The eLearning Coach Podcast #55

What You Need To Know About Cognitive Load

http://theelearningcoach.com/podcasts/55

Connie: Hello learning people. Welcome to episode 55 of The eLearning Coach Podcast. If you are a learning professional or if you think of yourself as a product designer, one of the most important aspects of your work is understanding how people learn. Yet, many of us in the field lack a basic understanding of the key theories that explain human cognitive architecture and how to best leverage it for learning. That's why I couldn't be more excited to present a conversation I had with John Sweller, PhD, the person who developed cognitive load theory, which uses which uses our knowledge of evolutionary psychology and human cognitive architecture as a basis for instructional design. Cognitive load theory is one of the most highly cited educational psychology theories and I can't wait for you to hear how John explains it.

John is an Emeritus Professor in the School of Education at the University of South Wales, Sydney. He is an educational psychologist and has authored over 180 academic publications.

Connie: Hi, John, welcome to the eLearning Coach Podcast.

John: Nice to talk to you.

Connie: So, it's been pretty exciting to have this conversation with the person who formulated cognitive load. So I think an appropriate question is, what are the components of human cognitive architecture that you think learning professionals should know about and should consider during their work?

John: Human cognitive architecture, as you might expect, is complex. It becomes a little bit easier because large aspects of human cognitive architecture are not relevant to instructional design, so we only need to concentrate on those aspects that are relevant.

John: Let's begin by looking at categories of knowledge. You can categorize knowledge in an almost infinite number of ways. And most of those categories don't matter all that much for instructional design because the instructional design is the same for category A and category B, then we don't need to worry about the fact that there's a difference between A and B. But there's at least one category that instructional designers need to keep in mind, it goes under a variety of names, let's call it primary and secondary knowledge.

John: Primary knowledge is knowledge where you have evolved to acquire. It can be very, very complex, but despite its complexity, we acquire it very, very easily. An example of primary knowledge is our ability to learn to listen and speak our native language. It's probably one of the most complicated things we do. Nobody needs to instruct us explicitly on it, we just pick it up automatically.

John: In contrast, secondary knowledge is knowledge which we're able to acquire, but which we only acquire with difficulty, consciously, and with a lot of effort. An example of that is learning to listen and speak to ... a second language, a foreign language that we're not familiar with. We don't pick that up automatically in the same ways we pick up a native language, a first language. Or learning to read and write, we don't pick that up automatically either. And the same goes for every single curriculum area, or every single area that's taught, by anybody from, all the way from little kids to adults.

John: We have specifically designed instructional systems and educational systems to deal with secondary knowledge. So, when I talk about cognitive architecture, I'm really not talking about primary knowledge, I'm talking about secondary knowledge. And when talking about secondary knowledge, the knowledge that we're concerned about when we're dealing with instructional design, there are several aspects of human cognitive architecture that we need to take into consideration. And those aspects apply to secondary knowledge.

John: The first aspect we'll talk about is working memory, which is critical to cognitive load theory. Working memory, the easy way to think about working memory is in terms of consciousness. What you're conscious of is what's going through working memory. And the major characteristics of working memory when dealing with novel information, firstly, it's very, very limited in capacity. We can only deal with somewhere between, maybe, two to four elements of information at a time. And by elements of information, it's anything you need to think about, or learn, or deal with when solving a problem.

John: Secondly, we can only hold information in working memory for maybe about 20 seconds or so. After that, almost all of it is gone. So, we've got a very limited capacity, limited duration working memory. How do we function under those circumstances? Well, we function because we've also got a long term memory. Long term memory differs from working memory in those two critical respects. It has no known limits of capacity, it has no known limits of duration. So it's very, very large.

John: Once information gets into long term memory, we can then take it back into working memory, we can transfer it back into working memory without effort. And again, without any limits. So, working memory is limited when dealing with novel information. Its effectively unlimited when dealing with processed, stored information that we get from long term memory. And that is critical, because it tells us why we're engaging in instruction. It tells us, we engage in instruction in order to get information into long term memory. Because once we get that information into long term memory, we're transformed. We become different people. We can do things we couldn't dream of doing otherwise.

John: So, that gives us the purpose and the function of instruction, and it tells us how our cognitive processes function. They function by getting material into working memory, processing it, transferring it into long term memory. And once its in long term memory, it can then be easily got back into working memory. It's not easy to get it into working memory as novel information from the outside environment, but it's very easy to get it into working memory once its been stored in long term memory. So, that's a major aspect of human cognitive architecture.

