be a hunter. They learn to recognize. They learn by their calls. They make a lot of noise when there are enemies around. And the young pick that up, and they imitate the parents. And they learn.

And this has been studied in the corvids. Crows in particular. There's experimental studies of this. Of the specificity of their learning. And the interesting thing about those studies of crows that have been done more recently, since the earlier studies, was that that knowledge can be passed on to the young that have never seen the particular face of the person that they've learned to avoid. They can somehow teach their young about this.

Songbirds-- we know they learn songs by mimicking adults. Scott mentions that in newly hatched chicks-- this was a surprise to me; I didn't realize that these old newborn chickens can learn by imitating adults. But the adults that they observe pecking food will avoid certain food because it's bitter, and feed on things that are not bitter. And the chicks learn that. The chicks don't have to learn everything by trial and error. They do learn just by watching. Just by watching adults.

It didn't fit the people that believed everything was conditioned learning. They learn by imitation and mimicry as well. And there's very good evidence for that. Experimental studies, for example, of the chickens. The little chicks that Scott mentions. And there's various other things you can find, too.

OK. So let's move to a slightly different topic, involving a lot of learning-- the topic of navigation. And I also gave you a reading from Niko Tinbergen, whose little book, *Curious Naturalists*, gives very interesting descriptions of some of his experiments-- not just with gulls, but also with insects. He studied landmark orientation in the female digger wasp. He studied several different species of these, usually sand wasps that make underground nests.

He studied them with his students. And not just by observing them. They observed them in great detail. They spend many, many hours observing details of how these wasps do it. But then they also did specific experiments.

And this is from the chapter that I asked you to read. In this case, the wasp had his nest over here. And they had observe the walls wasp bringing larvae back to the nest to feed-- they brought dead insects, rather, back to their larvae in the next. And the way they did the experiment was, they surrounded that nest with pine cones, and left those pine cones there for an extended period.

Because they wanted to see-- does the female wasp use landmarks like that, like the position of those pine cones, to actually find her nest? Or is she using other means? So then, of course, the test was simply to move the pine cones as you see here. And this is what happened, as he shows in his drawing. The initial place the wasp looks is right in the middle of those pine cones, where the nest used to be-- even though it's been displaced.

Now, of course, you could use other landmarks, too, to position it. But prominent landmarks are usually the ones being used.

This is from another study. I think we talked about this before. The colors aren't showing up over here. They're showing up a little better there. But anyway, here, in this particular wasp, he's got five different nests. At least, by the fourth day, he had five nests. These are the days here. And when they first observed this wasp, he was bringing food to one nest. Then she came back to the nest. And the open circle there indicates that all she did is open it and look.

And the next step was she closed it up. Close it up and left. She never came back again to that nest. Instead, she went to another nest here. This was all on the first day. And just opened it and observed. Then she flew away.

But then, rather than go get a larvae and bring it back, she didn't do that right away. She dug another nest. The third nest, here. She dug it, and then she got a dead caterpillar, and she brought it there and laid an egg, and left the caterpillar next to the egg.

It was at the beginning of the second day that she came with the caterpillar and laid the egg. She had dug the nest at the end of the first day. And then, later in that day, she came back to that nest two, observed it, and then she started bringing dead insects-- worms, or caterpillars-- to that nest.

And that's what the green dots here-- and if it's open, all she did is observe what is she observing? Well, she's observing if the caterpillar has been eaten. And also, she observes the state of the larvae. How developed is it? It is ready for metamorphosis, so she can close the nest?

And so here, she just observes it. Then she deposits three more dead animals. And then she closes it. Now there's enough. She's obviously able to make that judgment. And then she goes back to nest three, which she had only been to after digging-- she had only been two once, when she laid the egg.

She just observes it. And then she went and dug a fourth nest. Then she came back, observed that second one, and then, soon, came back with food and left it there. Then she went and dug a fifth nest, and laid an egg, and provisioned it. And then, she flew off to observe nest three.

You'd say, why did she have to come back? Well, she's obviously having to monitor. Is the larvae alive? Has it eaten the initial caterpillar that she left there? And if it needs more food, she starts bringing food to it. And that extended into the sixth day, where she brought food to it twice, and then she closed it up. And then she went to nest five, which had laid an egg at before. She looks at it, and then she brings a caterpillar to it. Here, she disappeared for a while; I think they don't know what she did. That led then to, finally, closure of that fifth nest.

