[SQUEAKING]

[RUSTLING]

[CLICKING]

JUSTIN REICH: Why don't we get started? Nice to see you all. So the thing I asked you to do in your learning journals for today was to do a little reflection on your personal edtech history. So what are tools, experiences, platforms, systems that either had some very positive effect on your learning trajectory, you, like, discovered a love of something because you did that, or were terrible and disastrous, they ruined your learning experience, they made you feel bad, they inhibited your learning in some way, or they were just interesting, or engaging, or useless, or some other way distinctive.

Go ahead and-- now, if you didn't do that assignment, you can probably still make up an answer to that question. So you'll be OK for today, but turn to the person or two next to you, get in small groups, and just share a little bit with each other about this little snippet of your personal edtech history. We'll do that for five minutes or so and then share out. Does that make sense? Ready, go.

Did anybody happen to be in a group where multiple people in the group mentioned the same thing? Were they all different? What were the most similar ones? Anybody?

AUDIENCE: There was a computer literacy class, like, similarity.

JUSTIN REICH: OK, good. So there's a class, which is called computer literacy, digital literacy. And it's kind of a pretty out, like, here's how to use PowerPoint, here's how to type, here's how to do pretty simple things with computers. Good. Yeah, that's a feature of a lot of different kinds of school systems. And was it successful in your environments or was it a drag?

AUDIENCE: A drag.

[LAUGHTER]

JUSTIN REICH: Yeah, good. Yeah, yeah. Yeah, there can be a bias that, oh, man, we got to put technology in the schools because the kids are going to be really fired up about the technology in the school. No, actually if you do boring things with technology, they're just as bored, sometimes even more so. Anybody have things that were really positive experiences for them, that somebody, you or somebody in your group talked about? Things you're like, yeah, that was great. That really made school better.

AUDIENCE: Well, I mean, I talked about the smart board, and just like the implementation of it in, like, every classroom, I feel like, especially as our generation was growing up. And like, I don't know, like, I made a little, like, side story about how that's the first time I felt like the student became the teacher because the old ladies at our schools didn't know how to recalibrate it or turn it on. So it was always the students coming up and showing the teacher how to use it.

JUSTIN REICH: Yeah. Yeah. Yeah. Good, a sort of inversion of expertise. How many of you were in schools where smart boards just kind of appeared in the middle of your educational trajectory? Yeah, you all are in the sweet spot. How many of you would say that they were-- anybody at schools where you felt like they were very effective, where they substantially improved learning? You thought yours? Good.

How many of you were at schools where this is like, no one is using these things, they could have just bought chalkboards or whiteboards or things like that? Yeah, so overwhelming majority of that. Yeah, smart boards are like a totally fascinating case study. A vast amount of money was spent on these tools that seemed to have essentially no positive impact at all on learning.

In lots of places, they really weren't used as interactive whiteboards. They were just like, oh, like, a whiteboard. Khalil and I were joking that they actually didn't work very well in his school, in particular, because they require slightly different lighting. Like, you need to dim the lights a little bit to be able to see them. But the classrooms that were built were not built for dimmable lights. They were built for lights that would be shined on whiteboards or chalkboards.

And so the rest of the-- there's a whole bunch of the rest of the infrastructure that didn't make smart boards work. Like, maybe your classroom didn't work, maybe there was no training for your teacher, maybe there were no curriculum materials that were available to integrate with the smart board.

And yet, how is it the case that we could be in a room full of people all across the country, all across the world, you all raised your hand and you're like, we all bought smart boards and they all didn't work? Why? Like, why didn't the first people who bought smart boards be like, team, this is not a good idea? Like, spend your money on something else.

Every dollar we spend on education technology has an opportunity cost. Every time you spend $1 in a school, you're trying to make people's learning better. And there are very few dollars that we spend that make learning worse. Does anybody feel like the smart board actively, like, made you dumber or learn slower or something like that? Probably not.

Like, probably you're smart-- like, it's probably the case that no one, none of you were harmed by smart boards. I think that's pretty unlikely. But every dollar your school spent on smart boards could have been spent on some other technology, on some other curriculum material, on art supplies, on professional development for teachers, on a zillion other things which could have improved your learning experience and not just been, like, a very expensive version of a whiteboard or something like that.

So there's a lot in that discussion of smart boards that I think is helpful for us to think about. How not just of the efficacy of any particular technology, the efficacy of any particular technology, how a particular technology is situated in different kinds of places, with at least one example from Dana who's like, no, these things are great. We bought them. Our teachers got trained with them. Students were able to help use them.

Like, it worked in that one context, even if almost universally, it didn't work in these other contexts. So those are all things we're going to have to unpack and figure out. And then we're also going to answer this question like, if there's this thing which overwhelmingly doesn't work, why did schools spend so much money on it? What were the incentives? What was the logic? Who was spending this money? All these kinds of things are questions we can keep asking.

Anybody have really disastrous stories? That's kind of like a banal weird, that thing didn't work. But anybody have, man, this was terrible kind of experience with technology?

AUDIENCE: At our smart boards, in one case, students were able to install, like-- so you can customize the brushes. And the students just chose kind of an awful picture of a teacher to put as a brush. And they used as a kind of a prank on them, which is technology allows you to do more creative pranks.

JUSTIN REICH: Good. Yeah. Yeah, so expanding our way of humiliating faculty, that could be one good way.

AUDIENCE: In seventh grade, my school decided to give everyone a Chromebook. But in seventh grade, I feel like everyone is not really mature. So in class, everyone would find a different game. It was like the snake I/O game, or the paint I/O. And the entire class would be on the game, like, while the teacher was teaching. So then they would have to block different websites, but a new website would appear every week, so they just had to keep blocking more and more.

