

Our Human Future with AI

Episode 5: University teaching and learning in an AI world

Christopher Newfield

in conversation with Thérèse Murphy

Thérèse Murphy (TM): Hello, and welcome. This is a podcast series by the Global Campus of Human Rights. It's about artificial intelligence, more specifically, it's about our human future with artificial intelligence. I'm your host. My name is Thérèse Murphy and in this episode, the question I'm asking is: how does AI affect universities? To help me answer that question, I'm delighted today to be joined by Chris Newfield. Chris was university professor in the United States for more than three decades. He's written a range of books all about life in universities, looking at metrics or quantification in universities, looking at public universities today, how we make them inspiring places. Currently, Chris is Director of Research at the Independent Social Research Foundation. It's a public benefit organisation that is dedicated to advancing the social sciences, and Chris blogs from his position as Director of Research at the Foundation. Chris, you're very welcome. Thanks so much for joining the Global Campus of Human Rights today.

Christopher Newfield (CN): Thanks for having me, Thérèse. It's a pleasure to be here.

TM: Shall we get started? To get us started, it would be great if you could set the scene for our listeners. So, what is it we need to know, Chris, about AI in our universities today? Is AI in our universities today? If it is, where is it? In what ways is AI in our universities? Over to you, Chris.

CN: Well, the first thing that I should say is that AI is an umbrella term for various kinds of technologies that we associate both with Artificial Intelligence, in the sense that it comes from programming in some basic sense, and also that has been represented as intelligence. We don't have to go through the whole history of how this term arose, but in my view, it's originally a marketing term for a whole range of specific technologies that aren't intelligent, you know, you don't think your thermostat in your house that you turn up and down, or that to regulate the heat, or that turns it on and off for you when you're not there to do that is intelligent, but it is reactive to its environment on the basis of a code that's been written to respond to do certain things under certain conditions. Okay, so the kind of AI, as it's come to be called, that we're going to be talking about in the context of the university is, for lack of a better word, a Large Language Model, and I assume that you've had other episodes that have laid this out, so I'm not going to say much more about how those work, but I think the important thing to realise, sort of at the outset, is that AI in the university is a Large Language Model that is inserted into an institution that is constituted by, operates through the production of language. That's not just the sense of management in which email constitutes people's relationships, it's also terms of teaching and research, so the core activities of universities operate through language in a way that is not the same case to say in the construction industry or the design industry, architecture, although lots of code is involved in that, also.

So, it's an institution that is vulnerable to as well as can take advantage of Large Language Models that essentially produce automated writing, that's the kind of the core offering. And so, in terms of the question of where is AI in the university, the answer to that is everywhere. So, it's in course management systems that for 20 or 25 have become the kind of the backbone of course instruction, where students can go online and get the readings for the course, they can get supplemental material, they're often this site where exams or quizzes are to take place, where teaching assistants interact with students. So those are now platforms for machine learning of various kinds that comes out of these models as LLMs. It's also gotten inserted into student advising, so, many of us picture students as going to a university where they go to the office hours of their professor, they make an appointment, and they can talk to their professor about their coursework, how their different courses fit together into an overall pattern, how they're succeeding or not succeeding in particular courses, how they could do better. In other words, there's a person they can talk to about their intellectual lives and how education is contributing to that, or not, and then do course corrections, or whatever needs to be done, and much of that has now been automated at a lot of different universities. You can go online and look at your transcript, but it's gone way beyond that. So, at places like Arizona State University, your curricular interaction is mediated by technology that tracks what you've done, makes suggestions about where you should go, tells you not only what required courses you haven't done yet, but also makes suggestions about if you liked this, you'll also like that. I don't want to just make it sound like an Amazon book suggestion, but that's the kind of the basic principle, and then it's more sophisticated than that.

Okay, so student services is been, as well as course management, are now sort of AI platform, administrative systems, office work, etc. but also the monitoring of the members of the university – students and staff and faculty members – to the extent that they now operate on these platforms can have essentially all the work that they do overseen or surveilled by folks that have access to these systems. At the University of California, just this last year, there's a company called Trellox, whose adoption is kind of an antivirus technology, but it also is connected to other kinds of surveillance services, and faculty tried and failed to prevent it from being forced as an installation onto their devices, including personal devices that they use to do their academic work, so that kind of... it's gone very deep into those systems. Then faculty are also using it in their teaching, faculty are using it in their research. Faculty are sometimes using it to help them write papers that they submit to journals. And then, last but not least, because this is the one in the news, students are using it to do their coursework, and there it just blends into plagiarism and cheating, which is a complex issue that is, we may or may not want to talk about. In any case, it's pervasive now, and it's happened just in a very short amount of time, and people are just, well, the last couple of years they've just been starting to figure out what to do about it, and to see if there are any options for backing away from some of the uses that people are all now taking for granted.

