INTRODUCTION: The following content is presented by MIT OpenCourseWare under a Creative Commons license. Additional information about our license and MIT OpenCourseWare in general, is available at OCW.MIT.edu.

PROFESSOR: Good afternoon here is the plan for today's lecture. There is out there somewhere, unless you're philosopher, a world, it's huge. A lot of stuff out there in the world. Their is inside use somewhere a long term memory, which is also huge. And 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 looks familiar to you by now. We can described 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 providing and 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 it that OK 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. And finally, it is of absolutely no use to you 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 is a world of 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 appreciates 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 like write down the answer on a test, or tell me who's running for president this year. Nobody.

AUDIENCE: Bush.

PROFESSOR: Yeah that's one of them. Bush yeah.

AUDIENCE: Kerry

PROFESSOR: Kerry OK. And a few other characters too. Right. OK. 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?

AUDIENCE: [LAUGHTER]

PROFESSOR: And on demand, you could retrieve that back out into what we will call and your working memory. If we were just thinking about this and 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 bi-directional 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 external -- oh no we can't see it there. Well, you know, 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 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 junk of this bottleneck.

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

But in trying to ask what's the limit of this bottleneck from the outside world into a more long term varieties of memory, one of the typical measures is what's called digit span. If you end up in the hospital because somebody bopped you over the head, one of the things you neurologist would 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 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. OK. So if I say, red, blue, purple, repeat. You say?

AUDIENCE: Red, blue, purple.

PROFESSOR: It never fails.

AUDIENCE: [LAUGHTER]

PROFESSOR: Right, [? Mara, ?] last year too? Always you get the repeat back too. And I'm not sure what it is about this that demands, that at least from some people, the repeats of repeat. But in any case good, good. You're not deeply brain damaged yet. Ready? We'll do some more of these. Green, pink, yellow, white yellow, repeat.

AUDIENCE: Green, pink, yellow, white, yellow.

PROFESSOR: OK. That was fine. All right, here we go. Oh, I forgot to mention. It turns out to be really easy to repeat back in essentially infinitely long list of these, if you write them down while I'm saying them.

AUDIENCE: [LAUGHTER]

PROFESSOR: There are a variety of demos that I'm going to do today that are of that form, and it's really boring. So, you know, OK we know you can do that, so don't do that, because that's boring. Be a nice honest person. OK ready? Where was I. I just did 5, right? Let's try 7. Ready? Oh and I told you when I'm going to stop. Oh what the heck. Pink, red, yellow, green, purple, pink, orange, repeat.

AUDIENCE: Pink, red, yellow, green, purple, pink, orange.

PROFESSOR: It's getting a little weak there at the end. But you will see, if you look at the handout, you will see that the answer to this question is given in the title of George Miller's classic paper on the subject. And is answer is 7 plus or minus 2, and some of you were on the minus side. Oh I should also say -- I should've said this at the beginning of the course -- people in intro psych classes, particularly by the time you get around 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.

AUDIENCE: [LAUGHTER]

PROFESSOR: You're probably fine. I mean if you, if it was the red, yellow, purple, you're saying what was the one after red. You know I admit there might be an issue there. 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.

AUDIENCE: Blue, green, purple, red, green, yellow

[LAUGHTER]

PROFESSOR: Yeah, yeah, yeah. You can hear. That was 10. And you could here it in the group, and you could probably feel it in yourself, that, you know, up to 6, 7, 8, you're yeah, we're good here. Oh man if I start taking in this next one, that I'm losing them here. Something bad is happening.

So there's this capacity limit on what you can get into -- well let's used this jargon for the time being -- into some sort of short term memory, and then spit out immediately to me. Now 7 plus or minus to what? Let's try another one. You ready? Same game. Red, red, red, red, blue, blue, blue, blue, green, green, green, green, repeat.

AUDIENCE: Red --

[LAUGHTER]

PROFESSOR: You know you can do this. That's 12 items. How come I all of a sudden manage to boost your memories so effectively? Yeah classification, or what Miller originally, and the field had 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 managed 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 somebody a 10 digit, you know, give 300 MIT undergraduates a 10 digit string, you know, a bunch of them crud out on you, but then there's a handful of people who just raddle it right off. And you say oh my goodness, these people have this amazing memory. And you ask them how'd you do it? And they say oh well that was obvious. It was the natural log of 16.

AUDIENCE: [LAUGHTER]

PROFESSOR: MIT students tend to junk numbers rather more readily than the rest of the population. They don't tend to junk color names particularly adeptly I've found, and so you get the right 7 plus or minus 2 answer. Now this sort of chunking, you should think of this more like a rate right. 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, you know, the first lecture shot your brain and that was it for the term. That's obviously not the case.