JUSTIN REICH: Yeah. Yeah. Yeah. Yeah, blocking websites is a good way to teach young people how to use VPNs. If you were like, I want to teach young people how to use VPNs, then blocking websites is pretty good.

But I wrote an op-Ed once years and years ago. Probably like just after I was teaching, probably 2006, or '07, or something like that. Where I describe them as knee-high fences around the internet that adults trip over and young people just jump over.

[LAUGHTER]

Good. Yeah. So here's this tool, which actually now is like Chromebooks are pretty ubiquitous in schools. There are plenty of places where there's pushback against these kinds of technology for people telling just the story that you described, our devices are canonically designed to capture our attention as much as possible, and to particularly capture our attention in ways that have us look at advertisements for people who are designing phones, and computers, and websites, and things like that, is like a central focus.

How do we get you looking at this thing as much as possible? And in particular, how do we then get advertisements to slide in front of these things so that people will pay you to look at it? That is not a set of characteristics that are good for a learning environment.

But on the other hand, we could-- I'm sure some of you may have mentioned them, but we can go to schools where Chromebooks are implemented and the kids are not playing snake I/O all day because they're doing, like, interesting learning activities that are well supervised that would not be possible without the Chromebook, without the apps, without the access to the internet and things like that.

So we could also think about what happened in your school that had this kind of purchase, and then there was clearly not the rest of the infrastructure in place to have that purchase be meaningful, whereas the exact same purpose in other places was supported by professional learning, and curriculum, and like, all the other things that might have potentially made that a valuable learning experience.

AUDIENCE: So I was talking to Herman, and I think we had a bit of a different experience from most of the people in the class, because we didn't do our basic learning in the US. So [INAUDIBLE]. I'm from Brazil. I think we use technology to a much lesser extent than here.

So I remember there was, like, one significant experience for me in terms of using technology growing up, which is going to a media lab to do English classes. And I think it comes to your point of having a purpose to having technology. And that was very clear. You can't really do listening and speaking exercises unless you actually have something to listen to and something to repeat. And like, each person would have their booth where they could practice by themselves. So I think the fact that it was very purposeful, and that it changed from day to day, because I think also there's an aspect of--

JUSTIN REICH: We're going to the computer lab. We're doing something different.

AUDIENCE: Exactly. You change floors. It was a different scene changes. Like, lighting. Like, the fact that it's different made it much more interesting. And when technology gets to this very common place, you take it for granted. You don't really value all that it can bring. So yeah, we discussed it.

JUSTIN REICH: Yeah, I mean, those kind of language classes are a great example of this. I know you probably all have taken language classes where the teacher goes around and you like each say hello to the teacher in the foreign language. And it takes like 10 minutes to let everyone say hello once. You all went to your individual booths and you could imagine designing technology where every person simultaneously says hello, and some algorithm evaluates whether or not you say hello correctly, and then it says something back to you and other kinds of things.

You're like, wow, wow. We each are getting our own teacher that maybe isn't as good as the human teacher, but all of us get one at the same time. That's pretty cool.

AUDIENCE: I would say, like, adding, I think, to what you said with purpose is also like, you went there to do that English practice. It wasn't used for any extraneous reasons. Something like this, like Chromebooks, I have a younger sister who's in high school right now. And all her homework has to be submitted on Canvas or something. And then she also just loves, like, watching YouTube and stuff.

And then my mom's problem is like, I can't remove her technology or limit her technology because that limits her schoolwork. So what do you do? Like, you can't control her, like, what she's doing.

JUSTIN REICH: Not only that, but the vast majority of how we learn to be parents is by observing our own parents. And this generation, parents of roughly my age, a little bit younger, certainly the next generation coming along, we can't turn to our own childhood examples to figure out how to navigate something like this. It's very difficult to come up with good parenting strategies from first principles.

A professional parent or a professional caregiver getting training and other things like that. What are you supposed to do? And no, it's a-- I think a lot of school-wide infrastructure has a bunch of unintended consequences. The idea that at night we're going to lock kids into their devices like that. And it can also differ by kids.

So for kids who have really good executive function who can be like, I'm going to do all my homework and then watch the YouTube videos. Like, maybe there's cool different things that we could do with the computers. But for kids who are-- yeah, it's weird. Like, do homework on the machine, which always has YouTube on it. Like, literally you're one click away from really fun things all the time. But do your quadratic equations, really.

It's nuts, right? Like, there must be millions of kids who cannot, are not good at that.

AUDIENCE: Yeah, I think to add on to that fact. I also went to an international school in Ethiopia. And it was a really interesting experience of how rules around technology would change almost like on a weekly basis. At some point, it was like, at some point, like, you can't use any social media. And then, like, phones were banned.

And then they're like, oh, I think it's call your parents, though, you can't call them. And then it was like limited time. And then like VPNs came. And then it was like-- it's always so interesting of how teachers and faculty would always try to make, like, new rules. And each time, they would have to get broken or something would have to change. And I think it was always, like, almost like a comedy show of how much it changed throughout the time.

JUSTIN REICH: When internet filters were first widely introduced in schools, one of the problems came up was the vice president at the time was a guy named Dick Cheney. And there are a lot of searches that were filtering out "Dick." But you need to figure out about Dick Cheney. So then like, someone sets a rule. The rule isn't right for the context.

GPT is another set of things right now that people are doing with tons of rules about. These are good things. We're going to keep talking about them. I'll take a break here. So that we can move on to some other things. But you all have valuable case study experiences from your time in schools that we can bring into our conversations together to think about how new technologies operate.