So she's aware of multiple nests at the same time. As you see here, she moves from two to three-- four, and then five here. Once she closes it, she seems to remember that's been closed, so she doesn't come back there anymore.

And there's these three clear stages in the way she deals with. She has to remember the location of the nest. She has to remember the contents after she inspects it. And obviously, she has to be able to remember the status of her offspring there in the nest.

So then he did more research on how she finds those nests. And in this case, this is a wasp-- a sand wasp-- that has the nest here. And I've colored that position red. They did experiments with the wasp. They would capture her and take her various distances in various directions from that nest, and release her. And you can see-- if they released her in this vicinity, except for one case here, she found her way back to the nest. In one case here-- they don't always know what landmark she's using. But she didn't go directly to it. But then she found it.

But notice here. If they went a little bit further, she generally didn't find it. If they released the wasp up here, they found that in spite of it being a greater distance, she was able to find her way. Not directly, but by this more circuitous route. And this is a very large pine tree here, in this ridge. She's able to find the nest.

So obviously, it depends on what they remember from their experience of the terrain. These wasps do fly up and spend time observing the landmarks. And in their earlier experience, if you take them to an area that they haven't observed, then they frequently won't find their way back. Or they'll at least wander around a lot before they will find it.

So this is another one of the specific experiments about the use of landmarks. Here, in the initial period, there was food here they gave the animal. She's carrying back something towards the nest. And there's a row of trees here. They're actually artificial, so they can be moved. And all she had to do was follow the line of trees to get to her nest. And over a period of time, she did that regularly.

So then they shifted the position of the trees. Now, what does she do? And you can see here. She pretty much follows those trees. You can see here-- she's wandering around. Here, she doesn't find the nest at all. And another time, they moved the trees over here. And you can see what she does. She expects to see the nest here, and so she starts flying back and forth. Here, she actually did find the nest.

Because obviously, it's adaptive, if they don't find it, to start a random search pattern-- which is what she's doing. So sometimes, they do learn to adapt to this moving of trees. But in many cases, they don't-- like here. So then, what do they do? They would have to build a new nest and start over again.

So this is the use of landmark orientation. And they found specific differences sometimes. In this case, Adriance and Berens-- they thought they were studying the same species, but when you compare the details of their notes, you see that they're actually bringing different kinds of food to the nest. One's bringing here the sawflies-- another type of larvae for their own larvae. And this guy observed only one nest at a time, whereas this one, like the one we talked about a minute ago, had multiple nests. Telescoped, indicating that they might be at different periods of development, but she took care of multiple broods at the same time.

They had slightly different breeding periods, and they had different details in the way they closed up the nest. The way they handled the sand. And what they did with sand when they opened it. Did they scatter it? Did they take it back to where they were getting the sand? And so forth.

So these details-- at least, Adriance, when he read the Berens result-- he realized that they must be different. And even though to Adriance and Berens, they looked the same, he decided, why not send them to a guy highly trained in morphology? Get him to examine these wasps. That was done, and the conclusion was, they're actually different species.

So there, they've been named separate species-- not because of, initially, the structure; that came later. Only after the behavioral differences were discovered. Because remember, the behavioral patterns, the fixed action patterns are inherited, just like structure of the body.

We gave an example of navigation when we talked about the Lawrence question about, can a fixed action pattern actually be maladaptive? And why would they become maladaptive? And he brought up the problem a full of photopollution in hatchling sea turtles. And this was done by Scott here.

And people have used knowledge of how these little turtles orient towards the sea to prevent the kind of photopollution that causes them to fail to reach the sea. And we know that the little sea turtle hatchlings-- I've pointed this out-- they tend to orient towards the sea because they're attracted to the blue light. And the sea obviously doesn't have trees impeding the view and the light, so that will be where the most of the blue light's coming from-- towards the sea, instead of towards inland areas. And they will crawl in that direction. Because during evolution, the strongest light was always seaward. Words

But then electric lights show up in a town. And the electric lights can be blue lights. Neon lights are often blue. Hot fires can make them start crawling towards the fire, even at night, just because of the innate behavior that they've inherited.

