JEREMY WOLFE: Good afternoon. Here is the plan for today's lecture. There is out there somewhere, unless you're a philosopher, a world. It's huge. A lot of stuff out there in the world.

There is inside you somewhere a long-term memory, which is also huge, and hopefully getting huger all the time. You cannot transcribe all of that world from the outside into your long-term memory, as you may have already discovered if you've had a test this term. And in a move that ought to look familiar to you by now, we can describe that as a bottleneck of some sort-- not the same bottleneck that I was talking about in the context of attention, but a bottleneck, nevertheless, that is involved in putting a severe limitation on the ability to encode the world into long-term memory.

The first part of today's lecture is going to be about this problem of encoding and about the nature of this bottleneck. The second part is that some stuff-- a mere trickle of the world-- managed to get through this bottleneck. It needs to stay there. It needs to stay there for a long time.

You've got things in here, in this long-term memory of yours that are 15, 16 years old, maybe more at this point. How they stay here, how they are made firm, is the problem of consolidation, which it would say underneath there, if it wasn't that I put a screen in front of it. And finally, it is of absolutely no use to you to have-- as you also may have discovered on some exam this term already-- to have a long-term memory that is absolutely chock full of marvelous material if you can't get it back out.

Unlike the story that we were telling before about perception where there was a world a bottleneck and then some construct that was your perception of the world, in this case, you've got a two-way street. That you cannot appreciate the full contents of your long-term memory at any given moment you need to retrieve-- that's what it says on the far board-- you need to be able to retrieve material from long-term memory and bring it back-- as we will see, in effect to bring it back into the bottleneck. Got to be able to get back out into this limited capacity bottleneck where you can do things write down the answer on a test or tell me who's running for president this year. Nobody's-- boy, that's very--

STUDENT: Bush.

JEREMY WOLFE: Well, yeah, yeah, that's one of them, Bush. Yeah.

STUDENT: Kerry.

JEREMY WOLFE: Kerry. OK, and a few other characters, too, right. Presumably, except for the most politically obsessed among you, that was not an active thought in your mind until I asked it. But Kerry and Bush live in your long-term memory. Isn't that a scary thought.

And on demand, you could retrieve that back out into what we will call your working memory. If we were just thinking about this in encoding terms, we might just call this short-term memory. Stuff comes in from the outside, it's maintained for a little while somehow in a short-term memory, and then somehow makes it into long-term memory in ways that we'll talk about.

But if we think about this in a two-way, in a bidirectional kind of way, it's better perhaps to think of this as your working memory. Think of it as your computer desktop. It's got access to all the stuff that's on your hard drive.

It's got access to the external world. You can tell it's got access to the-- oh, no, you can't see it there. Well, it's wireless. And so it's got access to the web. And so that's the external world.

But only a limited amount of it can be on the desktop at any one time. Similarly, only a limited amount of either the world or the contents of your long-term memory can be on your mental desktop at one time. That's working memory. That's an important chunk of this bottleneck.

Well, what does it mean to say that this is a limited capacity bottleneck of some sort? Let me illustrate the limit to that capacity. We saw an example, a particular example of that, in what's known in the trade as visual short-term memory last week when I was putting up four little, colored blobs and saying, now what changed? And there's a limit of about four objects that you can keep track of.

But in looking at the-- and trying to ask what's the limit of this bottleneck from the outside world into more long-term varieties of memory, one of the typical measures is what's called digit span. If you were doing-- you end up in the hospital because somebody bopped you over the head. One of the things your neurologist will do is read you a list of numbers and ask you to repeat them back. And how many you can do is an index of whether you are OK.

So let's see if you're OK. Except that I won't do this with numbers. I will do it with color names, and we can come back to that in a minute.

Where are my color names? Here are my color names. So what I'm going to do is read you a list of color names.

I'll then say, "Repeat," and you, in glorious unison, repeat this back to me. So if I say, "Red, blue, purple, repeat," you say--

STUDENTS IN UNISON: Red, blue, purple.

STUDENT: Repeat.

JEREMY WOLFE: It never fails.

[LAUGHTER]

Right, Mara? Last year, too. And always, you get the "repeat" back, too. And I'm not sure what it is about this that demands, at least from some people, the repeat of "repeat."

But in any case, good. You're not deeply brain damaged yet. Ready? We'll do some more of these. Green, pink, yellow, white, yellow, repeat.

STUDENTS IN UNISON: Green, pink, yellow, white, yellow.

JEREMY WOLFE: OK, that was five of them. All right, here we go. Oh, I forgot to mention-- it turns out to be really easy to repeat back an essentially infinitely long list of these if you write them down while I'm saying them.

[LAUGHTER]

There are a variety of demos that I'm going to do today that are of that form. And it's really boring. And so OK, you can do that. So don't do that because that's boring. Be a nice, honest person.

OK, ready? Where was I? Just did five, let's try seven. Ready? Oh, and I told you when I'm going to stop. Oh, what the heck. Pink, red, yellow, green, purple, pink, orange, repeat.

STUDENTS IN UNISON: Pink red, yellow, green, purple, pink, orange.

JEREMY WOLFE: It's getting a little weak there at the end. But you will see-- if you look on the handout, you'll see that the answer to this question is given in the title of George Miller's classic paper on the subject. And the answer is 7 plus or minus 2. And some of you are on the minus side.

Oh, I should also say-- I should have said this at the beginning of the course-- people in intro psych classes, particularly by the time you get around to starting to talk about cognitive things like memory, take demos very personally. And think 7 plus or minus 2, I got minus-- I'm doomed. Nah, you're probably fine.

I mean, if it was the red, yellow, purple, you're saying what was the one after red, there might be an issue there. But the advantage of a 300 person intro class is we're looking at the average. And what you're doing on any given single isolated demo trial is not actually deeply relevant to whether or not you're going to pass physics, for instance.

OK, one more. Ready? Blue, green, purple, red, pink, yellow, green, blue, white, orange, repeat.

STUDENT: Blue, green, purple, red, pink, yellow.

