

## The Case for Normalising GenAI use in Higher Education and Professional Practice

Interestingly, in my title, I talked about normalization and normalizing generative AI. And Nick asked me when he arrived, have we not normalized AI already? And the thing with technology is, once it's normalized, we don't talk about it so much. So the fact that we're talking so much about it will tell us that it's not been normalized yet. It's not something that we use every day, it's not something we think about. It's not something that we use without thinking about every day. It's still something that's in our consciousness. It's still something that takes time, takes effort, and we're still trying to understand. So we're in the process of normalizing, and at the same time, this technology is developing incredibly quickly, just like Nick mentioned, which makes that normalization process also more complex. But today I'm going to be looking at how we might understand normalization in the use of AI in higher education and also professional practices, in the sense that perhaps that's our goal, that we need to be seeing that as where we are going, so we can then manage that process for ourselves, but also for our students as well.

So I'm going to start with a question. Do you think you can detect your students' use of AI in their work? What do you think? So if you use AI to check AI, AI can tell you whether AI has been used, yes. So AI can detect AI, but can we detect AI? And that's something we also have to be conscious of as well. Only one of these is an authentic photo, and the concept of authenticity is also changing with the development of AI as well. So I asked several different models to generate a landscape of Hong Kong, and these are the images. The other one I took from a legitimate source, taken by a real human being. Can you see which one is a real image? See, you are unsure, right? So, can we detect AI in our students' work? We're struggling already, right? And this is an image of someone we know very, very well. You're right. It's C. C is the real image here. OK. But it's getting more and more difficult to tell. And if you ask, many different bots now to generate images, you could ask it to generate an image of a Hong Kong politician, for example, it will give you someone who looks very much like a Hong Kong politician, but isn't a Hong Kong politician, right? Equally, you could ask it to generate an image of Hong Kong in 1960 and it would very much look like an image of Hong Kong in 1960, but it's not Hong Kong in 1960, right? This is the complexity of the reality we're living in now, where this is blurring the lines and making things very difficult. And what's making it even more difficult is, I've generated these images, now I could easily upload these images to Google images as images of Hong Kong. So when you type in "image of Hong Kong", you'll get AI-generated images of Hong Kong, not real images of

Hong Kong, because the speed that these images can be created is way faster than real images are being put up. That's something I can talk with you more about later.

But of course, we can tell if they're used badly, OK? Here are some very famous cases of AI being used badly. These are two articles that have now been retracted from journals, one of which said, "Certainly, here is a possible introduction for your topic". This is an article that went through an editor, peer review, copy editors, and was published with that language in it. Another one said, "In summary, the management of... I am an AI language model." Right? "Real-time information." "Patient-specific data." This was an inclusion of a published article that went through editors, went through peer review, went through copy editing and was published. And this is an image that also appeared in an article that went through an editor, went through peer review, got published in an article. All of them were found to be AI-generated once they became publicly seen, and all of them were retracted. But we can see if they are used badly, like this, we can easily tell that AI is being used. However, if it's used well? As we mentioned, perhaps we don't have many opportunities. So we're seeing now that authors are using AI as part of their professional work, and they're using them in ways that they see as augmenting themselves and improving what they can do.

So this is an award-winning author that said she used AI to support her work. She used it in a minimal way, but she found that it enhanced what she was able to do. And a recent study that came out only about a month ago, less than a month ago, found that humans are finding it very difficult now to distinguish certain types of text as whether they're being generated by humans or AI. And this article even found that there was a preference for AI-generated poetry. OK. Maybe obviously, any of these studies can be picked apart. And, you know, how the poetry was chosen, and who the choosers are, but we can see now this kind of complexity of the situation we're now in. We're also in this situation where, yes, AI can detect AI, but in a very limited way. If you ask GPT4 to create a text, and they directly input that text into Turnitin, you're most likely to get picked up as AI-generated. However, if you manipulate that text slightly, adding, changing some of the terminologies, adding different aspects to that, then there's a good chance that it won't get picked up at all, and it will be seen as made by humans as well.