John: Instructional designers need to keep in mind that, when they're presenting people with novel information, they're going to have to take into account the difficulties that the learners are going to be dealing with when dealing with that novel information. And the instructional designer needs to keep in mind that the purpose of the instruction is to enable the learner to transfer that information into long term memory, so that then, they'll be able to take it back into working memory effortlessly. So, that's the basic aspect of human cognitive architecture that an instructional designer needs to keep in mind.

Connie: That's excellent. One interesting thing, an important part of your explanation is that it implies why it is so important to connect novel information to previously learned information, because then you can try to fit it into an existing structure of framework.

John: Yeah. We will automatically, when faced with any new information, try to use old information to assist us in processing it.

Connie: Mm-hmm (affirmative).

John: Which is another way of saying, the more you know, the easier it is to know even more again. It's a positive feedback loop. If you know very little, then just about everything you deal with is new. And we know that as soon as you go to a new area that you're not familiar with, everything is difficult.

Connie: Right.

John: But when you've been in the area a while, not only does it become easier because much of the area is stored in long term memory, it's also easier because when you come across new information in that area, you can use the old information to assist you in processing the new information. You can reduce the effects of that, those working memory limits, when dealing with new information by using old information. So that's absolutely correct, yeah.

Connie: I also like the fact that you mentioned two to four pieces or bits of information in working memory ... Doesn't seem to be any way to increase working memory, is there?

John: No, but keep in mind that, in one sense, if you think about the cognitive architecture that we've just been discussing-

Connie: Yeah.

John: You don't need to, because working memory automatically massively increases just as a consequence of having information in long term memory.

Connie: True.

John: The information is in long term memory, then working memory not only increases, it seems to become almost infinite. It's only novel information for which working memory has those limits. Old information, processed information, stored information, information from long term memory renders your working memory enormous.

John: So, the limits of working memory, in that sense, don't matter.

Connie: Yeah, that's a good way to look at it. It's more positive.

John: Yeah. Yeah, providing you put in the effort to learn the material, get it into long term memory, at that point, working memory becomes infinite. Can I just say ... why there'd been a change from seven, plus or minus two, to two, three or four?

Connie: Please.

John: It's still seven plus or minus two, but only when you're dealing with material that you're just memorizing little bits of information, okay?

Connie: Right.

John: If you got a new telephone number and you're trying to remember it, then it's seven plus or minus two to remember the digits of that telephone number. But, if you're trying to process something in some way, if you're trying to deal with it, work on it in some way, combine it, contrast elements, which is not what you're trying to do when you're memorizing a telephone number, then you're talking about two, three, four elements of information.

John: There's the difference between simple memorization as opposed to processing. And of course, when you're dealing with instructional design, you're primarily concerned with processing, which is why it's two, three, or four.

Connie: Right, right. And a simple example is multiplying a few digits by another few digits, you reach your capacity pretty quickly.

John: Very quickly, yeah, that's an excellent example. If all you have to do is memorize those digits, yeah, you can memorize about seven of them. When you have to start processing them, then it falls dramatically because the act of processing itself takes up working memory capacity.

Connie: It's fascinating. So how do you define cognitive load?

John: Okay, so, cognitive load is imposed on working memory when we're dealing with multiple elements of novel information, okay? So, if you've got difficult material, what do I mean by difficult material? I mean, material that requires you to simultaneously process multiple elements of information in working memory. When do you have to process multiple elements of information simultaneously? It's when element interactivity is high. By element interactivity, I mean that the elements interact with each other in some way. You don't have to just memorize them, you have to process them.

Connie: Mm-hmm (affirmative).

John: And in order to process them, that means you have to deal with them, put them together in some way, think about them simultaneously. And that's what increases element interactivity. Increasing element interactivity imposes a heavy working memory load, and that's what cognitive load is.

John: I'll give you some examples of that, because element interactivity is central to cognitive load theory. We have low element interactivity in a lot of areas where we just have to memorize a series of connections. If you're learning a foreign language and you have to learn the vocabulary of that foreign language, you can learn each vocabulary item independently of every other item. So, you can learn the translation of the word cat independently of the translation of the word dog. They don't interact with each other, you can learn each one independently.

John: Learning the vocabulary of a foreign language is an extremely difficult task, but it doesn't impose a working memory load. The working memory load is low, so cognitive load is dow, despite the fact that it's a difficult task. It's a difficult task because there are a lot of elements, but the elements don't necessarily interact. If you're learning how to put the words together into meaningful sentences, then they start interacting, and then you start getting a cognitive load.

