Jeff Cobb (00:00):

AI has the potential to do much more harm than a single biased teacher, but it also has the ability to do great good. If we can scale unbiased AI, it can help learning be the great equalizer that it's often been held up to be in the past.

Jeff Cobb (00:17):

I'm Jeff Cobb.

Celisa Steele (00:18):

I'm Celisa Steele, and this is the Leading Learning Podcast.

Jeff Cobb (00:27):

Welcome to episode 268 of the Leading Learning Podcast, the halfway mark in our series on the frontiers of learntech. Celisa and I set up the series in episode 265 talking about the big-picture frontiers of learning technology and touching on artificial intelligence, virtual and augmented reality, data, the learning ecosystem, personalization, and the increasing need for speed. Then we aired two interviews.

Celisa Steele (00:57):

In episode 266, I spoke with Donald Clark, learntech entrepreneur, CEO, professor, researcher, blogger, speaker, and author of *Artificial Intelligence for Learning*. Donald draws on his almost 40 years' experience working with and on learning technology in our conversation.

Jeff Cobb (01:16):

And, in episode 267, I spoke with Sae Schatz, director of the Advanced Distributed Learning Initiative and editor of the e-book *Modernizing Learning: Building the Future Learning Ecosystem*, published by the ADL Initiative.

Celisa Steele (01:31):

Still coming in the series are conversations with Sam Sannandeji, founder and CEO of Modest Tree, a simulation company that develops augmented reality and virtual reality training, and Ashish Rangnekar, co-founder and CEO of Bench Prep, a learning company that makes a platform designed for credentialing bodies, associations, and training companies. Before those conversations with Sam and Ashish, though, we want to focus this episode on some themes that emerged from the conversations with Donald Clark and Sae Schatz. Specifically we want to focus on bias and equity.

Jeff Cobb (02:13):

Bias and artificial intelligence or, more specifically, bias *in* AI is not a new concern. But it's one that's been garnering more and more attention in recent years, and it feels appropriate to focus on bias and equity in learntech at this moment because that focus is in step with many social justice movements we're seeing in the United States.

Celisa Steele (02:34):

We just watched *Coded Bias*, a 2020 documentary on Netflix, and it's very social justice-focused. The film investigates bias in algorithms, and it features the work of MIT Media Lab researcher Joy Buolamwini, who uncovered flaws in facial recognition technology—meaning the technology was really good at recognizing the faces of white men. It was less good with the faces of women and people of color. Because of her work, Google and other tech companies have worked to improve their AI, and it's gotten better at recognizing faces of all types.

Celisa Steele (03:13):

Joy founded the Algorithmic Justice League, which is a really interesting project that "combines art and research to aluminate the social implications and harms of artificial intelligence. AJL's mission is to raise public awareness about the impacts of AI, equip advocates with empirical research to bolster campaigns, build the voice and choice of most impacted communities and galvanize researchers, policy makers, and industry practitioners to mitigate AI bias and harms." That's from the AJL Web site, which we'll be sure to link to in the show notes for this episode.

Jeff Cobb (03:53):

And Cathy O'Neil is also featured in the *Coded Bias* documentary. Cathy wrote *Weapons of Math Destruction: How Big Data Increases Inequity and Threatens Democracy*, which was published in 2016. In the documentary, Joy and Cathy focus on the use of AI in policing, surveillance, credit and lending decisions, insurance, advertising, and more. At one point in the film, Cathy O'Neil says, "People are suffering algorithmic harm." And that's a turn of phrase that just really struck me.

Celisa Steele (04:27):

Yeah, both Cathy and Joy are focused on the harm. *Harm* is in that quote from Cathy, and the mission of AJL also mentions harm. So they have real concerns, and there is real reason for their concerns. Lost opportunities in accessing money through lending, greater likelihood of being stopped by police, higher interest rates, etc. There's enough harm and enough real instances of it that many are clamoring for legislation and regulation and standards. And in fact, as we're recording, the European Commission is expected to unveil a proposal on artificial intelligence regulations in the European Union this week.

Celisa Steele (05:12):

One concern covered in *Coded Bias* involved a teacher in Houston who'd won numerous teaching awards over many years, but then he received a poor evaluation when the district implemented an algorithmic approach to assessing teachers. And he and some other teachers sued and part of their argument—and, by the way, they won the case—was that they didn't know why they'd gotten the poor evaluation. The algorithm was a black box that they couldn't question, and so they couldn't contest the result because the premises for the result weren't known.