At the same time, when we say it can detect AI, it's detecting certain kinds of AI use, and the way we use AI is much more complex than just asking an AI tool to write an article for us. In fact, most of us in this room would rarely do that, but we will use it in certain ways that support what we do. So when we talk about AI and we say, "AI is used in this way", then we are forgetting all the other ways that it might be used as well. So in response to this emergence of these tools and the kind of complexity of their use, publishers, journals and universities have been advocating for the transparency principle, OK? So we can't detect it. It's being used in multiple different ways. So most of us, and the same with my institution and most publishers, and I'll show some evidence of this later, have been arguing that, okay, what we want from users is for them to declare that they're using these tools. So when they submit an assignment, they declare which tools they used, how they used them, why they used them. If we write for an article for a journal, most publishers now and journals say, "Did you use AI tools? How did you use them? Why did you use them?" to leave up to the judgment of the professor or the judgment of the editor to see whether those uses are legitimate or not. OK.

So sorry, the text is a bit small, but a study by Tang earlier this year found that, of 125 journal policies in the nursing field, about 27.6% explicitly require their contribution to declare the use of AI. My understanding now is that that's a much higher percentage. At least in my field of language education, most of the journals require some kind of statement of AI use. Similarly, in a study I did last year, we talked about a year ago now, actually more than a year ago when I did this study and collected the data, we found that more than half of universities that had policies around the use of AI in assessment had a requirement that students would declare their use in some way. So this is University of Toronto, saying that they need to say what tools were used, how they were used, and how the results from AI were incorporated in this work. And they've got UCL even asking students to provide the prompts they used, which, for me, suggests a particular kind of use of AI again. That we would go into an AI tool, prompting it to give us something, and then submitting that work, not thinking of the kind of complex array of tools that we are using when we're maybe engaging in a task, which is different from, "I use ChatGPT, and this is what I got", which is not what most kind of sophisticated use of AI would look like, you know, Nick showed a huge array of tools, they don't look like we interact with them in this way, right?

Guidance which most journals follow when it comes to the kind of ethical standards and when it comes to publishing and writing, has the same idea of this transparency and disclosure. So they argue that if we're using these tools, we should declare their use in the Methods section, how we use them, why we use them, what they add to the study, and also this idea of responsibility of the author for whatever they submit. OK? Our guidance at Baptist, these are the ones that Christoph mentioned that I was responsible for helping to draft and leading the task force that developed it. We also talk about the transparency needed here. OK. So I'm kind of driving this home, trying to drive this home as something that is kind of normal now. Okay, so we want to talk about normalization. Transparency is becoming normalized as a requirement.

So in our guidance, we talk about these four principles which guide our kind of AI policies and practices at Baptist. This idea of AI-Empowering that we want our students to be able to use these tools effectively, and what do we talk about effectively and capably in their knowledge creation and teaching and learning? AI-Critical is that they have healthy skepticism of generative AI tools and understand the strengths and drawbacks. AI-Ethical refers to maintaining integrity, transparent, ethical use. And then, I guess, what Nick mentioned about people-centric. We also talk about building and sustaining human uniqueness, so this idea is that we are central to this. So these are our principles, but transparency was a key aspect of this as well. So we require our students to declare their use and acknowledge their use. So in every assignment task that they need to submit, if it's done outside of proctored conditions, then they need to submit a declaration like this, very much like our old-style kind of plagiarism statements that we used to have, where they say, "I did or did not", and then there's a statement about it. Then they have to declare how they used it within the process, and then they have to acknowledge that they've kept record of it. There was debate about whether we asked them to submit all those things that they used AI for, but I was an advocate to say that that's not going to be workable, considering how they're using tools. But at the same time, do we want to store all that on our university servers? All those images and images and images and PDFs and PDFs and PDFs of their interactions, which is, you know, would have been a big data storage issue. So we argued against that, but ask the students to keep a record of it. If we pull them into our offices and say, "Did you really write this?" Right? OK.