To do a good job of making sense of technologies, we need to understand a little bit about how people learn. So we're going to spend two classes looking at two pretty well established contemporary philosophies, pedagogy, sciences about learning. Today you read a bunch about cognitive load theory. I'm going to ask you to do another exercise, actually, in the same groups that you were in would probably be fine for this one.

I want you to draw a visual representation of what the authors of the cognitive learning-- of the cognitive learning theory, articles you read in other things, called human cognitive architecture. What are the key components of the brain of thinking of memory in order to have learning happen?

Another thing that you could draw in this visual representation is based on this human cognitive architecture, there follows a series of instructional principles. There follows a series of better and worse ways to teach people. And so you could also try to illustrate how this human cognitive architecture leads to different kinds of instructional principles.

So groups can either go and take a chunk of whiteboard and use some of those whiteboard markers. Or part of why I'm doing this is just, I really like this paper, but it kind of got folded. So I just want to get rid of it. But you can take one of these and huddle around on one of the tables and draw. So if you're a paper person, you can draw on these and then stick them up there in a few minutes. If you're a whiteboard person, you can do that.

We still have all these markers here. I haven't given you totally clear instructions, in part because what I want you to do is to figure out how to think through some of this on your own a bit. But get in your groups, talk with each other about what you think are the key components of cognitive load theory and the human cognitive architecture that leads to cognitive load theory. And don't go more than like two minutes before starting to put pen to paper.

You're not going to be graded on whether or not you're right. But let's see if we can get a couple of different representations of what these folks are talking about, how people learn and based on how people learn how should we instruct them. Does that give you enough to start messing around for a little bit? Ready, go.

All right. You've seen a couple of these now. What are the main building blocks of this model of a brain? What are the key features of human cognitive architecture?

AUDIENCE: I feel like almost everyone has working memory and long-term memory.

JUSTIN REICH: OK, two big pieces here that are important, working memory and long-term memory. What are some of the characteristics of long-term memory? It's cool.

AUDIENCE: Different schema stored inside of long-term memory.

JUSTIN REICH: OK, so we store stuff in blocks. We want to talk some more about what these schemas are. That seems important.

AUDIENCE: It's kind of limitless.

JUSTIN REICH: It's limitless. That is bonkers. No one has discovered the boundaries of the amount of stuff you can store in the human brain. You have a finite amount of space, you have a finite number of ounces, you have a finite number of neurons. There's a lot of them, but a finite number. And as far as we can tell, there are no known limits to how much stuff you can store inside your brain.

That's amazing. That is an amazing feature of your brain. Anything else that seems important about long-term memory? All the stuff can be there. How cool.

AUDIENCE: It's also unconscious.

JUSTIN REICH: Yeah, you cannot think about all those things at the same time. It is unconscious. How long does stuff stay in there? It falls out, right? You forget stuff. If you're interested in learning, the fact that you forget things is really important.

So stuff goes in the long-term memory, it leaks out into long-term memory. Some of the very earliest experiments in psychology done by a guy named Ebbinghaus were about the forgetting curve. It was a psychology experiments that were mostly, as I understand, he mostly did by himself. He basically, like, tried to remember stuff. And then tested himself, and then saw how often he forgot things, and noticed that they were mathematical properties, the forgetting curve and things like that.

All right, long-term memory. It's functionally limitless. It's unconscious, so we can't directly access it. We can put things of different levels of complexity in there. What schemas do is allow us to aggregate information elements into more and more complex kind of things. What are some of the key features of working memory?

AUDIENCE: Very limited.

JUSTIN REICH: Very limited. About how many things can you hold in your working memory at any given time, give or take? Maybe this article didn't emphasize it. Most people know this. It's kind of like a fact that is out there in the world.

AUDIENCE: Like, seven?

JUSTIN REICH: Yeah, for a long time, people thought it was seven. A better guess seems to be even less than that, three to five. But somewhere in the very-- so here's another beautiful thing about-- you all come from all different parts of the world, you can look around this room and you can see that you're different in all kinds of ways. You're different genders. You have different skin color, different hair color, different eye color, all that kind of thing.

All of you have essentially the exact same working memory. All of you can hold three to five things in your head. It doesn't seem to decline as you age. Like, there's other problems with aging, things like that. But people of all different ages, everyone, every human you'll ever meet, the eight billion people on this planet, they can all hold three to five things in their working memory at any given time. It is like a beautiful thing that unites all people on Earth is this common feature of cognitive architecture.

How do you get or what are the things in working memory? How do you get them, or access them, or do something with them?

AUDIENCE: Well, I was going to say like as much as we can only hold like three to five things, like, those things you can say, like, are like schemas. And they can be super complex. And that's how you're able to solve really hard problems. So as much as the-- like, the example of the restaurant, you have the chairs and the tables, that's like the basic things. And then the food that you eat and the workers, and like, all these things make up a restaurant, which are super confusing. But your working memory just knows it as this really large, like, one object. And so in your working memory's mind, that's only one thing, but it's really a compilation of hundreds of things.

JUSTIN REICH: Yeah, good. So we can organize some things into larger things. And the reason why we're using imprecise words, I would say, is in part like, this is not the revealed word of the creator about your brain. This is like the best models we can come up with. And our models are imperfect.

What exactly is a thing in working memory? I think it's a little bit sketchy how much our understanding of that is. However, even with sketchy models, like, we can still do useful things with.

Let's do a quick working memory exercise. Do any of you speak Sanskrit? Oh wait, I can't show you this just yet. Well, here's-- actually, I'll do these first, and then I'll show you this next thing. Grab a piece of paper and a pen. If you don't have a piece of paper and a pen, you can pass that around. Here are two that I really like.