Jeff Cobb (05:45):

And I find that black box argument really interesting and maybe a little bit scary as well. We use a lot of technology we don't understand. My laptop, my smartphone, Google Search—I have a pretty crude understanding of how all of those work, but I certainly don't want to give any of them up. And now, recently, I came across what I found to be a really helpful distinction from Christopher Penn, and this was in a Marketing Over Coffee podcast episode. He says when we want to understand how software arrived at a particular outcome, we choose between explainability and interpretability.

Jeff Cobb (06:24):

And he says, "Interpretability is the decompilation of the model into its source code. We look at the raw source code used to create the model to understand the decisions made along the way. Explainability is the post-hoc explanation of what the model did, of what outcome we got, and whether that outcome is the intended one or not." And that's quoting Chris in saying that. And he uses an analogy to make explainability and interpretability a little more digestible. So explainability is tasting a cake. We can taste it and get a general idea of what went into making it.

Jeff Cobb (07:05):

We might not get 100% of the ingredients right—is that vanilla extract or almond extract that I'm tasting?—but it's a fast and easy way of testing. Now interpretability, on the other hand, is looking over the recipe for the cake. We look at the list of ingredients and the steps, and that allows us to verify that the recipe makes sense and the ingredients were good. And this is a more rigorous way of validating our results, so it makes sense, particularly in high stakes situations. If someone has a severe allergy, for example, if harm could come from eating the cake, then we want interpretability not just explainability.

Jeff Cobb (07:46):

But if the stakes aren't that high, explainability usually is the go-to. Interpretability can, well, really slow things down. It's costly. It's operationally difficult to do a thorough review. "For more complex systems like neural networks, interpretability is a massive, expensive undertaking. It slows down systems like crazy, and in most complex models we might never fully unravel all the details. Interpretability also reveals the secret sauce, to the extent that there is any, in the process."

Jeff Cobb (08:20):

And, again, I'm quoting Chris there. So AI software makers don't really want interpretability, at least not publicly available interpretability.

Celisa Steele (08:29):

But if you're that teacher in Houston, whose job is suddenly in jeopardy, you want interpretability not just explainability. Or if you're not getting job interviews because of AI screening, or you can't get access to a loan because of your zip code. I mean, these are places then where you're going to want that interpretability. And there's a power differential that comes into play, and this came up in *Coded Bias*. I think it was Cathy O'Neil who made the point that it's really, really hard as an individual to push back against large-scale, AI-driven decisions because a lot of them are invisible. They're happening in that black box.

Celisa Steele (09:10):

And so many people who are concerned about bias in AI are really focused at this point on simply making the issue known. They're wanting to call our attention to these often invisible systems. They want to raise awareness of the potential for misuse, whether that misuse is intentional or just incidental.

Jeff Cobb (09:39):

Coded Bias and Joy Buolamwini and Cathy O'Neil focus on the potential of AI and algorithms to create or continue harm. But in your conversation with Donald Clark, Celisa, he said he feels like the bias and AI discussions lean too hard to that harm side.

Celisa Steele (09:57):

Yes, Donald makes the point that calling out bias in AI is problematic because it's not like the alternatives are bias-free. I mean, if we take the case of learning, human teachers and facilitators are rife with bias. So eliminating AI does not, will not eliminate bias. And I actually think Donald and Joy agree on this point. Joy says in the documentary that "The past dwells within our algorithms." So they both acknowledge that past and current biases are reflected in AI and its algorithms. Donald just wants to make sure that the baby doesn't get chucked with the bathwater because there's bias everywhere and because the current educational system isn't working.

Celisa Steele (10:44):

In fact, he calls the current educational system "far too expensive, clumsy and slow." So Donald actually sees hope in AI because it can be taught and audited, perhaps more effectively than humans so that, over time, hopefully we catch and remove biases.

Jeff Cobb (11:02):

And in the meantime, of course, we have to act carefully since AI can scale bias. AI has the potential to do much more harm than a single biased teacher, but it also has the ability to do great good. If we can scale unbiased AI, it can help learning be the great equalizer that it's often been held up to be in the past though, as Donald points out, it's often actually not an equalizer, but something exclusive. It's not cheap enough or fast enough to be equally useful and accessible to all.