But interestingly, even though that's our policy in our university, we're getting a lot of cases where students are not declaring their use and this seems to be something that's quite common within higher education sector and also academic publishing sector as well. There are parallels between the two. So non-compliance seems commonplace. A study done very recently by Gonzalez at King's Business School found that 74% of students failed to declare their AI usage, despite the declaration being a requirement. OK? And the reasons for this were fear of academic repercussions. They believe that if they declare their use, their instructors or their professors would penalize them for that. As some guidelines are ambiguous, they don't know what uses to declare. So if they go to ChatGPT and say, "I'm going to need to write an article about blah, blah. Can you give me some ideas?" Is that AI use that needs to be declared? Or is it, "I have to write this article. Can you write it for me?" Is that the use that needs to be declared? It was unclear in most institutions, for most students, about what kind of use to declare, how we talk about it, and how we use these tools to declare them. Like Nick mentioned, do we declare that we use spellcheck or we turned on Grammarly when we were writing and it was giving suggestions of the next word? Do we declare that or not? How much of this text is written by me and the AI at the end of it? How do I justify myself? And students will justify themselves in different ways. They'll say, "This is my idea, therefore I don't need to declare". Or they might say, "I just use it to proofread, therefore I don't need to declare". They will justify it because the guidance was unclear. Inconsistent enforcement is another barrier to compliance. You know, the AI detection tools are not fully utilized by everyone, and they're not, as I mentioned, they're not reliable. So some students might get pulled into the professor's office, and some don't. So I'll take the risk that I won't get pulled in because he won't pull in everyone. That's impossible.

Next, peer influence. Others are using it in certain ways. Therefore, why would I not? And we have cases. We did lots of town halls with our students about their use at Baptist over the last year or so, and there is a real issue where they do think it is a fairness issue here. So my classmates are using it, they are "getting away" with it, taking the negative view of it, which is not my view, but anyway, that negative view that they've built with their discourse, and then they're getting better grades than me, therefore, why can't I use it too? That seems unfair. So the peer influence, others are doing it this way, therefore I should be able to do it this way too, and that kind of peer influence. And there's generally academic support for these issues of these points, fear of academic progression seems valid.

“The human evaluators tend to judge content they think is generated by AI more harshly than content that is they think is generated by humans.” So if you tell someone AI wrote this, they’re harsher than if you say a human wrote this. We have higher expectations. And if a human used AI to write it, others would tend to be harsher. So that can cloud our judgment when we look at a piece of work, and it’s important that we’re conscious of that and then we can understand why students may have that concern. And there’s been literature looking at authors of academic papers who have the same concern that they are being judged more harshly if they are doing it, and there’s risk that they might get rejected by the editor if they say they used AI. There’s risk that the peer reviewers might not be happy that they’ve used AI in a certain way, or there’s risk that eventual readers will judge them as well. So this kind of issue of judging for use of AI is a real issue that’s happening across different sectors within the higher education field. And the big issue that I’m also having, which I told Christoph not to watch this video earlier, is that the technology companies are also selling the same narrative. So I’ll just share this video with you first, but you can see down the side why I might ask you to watch this.

I also shared this, I also teach a course called “Critical AI Literacies”, and I used this video. This just came out less than a month ago, and I used this in my class as well to discuss this issue around how technology companies sell their products to us and the kind of people they think we are that buy their products, and they also suggest the same narrative that there's a risk that we might be judged on our use, therefore use them deceptively. OK? If you go to YouTube and type in “Apple Intelligence”, this is their new marketing campaign for their new feature, Apple Intelligence. Instantly when I saw this, I thought, “Apple thinks we are lazy, could do it by themselves, but decide not to. OK? We’re incompetent, can't do it themselves so give it to AI to do. Or dishonest, want to find a shortcut.” And if you look at each of their cases, it's something similar. There's no emphasis on transparency. It's all kind of about deceptive use that I want to show I can do something, but I can't do it, or don't want to do it, or don't feel that I want to use my time to do it, therefore I ask AI to do it. And this is the kind of narrative that's being spread. So when a student submits a piece of work, we might have that narrative too in our mind, and the students themselves might think we have that narrative in our mind, therefore they will take that into consideration when they are declaring their use, right? They

will use that when they're declaring their use with us. So my argument is that this will lead to a deception and breakdown of trust as a big, core and one fundamental thing that we have in higher education is a trust system.

(Yeah, sorry. I'm a mover. I do apologize. I'm not someone who stays still very well, so I generally move around a lot. Eye tracking, I want to keep your eye tracking me as I move. Just to keep your attention.)