[LAUGHTER]

JEREMY WOLFE: Yeah, yeah, yeah. You can hear that was 10. And you could hear it in the group and you could probably feel it in yourself that up to six, seven, eight-- yeah, we're good here.

Oh, man, if I start taking in this next one, I'm losing them here. Something bad is happening. So there's this capacity limit on what you can get into-- well, let's use just jargon for the time being-- into some short-term memory and then spit out immediately to me.

Now, 6, 7 plus or minus 2, what? Let's try another one. You ready? Same game. Red, red, red, red, blue, blue, blue, blue, green, green, green, green, repeat.

[LAUGHTER]

Never mind. You know you can do this. That's 12 items. How come I all of a sudden managed to boost your memories so effectively?

The classification, or what Miller originally-- and the field has stuck with-- the notion of chunking-- if you can break things into meaningful chunks, it's 7 plus or minus 2 chunks. And what's important is how you manage to cut the world into chunks.

This, by the way, explains why I don't do this demo in the traditional way of giving you digit strings. Because years ago, I discovered that if you give 300 MIT undergraduates a 10-digit string, a bunch of them cut out on you. But then, there's a handful of people who just rattle it right off.

And you say, oh, my goodness, these people have this amazing memory. And then, you ask them, how did you do it? And they say, oh, well, that was obvious. It was the natural log of 16.

[LAUGHTER]

MIT students tend to chunk numbers rather more readily than the rest of the population. They don't tend to chunk color names particularly adeptly, I've found. And so you get the right 7 plus or minus 2 answer.

Now, this sort of chunking-- I mean, you should think of this more like a rate. It's not 7 plus or minus 2 things ever, ever, ever. It's 7 plus or minus 2 things in some unit of time that you can manage to deal with.

Otherwise, as soon as I had presented 7 plus or minus 2 units of information in this lecture, you'd be like dead for the rest of the day. Or maybe the first lecture shot your brain and that was it for the term. That's obviously not the case.

The question is, how big a chunk can you manage to cut things into as it's coming in? One of the places you can see that is in language as you're trying to keep track of, for instance, what I'm saying. So if I ask you to repeat this sentence, "Memory is a fascinating topic that will be on the final exam," you say.

STUDENTS IN UNISON: Memory is a fascinating topic that will be on the final exam.

JEREMY WOLFE: But if I lettuce, quadrangle, forms, only, only, column, bug, dela, forms, aft, tunic, reply, quintillion, you say--

[LAUGHTER]

Roughly the same number of syllables, but you can't repeat that back because you can't chunk it. Well, you can, but the chunks that you were making were too small to survive the rate at which they were being presented. So everybody here, if forced to, would have produced some little bit of it, but people would not be able to produce the whole thing because you can't chunk it into meaningful chunks that work.

This is also the route, presumably, of the near-universal experience that when you go to some other country where they speak another language that you nominally speak because you took four years of it in high school, you discover that those Parisians, or those Spaniards, or the Japanese, or whoever, talk really, really, really fast. What's the chance that everybody else in the world talks really, really, really fast? That's not true, of course.

And those of you who are non-native English speakers, and have now arrived at MIT, are thinking, no, back in Kazakhstan, they all spoke really slowly or whatever. But these English speakers are wackos. The problem is in your native tongue or in any tongue that you are really fluent in, you are very good at cutting into meaningful chunks rapidly.

If you're busy trying to say, what was that word? Damn it, that was the week I was asleep in high school, or something-- no, it was the year I was asleep in high school. Anyway, the chunks go by too fast like in the nonsense that I was producing for you.

So the ability-- you're set up to get chunks through in some domains much better than in others. And in fact, let me illustrate that here by-- where did I hide it? That looks like it's probably it.

All right, got to memorize some more stuff here. Come on. You want to go, and there we go. Too much light. Let's kill the stage lights. Good.

I want you to memorize all these pictures. You ready? There's a picture, and there's a picture, and that's a picture, and there's the back end of a horse or two, and there's a Christmas tree or something. And I don't know, I must have raided a housing site.

But these are definitely pictures. And that's a picture. And that looks like a picture. And oh, look, there's tomatoes. OK, got them all?

OK, now what I want you to do is I'm going to show you some pictures that are either old or new. And you just say yes if you've seen it before, no if you haven't.

STUDENTS IN UNISON: No. No. Yes. Yes. No. Yes. Yes. No. No. Yes-- no.

JEREMY WOLFE: OK, the important point is that you're very good at this. The interesting side point is that this particular one is flipped, left/right reverse. So if you were sitting there saying, why are these people next to me changing their mind, it's because some people noticed that it was left/right reversed.

But if I had lots of time, I could have done this with a lot of pictures. The largest report in the literature is 10,000 pictures shown to the usual collection of college undergraduates. And some days later, the students still remembered, I think it's about 80% of those, about 80% accuracy on that.

If I read you 10,000 color names-- well, there aren't 10,000 color names-- if I read you 10,000 words or almost any other material that I would care to present to you, you would do nowhere near as well. We are not entirely clear on how you do this. This is quite a remarkable ability, but it does speak to the ability of picture information, familiarity with a picture, getting through this bottleneck into some sort of long-term memory with remarkable efficiency-- that you are somehow specialized for doing that.

There's also a little bit of a cheat here. And it's only a little bit of a cheat, but it's an interesting cheat. This was a recognition test.

The color one, the previous demos, were recall tests. I was asking you pull out of your memory what was there. Here, I was saying, was this there?

And you know this from your test-taking experience. You know that if I give you a choice between fill in the blank and multiple choice, which do you want to do?

STUDENTS IN UNISON: Multiple choice.

JEREMY WOLFE: Multiple choice because that's a recognition test. Fill in the blank is a recall test. The distinction is-- OK, here's your long-term memory. Here's the target that you're trying to find in long-term memory somewhere.

If I'm doing a recall test, I'm saying, go find this thing in all the vast halls of your long-term memory. And you send some little probe out into long-term memory that may or may not find it. If I say, is this in there, that's an obviously simpler matching kind of task, and you are typically much better at doing recognition than you are at doing recall.