What 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 for instance, what I'm saying. So if I ask you to repeat this sentence, memory is a fascinating topic that would be on the final exam, you would say?

AUDIENCE: Memory is a fascinating topic that would be on the final exam.

PROFESSOR: But if I lettuce, quadrangle forms only only column bug de la forms aftunic reply quintilian. Kelly You say?

AUDIENCE: [LAUGHTER]

PROFESSOR: Say it roughly the same number of syllables. But ouughh. You know, 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've 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. You know, it's not true of course. And those of you who are non-native English speakers, and have now arrived at MIT, are thinking no, you know, back in, you know, Kazakhstan they all spoke really slowly or whatever. But these English speakers are wackos.

The problem is you are in your native tongue, or in any town but 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. I know. I damn it that was the week I was asleep in high school or something.

AUDIENCE: [LAUGHER]

PROFESSOR: No that 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. 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? 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. OK.

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've rated a housing site. But, you know, 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?

AUDIENCE: Yes.

PROFESSOR: 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 havent.

AUDIENCE: No. No. Yes. Yes. No. Yes Yes, No. No. Yes. No.

PROFESSOR: 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, reverse. 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% 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're 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. 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 test. I was asking you what, you know, pull out of your memory what was there. Here I was saying was this there. And you know this. But, from your test taking experience you know that, OK if I give you a choice between fill in the blank and multiple choice, which do you want to do?

AUDIENCE: Multiple choice.

PROFESSOR: Multiple choice, because that's a recognition test. Fill in the blank is 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, you know, 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 case of 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 over 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. Go to sleep. Why don't you want to go to sleep.

AUDIENCE: [LAUGHTER]

PROFESSOR: It's like your children. Mara, what do you do? I don't want to know OK. You mean I'm supposed to hit it?

AUDIENCE: [LAUGHTER]

PROFESSOR: Oh 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. He's 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 your 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 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 like 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 Mali 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 --

AUDIENCE: [LAUGHTER]

PROFESSOR: -- they looked at my list and they just crossed these things out. Anyway. So 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. Moche. Yill. Saz. Tirt. Varl. Bince. Jucks. Gouf. Zauce. Rab. OK. Write down everything you can remember.

AUDIENCE: [UNINTELLIGIBLE]

PROFESSOR: Now don't copy off your neighbor. On your hand out on the second page I believe, you will find the axes for what's called the serial position curves. 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 in sort of number of hands, as a function of position in the list. Where in the list that I read did the word show up. And I will tell you the answer at 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 affect, 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 fib? Alright. Data point turns out to be about there I think. How many people got dut? Few, fewer. How many people got -- well you see your falling a little below the line here. I am worried about you guys. How many got tirt? Ouhh how interesting. How many people got varl? Ouhh that's pretty pathetic. How many people got zauce. Oh coming back. Oh that's pretty good actually. That's above the line here. And how

many people got rab? That's sort of thereish.

OK. What's with tirt?

[LAUGHTER]

You guys did bet it. Well you know, 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 means --

AUDIENCE: [UNINTELLIGIBLE]

PROFESSOR: What's a tirt? What?

AUDIENCE: [UNINTELLIGIBLE]

PROFESSOR: Tert. Oh oh tert like as in tertiary or something?

AUDIENCE: Yes.

PROFESSOR: Oh, see I'm spelling it t-i-r-t, and it never occurred to me. All right, all right. We're going to have to do something about tirt. Anyway. OK. 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 sort of 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 sort of by analogy to immigration. You know here's America, and here are all of our various forebearances and stuff trying to get into the country. And there's this choke point at immigration. You know, if you got your Uncle Charlie already in the country, he can pull you 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. Sort of 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, Yeah, OK.

AUDIENCE: I was writing down this fip and dut.

PROFESSOR: Oh OK. Apart from the fact that you were -- no, no, you presumably were being a good person and not writing them down when I read them first though right? OK. So what were you doing while you were waiting to dump them back out? Yeah.

AUDIENCE: I mean I was going over them in my head. Rehearsing them.

PROFESSOR: 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 auditorium material sometimes called a phonological loop. Does it say phonological loop on the handout anywhere? No. Well. You know, 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 sort of keep it alive.

Anybody got a different experience for zauce and rab at the end of there? I'm sorry I heard something that sounded promising. Yeah?

AUDIENCE: They were just there. PROFESSOR: They

were just there is actually the right intuition. 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. You know all in beautiful unison. Right got that? That would make a lot of noise, so when I do this, keep an eye on me. When I do this, write everything down. OK? OK. They kind of got the instructions. All right, where's the rest of my nonsense here?