So the perceived risk of using AI without declaring is lower than the perceived risk of declaring the use of AI and being judged more harshly as a result. I think that's a fundamental issue we're facing right now, at least in this cat and mouse chase, where instructors' jobs become about trying to find evidence that AI has been used, and students finding ways to cover up the use of AI. So it creates this real kind of challenge between these two people, which breaks the trust. We can't have an honest conversation if you're hiding what you do from me, and I'm trying to find out what you do. There's no way we can have trust in that. Ok?

AI use may lead to skepticism, was this produced by humans, machines or both? Right? Because we don't know, because there's no evidence to show either way. Increased risk of unfair treatment and the critical use of AI and erosion of institutional academic integrity, so rather than open, honest and transparent dialogue around the productive, responsible and ethical use of generative AI to augment human ability and intelligence. And we see that here. I mean, I saw this article from the UK Times Higher Education, and it made me very frustrated, because it created this new term that I think is highly problematic, "AI cheating", which is a very unusual term because we don't generally use a technology to talk about cheating. It's either plagiarism or it's not. There is no "AI" in it. It's taking someone's ideas without citing them, citing the original source, while they used AI to help do that, or they did that without the help of AI. It's the same thing. So "AI cheating" for me is highly problematic, because, again, it puts this AI in this position where using it is a deceptive practice that I'm using in a negative way, and it doesn't reflect how we use it.

So I was thinking about my own use of AI. You know, I'm an early thinker about AI, but I'm not an early adopter. I'm quite critical with how I think about using AI, but I still use quite a number of different AI

tools when I'm constructing an article, an academic article, or something like this. So I will use things like, we're in Hong Kong BU and we have our own platform, but I also use Poe. I might use that to think of some ideas, or think of some different theoretical frameworks. You know, I have the ones in my mind that I go to regularly, but there might be others that might be useful. So I might type in, "I'm looking at this research question, what theories might help me understand this?" And it gives me a list of different ideas that I can then research myself. Very useful for a researcher, because I end up talking to myself otherwise. It gives me ideas I wouldn't have thought of before. I might use tools like Consensus, which, you know, we subscribed to at Baptist. My library is very progressive and really, really trying their best to kind of facilitate our understanding of these things. So they subscribed to this tool for us. So I might use this to look at what literature is available and what suggestions it provides to me, as well as supplementing, using other databases like Scopus, which now has Scopus AI, we have that feature as well, right? And also Web of Science and others. I might use Grammarly as I'm writing, and I realize that I do now. I have Grammarly switched on, and I'm constantly going through and selecting, using my own language or what Grammarly suggests to me. OK, I'm the one making the decisions, but Grammarly is there giving me suggestions. Yeah, MAXQDA has AI-assisted coding processes now. So this is something that we might implement in our coding, right, where we've got tools to help us, giving suggestions and ideas, and supporting our processes. Maybe using things like Canva to generate materials and multimodal representation of things, giving me suggestions, ideas, and things like this. So actually, the way we're using AI is very complicated. It's not straightforward, as the companies suggest, or as we think it might be. So where do we need to go from here?

So these are my two thinkings. One is that we need to improve the transparency principle. So we need to improve on this, right? We need to think about how we make this more effective, because it's not working in this current state in any domain. Or we need to normalize AI and adapt our practices and assignments to match that new reality. So these are the kind of two directions I think we can go in. In the talk I gave for the library forum on Friday, which is the JULAC Library Forum, I only argued for the second one for them. And for us, I'm arguing for both, because of perhaps where we are in this level of development. Oh, this is AI-generated. OK. So when I say about improving the transparency principle, I think this is a, maybe a short-term thing that we can think about doing, but

one thing we can do is to state precisely how generative AI tools can be used in an assignment or a task. So rather than saying they can or can't be used, we need to be more fine-grained in how we discuss what uses are suitable in a specific assignment task. OK. These could be related to our course learning outcomes or program learning outcomes. So what does our course really want our students to demonstrate? What do they need to be able to demonstrate? If AI does those things, can they demonstrate them or not? If not, then AI can't be used for those things. However, if there are other things that an assignment requires students to do that aren't part of this course, learning outcomes or program learning outcomes, then they can use AI in those ways. OK, so we try to do a more fine-grained understanding of what AI tools can be used in those assignments.