So you have, presumably, special-purpose mechanisms in your system that are really good at doing some tasks. What do you do in the case of those unfortunate tasks-- actually, I don't need that anymore at all-- in those unfortunate tasks which you were not well designed to do? For instance, the memorization of neuroanatomical terms for a midterm in intro psych, which is only the beginning.

How many people here are pre-meds or think they might be pre-meds? You guys are the ones who are going to devote years of your life to the memorization of things that you were not built to memorize, like where all the bones are and stuff like that. So how do you do that?

Let's try memorizing some nonsense. And go to sleep. Why don't you want to go to sleep? It's like you're children.

Mara, what do you do? You don't want to know. I don't want to know. You mean I'm supposed to hit it? OK, all right. It went to sleep. It happens from time to time anyway.

Where were we? Oh, yes, nonsense-- this is a line of work that starts, really, with a guy named Ebbinghaus in Germany in the late 19th century who wanted to understand memory processes in isolation from these issues about meaning. And so what he did was he got himself, actually for starters, to memorize nonsense syllables that allegedly had no meaning.

It's extremely hard to find material that has no meaning. But I will attempt to read you some. Now grab a hunk of paper. I don't want you writing them down while I read them to you, but I do want you to be in a position to write all the words down when I tell you it's time to write.

Oh, and while you're doing that, let me make my annual disclaimer. MIT is a hugely multicultural institution. You speak a vast range of languages. To my knowledge, what I am saying is meaningless.

If it turns out that it is meaningful, particularly if it turns out that it's deeply profane, do let me know later and I will change my nonsense yet again. I have, over the years, said a number of absolutely remarkable things, including something in Malay so good that the young women who told me I couldn't say it wouldn't tell me what it was I had said. But they looked at my list and they just crossed these things out.

Anyway, so I don't think-- it's mostly gone. So all right, you're ready to write? I'm going to read you a bunch of nonsense syllables, and you're going to write them all down. Ready? OK.

Fip, dut, moke, yil, saz, tirt, varl, bince, jucks, goff, zos, rab. OK, write down everything you can remember.

[SIDE CONVERSATION]

Now, don't copy off your neighbor. On your handout, on the second page, I believe, you will find the axes for what's called a serial position curve, which is the way we want to represent the data from this experiment. We want to represent percentage recall, which in this case will be a certain number of hands, as a function of position in the list, where in the list that I read-- where did the word show up?

And I will tell you the answer ahead of time and then we'll check if it worked. The answer is that the data will look like this. There will be what's known as a primacy effect, where you'll remember the first words in the list quite well, and a recency effect, with some preservation of the last words in the list. That, at least, is my assertion.

Let us actually check the data here. How many people got fip? Data point turns out to be about there, I think. How many people got dut? A few fewer.

How many people got-- well, you see you're falling a little below the line here. I'm worried about you guys. How many people got tirt? How interesting.

How many people got varl? Well, that's pretty pathetic. How many people got zos? Oh, coming back-- that's pretty good, actually. That's above the line here.

And how many people got rab? That's sort of there-ish. OK, what's with tirt? You guys did better. Well, I discovered-- teaching concourse this morning, I discovered that one of my words has currently acquired a slang meaning that I'd never heard of before. So tirt you just got lucky on, or tirt suddenly mean--

[INTERPOSING VOICES]

JEREMY WOLFE: What's a tirt?

STUDENT: [INAUDIBLE]

JEREMY WOLFE: What?

STUDENT: Tirt beetle.

STUDENT: Tertiary.

JEREMY WOLFE: Oh, oh, tert like in tertiary or something? Oh, see, I'm spelling it T-I-R-T and it never occurred to me that it's-- all right, all right, we're going to have to do something about tirt. Anyway, the question here is going to be, how did you do it?

Part of the answer is clearly any word that happened to make some meaningful association to you is going to have a preferential chance of getting into long-term memory. Where did my beautiful, little map of the world go? It must be underneath there. You go up.

One way to think about this is by analogy to immigration. This here is America. And here are all of our various forebears and stuff trying to get into the country, and there's this choke point in immigration.

If you got your Uncle Charlie already in the country, he can pull you in, in ways that if you have no relations here, it's harder to get in. So if tirt reminded you of tertiary, that's Uncle Tertiary in long-term memory pulling you through. But the question is, what do you do if you don't have that?

So the things at the beginning, anybody got any intuition about what they were doing with fip, and dut, and stuff like that? Yeah, OK.

STUDENT: I was writing them down as [INAUDIBLE] and [INAUDIBLE].

JEREMY WOLFE: Oh, well, apart from the fact that you presumably were being a good person and not writing them down when I read them first, though. OK, so what were you doing while you were waiting to dump them back out? Yeah.

STUDENT: I mean, I was going over them in my head, rehearsing.

JEREMY WOLFE: Yeah, OK, you were sitting there rehearsing. How many people had the experience that they were rehearsing and discovering that, oh, man, there's more stuff here than I can rehearse? That's the capacity limit again.

But there's a rehearsal loop for auditory material called, sometimes called a phonological loop. Does it say phonological loop on the handout anywhere? No? Well, phono, like phonograph, phonological, where you repeat this stuff to hold it in this working memory, keep it alive in working memory, and in that way, keep it alive. People probably were not-- well, anybody got a different experience for zos and rab at the end there?

STUDENT: No.

STUDENT: [INAUDIBLE]

JEREMY WOLFE: Sorry, I heard something that sounded promising. Yeah.

STUDENT: They were just there.

JEREMY WOLFE: They were just there is actually the right intuition. They were somehow still-- it was almost like you could still hear them. They didn't need to be preserved yet. They were still fresh and hadn't rotted.

The claim though, is that the process that gives you the recency effect is different than the process that gives you the primacy effect. Let's give that a try. What I'll do is I'll read you another list of nonsense like the first one.

I'm going to ask you to write them down again, but this time at the end, instead of saying "write," I'll say "count." When I say "count," I want you to count backwards from 431 by 3's out loud, all in beautiful unison. Got that?

That will make a lot of noise. So when I do this-- keep an eye on me-- when I do this, write everything down. OK? They kind of got the instruction.