Celisa Steele (11:35):

So if we were to put this on a Pollyanna-to-Doomsday spectrum, it seems like we're somewhere in the middle.

Jeff Cobb (11:42):

Yeah, that sounds fair. I mean, artificial intelligence feels like a both/and at this point. It has dystopian possibilities, but it has utopian possibilities as well.

Celisa Steele (11:53):

Tied up in Donald's baby-and-bathwater comment is some idea of the risk involved. In fact, he used cars as an example. Tens of thousands of people die in car crashes in the U.S. every year, but we still drive cars. We're not talking about banning the use of cars. We've collectively concluded that the good outweighs the bad. The same is likely true for AI. We won't ban it. I mean, who wants to get rid of Google Search? But what rules, what regulations, what speed limits do we need in place to make it as safe as possible?

Celisa Steele (12:33):

I'd argue that AI for personalizing learning seems at the lower end of the risk scale, especially if the AI is not a gatekeeper giving access to content to some and keeping others out but is more of a guide on the side. It's recommending and trying to help learners find the more relevant and useful content.

Jeff Cobb (12:58):

This reminds me of part of my conversation with Sae Schatz. She homed in on the fact that we can't be satisfied when we can say everyone has Internet and a computer—that is, the tools for access aren't enough, in and of themselves. For equity in learning, everyone needs access to high-quality opportunities and experiences. And that's something that AI can help with if it's done right. If it's not, then we run the risk of exacerbating existing inequalities and creating what Sae called "hidden haves and have-nots."

Celisa Steele (13:33):

In his book *Artificial Intelligence for Learning*, Donald Clark points out the risk in AI, specifically for learning—because so much of the focus around the risk in AI is in these other domains, in policing or in lending and those areas. But this is what Donald Clark has to say about the danger of AI. He says, "The danger is that AI could deliver narrow, deterministic, prescribed pathways, not allowing the learner to breathe and expand their horizons, and apply critical thought." So, he continues, "We need to be careful that the learner retains the curiosity and critical thinking necessary to become an autonomous learner."

Celisa Steele (14:16):

But he also points out that "The degree to which human agency is included in AI-driven systems is a design issue." So it comes down to the human designers. Do we design AI as a guide and a nudge with lots of room to still explore or even ignore recommendations? Or do we make AI more of a gatekeeper with tight control over access to learning resources? As long as AI is a guide and not a gatekeeper, the risk of getting AI wrong feels minimal, and the potential for it right seems huge.

Jeff Cobb (14:53):

And if you're looking for a partner to help realize your learning business's potential, check out BenchPrep, our sponsor for this series.

Ashish Rangnekar (15:01):

BenchPrep is a pioneer in the modern learning space. We have been digitally transforming professional learning for credentialing bodies, associations, corporations, and training companies for over a decade. With an award-winning, learner-centric, cloud-based platform, BenchPrep enables learning organizations to deliver the best digital experience to drive learning outcomes and increase revenue. The platform's omni channel delivery incorporates personalized learning pathways, robust instructional design principles, gamification, and real-time analytics that allow organizations across all industries to achieve their goals.

Ashish Rangnekar (15:42):

More than 6 million learners have used BenchPrep's platform to attain academic and professional success. BenchPrep publishes regular content sharing the latest in e-learning trends. To download our latest e-books, case studies, white papers, and more, please go to www.benchprep.com/resources.

Jeff Cobb (16:12):

In *Coded Bias*, Joy Buolamwini says, "The past dwells within our algorithms." The good and the bad of our past are on display in the algorithms and in AI. And the past is there because AI needs data to focus. That means the datasets used to teach AI can be troublesome. Gaps in the data or over-representation by particular groups can skew results. So even if the algorithm is unbiased, the data might be biased. We have to look for and audit for bias in the data and in the AI processes.

Celisa Steele (16:48):

If you think about the term *systemic racism*, and the fact that learntech is made up of systems, like learning management systems, right there you see the potential for bias and injustice in learntech. I'm also struck by how we've heard this one point from really everyone we've talked to for this series. And that point is that data is the lifeblood. Data is what makes artificial intelligence, personalization, recommendations, it's what makes all those things work. As one example, here's what Celeste Martinell, vice president of customer success at BenchPrep, says about the importance of data going forward.