So this is an example of mine that I used this semester for one of my courses where I have my assignments, and I mentioned the uses that can be used in those specific assignments. So then the students can declare using these uses, and they know which uses can be used. And it will be strange if they didn't use them in these ways, because these are sanctioned uses. So they say I didn't use AI, then they probably have a moral imperative for why they don't, but they can use them in these ways. And what it allows for is open communication. In the courses, I can say to them, "I've thought really hard about AI and its capabilities, but I also want you to be able to demonstrate and also learn these specific skills and knowledge in my course, and this is the aim of my course. So if you use AI in these ways, you're not going to demonstrate those things. However, if you use AI in these ways, I don't mind, because they don't validate these things. However, you are responsible for what you submit, and you should validate what you get from the AI." And I give them a way to validate using multiple large language models, such as checking references to ensure they are legitimate sources of information, and talking about plagiarism and the issue around plagiarism; so they do cite sources that come from AI. That's actually still theoretically plagiarism, because they didn't read the original sources. So they should engage with the original sources, having that kind of conversation with them, so they understand the processes of going through and understanding what to get from them. But because we can have that open, honest conversation, it changes the kind of classroom context from "ban or not ban", right, "allow and not allow", and then we can model acceptable uses in the classroom as well. OK. That's the option that I apply in most of my courses at the moment, because it fits, aligned best with where we are right now in the development of these tools.

Our courses have not been developed in the time of generative AI. Our assignments are not reflective of this new development. So this kind of aligns with current practices. But I did have the chance to develop a new course this semester, and I took this more kind of normalized AI approach. So in this approach, and this course was a critical AI literacy course. For me to ban AI in an AI literacy course would be really weird, because it's about engaging with AI critically. So I can't ban AI in that course, because I would be saying that I'm being uncritical about the use of AI. I need to be critical with how I understand how to use it. So in this approach, we recognize GenAI as part of our reality. As Nick mentioned, there are so many now different uses in the legal profession, and that's true of other professions as well. I work a lot with frontline school teachers, and they are using AI. OK. Of course, they are, but they're not necessarily using it critically, because they haven't been engaged with it well. Sometimes they're using it deceptively, because they're scared of how they're being judged if they use it. If their principal knows or if others know they're using it that way, they might get judged, which is highly problematic, because if they're not having conversations, that means they might be uploading students' work to a large language model, and then that's being used by that model in certain ways, right? Okay? Or they don't know what happened to it next, because they don't know how these tools work. Okay? So this one recognizes that GenAI is part of our reality. Students who use AI will need the competencies to use generative AI effectively and responsibly in their professional practice, to augment their intelligence and abilities. And therefore, we need to modify assignment processes, tasks, and guidance to address this reality.

So we have to change what we do basically to address that this is a new part of how we do things. Our students will need these skills and knowledge when they go into the law profession or any other profession, because when they go into that profession, they will be using them. So we need to prepare them for that reality. So, in this course, this was my guidance on the use of AI. Given the nature of the course, AI can be used to assist you with any aspect of a certain process and product creation. The goal is for AI to enhance your capabilities and not to replace your capabilities, so try and make them aware of that. And at the same time, they are responsible, and they need to check factual information. And we discussed in class what that means and how they do that through a process. So normalizing AI helps us better understand the competencies needed, because we can have real conversations

about how they're used in certain ways, and then we can feed them back into our course design and our assignment tasks. Right?

We can communicate with students about the value of learning from the process of doing assignment tasks and model new processes that integrate generative AI. But at the moment we might struggle to do this, because they become very pragmatic. It's a gap about getting a good grade, not about the learning. So if AI assists them in getting a good grade, they will use it. However, they may neglect what they learn from doing the task in the first place, but having that conversation with them, we can do that, knowing that those tools are there. So you can use the tool, but by using the tool, you may not learn from the process of doing it yourself. So they've become more self-reflective and aware of themselves as they engage with these tools. Okay, and create space for open and honest discussion about AI use. My thinking is that when we talk about this, when we say how we use them, how we engage with them, how they help us to do something we couldn't do before, then they can understand from how they might use it, and we can have more open conversation and more critical use of AI. OK, thank you. That's everything I wanted to say.