John: An example of high element interactivity is something like multiplying together two digit numbers. You can't do anything without thinking of the entire process.

Or, a simple algebra equation, A + B all over C = D, solve for A. You can't change anything there without considering the entire equation, because as soon as you make one change in one part of the equation, it has effects throughout the entire equation because of that equal sign. So, that's high in element interactivity.

Connie: What about the case of a typical classroom training with PowerPoint for adult learning? And people will run through 30 slides, reading text on the screen, although the individual items ... I'm not sure, do they interact? Or, does cause high cognitive load just because it's a steady stream of information and people don't have the opportunity to process it and integrate it with their long term memory?

John: Yeah. Look, it does cause a high cognitive load. And we're all aware of it ... well, I hope most instructors are aware of it, but certainly the learners are aware of it. Which is precisely why, when somebody's making that sort of a presentation, if you haven't been handed out notes of some sort, you busily sit there taking notes. Why do you take notes? Because you're fully aware of the limitations of your working memory. I mean, intuitively, we're all aware of those limitations. And you know quite well that by the time a few seconds have gone, you'll have forgotten what was just said a few seconds ago. We're all aware of that.

John: So, you need some sort of a permanent record so you can go over that later, and over that again, and think about it, and over it again. And after a while, it'll all go into long term memory. Once it's in long term memory, then you're okay. But before then, yeah, there will be a high cognitive load. And when presenting information in any presentation, instructors need to be acutely aware of the fact that the cognitive load for them, as instructors, is low. It's low because it all comes from long term memory, right?

Connie: Right.

John: They know it all. And it's very easy to forget, for the learner, you put something up on a slide or you say something, and you automatically connect it up with what was said two, three, four, five minutes ago. The learner can't do that.

Connie: Good point.

John: If they've gotta connect it up, they're gonna fail. And because of that, you need to present them material in such a way to keep in mind that those working memory characteristics that the learner has, which are different from the working memory characteristics that you have because you are an expert in the area.

Connie: Right. So, I guess my question is, is that a case of element interactivity that's causing high cognitive load? Or is it often cause, also, by just mounds of information being thrown at someone?

John: No, it's the first. It's element interactivity.

Connie: Okay, so that is a case of element interactivity.

John: That is a case of element interactivity because, if you've got a whole lot of information and a whole lot of slides, and the slides don't really relate to each other. In other words, when you're talking about slide 10, it really doesn't relate to slide nine. And then you go onto 11, and it doesn't really relate to slide 10. That's not a high cognitive load. People will be able to understand what you say in slide 10 without remembering what's in slide nine, if they don't relate.

John: If they do relate, and if what you say in slide 10 requires your learner to have transferred the information from slide nine into long term memory, they can't do it because it's not in long term memory yet. So, it's element interactivity if the elements of slide nine interact with the elements of slide 10, and then you run into problems, and then you run into people who sort of hold their heads and say, "Oh, I'm having difficulty understanding this," that's when you have element interactivity.

Connie: Okay, great. One thing I wanted to discuss was how cognitive load affects how people understand and make use of user interfaces?

John: When we're talking about user interface design, we're really talking about, how do we design the interface in such a way that cognitive load, working memory load, they're essentially the same things, is not excessive. How do you reduce it? And we've developed a whole series of cognitive load effects which are designed to show us reasons why working memory load might be excessive when you're presenting information. Whether it's on an interface or whether it's in a talk or any other way, the same rules apply. And they all go back to that cognitive architecture, which we briefly outlined a little bit earlier.

John: One of the effects is called the redundancy effect. Let me give you an example. Often, when you go to a webpage which is designed to show you something, you'll find that it's busy. There are ... there's stuff everywhere that you've gotta attend to, that you've gotta think about. And you don't really know, well, should I be looking at this now? Should I not be looking at it? Don't put so much ... well, a lot of instructional designers take the view, there's no such thing as having too much on a webpage because if the learner doesn't want to look at something, they don't have to look at it, okay? They can just look at the stuff that's important.

John: Well, that's all very well, but the instructional designer knows what the learner should be looking at. The learner himself or herself does not know, they don't know what they're supposed to be looking at. The material is new to them. What should I be looking at? We really need a webpage that doesn't have a great deal of information, and make sure all of the information is essential information. Anything that's there, be wary of putting stuff on a webpage that's just there for, I guess, attractiveness reasons, just there to attract attention, because it will attract attention.