All right, where's the rest of my nonsense here? OK. OK, ready? Sep, forf, lig, vop, hearn, mope, jik, tind, mez, wamp, flob, gan. Count 431--

STUDENT: 428.

[LAUGHTER]

JEREMY WOLFE: All right, OK, write them all down. The giggles will probably do just fine.

[SIDE CONVERSATION]

OK, the claim here is that what that should do-- did you feel yourself rehearsing the first ones again?

STUDENT: Yes.

JEREMY WOLFE: The claim is that the primacy effect, which is due to some sort of rehearsal into long-term memory-- that primacy effect should still be there. But the recency effect due to some it's just thereness should have been disrupted by the counting, and the chortles, and stuff like that. OK, so let's check here.

Where's my list gone to? Here, list. I should remember this by now. OK, how many people got sep? Look, that's almost exactly the same number.

How many people got forf? Oh, a lot of forfs. How many people got jik? Ooh, deeply pathetic.

How many people got mez? Oh, a few more. How many people got flob? Oh, that's pretty pathetic.

How many people got gan? No, gan was pretty well gone. So these were supposed to be circles so there'd be a different data set. So anyway, whoop-- nice primacy effect. Recency effect gone.

STUDENT: What about [INAUDIBLE]?

JEREMY WOLFE: Oh, yeah, there are other words in there. This is the phenomenon of MIT students really wanting to do well on these things. Man, he didn't ask me.

It's like you have the same experience for keeps on the exam. You sit there and you study chapter 3 till you're blue. And oh, man, they only asked about chapter 2. Ugh.

Yeah, yeah, I'm glad you got whatever it was, wamp or-- the worst thing about this, you realize, is that the stuff that you rehearsed into long-term memory may stay there for a surprisingly long time. I had a guy come up to me on the subway some years back, say to me, "You don't know me, right? I said, "That's right, I don't."

He said, "Fip, dut." I said, "You are a very strange person, but I now know where I know you from, or don't know you from, or something like that. So you may have now warped your brain on an ongoing basis.

No, no, I don't want you to go up. Get back here. Stop. Come back down. You go up.

All right, so back on the first page of the handout, you have what was sometimes called the standard model for getting from the outside world-- whoops, my bottleneck disappeared again-- into long-term memory, which was stuff comes in to a short-term memory. It needs to be preserved in short-term memory by something like rehearsal. If it lasts long enough in short-term memory, it gets into long-term memory, where it manages to stay.

There are certain problems with this kind of a model, which you ought to be able to anticipate now, but we can illustrate easily. Enough how many people here had breakfast? OK, person with the MIT shirt on there, what did you have for breakfast?

STUDENT: Bagels.

JEREMY WOLFE: Bagels, anything else?

STUDENT: Cream cheese.

JEREMY WOLFE: Bagels and cream cheese. OK, that's good. We want a little bit of a memory load here because I need to be able to explain that the reason that she knows that she had bagels and cream cheese for breakfast is that on the way out of her living arrangements, she was going, "Bagels and cream cheese, bagels and cream cheese, bagels and cream cheese, bagels and cream"-- right?

STUDENT: [INAUDIBLE]

JEREMY WOLFE: OK, that's good. No, she wasn't doing that. You weren't doing that, but you remember what you had for breakfast, too. Obviously, this rehearsal thing, while it's important for getting nonsense from the world into long-term memory, is not the only way to get through this bottleneck.

Again, things with meaning manage to get there more readily. And we can see something about-- that fact tells us something, perhaps, about the way that our long-term memory is put together. Now here, what we're talking about is what's known as explicit long-term memory is distinct from implicit.

Explicit long-term memory are the memories that you can get to if I ask you about them, that you can recover. So if I say, who's running for president, you have an explicit memory that the answer is Kerry, Bush, and whoever else you can name down the line there. If I ask you, how do you say the word "president," what are the tongue motions that are required, if you say "president," you can now monitor it and figure out what they are.

But you have no notion ahead of time of how you do that. You certainly remember how to do it, in the sense that you can produce it on demand. That would be an implicit memory.

And how do you ride a bicycle-- there are a whole slew of things like motor memories that are implicit. But here, what we're talking about is getting into an explicit long-term memory. It's useful to divide explicit long-term memory into episodic memories-- memories for the episodes of your life-- and semantic memories-- your body of knowledge, and facts, and things of that sort.

They're not walled off from each other. If I say, what's the capital of the United States, you dig out of your semantic memory the answer Washington DC. If you've been to Washington, you might then start digging out associated episodic memories of being there, and so on.

But within semantic memory, the notion here-- remember the Uncle Charlie or whoever he was notion that things can reach out to the world and help pull things into long-term memory-- is related to the way that we think that your long-term memory might be organized. And one of the popular ways of thinking about this is a giant network of associations called a semantic network, often, where if I say, "cat," the first word that comes to your mind is--

STUDENT: Meow.

STUDENT: Meow.

JEREMY WOLFE: Meow, dog, mouse. So there's a node in your long-term memory that is "cat" in some sense in this story. Its close neighbors are things like dog, and meow, and-- what did we have-- mouse.

And now, it's important to realize that this is not some Linnaean taxonomy of species, and genus, and all that good stuff. Even though there's going to be connected with animal up here somewhere-- animal, and fur, and stuff. But you might also have connections that go like, meow, Mao--

[LAUGHTER]

Actually, that goes back here. I actually have a cat whose name is Chairman Mao by roughly this association. And anyhow, China, plates. Nobody has the notion that you could somehow map this out for any given individual. It's going to be a vastly complex and continuously changing set of associations.

But there is evidence for this notion of proximity by association, by meaning, in long-term memory. And the way you might find that out, one of the ways you might find that out, is by doing what are called priming experiments, which I might as well mention because I see that's-- is that on the handout properly? Well, if it isn't, it should be. It probably says priming somewhere.

Priming experiment-- what you do here-- let's give as-- evidence from priming, look at that. Suppose I put up on my computer screen here, I'll put up a string of letters. And you tell me whether or not it's a word. So I put up, and you say--

STUDENTS IN UNISON: No.