So, the way conservationists have solved the problem is, after studying it and realizing what very simple cues these turtles are following, they studied their reactions to other colors and found that if they used yellow lights or red lights, they're not as attractive as the blue lights. They also don't like flashing lights. Generally the sun doesn't flash, and the sky doesn't flash.

So they started using flashing orange lights. They talked to the people at the companies that were using the neon and doing the advertising. Putting up the billboards with neon signs. Putting it at their stores.

If they were near beaches, they persuaded them to change the color and to use more flashing lights, especially orange lights. And also, local towns made rules about when beach fires could not be used. The period of hatching of the little sea turtles. And it's led to a great increase in the survival of these sea turtles. Just by taking these very simple stimuli that they used to get to the ocean into account.

Let's talk about dead reckoning. It's a funny term. It actually comes from the phrase deduced reckoning. If humans leave one place and walk a certain distance away-- say, in a forest-- but then they want to get back to where they started from, and they haven't left a string behind them-- as we know from fairy tales, that can be done-- how do they get back? Well, humans can often remember directions pretty well. They keep track of directions. So when they turn, they know they've turned. And they know roughly how far they've walked.

Some people are a lot better at this than others. But that's a kind of deduced reckoning, using internal cues. Especially vestibular cues. But ants that have been studied are generally using external cues for direction. The position of the sun, and visual landmarks.

Let's skip down, because we'll talk about the Sahara Desert and desert ants in a minute. Scott gives the example of wood ants that use skyline features as visual landmarks. There haven't been many studies, but there are studies actually more extensive than the ones I did with Syrian hamsters, where you showed that Syrian hamsters can find their way back from a foraging trip. They go some distance from their tunnels. So it was a problem. How do they get back? Are they just using olfaction?

And it turns out-- no. They use visual landmarks when they can. They use a long horizon. Rabbits do the same thing. In fact, rabbits have denser receptors in their retina, in the part of the retina that represents the horizon. I have not noticed this in studies of the hamster, but hamsters definitely use these horizon cues. And probably, many-- some of the terrain they come from is mountainous, so probably they need to be able to use cues well above the flat horizon as well.

But anyway, they use visual cues, unless the levels of light get so low it's really almost completely dark. Then they will switch and use olfactory cues. So they're capable of using multiple cues for their orientation. Scott also uses the example of another wood ant. He uses vision to judge distance, by using landmarks at a certain height. And he will adjust his distance from the landmarks. So the angle is at a certain point. And by varying the height of that landmark, they can change the navigation of the wood ant. Like, the elevation of the top of a wall on one side.

The Sahara Desert ants are particularly interesting. And this is especially because there are some species of ant that go out in the hottest part of the day, which by itself is very surprising. But they're in an area with basically no landmarks. And yet, they go out from the nest. They wander around. They seem to meander around. What they're looking for? Dead insects, especially. Small dead animals that they can use as food.

And if they find it, then they take it. And they don't follow their trail back. And some ants are known-- other species of ant; forest ants, for example-- will often even olfactory trail. Or they will remember by other cues how to get back, but then they'll leave a trail when they come back. And other ants will leave the nest and follow that olfactory trail.

But not these desert ants. They also seem to know exactly how to get back. I made up this slide based on what I heard Rudiger Wehner talk about here in a colloquium talk he gave in October of 2011. Just two years ago. And if you search under Wehner, you will find a number of these studies that he's done with students on how these ants do it.

He studied-- this is called path integration. And also, studies of landmark orientation that we've been talking about. Now let's talk about path integration. It's like the humans use, that are just using vestibular cues and their strides to know about how far and in what direction they've walked.

They chose the study of this desert ant because there's just no landmarks in the Sahara, in the area they were studying, that these ants could be using. Just sand. And he found that they will walk a little ways and they will stop. Look around. Then they will walk further. So they could be stopping here, here, here, here, and so forth. They walk and stop. Walk and stop.

What they're doing is, they're checking direction. How do they know direction? Basically, from the position of the sun. And they don't actually have to see the sun itself; they can see the pattern of polarization in the sky. We can't see that, but some birds can see that, and ants can-- many insects can see it. They're sensitive to polarity of the light coming from various parts of the sky. And it has a particular pattern that varies with how high in the sky the light is.