This one, I think, illustrates really nicely the idea that the working memory is the bottleneck. If you are an educator, one of the main things you're working against is people's limited working memory. If people had huge working memories, you could tell them all kinds of complex things.

You'd be like, let me just explain everything. And they would be like, great, now it's all in my long-term memory. But you actually can't do that. People can only attend to a small number of things at any given time. And for novices, those things that they attend to are necessarily going to be simpler than for more complex people.

So for the people in cognitive load theory, as they think about instructional environments, the bottleneck of those instructional environments is the limits of your working memory. Anybody read the Sweller article? Anybody remember how cognitive load theory got started to begin with? What was the canonical experiment that got John Sweller kicked off onto cognitive load theory?

AUDIENCE: Wasn't he, like, making his students do some kind of weird math, like, adding and subtracting, and multiplying multiple things?

JUSTIN REICH: Yeah, yeah. So you solve a math problem, you're given number A and number B, and there's only two operations you can do. It's something like you can multiply by 3 or subtract 69. Those numbers aren't exactly right. But those are the only two operations you can do.

And somehow you have to get from the first number he gives you to the second number he gives you. People do multiple forms of this problem. It turns out that solution is always alternating between these two things. Add 69, multiplied by 3. Add 69, multiplied by 3. Add 69, multiplied by 3. Nobody notices this.

None of the people. Maybe not none. Very few people in the experiments notice this pattern. This is deeply troubling for John Sweller, because John Sweller is in a moment in math education where people are really big on problem solving, really big on discovery, really big on intuition. And he's like, my guys, people are not-- they're not inducing the key principle that's here.

Like, they're doing a kind of mathematical operation, but they're not figuring out the important thing here. Like, think about how much consequences that has for math instruction. They're solving the problems correctly without understanding some key feature about how you solve them correctly. That's bad for mathematics instruction.

And then he comes up with this theory. Like, why weren't they able to do that? And he describes it as that people do this thing he calls means end analysis. They're constantly looking where they are and thinking about what would get them to the next step. If you're constantly thinking about where you are, where you want to be, and what might the next step is, those are like three things to block up your working memory. And so now you don't have a whole lot of working memory left.

And his theory was that people, like, essentially filled up their working memory, their working memory was at capacity. It had overwhelmed. There was too much cognitive load, too much cognitive load on the working memory. And so they couldn't figure out this underlying principle.

And so Sweller-- and then he just gets fascinated for the rest of his life about thinking about what is the human cognitive architecture? Like, where is the cognitive load? Where are we overburdened? Where are there bottlenecks? And what would be instructional designs that would help us relieve those bottlenecks, work around those bottlenecks?

In particular, really skeptical of this idea that you just give a bunch of novices stuff and they can figure it out on their own, with enough time and enough exploration, things like that. So that's the bottleneck here that I really like. And I really like this picture that has of infinite long-term memory. I mean, in theory, you should just go up forever, but it's big enough. That works.

A nice thing about this one on the side is it notes that where does stuff get into working memory? Kind of from two places. You can pull it out of long-term memory, you can remember stuff and then you can put it temporarily in your long-term memory, or you can attend to it. You can see things in your environment. You can watch stuff. You can read things.

You can listen to people and talk to people. As I'm explaining things to you now, you're taking these constructs, and to the extent that you're listening and attending to them, you're putting them in your working memory. If we mess around with them enough, if we get you to apply them, think about them, if you write them down, some other stuff like that, it will eventually get stored in your long-term memory.

And then there's this nice piece here that stuff sort of not only goes from working memory to long-term memory, but things kind fall out and get forgotten. And yeah, I think pairing these words together, when stuff in working memory gets stored in your long-term memory, they call that learning. When stuff in long-term memory gets put into your working memory, when you become conscious of it, we call that remembering. That's kind of a nice way of summarizing that.

There are, in this model, everything that you learn essentially is kind of a fact that you remember. There's a great quote by Daniel Willingham, who's a psychologist, that "memory is the residue of thought." When you think about stuff that you remember. And in this model, people tend to think of the things you remember as discrete facts. You can probably think about-- and the accumulation of those discrete facts is what leads to expertise.

That understanding is, in a sense, the accumulation of lots of these facts. There are probably some domains that you could think about where that model feels really good and makes sense. Oh, man, that person is a really good economist because they just, like, know a ton of stuff about economies, and they know a ton of facts about economics, and they remember lots of different calculations.

Would you say that Maya Angelou was a great poet because she had way more facts than other poets do? That seems like a unsatisfying description of how people become great poets, but something to wrestle with and think about a little bit. OK, you have your piece of paper with you?

I'm going to show you a phrase. And I want you to copy down-- I'm going to show it to you for about 10 seconds. So hold your pens up in the air. You're going to look at this thing for 10 seconds, and then I'm going to take it off the screen, and I want you to copy it down. OK, ready?

I want you to copy this down. Look at it. Memorize it. Look at it. 5, 4, 3, just this row. You don't have to do that one. Just this row in the middle. 2, 1. All right. Copy down as much of that as you can.

All right. And 3, 2, 1. Now, everybody hold up what you did. Show it to each other. How'd you do? Oh, man. These are a little bit disappointing.

All right. Let's do a second one. Let's do a second one. Let's see if you can do any better than this. All right, I'll help you. You might feel better about this one. I'm going to give you a set of writing now that has roughly the same number of marks in it.