JEREMY WOLFE: No, even though it's now in your long-term memory. But if I put up--

STUDENT: [INAUDIBLE]

JEREMY WOLFE: Yes. OK, now, if the next word was something like truck-- unless you have a dog in your truck or something like that-- well, all right, let's compare this. The next word might be truck or the next word might be cat. You'll be faster to confirm that cat is a word than truck is a word.

Why is that? Well, seemingly, when you go in to your semantic network to discover that "dog" is a word, you light up the dog node. And the activity spreads-- it's called spreading activation-- away from that node to the neighboring nodes.

As a result, when you get the word "cat," cats are already a little activated. The bell's been rung a little bit already. And you're quicker to confirm that "cat" is a word than "truck," which is down here in the next county somewhere.

You can use-- I mean, in principle, I suppose, you could map out somebody's semantic network with a technique like this. But it would be a lifetime's endeavor and not clear what the point is. But you can use this sort of a technique to show that some things live closer to each other than other things.

And by the way, when you go and reach in to retrieve-- this is jumping to the retrieval part-- but when you go reach in to retrieve something from one node, you might goof and pull something from the neighboring node. You know this, or you have experienced this if you are not an only child, perhaps, even if you are.

How many of you have siblings? Keep your hands up. How many of those have ever had your parents call you by the sibling's name? Almost everybody, including some people who didn't have their hands up before, who didn't know that they had siblings until that happened. Or worse, you get called by the name of the dog or something like that.

This is not because-- whatever you may believe-- that your parents are unusually dim people. It's because all those terms presumably live very close to each other in this semantic memory. And when particularly under some pressure, you reach in to grab kid number-- all my kids have the same name. Their kid, then the name is Ben-Philip-Simon, whatever your name is. And it works for all of them.

Now, there are certain drawbacks to this. It turns out to be decidedly unfortunate if in a moment of passion, you murmur the name of the last girlfriend, no matter how close they live in your semantic network. We can explain it, but we can't excuse it, if you know what I--

So big world, small desktop, small working memory, big long-term memory-- how does stuff-- what's going on when the stuff in long-term memory gets turned into something long term? It's really there for a long, long time. If you get a chance, David Pillemer's article that's in the website under the Memory category, about very long-term memories, is fun to read as he goes and probes people's oldest memories from way, way back.

How does stuff get consolidated? Well, we know some things about this. We know quite a lot about this, but I want to tell you one bit, a couple of bits, that are important on a physiological front.

So here's the brain again. Underneath the temporal lobe-- not in the cortex of the temporal lobe, but underneath the temporal lobe-- are-- struck the hippocampus, and to a lesser extent the amygdala-- I'll just mention that but really, we're talking about hippocampus here, parts of the so-called limbic system that are vitally important in this process of consolidation. And we know that in the first instance, from perhaps the most famous patient in the neuropsychological literature.

He's a patient known in the literature as HM. He's the inspiration for the film*Memento,* if you saw *Memento.* When he was a young man, he had a bike accident that was probably the cause of his epilepsy.

So he was an epileptic. You may recall I said earlier in the course that epilepsy is a sort of an electrical storm, often started from a damaged piece of tissue in the brain. If it's not responding to drugs, which this was not, one of the treatments is to go and try to excise to take out the generator.

The evidence in HM's case pointed to structures deep in the temporal lobe, particularly the hippocampus and amygdala. And what was done in the late '50s or early '60s-- anybody remember? Anyway, a long time ago at this point, up in Montreal, was that these structures were largely destroyed on both sides of his brain. That's important. This was a bilateral lesion.

And this did have the effect of largely controlling his seizures. The seizure problem went away. However, he is a one-of-a-kind patient because the side effects, the unintentional consequences of this lesion, were so devastating that nobody could ethically ever do this again, at least not to a human patient.

So he's still alive. And in fact, he's in a nursing facility near Boston because he was brought to Boston essentially as a psychological subject, been unable to live on his own since the surgery. What's his problem? Well, he really has two, but the most relevant one for the present purposes is that he is simply unable to form new, explicit, long-term memories.

So take his interaction with a once upon a time TA in this course, named John Gabrieli, now a professor at Stanford, soon to be a professor back here at MIT, but who I think he did his doctoral work studying HM. But he certainly studied him as a graduate student.

He'd come into the lab and there to meet HM. Hi, have we met before? HM said, no, I don't think so.

He has perfectly intact short-term memory, which is fine for conducting conversation, also perfectly intact or decently intact memories, long-term memory from before the operation. So it's not that we've somehow wiped out his storehouse of memories. It's the ability to put new stuff in here that's gone.

So in comes Gabrieli. Have we ever met before? No. Hi, I'm John Gabrieli. Nice to meet you. I'm HM-- well, I suppose he didn't call himself HM, but anyway you use initials to protect the privacy of the patient, not necessarily the initials of the patient or something to label them. So anyway, fine, have a nice conversation.

Gabrieli goes out, comes back in, says, hi, have we ever met before? No, I don't think so. Well, hi, I'm John Gabrieli. Nice to meet you. Have the same conversation. You can do this over and over again because he's got no--

So in a more controlled kind of way, do a digit span kind of experiment. Give him a few digits to remember. No problem. He's got the same 7 plus or minus 2 that the rest of us have, pretty much. Distract him momentarily, and it's just gone.

You've had this experience yourself if you tried to get from a phone book to the phone with a phone number, and somebody says, did you do x, y, and z? Oh, there goes the phone number. That's his continuous experience and has been for 50 years at this point.

A few things have gotten in. He happens to know that men have landed on the Moon for instance, perhaps because there are little bits of the structure that are maintained, perhaps from another route. We're not quite sure. But basically, no new episodic memories.

What this is telling us is that the hippocampus in particular, other research indicates, is critical for the act of consolidation. It is not the memory itself. It is the mechanism that allows you to make those memories solid in some fashion. Yes.

STUDENT: Does he know he has a disorder?

JEREMY WOLFE: He knows he's got a memory problem. And this actually is a good point. It ties into his other problem, which is, in some sense, an offsetting benefit.

