Brenna Clarke Gray:

Hello and welcome to You Got This!, a podcast about teaching, learning, community, conversation in your digital life, made for everyone at Thompson Rivers University. I'm your host, Brenna Clarke Gray, Coordinator of Educational Technologies, and this podcast is a project of your friends over at Learning Technology & Innovation. I record this podcast in Tk'emlups te Secwepemc, within the unceded, traditional lands of Secwepemcú'ecw, where I hope to learn and grow in community with all of you. Today I'm thinking about robots. I'm sorry, I swear last time this year, I promise. But robots, it is. Let's get into it.

So before we start, friends, apologies for my voice, which sounds much cooler than usual. I've been down with a head cold. The tests say it's not COVID. I don't know if I believe them, whatever, it is what it is. I am poorly. I canceled my interview this week on account of I didn't have a voice for several of the days and I was hoping that by the time I sat down to do this minisode, which is what this is today, by the way, a minisode, I would sound more human and instead I still have these kind of bullfrog vibes. So I hope you're enjoying them.

It's been a busy couple of weeks thinking about AI for me. In the last two weeks I've done a classroom visit about AI, a session on campus introducing our AI guidance document. I got to do a keynote about AI at Tech for Teaching Days for college in Ontario. I did a CELT Talk on AI and I've been planning programming like crazy for the fall term AI, AI, AI. It's all I'm thinking about. I increasingly feel frustrated about the quality of the discussion. That's not to say like I'm so great at having complicated discussions, but what I do think I offer as a strength in this space is a willingness to be a buzzkill. And if you came to my CELT Talk last week, you may have gotten those buzzkill vibes. Sorry. Not sorry.

I think there are a lot of really important conversations that we're not having about this technology. The conversation that we are having lots and lots and lots is academic integrity and why that's important. And you know what? I get it. I get why that's top of mind. It seems to be the thing that is changing our practice the most. We're all really freaked out about what this means for assessment design, what this means for the kind of work we want to do with students. I get it. I genuinely get it. I'm not naive, but there's not really an answer to that question that's going to be easy.

So something I've been talking about a lot at the sessions I've been at is that there's not really a technological solution to this. There's not like a Turnitin equivalent. I mean, Turnitin has developed an equivalent, but it sucks. Independent tests suggest it could be as much as like 90% false positives according to some analysis. There's just no way to ethically hand student work over to a for-profit tool like a Turnitin when we don't know if the results it's giving us are honest. That's a super troubling thing to do.

And I hate to be perceived as flippant when I say that the solution is assessment redesign because I know how not minor that is as a thing to say because assessment redesign takes resources and time and you should be paid for the time that you spend redesigning your assessments. And I'm not sure there's an appetite for that administratively anywhere in the sector. We need to be investing in those practices and I don't know when that's going to come.

I think that the best solution to the academic integrity "problem", I'm putting it in quotation marks, is a real back to basics approach. Like make sure the questions are really specific to the classroom experience, include reflection when you can. Have a good rubric that holds students to account to answer the questions that you are actually asking. Something I think about a lot when I see ChatGPT-generated essays in particular is just how easy it is when you're marking stacks upon stacks of essays. And this is a world I know so well. When I was a community college English instructor, it was like there were times in the semester when I was marking 120 essays a week or 120 essays every two weeks. So I get that pressure, I a hundred percent get that pressure.

I also remember how seductive good grammar is in that context, because you're working so hard, you're reading through difficult pros. Pros when you're learning how to write is hard. It's hard to write and it's hard to read. And then you come across a paper that was like grammatically perfect and it could potentially take me a few reads to realize like, "Oh, this student happens to have really beautiful grammar, but this is empty or this isn't really answering the question I asked." Or, "This is a generic assignment. Oh, I think this assignment's been handed in more than one class." That back to basics approach and trying to slow down so that we're not seduced by form because that's what ChatGPT and these tools are good at. They're really good at form, so we need to not be seduced by form. I get how hard that is, but I think that's ultimately all we can do with academic integrity for now. Because creating new policy around the tools that exist in this moment is not going to help us very much when the tools shift down the road.

The thing I talked about at my CELT Talk this week is that there are real questions about our values that are all wrapped up in generative AI. Issues like sustainability, like our understanding of intellectual property, like what it means to consent to have your work used, like how much fresh water these tools use and how much their carbon footprint is. What about the labor situation, the underpaid labor in the global south that makes these tools usable? This is why I'm a buzzkill because all of these things are so much more urgent to me than the unsolvable question of academic integrity. Because we have certain values as a campus community and I want us to talk more about how we reconcile using these tools that maybe actually contravene a lot of our fundamental values as a campus community. How do we reconcile that? I think that's a conversation we actually have to have.

If for example, sustainability is important to us, what does that mean when we look at the water use of a tool like ChatGPT? What does that mean when we look at the carbon footprint of a search done with AI versus without AI? And can we have an honest conversation about when we need AI and when we don't? Like you don't need AI for the average web search. It's just this glitzy thing that we're all head up about right now. But the standard search engine was really great at telling you whether or not you could bring a dog to a provincial park. You don't actually need AI to solve that for you, and you shouldn't look to generative AI to solve that for you because if it doesn't know the answer, it's just going to make it up. That is what generative means.