AUDIENCE: [UNINTELLIGIBLE]

PROFESSOR: OK. OK ready? Sip. Forth. Lig. Vop. Hearn. Mope. Jick. Tinned. Mez. Wamp. Flob. Gone. OK. Count 431.

AUDIENCE: 431. [UNINTELLIGIBLE]

PROFESSOR: OK. Write them all down. The giggles would probably do just fine.

AUDIENCE: [LAUGHTER]

PROFESSOR: OK. The claim here is that what that should do, did you feel yourself rehearsing the first ones again?

AUDIENCE: Yeah.

PROFESSOR: The claim is that the primacy affect, which is due to some sort of rehearsal into long term memory. That that primacy effect should still be there. But the recency effect due to some sort of it's just thereness, should have been disrupted by the counting and the chordals 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 sip? And look. That's almost exactly the same number. How many people got forth? Ooh. A lot of forths. How many people got jick? Ouhh deeply pathetic.

AUDIENCE: [LAUGHTER]

PROFESSOR: How many people got mez? Few more. How many people got flob? Ouhh that's pretty pathetic. How many people got gone? No gone? Gone was pretty well gone. These were supposed to be circles, So they'd be a different data set. So anyway, nice primacy effect. Recency effect gone.

AUDIENCE: [UNINTELLIGIBLE]

PROFESSOR: Yeah. There are other words in there. This is the phenomenon of MIT students really wanting to do well on these sort of things right?

AUDIENCE: [LAUGHTER]

PROFESSOR: Man, he didn't ask me. It's like, you know, you have the same experience for keeps on the exam. You sit there and you study chapter 3 until you're blue. And oh man they only asked about chapter 2. So yeah, yeah, yeah. I'm glad you got whatever it was wamp or --

AUDIENCE: [LAUGHTER]

PROFESSOR: 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 know you. He said fip, dut.

AUDIENCE: [LAUGHTER]

PROFESSOR: I said you are a very strange person. But I now know where I know you from, or don't you from or something like that. You may have now warped your brain on ongoing basis.

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 a standard model for getting from the outside world -- whoops my bottleneck disappeared again -- into long term memory. Which was stuff comes into 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?

AUDIENCE: Bagel.

PROFESSOR: Bagels. Anything else?

AUDIENCE: Bagel and cream cheese.

PROFESSOR: Bagels and cream cheese. Oh 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, Bagel and cream cheese. AUDIENCE: [LAUGHTER]

PROFESSOR: Right?

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 that 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 know as explicit long term memories as 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, you know, 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 asked 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 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 for the episodes of your life, and semantic memories, your body of knowledge and facts and things of that sort. There not like 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 up memory might be organized. And one of the popular ways of thinking about this is sort of a giant network of associations called a semantic network often. If I say cat, the first word that comes to your mind is?

AUDIENCE: Meow.

PROFESSOR: Meow.

AUDIENCE: Dog.

PROFESSOR: Dog.

AUDIENCE: Mouse.

PROFESSOR: OK. So there's a note in your long term memory that is cat in some sense in this story. It's close neighbors are things like dog and meow and -- what did we have -- mouse. And now it's important to realize this is not some sort of Linnaean taxonomy of, you know, species and geneus and all that sort of good stuff. You know, even though it'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, mow.

AUDIENCE: [LAUGHTER]

PROFESSOR: Actually that goes back here. I actually have a cat whose name is Chairman Mao by roughly this association.

AUDIENCE: [LAUGHTER]

PROFESSOR: You know, then, you know, China.

AUDIENCE: [LAUGHTER]

PROFESSOR: Plates.

AUDIENCE: [LAUGHTER]

PROFESSOR: So, you know, 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 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. Evidence from priming. Look at that. On my computer screen here, I'll put up a string of letters, and you tell me whether or not it's a word. Right? So I put up. And you say?

AUDIENCE: [UNINTELLIGIBLE]

PROFESSOR: No. Even though it's now in your long term memory right. But if I put up. 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 say 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 activities spreads -- it's called spreading activation -- away from that node to the neighboring nodes. As a result, when you get the word cat, cat's already a little activated. The bell's been rung a little bit already. And your quicker to confirm that cat is a word than truck. Which is you know, down here in, you know, in the next county

somewhere.

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 ever, and not clear what the point is. But you could you can use this sort of a technique to show that some things live closer to each other than other things. Oh and by the way, when you go and reach in to retrieve -- this is sort of 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. And you know this, or you have experienced this if you are not an only child perhaps, even if you are. How many do 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.