John: Now, if you have something moving on a webpage or something that's very colorful and really doesn't have anything to do with the basic message that the learner has to acquire, it's going to take up those limited working memory slots, which are no longer going to be available for learning.

Connie: Mm-hmm (affirmative).

John: So, don't have people looking at something that they don't need to look at if what they're looking at is complex, difficult to understand, they don't need a cartoon as well.

Connie: Right. And that's exactly what I do recommend, because there are ways to guide atn, visually, if you don't have-

John: Exactly.

Connie: Right, use that and remove anything extraneous.

John: Yeah. So, redundancy is one issue. I use the term redundancy in a broad form, but redundancy is anything that's there that you don't need to be there, that doesn't need to be there, anything at all. Sometimes it's redundancy because you're saying the same thing in different ways. Again, if you do that simultaneously, in other words, if you, for example, provide something in written form and also in spoken form, people will learn more if it's just in written form or just in spoken form, not in both. If it's in both, then, they're trying to coordinate the written and spoken material, that takes up working memory and that's useless. That's not the purpose. The purpose is to acquire the information.

John: Attention is another major issue. If you've got information on one part of the screen and you've got information on another part of the screen, and you can't understand either sets of information until you've put them together, make sure that you put them together physically, okay? Because if the learner sees a list of things and then has to find one member of that list in another list, and has to search the second list while remembering what was in the first list. Or has to look at a diagram and then look at a series of statements and, while they're looking at the diagram, they have to remember the statements or, while reading the statements, they have to remember the diagram. That's working memory. That takes up working memory.

John: Put them together, in other words, things that ought to go together, put them together. Put them together either by putting them adjacent to each other, or use arrows so that when you talk about such and such, people know exactly where to find such and such rather than try to search, now, where's that? I know it's here somewhere. While they're searching for it, working memory is ... well, it's blowing.

Connie: Right, that's avoiding the split attention effect.

John: That's avoiding split attention.

Connie: Right.

John: A lot of other such effects, but the last one I'll talk about is also important, and that's the transient information effect. Transient information applies to some ways of presenting information, but not others. If you've got a static page, no movement, no change, then you don't have to worry about transient information. If you've got animation, if you've got a video, if something is moving, then what you had a few seconds ago has disappeared to be replaced by what you've got now. That's transient information.

John: Given the characteristics of working memory that we talked about before, you only hold things in working memory for about 20 seconds. If you've got a stream of information and it's transient, it changes, you need to keep in mind that, while people are processing the current information, they've lost the previous information. Now, that may not matter because the previous information may not be relevant. But if it's relevant to them, you need to put it in some sort of static form so that they can, for example, go backwards and forwards.

Connie: The learner needs some type of control to be able to stop the flow of information until they've processed it.

John: Yeah, exactly. And incidentally, speech is also transient. The reason we invented writing is in order to transform transient speech into something that's permanent, that we can go back to, that's there all the time. What I said a few minutes ago, that's gone. What I said a few seconds ago, that's gone. Speech is transient, so you need to keep that in mind.

John: As technology developed, a lot of instructional designers took the view, oh, look, we can do wonderful things we couldn't do before. We can come up with animations, we can use videos, we can use speech. They're all transient, and that's a problem. There's something to be said for the old textbooks, you know? It was there.

Connie: Right, good point.

John: It didn't go away.

Connie: Right. So, John, can you explain the difference between intrinsic and extraneous cognitive load? Because that has so much to do with people who design learning experiences, it's so relevant to them.

John: Yeah, it's directly relevant, and it goes back to element interactivity can be due to two different factors. You can have high element interactivity because the material is very, very complex. In other words, it's intrinsically complex. It's not something where you just learn one thing after another thing after another thing, which are not closely related.

John: If it's a high intrinsic cognitive load, it's because you're dealing with material where every single element is related to every other element, and you really can't think about them. You can't process them. You can't learn them without thinking about all of them simultaneously, all of the elements simultaneously. So, if the natural intrinsic element interactivity is high, you have a high intrinsic cognitive load.

John: Now, element interactivity can also be affected by the way you present the material, by the activity you get learners to engage in. So, some types of activities result in a very high element interactivity. It's not a natural part of the material, it's just the natural part of the way that you happen to be presenting the material. That's extraneous cognitive load, because you can control that. You can't control intrinsic cognitive load ... well, you can control it just by not teaching the difficult parts. If you don't teach something that's difficult, then the intrinsic load is not gonna be high. But if you wanna teach that, you can't change the intrinsic load.