It's not the exact same number of marks, but one of the things we could do to calculate how difficult this is to just count the number of marks on the screen. I'm going to give you another thing like this with roughly the same number of marks. Let's see if you can do any better.

I want you to look at this and then get ready to copy it down. This English phrase in the middle. Ready? 5, 4, 3, 2, 1. OK, copy it down.

Oh, you're finished already? You're really good at that. Nice. Good. Why were you so much better at this one? Yeah, yeah. You have those squiggles are organized into schema.

When you were very, very young, you observed these English letters. Actually, if you learned English in the middle of your life, whenever you started learning English, like just now you observed these Sanskrit letters. For those of you who are native English speakers, you can't remember the time, probably, in which there would be no distinction between English letters and Sanskrit letters.

So when you were copying down Sanskrit, none of this has any meaning to you. So you're like, literally memorizing the order in which these lines appear. Which is really hard because there are lots of those lines. There's about the same number of lines, I haven't actually measured this, but it's close. Take for granted that it's close.

But you know that that two combinations of line is an L. And you know that an L, an E, and a T is a word, "let," that has some kind of meaning. And you can actually, "let the entire world be happy," even bundle that into a phrase that you can memorize in your head. And so copying down that number of things is trivial because this is organized in your head in, like, a tightly bundled schema in a way that you have no schema for understanding these kinds of things.

Probably like the first thing you started putting in your schema was like, well, there's a line that goes across there, and then a bunch of things that drop down after that. I don't know anything about Sanskrit, so I don't know anything about it. Does that give you, is that an illustration of how schema operates in your mind?

This completely overwhelms your working memory efforts because you have to recall the orientation of each individual line to be able to copy it. This does not at all do that because you've already bundled those lines into schema. Another famous example, which I think is written about, is that chess masters. When chess masters look at a chess board, like, all of the pieces, they can remember easily all of the pieces in their arrangements.

Because if you know nothing about chess, like me, you have to be like, OK, the pawn is in the second row, black square. And the king is in the first row, white square. And the bishop kind of moved over here. But that's not what chess masters are doing.

Chess masters are going like, oh, that's the Turkish Armada Gambit or something like that. And they've bundled all of those positions together into one thing. So even though we can only attend to three to five things at any given time, those three to five things can be enormously complex. And that's what lets us think about complex and difficult things.

OK, that piece of schema makes sense for you? All right. So these are some pretty good fundamental models of human cognitive architecture kinds of things. Anybody want to ask any questions or make any observations about some of these key elements? How our brains work?

AUDIENCE: I have a similar thought to that. There's a weird alphabet that's just, like, alternative to writing in Latin, in English. So it's just one-to-one mapping between English or Latin characters to that alphabet. And as a fun thing, I just wanted to learn it.

So I did all of that, like, remember the mapping? And even though, hey, if you show me just like a glimpse of it, yes, I can tell you. So it's not, like, a process of how long will it take me to remember, but it was very hard to actually transcribe things fast from one thing to what I know it is.

But also, as I did that more, I started to-- I started kind of like to remember how I learned English as well with some of the squiggles, for example, at the end like that says I-N-G, I started kind of internalizing that. So whenever I saw like that kind of structure of whatever shape these characters, I kind of, oh, yeah, I know that's I-N-G. I didn't have to think individually about each character. I started to chunk them more and more into specific components that made transcribing things much easier.

JUSTIN REICH: Yeah, yeah. I-N-G is a great example of what it probably requires five different line strokes or something like that. Line, dot, 1-N-G, something like that. And then eventually you bundle it as an ING, and just memorize those characters together. And then eventually you realize it's a gerund. And it has a particular kind of modification to verbs and things like that. And you're not even thinking about it now when you look at the word "jumping," but you went through a process of learning all of those kinds of things.

OK, good. So based on these models of human cognitive architecture, what are some of the key instructional consequences that come out of this model?

AUDIENCE: If you overload the working memory, students can't store the necessary information in the long-term memory.

JUSTIN REICH: Good. So what teachers should be thinking about, in this model, which is thinking about all the time, is how much cognitive load, how much stuff is in people is, like, could potentially be overwhelming people's working memory? What are some ways to conduct instruction that respects this bottleneck of working memory? What are some effective strategies for dealing with that?

AUDIENCE: Guided practice is one of them that I saw, where it's like a completion exercise. Because then it's not completely novel to you and you have more-- you can basically abstract more ideas into schemas that are kind of bigger and then execute them much faster, instead of just focusing on the smaller aspects of that problem.

JUSTIN REICH: Good. So, guided, one version of guided practice. Probably one of the most well-known phenomenon in education psychology is called the worked example effect. And the idea is if you show people how to solve a problem, for novices, that is a more effective way to learn it, than have them just try to solve the problem. That a pretty typical instructional sequence that works pretty well is first you do a worked example.

You say, this is how you solve this kind of problem. Then you would take parts of a problem and have students just do those parts. That might be a kind of guided practice kind of thing. OK, you just factored this polynomial and it took you four steps. And the first thing I did is I just showed you those four steps. The next thing I'm going to do is I'm going to show you three of those steps and have you do one of them.

I just showed you a canonical way to write an introduction paragraph. Here is an introduction paragraph that's 3/4 written. It doesn't really have a good hook. Write a sentence that would fit in this paragraph that would be a hook that engaged the reader. You don't have to do everything else that an introduction has to do.

You don't have to introduce the topic because we've already done that. You don't have to lay out an argument. You don't have to establish your expertise. Just do this one thing.

And then you can imagine, all right, instead of if it's a four step process, then you'll do two of the steps, then you'll do three of the steps, then eventually you'll do all of the steps. And you might alternate which of those steps, do a bunch of one step, missing problems where the one step that's missing is different phases in there. But we're going to decrease the scaffolding.