AUDIENCE: [LAUGHTER]

PROFESSOR: Who didn't know that they had siblings until that happened. Or worse, you know, you get called by the name of the dog or something like that. This is not because whatever you may believe, you know, that your parents are unusually dim people. It's because all those terms presumably live very close to each in this semantic memory. And when particularly under some pressure you reach into grab kid number -- all my kids have the same name. Their name is Ben Phillip 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.

AUDIENCE: [LAUGHTER]

PROFESSOR: 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. 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 sort of 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 struc the hippocampus, and to a lesser

extent the amygdala. I'll just mention. 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 treatment 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 50's, in the late 50's or early 60's, 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 is a bilateral lesion. And this did have the effect of largely controlling his seizures. The seizure problem went away. However, he's a one of a kind patient, because the side effects, the unintentional consequences of this lesion, was 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.

Take his interaction with a once upon a time TA in this course name John Gabrieli now a Professor at Stanford soon to be a Professor back here at MIT. But who did, I think he did his doctral work studying HM. But he certainly studied HM as a graduate student. He'd come into the lab to meet HM. Say hi, you know, have we've met before. HM said no. I don't think so. He's perfectly intact short term memory, which is fine for conducting a conversation. Also decently intact memories, this 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, you know, it's 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. 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. Have the same conversation. You can do this over and over again.

In a more control 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 try to get from the phone book to the phone with a phone number. And somebody says, didd you do x, y and z? Ahh. There goes the phone number.

AUDIENCE: [LAUGHTER]

PROFESSOR: 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 they're a 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?

AUDIENCE: But, does he know he has this disorder?

PROFESSOR: He knows he's got a memory problem. And this actually is a good point. He ties into his other problem, which is in some sense an off setting benefits. The Limbic System, certainly the amygdala, very important in the mediation of emotion. His experience of emotion is greatly flattened. Flattened affect in the jargon of the trade. So. 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 women as the case may be. I suppose it doesn't particularly matter. It's going to be a disaster under either circumstance right? You're going to think I'm in a sci-fi movie, and I don't have the part. You know the good guy with the laser gun part that I really wanted. You know 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 has been very severely blunted. So much so that he needs to be asked whether or not he's sick. You know do you have a stomach ache? No that you mentioned it, yes I do. It's not that he couldn't feel the pain, it's just without the emotional overlay so, you know. I see blue. I have pain what's the big deal?

AUDIENCE: [LAUGHTER]

PROFESSOR: Yes?

AUDIENCE: [UNINTELLIGIBLE] motor memory

PROFESSOR: Yes. Nice point. So for example, I don't know how many of you have hung out at -- you know get rid of the brain here. Who needs the brain? So any you ever done a mirror maze? Some people may have actually built one of these once upon a time. So imagine you have like a star. Oops. Imagine you could actually draw a star.

AUDIENCE: [LAUGHTER]

PROFESSOR: All right. Well. So you make this out of like 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 right? Anybody can do this. What people can't do regular readily is do this looking at the image in a mirror. Right? Because then think they got [UNINTELLIGIBLE].

AUDIENCE: [LAUGHTER]

PROFESSOR: But you learn to do this. And HM has been trained in one set of experiments to do this. In multiple sessions he got better. If you asked him have you ever done this before on the nth time that he did it. The answer is no. I don't think so. And he goes [UNINTELLIGIBLE]. Hey, you know, that was really good. Most people asked how did you do that? Yeah. I've always been good at that kind of thing. Or, you know, just lucky. No episodic memory, no explicit memory of ever having done it. But yes a motor memory that allows them to do it.

Similarly, your sense of familiarity runs on different pathways. So if I bumped into you wandering around the halls, I may well have the sense that I've seen you somewhere before. And I might guess that that's, you know, 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 looks familiar. Do you know my name? Know you're name? Forgotten 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. It's yeah maybe from high school, or something like that.

We're about the same age as the 70 year old HM to the 20 something year old Gabrieli. So those sort of things he does learn. It's that 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 were a rat on a small pedestal not too far above the floor, what are you going to do?

AUDIENCE: Jump off.

PROFESSOR: Jump off, right. I'm a rat. Man I got to go 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. So when I jump off my little toes get an unpleasant, not damaging, but an unpleasant shock. So now the guy puts me back on here. What do I do?

AUDIENCE: Jump off.

PROFESSOR: 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 know you can kill people or animals that way. Obviously that's not the point here. This is to disrupt the on going electrical activity of the brain. Not to scramble the structure. Not to cook proteins. But just to disrupt any ongoing electrical activity.

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 a 100% of the time the rat will step down the next time to. Why? Not because the ECS is like 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.