TM: Okay, Chris, thanks for setting the scene. So, AI, Artificial Intelligence, is pervasive, I heard you say, in the university. It sounds like the university is becoming the automated university. AI is there, or in the form of, as you say, Large Language Models, is there cascading through university life, through student services, through course design, delivery, assessment, through the monitoring, or to use a stronger word, surveillance of faculty, administrative staff, students. They're also in the research, potentially of faculty. So, the automated university seems potentially to be a real thing, and I got a sense in your response that there is reason, or reasons to be concerned about this. So, maybe we could move on to that. Now we have a sense of the pervasiveness of AI in contemporary universities, acknowledging that it's probably not the same everywhere in the world at the present time, but given it is pervasive in universities in some countries in the world, let me ask the obvious next question: what's your view on this? Is the automated university the

university-AI interface, is that an inspiring place? Is that a place we want to see replicated around the world, or not? Over to you, Chris.

CN: Well, automation in the mechanical sense, no. I'm really a critic of that. I mean, this is the point in any conversation where a critic of AI says I love technology, I use it. I mean, this is true, but it's like any relationship, right? If there's, if there's mutual respect, if there's mutual understanding, if there's kind of equality over decision making, if there's parity, so that you can work things out, then if you feel like you're in control, if in educational terms, if you are in control, then it's fine. So, using AI for, you know, boilerplate kinds of writing where you don't really need to use your brain, and there's a lot of repetition, that's great. Using AI to find materials on a large scale, much more quickly than you can, that's really good. Using AI in the way that we use computers to do volume calculations, I think that's really useful, although bear in mind that Large Language Models are probabilistic systems, and so they're looking for the most likely continuation of the sequence, and that can often be wrong, right? What we euphemistically call 'hallucinations', and which we could just call mistakes or errors, are not going to be programmed out of these, and so using them in research or in student writing, or whatever it is, has to be done very carefully. It has to be checked.

Okay, so there's a lot of things in which automation really extends our capabilities, and I'm all in favour of that. That's really my criterion. Does it extend our capabilities, or does it interfere with them? If it extends, great. If it interferes, bad. So, I guess there are four areas of concern that I have in particular: when the use of artificial intelligence, in for example writing, encourages the replacement of thinking of the thought process, because it's easier, that's bad. When AI intervenes in relationships between people, that's bad. Third, when it damages cognitive agency, that is when it keeps people from engaging their faculties, including their desires, their moral frameworks, their understanding of the world around them, when it discourages the use of context, all those kinds of things, that's really bad, and when it steals from the intellectual commons, you know, mass downloading of millions of books, for example, which is how these LLMs got started. When it steals from the intellectual commons, and when it does not replenish the intellectual commons, that's very bad. So, there are four areas of real where it's damaging the world, everything from students' brains to the general public commons of knowledge without replenishing it, and that has to be watched and stopped in the cases where it's really disruptive.

TM: Understood, Chris. So, clear list there, four areas of deep concern for you, and then you articulated for us a general benchmark looking at AI in use in universities. Benchmark should be: does it extend our capabilities, or does it interfere or compromise them? And then, applying that benchmark, you determined on four contemporary areas of concern in relation to how artificial intelligence is being used in universities today.

I want to pick up on that idea of human capabilities, but I'm not going to phrase it in as creative and inspirational a way as capabilities but instead deploy the term competences or skills. Like you, Chris, you know, I have a lot of experience working in universities. The same might be true of many of our listeners. And one thing I've noticed, certainly over the last decade or so, from a base in a British university, is that there is a lot more talk about skills, about skills that economies need, skills that societies need and those skills that we need as individuals or people in societies, and I've also noticed that two skills in particular are being mentioned more and more, and they seem to run in opposite directions. So, the first skill I hear mentioned a lot is, not surprisingly, digital competence. So, today everyone in universities, and indeed more broadly, is expected to be digitally competent. This is a good thing. It makes the world more efficient, because apparently it makes all of us more efficient. The second skill I hear mentioned a lot is probably better described as a bundle, or a cluster of skills, and it's a bundle of skills that AI doesn't have. And interestingly, it's a bundle of skills, at least when I hear it discussed by the CEOs of AI tech companies, it's a

bundle of skills linked not exclusively but strongly to university disciplines in the arts, humanities, and social sciences. So, to make that concrete, you know, it's a skill like being able to tell a story, a skill that you might develop if you're a student or a professor of English literature, or let's say of history. And apparently, according to the CEOs of the AI tech giants, AI isn't as good as humans at telling stories, so if AI is reducing the number of graduate entry jobs, if it might be coming for other jobs too, it seems that a skill we should all be honing is storytelling, and then more broadly, the sorts of skills that traditionally we have said are cultivated by arts, humanities, and social sciences disciplines.