And the idea is that you can get people in their working memory just to attend to that particular thing they're trying to solve, to store that thing in the long-term memory. And as they have that more in their long-term memory and access it, as they start seeing the relationship between different steps, as their schema for an introduction is not, like, wait, what? It's the hook and then the thesis, and then my expertise?

Is that you start, like, recognizing those things as an introduction. And that takes up less space within your working memory.

So they call this the completion problem effects. Replace conventional problems with completion problems, providing a partial solution that must be completed by the learners. This list came out of "Sweller, 1998." All of these things get constantly tested and improved.

So one thing that people realized is in effect, which I think is not on this list, called the expert reversal effect, but I think it shows up in some of the readings that you did, which is that for novices, the worked example effect seems, to generally speaking, hold true. For experts in a domain, open-ended problem solving does seem to work better than just going through worked examples.

So once people get good enough at a thing, it's a good idea to let them solve open-ended problems, presumably because their working memory is not overwhelmed. They're not overwhelmed by cognitive load, figuring out what they're supposed to do next. And so it's a chance for them to refine their ability to apply things in different domains, to get better particular portions of things, stuff like that.

Now, what is a novice and what is an expert? Like, that is different in every domain. And you'd have to-- it would have to be-- it'll be different in nuclear engineering than it is in poetry or whatever else it is. But generally speaking, these are things that seem to hold.

How do these cognitive load theorists conduct their research? If you're a cognitive load theorist, how do you establish that an effect like this is true? What kind of research methods would you use?

AUDIENCE: Well, you could do one where you just directly ask the person. So they have their own, just, understanding of what they're able to handle. Or you could set up the problem so that you can just see directly from the work that they're able to do, like, how well they're cognitive load is being held.

JUSTIN REICH: Good. So one question that cognitive load theorists will be interested in is like, when do you feel overburdened by cognitive load? It's kind of interesting that people are actually quite lousy at self-assessing their learning in all kinds of ways. There's a classic study, I don't think we'll read about it, where people assess their baseball knowledge and then take a test about baseball. And the results are, like, the most perfectly spread scatter plot you've ever seen.

There's basically no relationship at all between your self-assessment of your knowledge and your actual knowledge of things. So this is one of the real challenges in cognitive load theory is that we're not really aware of our own cognitive load. We're not really aware when we're encountering desirable difficulty, which is giving us a good amount of challenge, versus undesirable, where our cognitive load is-- our working memory is overwhelmed and we can't feel a thing.

So people have tried to do work along those lines. What are some other ways of assessing whether or not an instructional effect is better than another instructional effect? Yeah.

AUDIENCE: I think in one of the papers that talked about physiologically measuring the cognitive load just by using heart rate and things like that.

JUSTIN REICH: Skin conductance, other kinds of things.

AUDIENCE: It also said that this was very unreliable as a measure of cognitive load as well.

JUSTIN REICH: Yeah. Yeah. Yeah. Knowing when people are working memory-- it's kind of funny, right? This is like the heart of the theory. How do we tell when you're working memory is overwhelmed. And it's pretty hard actually to tell. We don't have a great grasp of how many-- there are some kinds of exercises we do, which especially really simple problem solving or puzzle kinds of things, where you can really tightly control how many elements are in working memory.

But if you're like, actually trying to teach something, someone in the real world, like the Krebs cycle or something like that, how many things are in working memory at a time. It's a lot harder to check. So a pretty common set of methods that they'll use are experimental designs in which you have one group be a control group, and then you have another group be an experimental group. You measure people's skill at some-- in both groups, you measure people's skill in some kind of baseline.

So you say, we're going to design-- we're going to go into a math, a school where they typically use a lot of open-ended problem solving. And we're going to instead have the teacher just do stuff using worked examples. And we're going to see whether the worked example condition is better than the business as usual condition. And in fact, to make sure that there's no difference between the kids who get the regular condition, the business as usual condition, the kids who get the work example condition, we're going to randomly assign the teachers.

We'll get 20, or 30, or 40, or 50 teachers sign up for the study. Half of them will do the old thing, half of them will do the new thing. We'll measure all the kids' skill at baseline, and then we will measure the skill from the groups afterwards. So if this kind of model works, like, the two groups at baseline, in any group there's going to be a distribution. We're going to teach you how to factor polynomials, or add single digits, or conjugate Spanish verbs, or whatever it is.

But in any given classroom, there should be a distribution of skill, right? Some people are having a hard time with it, most people are doing fine, and some people are doing really good. At baseline, these two groups should be identical, right? Because we randomly assigned your teachers to two different groups, and random assignment is magical. And like every once in a while, it doesn't work, but for the most part, like, the two groups should be the same.

What should happen afterwards is that-- now here's something that makes learning research really hard. What kind of learning will-- like, learning will happen in both conditions, right? Almost everything we do makes people learn somewhat. Like, there are very few things that we do that we're like, boy, that was completely ineffective. Even if we're not using the most optimal techniques, students still learn stuff.

So in the control condition, let's call orange control, they're still at Time 2 going to be improvement, T1, T2. You're going to see of distribution move. What we want to see, though, is something like this. Like, oh, we learned more in the experimental condition, when we went from T1 to T2. And in particular, we're interested in-- the way cognitive load theorists do this, they're particularly interested in averages.

Like, in general, are kids more or less better off having to do this? Like, this has 10 teachers, each of whom has 20 kids. That's 200 kids. This has 200 kids. Is one group better than the other?

