

TOPcast Episode #139: Automating Digital Teaching: What Can Go Wrong?

Narrator: The University of Louisville empowers students with over 50 fully online degrees and certificates in areas like business, public health, social work, engineering, and more. Flexible coursework allows time to focus on all of life's priorities. Learn more at louisville.edu/online.

(Intro music)

Tom: From the University of Central Florida's Center for Distributed Learning.

Kelvin: And the University of Louisville's Delphi Center for Teaching and Learning.

Tom: I'm Tom Cavanagh.

Kelvin: And I am Kelvin Thompson.

Tom: And you are listening to TOPcast, Teaching Online Podcast. Or potentially watching.

Kelvin: Or potentially reading.

Tom: Or reading. Yeah, that's true. Our awesome team here at CDL does put together the transcripts. So yeah, you could be reading it.

Kelvin: So, if we are just coming to you through the written word, imagine warm tones in these words that you're reading. [Laughter]

Tom: Or if anybody happens to be translating this into semaphore or interpretive dance or something, let us know because we'll add that to the show notes. [Laughter]

Kelvin: That's right. We're all about access. [Laughter]

Tom: That's right. In all forms. Yeah. All right, Kelvin, I see you sipping there. Oh my gosh. That is a mug. That's a throwback. So, what is in your thermos today? So, for the listening audience who may not be watching, Kelvin's got a WebCT mug. That's a deep cut right there.

Kelvin: It is, it is. I thought it was appropriate for today. It's not my oldest mug in the collection, but in our circles, people grasp the datedness of it rather quickly. [Laughter]

But my coffee today, Tom, in my mug, is also a call back to the past. We have personally shared it previously in person. It is a single-origin El Salvador, (with Spanish flair) “La Siberia.” I might have said La Siberia from Beacon Coffee in Camarillo, California. We drank it previously three years ago in the show that aired in January of 2020. So, I might have mentioned this at some point in the past. I won't get into it a lot, but I got some maybe what was good information several years ago about if you have an abundance of coffee, there's people who put coffee in the freezer. Then there's people who know enough not to put coffee in the freezer. Then apparently there's people who say, “Well, you could put coffee in the freezer if you did it in this way,” which is what I've been doing when I have an abundance.

So, this is coffee from the freezer, and it is surprisingly good given its age, and I've shared it with a colleague, and she also, I gave her the warning. I'm like, “Now, this old coffee.” She says, “No, this is actually good.” Okay, there you go. So, that's the coffee. How do you find the connection to today's topic?

Tom: Well, I'm not sure I'm finding a connection to the coffee, but I'm finding a connection to your mug, potentially. [Laughter] The technology aspect of the mug.

Kelvin: That's right, that's right.

Tom: So, I don't know. Maybe elucidate me on what you're thinking.

Kelvin: Yes.

Tom: As far as the coffee connection.

Kelvin: Well, I threw in the mug for bonus just for that reason. Callbacks to the past, and pros and cons, positive benefit of something like your freezer coffee tactic. But people might be questionable, and it might not work. So, that was my shot at it.

Tom: Okay. I think I get it. So, we're going to be talking about pros and cons, more cons, today of certain kinds of assistive ed tech solutions. I get it. I think I get that connection. So, you want me to talk about how this connects to some previous podcast episodes that we've done?

Kelvin: Sure, yeah. [Sips coffee]

Tom: So, as you've hopefully noted in our prep notes here, in the past, we've talked about different kinds of technology tools that bring different kinds of automation to the work that we do in digital teaching and learning. So, for instance, here are just a couple episodes if you wanted to check those out, different aspects. I don't think these are repetitive. Episode 81 was “Technology Adoption as Enabling the ‘Right Thing,’” and we talked a lot about proctoring in there. Episode 117,

“Online Discussion-Enhancement Tools,” Episode 119, “Chatbots in Online Education,” and then a fairly recent one, Episode 131, everyone's favorite subject: “ChatGPT: Friend or Foe?” Yeah, it's a tool. Can be used for good or ill, right?

Kelvin: That's right, that's right.

Tom: So, we'll put some links in the show notes for that. But you want to explain how we're going to evolve from those previous conversations into something maybe a little bit different?