Celeste Martinell (17:31):

When you think about trends, it's hard to talk about trends in learntech and not talk about data intelligence and the positive impact that will have on learners and learning organizations. I think we'll start to see learntech move beyond just pulling user and content data together to predict learner outcomes. And we'll really start see the entire learning ecosystem brought together to tie learners' in-platform performance, if you will, to its out-of-platform outcomes. So really looking at a learner's performance and evaluating it against their job or career trajectory, for example.

Jeff Cobb (18:05):

Data was a topic you and I spent some time on in episode 265, Celisa, but data was a refrain that both Donald and Sae returned to again and again. Sae mentioned the conventional wisdom summed up in that cliché that data is the new oil. Now oil is arguably a 20th-century point of comparison. A more 21st-century view might argue that data is the new solar or the new wind. But the point is that data is so important and valuable, and it's necessary for powering and enabling other types of activity.

Celisa Steele (18:40):

Donald Clark outlines four levels of use of data. The first is describe, then analyze, predict, and prescribe. The levels move up in terms of difficulty. Using data to describe who completed which course and when is relatively easier than getting to that fourth level, prescribe, where we're using data to help us understand not what's happened but what *should* happen. What should this learner study? So getting into recommendations and true personalization.

Jeff Cobb (19:14):

And those levels of data remind me that AI is playing quite a significant role in marketing today. Marketing is pretty good at getting beyond those lower levels of describe and analyze and into those upper levels, predict and prescribe. I feel like many learning businesses are still in the describe and analyze levels with their learning even while marketing is further along. And I know our listeners have to market their learning offerings, so marketing's use of AI is important and relevant in its own right but also because, as we've pointed out before, what happens with martech is so often a bellwether for what happens with learntech.

Celisa Steele (19:55):

I think that's a great point, and it fits with Donald Clark's assertion that consumer tech drives learntech. It's also a point that Joe Miller made when I spoke to him. Joe is vice president of learning design and strategy at BenchPrep, and he says, he likes to look at the other "x" techs and what they're doing—so fintech or retail tech—and look at how they're approaching problems and opportunities because that can be instructive for organizations looking to get the most out of their learntech.

Jeff Cobb (20:34):

Bias and equity are huge societal issues, and learning has a role to play, hopefully on the side of reduced bias and increased equity. To help make sure you're on that right side, think through where and how you're learning business uses or would like to use artificial intelligence and other automated technologies. How risky are the ways in which you use or plan to use AI and automation?

Celisa Steele (20:59):

More risky ways would be tied to when the technology is a gatekeeper, when it allows access to some and keeps others out. Less risky ways would be where these technologies make suggestions and recommendations to learners that can be acted on or ignored. For each of your uses of automation and AI, is the technology interpretable or explainable? And does that match the risk? For riskier uses that might be tied to job promotion, salary increases, etc., you'll want to lean towards the interpretability side, being able to check out the cake recipe, so to speak.

Jeff Cobb (21:40):

These are definitely difficult questions to answer, but they're important to explore and reevaluate as your learning business gets more into AI and other aspects of learntech that automate more.

Celisa Steele (21:52):

You can find show notes for this episode at leadinglearning.com/episode268, along with a transcript and a variety of resource.

Jeff Cobb (22:00):

At leadinglearning.com/episode268, you'll also see options for subscribing to the podcast. To make sure you don't miss the remaining episodes in this series, we encourage you to subscribe, and subscribing also helps us to get some of that all-important data on the impact of the podcast. Data is, of course, the lifeblood of podcasts too.

Celisa Steele (22:21):

We'd be grateful if you would take a minute to rate us on Apple Podcasts. Jeff and I personally appreciate it, and reviews and ratings help the podcast show up when people search for content on leading a learning business. Go to leadinglearning.com/apple to leave a review and rating.

Jeff Cobb (22:37):

Lastly, please spread the word about Leading Learning. In the show notes at leadinglearning.com/episode268, you'll find links to us on Twitter, LinkedIn, and Facebook.

Celisa Steele (22:49):

Thanks again, and see you next time on the Leading Learning Podcast.

[music for this episode by DanoSongs, www.danosongs.com]