AUDIENCE: [LAUGHTER]

PROFESSOR: It sounds very odd. It sounds very odd. And those of you who've seen the now quite old movie to kill, *Not To Kill A Mockingbird*, *One Flew Over The Cuckoo's Nest*. One of those bird movies. *One Flew Over The Cuckoo's Nest* probably have a dim view of electroconvulsive therapy. But that's badly earned.

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

Like most medical treatment, 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. This is in studies for instance of memory for TV shows seen only once. You know sort of thing that you might dimly remember under the best of circumstances.

These rather limited kind of memories can be damaged by a 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 may be something like how important 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 *Bridgette Loves Bernie*, or Benji. All right that kind of good. You know 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 the sense of sort of 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 WITH STUDENT]

OK. 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.

We won't do this is 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 a Thursday

AUDIENCE: [LAUGHTER]

PROFESSOR: When she came to, she might not remember having actually been hit, because the set of memories from immediately before that whop 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 say, I'm fine, I'm fine. But have no recollection of that either. That would be an anterograde amnesia. 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 if any rich sort of memory is not the recovery, the replay, of the tape for 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 a notion that any memory is probably stored multipley and in different bits of brain. Different aspects of 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 asked 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 hit with the bat and stuff like that. But stories that you remember is sort of shaping incidents in your childhood. But that also get retold every time the family gets together, right.

So in the Jewish tradition is a holiday called Passover. The 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 conching out at the table. And I know whose house we were at. This is a tale that has been retold to my sister's no doubt the light every year for the intervening 40 plus years. And it is absolutely unclear to me given what I subsequently learned about memory, whether what I am remembering is the original event at this point or 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. OK. I'm going to read you a list of words, and then we'll do a recognition test. 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 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? Looks like my test list. OK. Did you hear the word fight?

AUDIENCE: Yes.

PROFESSOR: Science?

AUDIENCE: No.

PROFESSOR: Rage?

AUDIENCE: Yes.

PROFESSOR: Fear?

AUDIENCE: Yes.

PROFESSOR: Emotion?

AUDIENCE: Yes.

PROFESSOR: Wrath?

AUDIENCE: No.

PROFESSOR: Club?

AUDIENCE: Yes.

PROFESSOR: Anger?

AUDIENCE: [UNINTELLIGIBLE]

PROFESSOR: Well. Road?

AUDIENCE: Yes.

PROFESSOR: Ire?

AUDIENCE: Yes.

PROFESSOR: Hatred?

AUDIENCE: No.

PROFESSOR: OK. How about enrage?

AUDIENCE: Yes.

PROFESSOR: OK. so there were a couple in there where you could hear people hesitating, or you could hear a diversion 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 it 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 what's going on there. This is in 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 did 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. 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, but all of these are pointing towards cat. And then later I say well was cats on the list? And you go and look at that note 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 other, 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 memore. You have the 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 40's 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 40's it was done as a hand drawing cartoon. It's been replicated with, you know, snazzy video. It doesn't matter. Two guys get into a fight, or get into a sort of shouting kind of argument. At some point one guy pulls a knife, and then sometimes subsequent we'll ask you about this, OK. 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 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, you know, white guys are all biggits 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, you know, our nice clear memory. These people weren't sitting there saying, you know, oh yeah I'm being paid, you know, 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. 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. 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. All right. OK. You see this. 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 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 smashed into the blue car? That's the only manipulation here. You all saw the same video. So what do you figure the difference is? What's the result look like? So who gives the higher speed responses?

AUDIENCE: The smashed guy.

PROFESSOR: The smashed guys Lawyers know this. This is why courts attempts to avoid leading question. But it's very hard to avoid. But, you know, that's a pretty subtle kind of lead. But it's not a subtle kind of effect. It's about as I recall a 15 MPH kind of effect, which is a sort of 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 were screaming through the intersection at 45 MPH that's worse than if you're screaming through the intersection at 30 MPH 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 sort that suggests that this might not be true. Right? 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 18 or whatever it was, and you know, 10 cool. Oh man it's not 10 out of 10 is it?

AUDIENCE: [LAUGHTER]

PROFESSOR: You know 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, you know, bad, horrible individuals. And you should remember this is one of these sort of context effects. That may be a little over stating 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 problems that did you get done because you fell asleep in there, and there are now drool marks on it. And not only is there drook mark, there's a little heart shaped note next to it saying, I'm leaving you for your roommate.

AUDIENCE: [LAUGHTER]

PROFESSOR: Right? So if I asked you at that point how was your childhood? Not how do you feel, but what kind of childhood did you have? You had not maybe a miserable childhood. But to exaggerate, you know, you grew up in a closet. 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.