Kelvin: Today, we thought we'd maybe zoom in on discussing the potential of unintended negative consequences of these kinds of technological automations within digital teaching and learning. In fact, back in one of those episodes that Tom just mentioned, Episode 119 about chatbots, we quoted Dr. Jutta Treviranus from Ontario College of Art and Design University, who was the recipient of the 2022 Women in AI Award as saying on a call that I had had as a board where we were both present. She said, "Disability can be defined as 'a divergence from the average.' More automated approaches tend to amplify the average," with the implication that that could then disadvantage folks who are on more of the trailing edge of the bell curve, so to speak.

So, in that episode, we also then noted, “the greatest good for the greatest number” is an understandable principle that is a hallmark of our digital teaching and learning work. But we got to guard against the systematic exclusion of any human differences as we carry out our work. And that's really the direction we want to explore in looking at the potential unintended negative consequences of automations within digital learning. There's all kinds of automations.

Tom: Obviously, there's an awful lot of talk about ChatGPT, and I've read an awful lot about bias built in there, whether it's perspective, just the data set it's pulling from is limited. Political bias I've read about in ChatGPT. So, there's all kinds of bias that's in there. So, in the work that we do with students, how do we ensure that we are doing the greatest good for the greatest number, but we're not leaving behind the individuals who might be a couple of standard deviations off that mean?

Kelvin: That's right. Yeah. I hope this comes through as a theme. It sure feels like it has been. When we talk about things anywhere close to this, and maybe even in online teaching or learning in general, Tom, I feel like we've tried to lean into the human that this is not about, we even joked about in that ChatGPT episode, you could have the discussion forum created by generative AI, and then feedback given by generative AI, and how horrible and dystopian that would be. But we will lean into human affordances, right? [Laughter]

Tom: Right. We just get our popcorn and watch the machines talk to each other, I guess. Yeah.

Kelvin: That's right. [Laughter]

Tom: Yeah.

Kelvin: That's right. Back in one of those episodes that you mentioned, Episode 81, I threw in something there. I went back and looked. "One way of examining critically the tools available to us as online learning leaders is to consider what 'right thing' those tools enable or prevent. So, part of our job is to identify and promote right things from both... a strategic and a moral perspective." So, that's just the flip side. What we're talking about today is just the flip side of the ups, of the benefits. We're looking at how do you mitigate the negatives.

Tom: Yeah, it puts me in mind of something that we've worked on here when it comes to our learning analytics initiatives. I even think back to the work we've funded for a senior researcher from our data mining institute to look for patterns and create a predictive model, so that we could anticipate the students who might need extra supports, supplemental instruction, whatever it might be. At the end of the day, and I think I've talked about this in the past on this podcast, but what he found was that the biggest predictor of success was actual performance. So, it was like GPA was a predictor of success. How you performed on graded activities in the course was a predictor of success. What was a less strong predictor, and in some cases not a predictor at all, were things like ethnicity, socioeconomic status, transfer versus first time in college, gender. All of that, even number of logins. None of it was all based on actual performance.

Which was somewhat reassuring to me because I've been part of groups like a committee here that was commissioned by the provost at one point to talk about the ethical use of student data. It was run by our institutional research office or chaired by them. We talked about things like, we want to help students, and we know that students from lower socioeconomic backgrounds tend to come in less prepared or struggle. So, let's help them by putting in predictive kinds of solutions. But you're making assumptions that are based on a bias before the student even has a chance to show you what they can do, and it could potentially be detrimental to some students if you're not really careful. You don't want to make assumptions or profile students based on some ethnic, or demographic, or socioeconomic characteristic. You want to judge them based on what they did. So, I was gratified by that research that was resulting, but it's a snarly thing is, I'll use a Kelvinism, because you want to help, but you also want to be really careful about how that data could potentially be misused intentionally or not.

Kelvin: Yeah, I think that's well said. So, the fact that that researcher went looking to see, okay, here's what we found, but we went looking because there's hypothetically vulnerable subpopulations, so let's see if they've been disadvantaged. And that's true. I think even back in the early days at UCF with Dr. Chuck Dziuban, Dr. Patsy Moskal, early on in that impact evaluation work for distributed learning, are any vulnerable student subpopulations disadvantaged by online courses, let's say?