The limbic system, certainly the amygdala, is very important in the mediation of emotion. His experience of emotion is greatly flattened-- flattened affect in the jargon of the trade. Which is good because he does know that there's something wrong, but it doesn't bother him that much.

You know what bothers him? What bothers him is looking in the mirror. Now, why should that be?

Well, imagine this-- his lesion happened when he's in his 20s. He's still got the long-term memories that he had in his 20s, basically your age. Now you go back to your room, you look in the mirror, and what looks out at you is a 70-year-old guy-- or woman, as the case may be.

I mean, I suppose it doesn't particularly matter. It's going to be a disaster under either circumstance. You're going to think, I'm in a sci-fi movie and I don't have the part, the good guy with the laser gun part, that I really wanted. I'm in big trouble here.

And this is disturbing to him. Not as disturbing as it would be to you, again, because his emotional responses have been very severely blunted. So much so that he needs to be asked whether or not he's sick.

Do you have a stomach ache? Oh, now that you mention it, yes, I do. It's not that he couldn't feel the pain, it's just without the emotional overlay. So I see blue, I have pain, what's the big deal? Yes.

STUDENT: [INAUDIBLE] motor memory?

JEREMY WOLFE: Yes, nice point. So for example, I don't know how many of you have hung out at-- get rid of the brain here, who needs the brain. So any of you ever done a mirror maze? Some people may have actually built one of these once upon a time.

So imagine you have a star. Oops. Imagine you could actually draw a star. Well, all right, well, so you make this out of aluminum foil and you go and trace it with a metal stylus. And if you go off the aluminum foil it makes a nasty noise because you're breaking a circuit or something like that.

Anybody ever build one of those? Oh, what's the world coming to. It's easy. Anybody can do this. What people can't do readily is do this looking at the image in a mirror. Because then, they [GIBBERISH] think they got-- oh [MONKEY NOISES].

But you learn to do this. And HM has been trained in one set of experiments to do this in multiple sessions. So he got better. If you ask him, have you ever done this before, on the n-th time that he did it, the answer is, no, I don't think so.

Well, he goes [MAKES NOISES]. Hey, that was really good. Most people don't do-- how did you do that? Yeah, I've always been good at that kind of thing or just lucky-- no episodic memory, no explicit memory of ever having done it, but yes, a motor memory that allows him to do it.

Similarly, your sense of familiarity runs on different pathways. So if I bump into you wandering around the halls, I may well have the sense that I've seen you somewhere before. And I might guess that I've seen you in intro psych because that's where I see a lot of people whose names I don't really quite know and stuff like that. So that would work.

And HM gets that sense of familiarity, too. So after a while, actually, with Gabrieli, Gabrieli comes in. Have we ever met before? Yeah, you look familiar.

Do you know my name? No, your name, forgot your name. Who do you think I am? I think you're probably somebody from my high school class.

Now, why would that be? Well, Gabrieli at the time was a young 20-something-year-old guy. This guy's got a 20-something-year-old long-term memory.

A reasonable guess of somebody who looks more or less familiar but isn't in my current collection of people I know well is, yeah, maybe from high school or something like that. We're about the same age, says the 70-year-old HM to the 20-something-year-old Gabrieli. So those are things he does learn. It's the new, explicit, long-term memories that require the hippocampus to be consolidated.

How long does this take? Let me say a quick word about that and then we'll take a quick break. You can measure the time course of consolidation by doing an experiment of the following sort. You can draw bad rats. OK, there's a rat.

He's on a little pedestal. If you are a rat on a small pedestal not too far above the floor, what are you going to do? I'm going to jump off. I'm a rat, man. I got to go and look around at stuff.

So I jump off. Well, what they didn't tell me back at home was that this particular floor is electrified. When I jump off my little toes get an unpleasant-- not damaging, but an unpleasant-- shock.

So urk. And now, the guy puts me back on here. What do I do? I don't jump off. I'm not a stupid rat, I'm just a rat. And so I stay there.

OK, well, I'm a rat who's in an experiment. And the question is, what percentage of the time do I jump off? And the variable here is after I jump off, and after I get my toes shocked, I'm also going to get a dose of what's called electroconvulsive shock, which is basically electrical current run through the brain.

You can kill people or animals that way. Obviously, that's not the point here. This is to disrupt the ongoing electrical activity of the brain, not to scramble the structure, not to cook proteins, but just to disrupt any ongoing electrical activity.

What you discover if you give-- so this is the delay to ECS. Rat steps down. When does he get his ECS?

If he gets the ECS immediately, put him back on the platform, basically 100% of the time, the rat will step down the next time, too. Why? Not because the ECS is fun or something like that, but because the ECS has wiped out the memory. The memory is not consolidated and the rat simply doesn't know that his little feet got fried. And he steps down again.

As the delay gets longer, you get an essentially exponential-looking curve with a time course, in this case, of seconds that represents the time of consolidation of that memory, at least to the point of supporting the behavior of not stepping down. But it's probably better to think-- it's not the case that all memories consolidate within seconds. There's evidence from human studies.

Now, why would you do human electroconvulsive shock studies? The reason is that electroconvulsive shock in humans is electroconvulsive therapy. And it sounds very odd. It sounds very odd and those of you who have seen the now quite old movie *One Flew over the Cuckoo's Nest* probably have a dim view of electroconvulsive therapy. But that's badly earned.

In the case of intractable depression, depression that's not being broken up successfully by medication or by psychotherapy-- severe depression of that carries with it a very serious risk of mortality from suicide. So this is not a casual disorder. It turns out that electroconvulsive therapy is often very effective in breaking up such depressions and is actually a useful part of the spectrum of treatments that are available in this case.

But like most medical treatments, it's not entirely cost free. It is the case that there are memory consequences of electroconvulsive therapy. It's part of what actually gives it a somewhat bad reputation.

And you can actually see those memory consequences. You can see traces of them going back years, as if-- so this is in studies, for instance, of memory for TV shows seen only once. It's the sort of thing that you might dimly remember under the best of circumstances. These rather limited kind of memories can be damaged by electroconvulsive therapy years later, as if this exponential has a very long tail.