So, what's my question for you, Chris? To the extent that what I'm describing is a phenomenon that makes sense to you, you recognise it, you know, what's your reaction to it? You know, if we're meant to be digitally competent, but then we're also meant to hold on to, to refine, to hone these arts, humanities, social sciences skills, how do we navigate this? Over to you.

CN: Yeah, that's a very good question. They are compatible when individuals are able to decide themselves what the relationship is. I mean, that's kind of my most general answer. And they're compatible when the digital doesn't dominate the sort of the art side in the way that they currently do economically, you know, where if you code, your income can be, you know, 5 to 10 times what artists can expect to make, or a web designer who's working as a freelancer. Okay, so I guess the first thing to say about the skills is that to use a chatbot, to use something like Chat GPT or Claude isn't something that you need to go to college to learn, right? I mean, the reason the adoption has been so rapid is because it's essentially, it's kind of foolproof just in terms of the basic use: you ask a question, you get an answer. So, the idea that universities have to teach digital competence in the sense of how do you use a Large Language Model for your job I think is greatly exaggerated. Let's say, it would take a week to learn how to incorporate an already existing model that you don't understand how it works but you know how to use it, to integrate that into the workplace.

So in other words, getting a college degree in order to use AI on a job doesn't really make any sense, it's all the other stuff that you're learning, all the various dimensions of intellectual development that an employer is going to pay you the so-called college premium to bring into the workplace, and that involves a kind of cognitive sophistication and interactiveness that is not something that you can learn just by using a Large Language Model.

In terms of storytelling, I mean, I think that storytelling is kind of a, is a one-word stand-in for a whole series of cognitive capacities that I see as fundamentally rooted in interpretive abilities, right? So, interpretation is non-reductionist, meaning it doesn't really believe in single variables, right? It doesn't think that there's one cause and one effect. I mean, the kind of stereotypical example of the single cause is the high street in my British town looks bad and consists mostly of dodgy cigarette shops and charity shops because of migration. And it's like, I mean, that's a kind of a common political position, and it's based on a single cause – which I don't actually think is correct. So, storytelling would bring in a lot of different dimensions, and the capacity to bring in a lot of different dimensions requires interpretive capability, that is something that universities teach. It involves having knowledge in different kinds of disciplines, so knowing something about sociology, the history of art, history in general, something about economics, some sciences, you know, how do you participate in climate change politics if you don't really understand at all the role of greenhouse gasses in creating global warming? And that then has a lot of psychological dimensions. Why do so many people not seem that concerned about the relation between fossil fuels and, you know, climate disorder that we're currently experiencing? Those kinds of things have to be synthesized in a story, and the person who can synthesize those things can't simply be the person that uses a Large Language Model or search engine, or whatever it is, to assemble just to have some data points. It requires a person that can take those data points and understand their relationships in all their complexity. Develop a multi-causal explanation

for how we got here from there, and then articulate a story narrative that brings in all of the factors that are involved in their just proportions and in language that people can understand, and that's the essential thing that universities contribute.

TM: Nice, I like that. I particularly like this idea of interpretative capability, and how there's a bundle of capabilities or skills inside interpretative capability. I wonder if I could push you to use your own interpretative capabilities for us to help us navigate what I think is a pretty common challenge. And it's this: it's quite hard to pause and to develop an explanation around the use of Large Language Models in universities, to develop an explanation that doesn't rush to the wholly negative or the wholly positive, to develop an explanation that, as you say, brings in the different factors, builds the arc of the story, adds the nuance, we might say, really gets into it. That's hard because it can feel overwhelming, it can be easy to be labelled anti-innovation. It can be easy to be labelled out of date, a dinosaur, you know, unwilling to move with the times. And then it can simply be difficult to find the time, and then, if that's the context, the so-called efficiencies that AI, that these Large Language Models bring, can seem even more attractive. So, what I'm trying to ask you it's: essentially, what advice do you have about how universities, as institutions and people inside universities who make up the universities, might navigate all of this? And do you have any – to go back to the idea of stories – do you have any stories of resistance or even rupture that you could share?