And indeed, no. Since then, in subsequent years, we've seen through publications from our colleague, Dr. Tanya Joosten and others, we see that there are benefits quite often to intentional design as enacted in online courses for everybody, but especially folks who come in maybe under prepared, or from those vulnerable subpopulations, they got built in scaffolding. But I think you're right. Don't just assume, and do no harm.

You mentioned online proctoring and some of the documented, now you mentioned, I think you mentioned generative AI and some of the documented bias there. There are some studies that I've seen. There's one that came out of a research group here at U of L with automated proctoring and some built-in bias there. We'll put a link in the show notes. But even things like you were talking about analytics, the almost embedded analytics inside of our tools that give us transparency as instructors about, "Well, what's going on?" Summarizing. There's an Oxford University affiliated research group that has a line of inquiry around the ethics of algorithms that's worth exploring. We'll throw that in the show notes.

And to that point, just yesterday, I kid you not, we have a faculty professional development cohort going on right now here, and I saw one of those facilitative messages from one of our team members, Program Manager, Robin Zahrndt, here, who was just so transparent modeling the thoughtful use of that embedded stuff. I'll just give you a taste of it. She says, "The LMS uses some AI algorithms. And while this information is helpful, I want to point out that language-based AI algorithms can present some equity concern, especially for English language learners and students who are not familiar with traditional academic speak." So, she's very transparent about what's available, and what she's doing with it. And I thought, well, that's helpful. It pops the hood metaphorically for the faculty who are coming into this experience about the support of teaching through these digital tools and modeling that. Nobody wants to be othered.

Tom: Yeah, that's interesting. Especially as we look at AI and generative AI is not going away. I think you and I are both of the philosophy, they lean into it, teach students how to use it effectively because it's going to be part of their professional workspace going forward. So, there's this whole new profession of Prompt Engineering, and apparently, it's highly paid because if you can save some corporation several millions of dollars a quarter, then you are worth several hundred thousand dollars a year to just figure out how to prompt the AI to maximize efficiency. But if you're not a native English language speaker, tuning that prompt and being able to extract all the possible efficiency that you can out of those AI whatever the results might be, it's going to be really hard. So, I can see that. That's interesting.

You mentioned proctoring, and we might as well go there. Because that is the area where there is so much, just get on Twitter and start #proctoring or something and look at the discussions. There is an awful lot of hatred for online proctoring tools, especially those that are automated. Because in the past, as the technology has

evolved, there have been some really unsavory sorts of results. I'm not saying any of it was intentional, but just because the technology wasn't up to speed yet. So, for example, there have been things where people with darker skin were unable to be recognized properly by the video algorithms. That's not good. And how do you mitigate just a normal somebody looking off screen, or scratching their head, or something, and now suddenly they're flagged for potential cheating. That's not cool. It creates such test anxiety for some students that they can be even debilitating, I know for some.

So, that's an area that I think we're going to have to get our arms around. I don't think online proctoring is going away. So as, I think, you said mitigate some of these issues to ensure that it's being used ethically and effectively.

Kelvin: Yeah, no, that's right. Well, you talked about prompt tuning, prompt engineering, whatever. I think one commonality to all of these automated, maybe we could even say, algorithm based, or algorithm influenced tools, is the training. What's the basis? This is back to Dr. Jutta Treviranus's quote that I shared, that we tend to look at the average or whatever, and then deviations are left out. So, how do we widen the scope of the training inventory database, whatever, so that we don't have that?

This research group here that did this study about online proctoring, they included recommendations, practical ones, like you should be dialed in, have better lighting, for instance. But then also, there were recommendations back to the designers to widen the scope of their training population, if you will. And yes, that study here, it was actually interesting. It was both darker skin tones, and the intersection of that with people who were female students, and dark-skinned were even more disadvantaged than darker-skinned male students. So, weird stuff. It was a very detailed study. It was really very interesting.

Tom: That's amazing. Was it isolated to a particular platform, or was it across multiple proctoring platforms?

Kelvin: They zoomed in on one platform at use at this university at the time. But again, then you run across similar reports with others. So, it's just tough to attend to. And I think that's maybe a principle, and I think we've talked about this a little bit in the past. Transparency, just be like, how does this all work? And then how do we see how we're doing? And if we're surfacing issues, then we've got to be able to address those issues.