And when you want to call it consolidated is perhaps a function of maybe something like how important the memory is. If the memory is, if I do this again, I'm going to be hurt right now, that memory can be made solid enough relatively quickly that you don't go and do it again. On the other hand, if the memory is was it Bridget loves Bernie or Benjy? All right, that kind of no-stakes kind of memory is apparently still vulnerable, at least still partially vulnerable, going back for a very long time.

The hippocampus seems to be behaving like something like scaffolding. It holds the bits of the memory in place until it's firm enough to stand on its own. And then, you can really call, it in a sense, a permanent memory.

Now, that doesn't solve the issue of getting stuff out of memory. And I will say a word or two about that momentarily. But first, you can wipe out your short-term memory by stretching for a second.

[SIDE CONVERSATION]

STUDENT: Could you teach HM to, say, play tennis or something like that?

JEREMY WOLFE: Presumably, but he would not remember the rules. He'd remember the skills.

STUDENT: That's really interesting.

JEREMY WOLFE: Oh yeah.

STUDENT: [INAUDIBLE] So I kind of got how they [INAUDIBLE].

JEREMY WOLFE: Oh, I didn't describe the details of the human experiment, which is beautifully elegant but I decided I was running out of time. Take a collection of TV shows that only ran for one season so you know what year they would have been experienced. And then, make up a test that says, tell me which of these are real TV shows and which are fake.

And so you get a curve. People are good for the most recent ones, and everybody gets worse as it goes back in time. If you had ECT here, so you were shocked here, your memory for recent stuff is depressed relative to normals. And it doesn't come back to the normal curve for about five years afterwards. So you can go back five years in time.

STUDENT: [INAUDIBLE]

JEREMY WOLFE: No, no, no, the memory was there. So I hit you on the head. I'm not going to just interfere with your memory for right now. I'm going to affect your memory going back into time.

And under the right circumstances, I can measure that effect back in time for years. If this is still unclear, we can talk about it later, but I'd better-- OK.

Just looking at the handout, let me just make sure that I explained the jargon here. So a loss of memory, as everybody knows, is an amnesia. A loss of memory for events prior to the point of injury would be a retrograde amnesia.

That's not what HM has. HM has an anterograde amnesia, a problem with his memory subsequent to the lesion, subsequent to the injury. If I-- we won't do this as a demo, but if I were to take a baseball bat and pop her on the head, well, apart from the fact that I probably wouldn't be back on Thursday, when she came to, she might not remember having actually been hit because the set of memories from immediately before that whomp would have been wiped out in a retrograde amnesia.

It is also possible that she wouldn't remember things for a period of time after she had regained consciousness, if I managed to knock her out completely. There might be a period of time after she regained consciousness before her memory would-- she might have been talking, and wandering around, and saying, I'm fine, I'm fine, but have no recollection of that, either. That would be an anterograde amnesia. An amnesia for, sorry, I didn't really mean to hit you on the head. You'll be fine.

So let me say a few words in the remaining time about retrieval from memory. Like any of these topics, there's a great deal to be said. But probably the most important thing to understand about retrieval from memory is that retrieval, particularly of any rich sort of memory, is not the recovery, the replay, of the tape or some true, guaranteed image of what you experienced, or what you learned, or something like that.

This sort of image here of a scaffolding holding together that memory is also an image of the notion that any memory is probably stored multiply and in different bits of brain, different aspects of it, in different bits of brain. And if you want to recall it, that what you're doing is reconstructing it. You can probably get some intuition about this if you ask yourself about famous stories about you in your childhood that are famous in your family, that you remember. Not just the stories they tell about you from when you got hit on the head with the bat and stuff like that, but stories that you remember as shaping incidents in your childhood, but that also get retold every time the family gets together.

So in the Jewish tradition there's a holiday called Passover. Tradition is that you drink four cups of wine. It's not really a brilliant tradition if you happen to be about four years old.

And I vividly remember my sister getting into the nice, sweet wine that you drink and conking out at the table. And I know whose house we were at. And I can see this.

This is a tale that has been retold, to my sisters no doubt delight, every year for the intervening 40 plus years. And it is absolutely unclear to me, given what I've subsequently learned about memory, whether what I am remembering is the original event at this point or a reconstruction based on all the family stories. I think I can still remember the original thing, but there's no real way of knowing. What you are remembering is a reconstruction.

Before I talk more about that, let me load you up with some more words. These aren't going to be nonsense unless I can't find them. There we go, there's some nice words.

Which words do I want to use? Oh, these are good words. I'm going to read you a list of words and then we'll do a recognition test. You don't have to write anything down this time.

I'm going to read you a list of words and you're just going to tell me, yes, I've heard this word. The word was on the list or the word was not on the list. Ready? So try to remember all of these.

Fury, rage, enrage, carpet, club, emotion, ire, mountain, fight, fear, mean, mad, place, hate, and road. OK, you got all those nicely stored in memory? OK, so let's see. Where's my test list here? It looks like my test list. OK, did you hear the word fight?

STUDENTS IN UNISON: Yes.

JEREMY WOLFE: Science?

STUDENTS IN UNISON: No.

JEREMY WOLFE: Rage.

STUDENTS IN UNISON: Yes.

JEREMY WOLFE: Fear.

STUDENTS IN UNISON: Yes.

JEREMY WOLFE: Emotion.

STUDENTS IN UNISON: Yes.

JEREMY WOLFE: Wrath.

STUDENTS IN UNISON: No.

JEREMY WOLFE: Club.

STUDENTS IN UNISON: Yes.

JEREMY WOLFE: Anger.

STUDENT: Yes.

STUDENT: No.

STUDENT: No.

JEREMY WOLFE: Well, road.

STUDENTS IN UNISON: Yes.

JEREMY WOLFE: Ire.

STUDENTS IN UNISON: Yes.

JEREMY WOLFE: Hatred.

STUDENTS IN UNISON: No.

JEREMY WOLFE: OK, so how about enrage?

STUDENTS IN UNISON: Yes.