Sometimes you can do some of this stuff. Some of this stuff happens in the field. Some of this stuff happens in laboratory kinds of experiments. One, let's see if-- so, this is a modality effect, which is kind of cool. Replace a written explanatory text and another source of visual information, such as a diagram, with a spoken explanatory text and a visual source of information.

It turns out, again, we don't-- I would say we don't totally understand this, that your working memory seems to handle visual and auditory inputs differently. And so you have a little bit more working memory to work with if you're hearing stuff and looking at stuff than if you're reading stuff and looking at stuff. Why? Nobody knows.

Evolution? Who knows? But it's kind of cool, right? So how might you study something like that? You take a textbook and you bring a bunch of students into the lab. And you have them do some kind of textbook learning exercise. You measure them beforehand. You measure them afterwards.

For some random half of the students who are walking into your lab, they probably are taking a freshman psych class, and they have to do this for extra credit or something like that. And you give them the, like, auditory enhanced version of this. And if the people in the auditory enhanced condition in the modality effect condition do better, then you didn't have to go all the way into the field. You didn't have to recruit a bunch of teachers in real classrooms and things like that. You could study little bits of this in the lab.

But again, the thing that these folks are particularly interested in is changes in average effects. Like, these models tend to assume that we're teaching stuff that everyone should learn. It would be good if all the kids in my math class learned to factor polynomials. And we want strategies that, on average, have more of those kids learn more about factoring polynomials and other things.

I'm sort of emphasizing this point, because a bunch of stuff that you're going to read for next class doesn't necessarily make those assumptions. It's still about learning, but thinks about doing research differently and thinks about goals differently.

OK, here are-- oh, good. We just did question number 3 here of things that I wanted to talk with you about. Number one, we emphasized, for cognitive load theorists, what is the bottleneck that they're most concerned with? Working memory. There's a limited number of stuff that you can operate with in working memory. So you got to think about, especially for novices, how do you reduce the complexity of a learning environment?

What are the goals of cognitive load theory? What are, like, the end states of proficiency? For cognitive load theorists, what does it mean to be good at stuff?

AUDIENCE: Would it be to have built up a bank of schemas, and then be able to retrieve them at any time, and use them to work more quickly through a problem?

JUSTIN REICH: It could be more quickly. So definitely, I think the first part, I would say you're spot on about m you have memorized a bunch of stuff. The stuff that you've memorized is organized into schemas of varying degrees of complexity. There are some things that don't organize well into schemas, but we just memorize them anyway.

Famous things that don't fit super well into more complex schemas are a bunch of Asian languages have characters that don't, like, have a lot of logic across them. Like, you just have to memorize these things. And there are a lot of them. And like, schemas will only get you so far.

But an important thing about, there is no relation-- think about it this way. There's no relationship between proficiency and the learning process. Like, what do cognitive load theorists think a really good surgeon is? Like, a really good surgeon is both proficient at certain kinds of well-established techniques, but also flexible, intuitive. A patient has cancer and you cut them open, and it's some set of physical problems they've never seen before.

They don't want doctors who are like, oh, I must memorize the procedure and just do what I'm told. They want doctors who are flexible. They want poets and playwrights who are creative. They want teachers who are engaging and effective. Like, their vision of proficiency is the same as almost anybody else who studies teaching and learning.

Just because they have methods that are, like, it doesn't mean that they want people the rest of their lives trying to reduce the amount of working load, load on their working memory. They want people who can do all of the things that we think experts should be able to do. But they say the way you get expertise is not to take a surgeon and be like, just get in there and start cutting people up, and like, see what happens.

You want to train them through a particular series of processes that has them memorize a particular series of facts. Those facts can be assembled in knowledge, or techniques, or other kinds of things. And then flexibly deploy those things they memorized to do good work in the world.

So don't confuse the way that cognitive load theorists want people to learn, which can be pretty structured, have this devolving scaffolding. That doesn't mean that that's what expertise looks like when you get to the end of that process. We might teach people to be really good computer programmers through a set of worked examples, but it doesn't mean that we want computer programmers who just copy things that they've seen in the world. We want those computer programmers to be flexible, creative, novel, inventive, all those kinds of good things.

You read two other pieces that I didn't talk as much about, but we talked a little bit about Sweller, "The Story of a Research Program," that this research comes out of a particular historical context, which is people in math pretty fired up about problem solving. And this guy, John Sweller, being like, I think there are some real problems with open-ended problem solving. I think it doesn't do what we want it to do, and there might be better approaches to try.

What did you take away from the Michael Person article? So Michael Person is a teacher in New York City, just teaches math to a bunch of different grades. And he's trying to make sense of all this research. What did you take away from the things he was trying to make sense of?

So the title of his piece is, "Not a Theory of Everything." What did that allude to?

AUDIENCE: I could give a response to say, applying cognitive learning theory to cognitive learning theory, and just restricting it to where it is, putting the boundaries, making sure we understand what it is and what it's not, to make sure that when we're discussing it, we're not stepping beyond it and starting, like, confusing things together and not knowing what we're talking about.

JUSTIN REICH: Good. So some of it is him saying like, all right, like, what exactly are these terms and how are they used? How do they mean at a time? When Michael Person says, it's not a theory of everything, what are some of the things that prominently it is not a theory of?

AUDIENCE: Some motivation.

JUSTIN REICH: Yeah. Why do you think-- so, Michael teaches math in a school in New York. Why do you think motivation would be something that would be a pretty keen interest to Michael?

AUDIENCE: I think that could affect how likely people are to remember the content right there. If you're really interested and really motivated, that could affect.