Tom: Yeah. Well, algorithms are written by humans, and humans have flaws and biases. Even unintentional consequences come that it's not something anybody wanted to happen. It's not an online learning or even an educational example. But the one that I've heard about in the past when it comes to self-driving cars is that some of the early technology, the video in those cars, and the radar, and everything that recognizes objects, and directs the cars didn't recognize people with darker skin as

a human walking across the road. Well, that's bad. Not only is it bad from an ethical... Just optics, again, it wasn't intentional, I'm sure, but it was bad. It needed to be fixed, but it's also potentially life-threatening, right? So, that's super bad, right?

Kelvin: Super bad.

Tom: So, that's just a really extreme example, but we are talking about people's educational lives here. So, how can we ensure that we have the same safeguards built in that we're not leaving behind how you began, the individuals in the service of the many?

Kelvin: Yeah. What we've been really talking about today is these inadvertent systematic problems built in. But of course, maybe we'd be remiss if we didn't acknowledge because I think part of what you said there takes me this way. Ill-informed use of these tools as well. Just this morning on Twitter, I collect various tweets and tweet threads related to generative AI. There's a lot of faculty doing really, I think, creative, thoughtful things. And I'm like, ooh, let's collect that and share that out when I run across colleagues. But then sometimes I collect the opposite of that, and I'm sure it's going to be in the news if it's not already. A faculty member, I won't beat up on their university, but a faculty member, perhaps ill-informed, took student work, as he said, ran it through ChatGPT, and ChatGPT said, "Yeah, I wrote this." So, he failed a group of students, kept them from graduating, and then a lot of folks have been taken his email, ran that through, ran that through ChatGPT, it says, "Yeah, I wrote that, too." [Laughter]

Tom: Yeah. Well, you know...

Kelvin: Ill-informed.

Tom: Chuck Dziuban and I have done some stuff where we put ourselves in like, tell me everything. "Tell me top 10 Chuck Dziuban quotes, or what he did. Tom, you have to put top 10 Tom Cavanagh quotes anywhere." [Laughter] So, I've done that just as a goof, and ChatGPT is a dirty, dirty liar in my experience.

Kelvin: A dirty, dirty liar. [Laughter]

Tom: It's true! So, some people know that I have a sideline hobby of writing novels, and I asked ChatGPT, "Who wrote these novels?" and I listed the novels that I've written, with the publication dates and the publishers, and it comes back, and it says, "Michael Connolly wrote those novels." I'm like, well, no, I wish I was Michael Connolly as my mystery career. But no. Then I say, "No, that's incorrect." I'm... like, "Oh, I'm sorry. Robert Crais wrote those notes." "No!" and it just... [Laughter]

Kelvin: There's a Jon Lovett skit in there somewhere. [Laughter]

Tom: That is so true. That is so true. Yes. Yeah, my girlfriend, Morgan Freeman. Yeah, Morgan Freeman. Morgan Fairchild. Oh my gosh. [Laughter]

Kelvin: That's even funnier. [Laughter]

Tom: That is even funnier. That's exactly what ChatGPT would say. ChatGPT would just lie about it. And it lies with absolute confidence.

Kelvin: Yes.

Tom: Morgan Freeman, where did that come from? That is so funny. But all of that to say that that is a really bad practice what that faculty member did, because ChatGPT isn't going to tell you if it wrote it.

Kelvin: No, no.

Tom: It'll tell you what you want to hear.

Kelvin: Yes.

Tom: And there are tools, and we talked about that in that episode that will give you a percentage of how likely was this thing written by a bot? But that's not what ChatGPT will do. So, I think that is really a bad practice, if I could be so judgmental, and not a good use of AI technology and certainly harmful for students.