OK, but what if the ant takes a long time before he finds an insect? So time has passed? Well, you could say, that's why he has to stop so often-- to recalibrate here. But in fact, you can transport them, and they'll still keep the direction clear from looking at the sky. And even if time has passed. You can trap them here. So you trap the animal here, and you keep him there. Let the sun move quite a bit. Then you let him go. Or you trap him just after he finds a dead insect. But now you keep him from going back right away.

He still finds his way back, because he's got a biological clock. His knowledge of direction from the sun is always re-calibrated according to time of day.

So anyway, the animal finds the insect, and he heads directly back to his nest. And he can do that without any landmarks. Now, if there are some landmarks, they do the same thing. But now they keep correcting for any cumulative error that occurred. Because they're not going to be absolutely perfect. But how do they figure out-- just knowing the direction back to the nest every time isn't enough. You have to know how far you've traveled, too. How do they know that?

So they did experiments. Some amazing experiments. They were really quite surprising. They know, because of how many steps they took and the length of their stride. How did they know that? They did experiments where they put little stilts on the ant. If you put the stilts on the ant after he finds food, and then you release him, he's actually moving a lot further with the same number of steps. And he will make errors. He'll go beyond his nest.

How can they vary the number of steps, but not the stride length? Well, they put ramps right in his path, so the ant has to go an additional distance. He has to make many more steps to get the same distance, because of those ramps.

Remember, these ants are-- if it were a human, you would notice. You might take some of this into account. But the ant is limited to his use of direction and step length. And number of steps. So he's constantly keeping track of that, so when he gets to his insect here, he gets back pretty accurately to that nest. There will be some cumulative errors if there's no landmarks at all he's using to recalibrate.

So now let's go to homing pigeons. And this is dealt with more in Scott. And pigeons, we know-- they know a direction. They've got some kind of compass. And at least when they're in the region of their nest, their loft, they do have a map of the region. So they use both a compass and a map.

And we know they use multiple cues. There's evidence now-- and it was very controversial. 30, 40 years ago-- even now, you see some studies of this, but now there's much more agreement that they are capable of using the earth's magnetic field to know direction. But they use the sun if it's available. If there's no sun, or if they're traveling at night, they use the magnetic field. Or, in the case of songbirds, they can actually use the stars.

When they're in the vicinity, they don't have to be perfect to do that. They just have to get to the vicinity of their loft, because then they have a map. Because they do spend time flying around their loft, and they do learn all the landmarks around there. They develop a map. And so, in the vicinity of their loft-- because pigeons can home from very long distances. So they might not get that close, just from their compass sense.

But they use not just visual landmarks. There is evidence that they use olfactory cues. And there's even some evidence that patterns of ultrasound can make a difference, and be used by at least some of these birds.

We have time to talk about the Morris water maze, used here in the lab for studies of mice and rats. What is the Morris water maze? You do this horrible thing to these little guys. You throw them in water. And the water's too deep. And they have to swim to stay alive.

It's usually a circular maze filled with water. In one place, there's a little platform. He can't see it. But if he finds it-- if he swims to it-- he can climb there, and then he can stop swimming and rest.

So if you put them in there repeatedly, they learn where that platform is. Well, how do they learn? They must be using specific cues. How do they learn the position of the platform that they can't see? You can do studies. I'm summarizing there, some of the studies that have been done.

We, of course, have the maze in a room. You can rotate the maze around. So all the visual cues in the room stay the same, but you're rotating the maze. What happens? He goes to the wrong place, because he's using those visual cues. Another way that that's been done, just to make sure he's not using the magnetic field or whatever-- you rotate the room. Meaning, you have sort of a room inside a room. You can rotate that, and leave the maze constant.

But the mistakes he makes always seem to depend on the visual cues. The overhead cues in particular. Our position with respect to that platform. So that's how they find it.

And we know that the hippocampal formation, so important in forming a cognitive map of the environment, is necessary for the pigeon-- not for his compass direction. They'll still go in the right direction, even without a hippocampus. But they won't have that map. That will be messed up by hippocampal damage. And similarly in the Morris water maze, they can't use to cues to form a map of that maze without the hippocampus.

We'll come back here to just finish the last few slides next time.