CN: The main form of resistance right now is professors opting out of the use of... not opting out of the use of AI tools because they don't actually control whether students use them or not, but setting up situations in which AI is not so useful, for example, oral exams. I mean, this is extremely time consuming, but you can gauge students' knowledge pretty more effectively – assuming they don't have an earpiece, something like that, where somebody's feeding them the answers – better with just a face-to-face attempt. Another thing that people are doing at a larger scale is moving from a sort of computer-generated, you know, where you're using your computer as a typewriter, basically to handwritten exams. So, there are various workarounds like that. That's the kind of the first level of resistance, I think it probably improves particular student experiences, because they know they're going to have to go brain only at some point in the course, so they continue to keep their brain engaged with the material, with the understanding that that's the destination that's going to determine their final grade.

The second form of resistance in this, I think is somewhat more effective is collective organisation around ground rules for teaching, and also to some extent research, but for teaching and learning. And these are in the United States, at least, are union-based – individual professors can't do this on their own. And that involves setting up parameters for requiring students to disclose the use and then establishing with the university some ground rules about AI policy. So, in the beginning, universities just kind of bought this hook, line, and sinker, and they said, you know, this is the future, and our students need to be prepared for the future, and our faculty have to get with the future, and so everybody has to use this. And over the last couple of years, as the research has come in about the cognitive offloading and brain rot and conceptual problems, I mean the de-education all effects of AI, there's been more of a sort of balanced approach, right? Well, you need to know how to use these things, but you really need to control your use, and, you know, it's like alcohol, right? Just don't get drunk, and don't drink while driving, and don't think while on AI all the time. I mean, that's been kind of the middle of the road position.

I think it needs to be stronger, and so I guess the, the next thing that I would say around resistance, or just setting up an alternative scenario, is just trying to create AI use in designated areas. And, I mean, every person that I've suggested this idea to says that's impossible, but it's not clear to me that it is, and the idea here is that we have courses about how Large Language Models work, their strengths, their weaknesses. I

mean, they are kind of miraculous programming achievements, you know. The technology itself is really interesting, very clever, very powerful, and can do all sorts of extremely cool things. So, teach people about how that works, and that includes non-computer science people, right? It includes literature majors. Everybody should have a, you know, understanding of this, you know, powerful and important technology, but for purposes of also not using it in the areas where it's doing real damage and where you need to be brain only and use, you know, develop your brain, especially in educational context. So, it's like no AI except in designated areas, right? It's we use this with trail bikes, you know. We should be able to use it with AI, and absolutely, you know, indispensable higher education contexts.

TM: I like this Chris, so I'm imagining the university as a space with brain only bubbles – if I can call it that – and then also these AI designated spaces, these AI safe spaces, and I suppose what I'm hearing you say is that this is all about trying to navigate rather than reaching for the extremes, thinking that we can, we should open the doors and welcome in AI wholesale, or we think we build a wall to try to keep it out by only having oral exams, for example. So, you're encouraging us to be more realistic, to decide about the human capabilities that we need to endure, we would need to ensure are enduring, and for those we need the brain only bubbles, but then also because AI can be cool, these AI designated spaces in AI designated areas, as you say.

Chris, I think we've got time for one final question, and given it's the final question, it would be great if you could look into the future for us. Do a small horizon scan, so take us to – you choose – take us to 5 years to 10 years into the future. What would you like to see in terms of the university-AI interface? Over to you.

CN: I love guessing about the future, so, thank you for this question. Let me just tell you my premises. One is that the future also always has to be fought for, right? There's no internal arc, and it's never the case that the most powerful people today are going to be the most powerful people in the future and are going to just decide how things are. That's generally not how it's worked. There's backfires, or there's unintended consequences, but there's also sudden eruptions of counter energies or counter narratives that then produce mass mobilizations, and that really do change things. So, I think, there's a lot of contestations now that can go in multiple directions.

The second thing I would say is that I think about this sort of along two axes. One is democracy versus platforms. We are in a period of tech oligarchy, there's high concentration of control over both technologies and the wealth that they produce. So, that's the platform side and the other side is that democracy remains a regulative ideal. There's been very well documented sort of episodes of backsliding in democracy like the one we're in, you know, my home country, the United States, is in a serious case of retrogression around that, but nonetheless, it remains a framework for thinking about how people should live with each other in societies. And then the other axis is what I call sort of 'high mass knowledge' versus 'low mass knowledge', and this is a slightly awkward way of talking about whether knowledge has been broadly distributed, and or has been democratised in the society. My assumption about humans is that there, you know, intelligence is absolutely everywhere, that people's baseline state is intelligent until it's squashed or repressed or distorted by our lives and the various factors in them. And to the extent that people are offered the opportunity to learn, they do, people love to learn, people love to have knowledge, and that the tendency without sort of hegemonic repressive forces is towards high mass knowledge rather than low.