Kelvin: Yeah. Do no harm. Do no harm. I remember years ago, we've had primitive rough-hewn kinds of, I don't even know that I really want to call them data. Information, something in the LMS for decades. And I remember working with some faculty colleagues like, "Oh, it'll tell me if the students have logged in or if they've read this thing." Yeah, don't trust that. Look at it if you want, but don't base grades on that. Don't take some indication on a dashboard, especially back then, and make any interpretations that are going to negatively affect students with it. No, no, no, no, no. By the way, here's a bonus one. You mentioned the generative AI detection tools, ran across; we'll throw it in the show notes. Pre-publication study, those tools disadvantage non-native English speakers. So, there are traits in the writing of non-native English speakers that give more false positives that like, "Oh yeah, this was generative AI."

Tom: Yeah. I've also heard several anecdotes even from faculty member here, that they put an old paper that they wrote, an old chapter, an old journal article or something in to one of these detecting tools, and it comes back, oh yeah, this was written by a bot, and clearly it was not. It was written 10 years ago by a faculty member for some academic journal. So, those are not a hundred percent reliable.

And got to be careful in the accusations you're leveling at students based on some of those reports. The technology probably will improve, but so will the AI, the generative AI.

So anyway, it's a new frontier and just like any new technology, I think we are going to have to make sure that we balance the ethical use of it along with the blue-sky opportunity that is in front of us, keep ourselves grounded a little bit while we continue to try to take maximum advantage of these tools.

Kelvin: Yeah, I think that's right. Maybe two quick additional shout-outs as we begin to think about wrapping up. One, our dearly beloved producer, Tim, shared with each of us this New Yorker piece, which I thought was quite thoughtful about AI. And there, I think, a big idea is we can sometimes look at these tools very naively, almost like magic. There, the author of that article, talked about genies, or it's like magic, and that's dumb. We have to understand a little bit about what's going on. As you said, the technical euphemism these days is "hallucinations" of the generative AI writing tools because all they're doing, I try to ground myself in this. We've been for several years, you're typing something, and you go, "I hope this note finds you well," because it's anticipating what you're going to write next. You're like, "Don't help me." But sometimes you go, "All right, fine. Help me."

That's all this is doing. So, it doesn't know anything. And with the exception of connecting generative writing AI to active web data, like Bing integrated with GPT4 or something, it doesn't know anything. It doesn't have access to anything other than just what you're typing in there. It's just really good at predicting; the algorithm is anticipating what you're going to write next. That's all it's really doing. It's just it's helpful. So, you have to understand that to know it isn't magic. "Run it through ChatGPT" isn't doing anything. It's not making any determination, really.

Tom: No, and the thing that just, I said it before, but I'll just underscore it that... So, hallucinate is a fun word, I call lie.

Kelvin: I know, I know. [Laughter]

Tom: So ChatGPT will not just lie, but lie with just such bold confidence that it's almost compelling, and endearing, and it's just a complete lack of any credibility. [Laughter]

Kelvin: Well, we're just anthropomorphizing the heck out of this whole class of technology.

Tom: We are, aren't we? Yeah. [Laughter]

Kelvin: That's right. The other shout out is when we listed these past episodes that we're doing callbacks to, I neglected one, and that was back in Episode 46,

“Empowering Humans Through Learning Analytics.” What makes that come to mind is we pulled together some resources back then about the thoughtful use of data and analytics. I think you knew something more about this than I did at the time, like the Asilomar Convenings, and so forth. Some of that stuff, useful to go back and look at, and frame our use and understanding of these automated tools, I think.

Tom: Yeah. Yeah, that was actually a really good discussion with a lot of really, I thought, thoughtful people across the space.

Kelvin: Yeah.

Tom: All right, cool. You want to land the plane?

Kelvin: I will try to line it up and get it on the ground. So, we might say that a key aspect of what we call digital teaching and learning is the use of technology tools to support effective teaching and learning. However, usage of these tools without understanding and mitigating the potential unintended negative consequences is there's a lot of words that generative AI could anticipate what I might write next. You can fill in the blank. What is it? Doing this stuff without understanding and mitigating is... We'll just leave it at unwise. It's unwise.

Tom: It is unwise. Yes. I think that's fair. We need to keep all that in mind as we go forward.

Awesome. Well, a very thoughtful conversation, and probably one we may revisit in the future as technology continues to advance.

Kelvin: Mm-hmm. Mm-hmm.

(Outro music)

Tom: So, Kelvin, until next time. For TOPcast, I'm Tom.

Kelvin: I'm Kelvin.

Tom: See ya.