JUSTIN REICH: Yeah, I think you could have a subtext of Michael's argument, which is like, my guy, like, you're clogged working memory is not the problem I have. My problem is my kids don't care. Like, if you're not actually interested, and engaged, and want to learn these things, who cares if you're stuck? It's not a problem to be putting seven or nine things in their working memory. The problem is they will put zero things in their working memory.

The problem is not the bottleneck that appears here. The problem is you can't get them to do this. You can't get them to attend to anything. I think there's part-- and for him as a teacher, motivation is just essential, like, a central problem of what he does every day. How am I going to convince these fourth graders in front of me that it's worth learning this math?

And that, I think, reading between the lines, I think that's part of what feels really important to Michael. And John Sweller says, this theory doesn't do that. This theory doesn't attend to issues of motivation.

Now, the good thing about doing that, and I think Michael comes to this conclusion in some respects, is like, yeah, it's probably not a good idea for every theory to try to do everything. It doesn't diminish the contributions of cognitive load theory to say it doesn't attend to motivation, but it probably should draw our attention to the fact that working memory is not like the single bottleneck that exists in all learning environments everywhere. There are all kinds of other bottlenecks that might be around.

Again, some of the stuff that we're going to read about next week thinks of motivation as a very central issue, as maybe like the central issue to teaching and learning, which doesn't mean that this other theory is better. But that the theories that we have are just models. They help us make sense. Nobody actually understands how the brain works. Nobody actually understands how people learn. That is beyond our ken.

But there's a lot that we can know about it. And we can build these models which help us make good predictions, but they don't help us predict everything.

The other thing that I think Michael's article is really good for highlighting is that these theories change. How many of you read in the original article about this thing called germane load? Yeah. Researchers in cognitive load theory have basically given up on germane load. It was an idea they had for a while.

They were working on it. They thought it would predict some stuff. They thought it would explain some stuff. And after a while, they were like, yeah, I don't think this germane load idea is all that useful. And people have more or less given up on it. So these theories are constantly changing and evolving too. The version of cognitive load theory that I might teach in this class 20 years from now is not going to be the exact same as I would have in the 1980s when John Sweller was just putting these things together.

Some parts of it have stood the test of time. I would say, particularly this idea, the idea of extraneous load seems to be enormously helpful. If you put too much stuff in front of novices, they won't attend to the right things and they won't remember the things that you really want them to remember, or they won't be able to attend to what you want them to do. But if you reduce some of that extraneous load, and there are a bunch of very specific strategies that we can use to reduce that extraneous load, some of which are very, very fine grained.

If there are multiple sources of information, frequently pictures and accompanying a text. Like, don't put the picture over here and the text over here. Put them in the same place and people will learn better. That's great. Think about how many people today in the United States, there are about 57 million kids in K-12 schools, and tens of million people in higher education.

Like, if all their textbooks got a little bit better by adhering to this principle, everyone would learn a tiny, tiny bit more, but amortized over 57 million kids going to school for 13 years. That could be a lot of learning. That's pretty cool.

All right. 2:20, what are the last things that I want to leave you with? You now have 50% of the ideas that you will need to do assignment number 1. So again, assignment number 1, which you should go to the syllabus and read, is you're going to pick an educational technology. You're going to explain if that educational technology is more inspired by instructionists. By people who are really fired up about direct teaching of things, the cognitive load theorists or instructionists. The worked example effect is an instructionist kind of idea.

There are lots of education technologies which are instructionist in nature, which are like, how are we going to tell people how to do this thing so they do it better? There are other pedagogies, education technologies, which are much more focused on experience, much more focused on apprenticeship. How do we get into the world and muck about and do things? And that is what you will read about for Monday, which is-- and both of these ideas are ancient.

I mean, I think I mentioned to you before this quote from Plutarch, that "education is not the filling of a pail, but the lighting of a fire." The cognitive load theories are like, no, pail filling is awesome. This is like super effective. We can measure how much stuff goes into the pail. We can measure how much stuff leaks out of the pail.

We can measure the optimal flow rate into the pail. Pail filling is cool. We're now going to find the contemporary versions of the flame kindlers, people who are trying to really get their minds around how do you get people to learn and naturalistic apprenticeship reactions?

So you're going to have to find a technology. You have to explain whether it's more instructionist or constructionist, more inspired by cognitive load theory or situated learning. Describe how the features of that technology reflect these pedagogical values. That's like the first couple pages of the assignment.

You have something like 1,200 words, but you can include pictures, you can annotate things. You can be like, look, this is like a little screenshot of the thing, and like, here's where instructionism is written all over it. Then in the last part of the paper, you have to take this same technology and imagine that it was redesigned or designed originally by people from the opposite pedagogical perspective.

So you might go and find a technology and be like, man, this is like straight cognitive load theory. What if the same technology with the same kind of goals was designed by situated learning constructionist kind of people? That is the task. So you can be thinking, as you're doing all of these kinds of readings, what are some of the key principles that I could use? How would I apply these key principles to education technologies?

Which education technology am I most interested in? How am I going to learn about the other pedagogical philosophies so I could apply in the same way? That should give you some stuff to work with.

OK, 2:25, we'll meet on Monday. You'll read stuff about situated learning. We'll do similar kinds of exercises to make sense of it. If you have questions, come, swing by. Take your books with them. Write your names in them so that you don't get them confused with other people's.

And you can expect another email from me hopefully in the next couple of days. I remember almost every week. I don't promise it comes out. It's like a courtesy service, but it's like a little summary of what you need to do for next week.

OK, I'll stick around if anybody needs anything. Otherwise, have a wonderful weekend. And I'll see you on Monday.