I'm nervous about how excited we are about the whiz-bang possibilities and the glitz and the glamour and not digging into the values issues that are at the core. People keep asking me what I think about academic integrity policy, and I don't think anything about academic integrity policy right now on this issue. I think the much more urgent component of artificial intelligence discourse at the university and in lots of other places is that we need to establish a disclosure policy. What are the parameters around which AI is a technology that we will use and how do we go about disclosing it? And that's not a policy that's just for students.

I want to know when Markham chooses to use AI. I want to know when a policy has been drafted through AI. I want to know what AI tools are at use in the registrar's office. I want to know all of those things. I want to know if folks are using AI to support their research. I want to know if they're using AI to support their teaching, if they're using it to generate multiple choice questions. And I want to know all that without a value judgment. I just want us to be able to have an honest conversation about this stuff.

Part of why I think disclosure is so important is because there's going to be really different disciplinary norms. The journal Nature has taken a strong stand AI cannot be your co-author. And if you use AI, you need to disclose it either in the methods section, in the acknowledgement section, but somewhere, you need to disclose how you're using AI. The reason why AI can't be a co-author by the way, is because AI can't be held accountable. Doesn't that seem really important to think about? As we lean into using

these tools, they can't be held accountable right now. That's wild to me when we're putting them in decision making situations.

I talk in my CELT Talk and I talk a lot about this thing I call the accountability gap. I keep waiting for someone to come up with a better name for it, but it's this gap that I see emerging where we look to AI to make difficult decisions for us so that we don't have to feel responsible for it. Most obvious example of this is police forces using AI facial recognition to try to catch suspects. You're putting so many layers of distance between the arresting officer and the decision to arrest. And I think that is really frightening, but I see that in education too.

Let's take algorithmic test proctoring as an example. The companies that make those tools, they're really clear that they don't take responsibility for the decisions. That they're just giving you information and professors are the final person who make the decision about whether academic integrity has been violated or not. And that's great language to use for sure, but I've seen how easily people defer to the suggestions made by machines. It's a function called automation bias. We really want to believe that machines know better than us and we really want to lean into their decision making. So how does that square with a tool's unwillingness to take accountability for the decisions that it makes, responsibility for the decisions that it makes?

I think this stuff is really important and I think talking about it honestly within the structure of the institution has to start with disclosure. There are ways in which we could mobilize AI to do great things for learners who need additional supports. Better transcription technology. Please, I would love AI to facilitate that. Automatic alt text on images. Yes, bring it on, bring it on and make it correct, but bring it on. I'm very sympathetic and curious about discussions where different kinds of processing disorders or just like anxiety, fear of the blank page can be overcome by some of these tools. I think those are all good conversations to have, but we don't have to pit those needs against the conversation about values. What we need to say is what's the use case where AI is worth the cost?

I actually think AI is worth the cost when it comes to water and carbon, if it's supporting a learner who needs support. I don't think AI is worth the cost to just mess around play. And you know me, I'm a big proponent of play, but did you know that a 20 to 50 query ChatGPT string costs 500 milliliters of clean water? That number blew my mind. I will share the reference in the show notes. I do not use or think about generative AI the same way since I learned that, and I hope you don't either. I hope it changes your perspective too. That's all I'm trying to do with these conversations.

I know it's not as fun as generating AI imagery and playing around with prompt engineering, but I actually believe that our values as a campus mean something. And I want to see how we take on this challenge because AI is a technology that fundamentally challenges a lot of what we say matter to us as a campus. The first step in addressing that is a disclosure policy where we talk about when AI is appropriate to use and when it's not, and what our expectations are around disclosure, as I say, not just for students but for the whole community. And from there we can start to talk about what use cases have value and which ones don't, and that's again, going to be a real conversation about real values. I actually trust us to have that conversation. We got to get over the whiz-bang first. Sometimes you got to listen to your friendly neighborhood buzzkill.

So that is it for season three, episode 23 of You Got This! Thanks for suffering through my voice with me. This is going to be an unpleasant episode to edit.

As always, if you want to write to us, you can email me. I'm bgray@tru.ca. And I'm also on Twitter @brennacgray. And in both cases, that's Gray with an A. All of our show notes and transcripts are posted at yougotthis.trubox.ca, and of course you can always comment on individual episodes there.

I'm going to leave you today with a Tiny Teaching Tip, and my Tiny Teaching Tip this week is that I want you to think about what language you're going to use about AI in your course documentation for next term, because I think you should have something in there. I'm not going to be prescriptive. It's going to be different for every discipline, whether playing with AI is appropriate or not, which assignments it would be appropriate for, which ones it won't. To what extent you see AI as a violation of academic integrity or not. These things are going to be different for everyone. But the critical thing is that just like disclosure, you've got to have the conversation. Start to think about your syllabus language and let me know what you come up with because I'm going to be working on some boilerplate to suggest, and I'd love to know what your starting places are.

Until next time, take care of yourselves and each other because the robots are not going to do it. We'll talk soon. Hopefully, I'll have something more like my normal voice next week.