John: Extraneous load, you can change that. You can change the way you teach. For example, I taught previous about split attention and redundancy and transient information. All of those increase the element interactivity because of the way you're presenting the elements. So, you can present the elements in such a way as to reduce extraneous cognitive load. And good instruction does just that, it reduces unnecessary element interactivity, it reduces unnecessary extraneous cognitive load.

Connie: Right. So, for example, in an eLearning course, if something is very complex, you might show how things are connected at first, then teach each small part, and then, at the end, show again how things are connected.

John: Yeah.

Connie: Is that a strategy that you would recommend?

John: Yeah, that's known as the isolated elements effect. In its simplest form, you teach people material which is connected, but you teach it in isolated form. Now, when you're teaching it in isolated form, they can't really understand, because you can only understand it once it's all been connected. They can't understand it anyway because the element interactivity is so high. So, you reduce the element interactivity by teaching one element at a time. Once they've acquired all those elements, then you can start putting them together.

John: So, from the learner's point of view, they use their working memory resources to first acquire these elements, pretty much learn those elements. And then obviously you don't want them to stop having just learned something. Then, they have to learn how to connect up the elements. And learning how to connect up the elements, that requires working memory resources. But that's when understanding comes in.

John: We know, if you're dealing with something that's really, really complex ... now, I really can't understand this, you know what happens. First you learn little bits of it and you say, oh, look, I can understand this bit, but I don't see where it fits in, I don't really understand what's going on here. You keep on working on it, and after a while, you get that aha moment. That's the point where all of those interacting elements, you finally got them into a long term memory, then back into working memory, and you can now hold and process all of them simultaneously in working memory. That's when you get that aha moment, when you can-

Connie: Right, yeah.

Connie: I was thinking that learning how to program is a great example of that.

John: It's a very good example, yeah. And a lot of areas of mathematics as well.

Connie: Mathematics, yes.

John: Yeah, it comes up very much there.

Connie: I don't want to take up too much of your time, so maybe I can just ask you one more question?

John: Yeah, sure.

Connie: I recently read an article, it was from several years ago, and it talks about why guided instruction appears to be superior, at least for novices, than other forms of instruction. Can you talk about why that is so in terms of our cognitive architecture?

John: Sure. Again, that goes back to element interactivity. One form of guided instruction is the worked example effect. And that occurs when you give people a problem with its solution, with a worked example written out so that people can follow it. There's difference between doing that and having people try to work things out themselves. Try to work something out yourself is very, very expansive of working memory resources, okay?

John: Essentially, what you're doing is problem solving. You're saying, okay, I'm here, now, what do I do next? Do I go to this other point? No, I'm not sure if that works. How about this other point, how does this relate to this? Doing all of that, that's the process of problem solving, of trying to work something out, very expansive of working memory resources.

John: You can reduce the required working memory resources by telling people, "Right, this is where you are, this is the next thing that happens, and this is the next thing." You tell them directly, they'd have to look at, this is where I am, do I go here or do I go there? Where do I actually go? What do I do here? You tell them directly. And by telling them directly, you've reduced extraneous cognitive load. You've reduced the element interactivity.

John: Instead of just having to think of, say, four or five elements, okay, I'm here, do I go to this place, this place, or this place? How do they all compare? You don't have to do that anymore. You're simply told, you're here, and I'm telling you, this is the next best thing you ought to be doing. That reduces working memory load. That's been demonstrated again and again and again.

John: It only works with novices for whom the material is difficult. Once you've got a reasonable knowledge base, then you can start problem solving. At that point, problem solving works.

Connie: And discovery, once you have a good knowledge base, discovery is okay. But for beginners-

John: Yeah, you're not in a position to discover-

Connie: Right.

John: If you don't know anything. There's simply no point discovering there. But once you've got a good knowledge base, then you can start exploring things yourself and finding out things that other people haven't told you or were not in a position to tell you.

Connie: That makes so much sense, yes, yes. John, I don't want to take up anymore of your time, but this has really been a highlight of my five or six years of podcasting. So, thank you so much.

John: Thank you very much for your great questions. It's been a pleasure talking to you.

Connie: Same here. Take care. Bye.

John: Okay, thank you very much. Bye.

Connie: Bye.