JEREMY WOLFE: OK, so there were a couple in there where you could hear people hesitating or you could hear a divergence of opinion. And those are the critical elements there. The most critical element in this case is the word anger, which a fair number of people asserted was on the list.

It was not, in fact, on the list. And even those people who were asserting that it was not on the list, there was a measurable [GROAN]. What's going on there?

Well, this is an effect known as the Deese-Roediger-McDermott effect, or on the handout, the Deese-Roediger-McDermott demo. Roediger and McDermott discovered it. And then, like all good phenomena in psychology, discovered that Deese had published it 10 years prior. And so it now goes by all of their names.

In any case, what they did-- they were deliberately looking for this. What they did was they-- where'd my semantic network go? Oh, here's my semantic network.

What they did was they took a word and they asked people to give the associates to it. So let's take the word anger. And so what words come to mind when you think of the word anger? Well, fight, rage, enrage, emotion, ire, this collection of words.

You then use a training list of words that has the associates but not the target word. And what you're in effect doing is saying, animal, meow, mouse, dog, fur. And each time you do that, the spreading activation gets all sorts of stuff. All of these are pointing towards cat.

And then later, I say, well, was cat on the list? And you go and look at that node. And lo and behold, it's glowing softly in your mind, and you say, yeah. Yeah, that was there.

And if you take one of the-- sort of a low-- it didn't particularly work in this case. Ire sometimes works, for instance. You take a word like ire, people are no more sure that ire was on the list than they are that anger was on the list.

People are thoroughly, in these experiments, confused about whether or not the target word was on the list because you don't have access to the truth about your memory. You have access to what you can manage to reconstruct about it. Now, that's fine. That sounds benign enough as a lab demo.

But let's suppose that you're not just sitting in some nice intro psych Class, But you're sitting in the witness box. And somebody is saying, where were you on the night of May 5th? And expecting you to tell the truth.

Well, of course you're going to try to tell the truth, but the truth in this case is going to get shaped by the nature of your particular semantic network, among other things. Let's talk about that part a little bit first. Famous experiment done first in the '40s, replicated many times since in many different ways.

But let me describe the original version. You-- and in this case, you are a white male-- are shown a depiction of an altercation on a subway. Two guys get into a fight. In the '40s this was done as a cartoon, hand-drawn cartoon. It's been replicated with snazzy video. It doesn't matter.

Two guys get into a fight or get into a shouting kind of argument. At some point, one guy pulls a knife. And then sometime subsequent, we'll ask you about this.

The variable of interest in this experiment is-- well, the question you're going to be asked is who pulled the knife. And the variable of interest is the race of the participants. So you've got a nice little 2 by 2 here.

They can both be Caucasian. They can both be African-American. Or it can be crossed.

The interesting and disturbing finding is the number of times when the knife is pulled by the white guy, it ends up in memory in the hand of the Black guy. Now, as I say, this has been replicated all sorts of different ways. This is not a white guys are all bigots experiment.

This is we all have biases. And the interesting and disturbing thing is that they can actually interfere with what we think of as our nice, clear memory. These people weren't sitting there saying, oh, I'm being paid-- 1940-- $0.10 an hour to be in this psych experiment. Why don't I see if I can stick it to an imaginary Black guy?

No, this is just they were presumably being as honest as they could be. But the structure of their memory, which included things like, who do I think might pull a knife, influenced what they recalled, what they pulled out of memory. And we can impose this structure from the outside, as well.

So suppose we do the following experiment. This is an experiment also done many different ways at this point. But Elizabeth Loftus is the name associated with this line of work.

You're going to see another film. This time you're going to watch a car crash. The red car is going to get into an accident with the blue car.

And you're going to be asked-- you see this, and you're going to be asked about it. We'll show all of you guys the same film. And now, we'll ask the question three different ways.

At what speed did the car did the red car bump into the blue car? At what speed did the red car hit the blue car? At what speed did the red car smash into the blue car?

That's the only manipulation here. You all saw the same video. The difference-- so what do you figure the difference is? What's the result look like?

So who gives the high speed response, the higher speed responses? The smash guys. Lawyers know this.

This is why courts attempt to avoid leading questions. But it's very hard to avoid that. That's a pretty subtle lead.

But it's not a subtle kind of effect. It's about, as I recall, a 15 mile an hour kind of effect, which is the thing that has an impact-- you should pardon the expression-- on what a jury is going to think about whether or not you, the driver of that red car, was at fault or not. If you're screaming through the intersection at 45 miles an hour, that's worse than if you're screaming through the intersection at 30 miles an hour or something like that.

So your memory can be influenced by the way that you are asked about that memory. This shows up not just on the stand. Well, actually, a version of this shows up beautifully in the experience of most people at some point who come to MIT.

Most people who get to MIT are smart people and they would do well to remember that. Because most people sometime, like within the first few weeks of being at MIT, receive a data point of some that suggests that this might not be true, like the first calculus test or something like that. You've never gotten a score below 99.8 on a math test in your life. You get back the first test in 180 whatever it was.

And 10-- cool, 10. Oh, man, it's not 10 out of 10, is it? Yeah, this is really bad.

And so that's a depressing data point. But a surprising number of people come to the conclusion from that that they not only are stupid but always were stupid, and probably unlovely, and generally bad, horrible individuals. And you should remember this is one of these context effects.

So that may be a little overstating it, but look-- you know, and this has actually been tested with clinical populations of depressive people who go from being depressed to being undepressed. Suppose you ask somebody who's depressed, clinically or otherwise. You wake up in the morning.

It's pouring cold rain out there. The problem set didn't get done because you fell asleep and there are now drool marks on it. And not only is there a drool mark, there's a little heart-shaped note next to it saying, "I'm leaving you for your roommate."

[GROANS]

So if I ask you at that point, how was your childhood-- not how do you feel, but what kind of childhood did you have, you had a-- maybe not a miserable childhood, but exaggerate. You grew up in a closet, fed on-- later on, if you're feeling better, I ask you about the same childhood. The childhood has magically improved. All right, we'll see you on Thursday, if I remember correctly.