Okay, so a very real possibility – I'm going to talk about 10 years from now instead of 5, because I don't think we're going to be ready in five, I think the 2020s is not going to be a great decade overall – is the intersection of high mass knowledge and democracy. And so if I think about that particular conjunction, here's what I see: first, majorities of many countries – perhaps even the majority of countries – will have rejected energy-

intensive and capital-driven development as the only form of development, and they will be looking more to human capabilities and democratised access to understanding as the mode of advancing.

Second, there will be more demystified understanding of rage politics, of the conversion of legitimate political grievances into anger-based scapegoating – the kind that's very common, especially in English-speaking societies right now – and they will have attempted to replace this on a broader scale with interpretative thinking, you know, with understanding the complex variables that we all live with, particularly when they're thinking about the other, about people that are who are different from them, and societies and cultures and religions that are not theirs. So, I mean, I think this is already starting to happen, actually, that people are getting tired of so-called populism and are interested in a more formative and a more developmentalist modality of politics that is also inclusive and that brings together people's different kinds of abilities into a form of common life or the modalities of common life.

Okay, that's the second. Then the third is I look forward to the time in the 2030s where we will have made advanced knowledge into a human right, access to that. So, the UN and other organisations already see education as a human right, and I would love to see the advanced version of that, tertiary, put into the same category.

Then the fourth. When you have exposure to it, and I say this again on the basis of my many years of teaching, almost everybody just wants it. When I would go into a class full of business majors, and when I taught a lot of distribution requirement courses – including a course called Detective Fiction, which was not majority literature majors, economists, engineers, every every kind of student was in there – I would say to them whatever you are outside of this class, in this class you are an intellectual, and you're here to think, and you are a thinker. You're not just learning how to, you are somebody who already does this and likes to do it, and the goal of this course is to help you be better as a thinker on these subjects. Nobody ever came up and said this is a bad idea. I really don't want this. I'm dropping the class because I'm not an intellectual. That just never happened. I mean, there may have been quiet quitters, but to the extent that I could track the people who left the class, it was not because of that, it was for other reasons. So, there's a popular desire for this, a mass desire for this kind of advanced knowledge, and I think we can tap into that.

Fifth, we will be asserting democratic rights over technology. This idea that technology unfolds its owners and its developers – you know, the Elon Musk figures of the world – the idea that they decide because they're close to the technology and technology cannot be resisted is a relic of, you know, the 1950s and really the 19th century high capitalist interpretation of industrialisation. It's not true. We, the people of all the nations, should and will make decisions over the technologies we want and the technologies we don't want. We have decided that we do not want nuclear weapons to proliferate, even though they currently are. We just have decided that we don't want leaded gasoline to poison children. You know, environmental protection remains a majority interest, even in relatively conservative countries like the United States, and we have every right to decide – and I think we will decide this – that the unmanaged use of Large Language Models for the purpose of cognitive offloading is not acceptable in school at every level, and it is not acceptable in the workplace, and it is not going to be used to drive prediction markets and politics, or any of that other stuff. And then finally in universities this will have gotten through to the folks who run them, and there will be a much more collaborative relationship within universities about how to use technology, how to define learning, how to set up curriculum in relation to these kinds of tech forces. So, I guess I mean, my feeling is humanity has the right to choose. We know this now, and 10 years from now we will have gotten much better at actually implementing that idea around AI.



TM: Thank you, Chris. Thank you for articulating the values, and thank you for telling us that yes, this future you've sketched for us does have to be fought for, but making it clear to us why it's worth fighting for that future, and how it includes access to higher education, and how... I guess this reminds me of a guest on another podcast episode, Audrey Tang, she talks about how the superintelligence is us. And as you talked about your experience in the US university classroom teaching this detective fiction course, and when you said to your students, you know, you are thinkers, I am a thinker, we're here to think about this stuff, you've never had any explicit pushback. Human beings, even though we might gripe about it at times, we like thinking about stuff, including hard stuff. And so, it's in this that you, you find some hope. Chris, it's been great. Thank you so much for joining the Global Campus of Human Rights today, to share your views on our universities and AI, and what the future holds. Thank you very much.

CN: Thank you, Thérèse.