

C O N V E R S A T I O N T H R E E

Waiting hopefully for the enshittification to come


Patrick Adler, Noel Castree and Mark Carrigan on the new AI policy at Progress in Human Geography, what counts as authorship when a model is in the room, and why the political economy of AI may settle the question before the disciplines do.

Recorded 1 September 2026 · Runtime 1:03

Speakers Patrick Adler, University of Hong Kong · Noel Castree, University of Manchester · Mark Carrigan, University of Manchester

Transcript generated by Zoom, edited by Claude Opus 5

Read the [PiHG policy statement](#) and Mark's [blog](#).



Generated by ChatGPT Images 2.0 (GPT Image 2)

Patrick talks with Noel Castree, editor of *Progress in Human Geography*, about the journal's recent AI policy statement. They're joined by Noel's University of Manchester colleague Mark Carrigan, Associate Dean for Digital Education in the Faculty of Humanities, who has written extensively on AI in scholarship.

They begin with the incidents that prompted the statement: a spike in suspected AI use, two published papers that readers felt were awry, and an investigation the journal had never had to run before. Noel

defends the policy on the principle of human authorship and is unapologetic about being strict, saying he doesn't mind being thought too suspicious. Mark praises the statement's transparency but disagrees about how restrictionist it really is.

Asked whether a policy this strict suits *Progress* in particular, Noel argues that review journals are a special case worth protecting, because interpretive reviews depend on immersion and close reading in a way other formats don't. Mark then offers a non-offloading use case: scholars interrogating their own archives, in the tradition of C. Wright Mills on intellectual craft. Noel accepts that this would be compatible with the policy, even at high volume. Patrick shamelessly asks both whether the disciplines and universities need permissive but transparent sandboxes to work this out. Noel calls transparent experiments like Disembedded worthwhile, but insists a social model of authorship should persist.

The conversation then turns to political economy. Mark argues that AI is bound up with a platform capitalism that preceded it, and predicts access will become too expensive for most academics, inverting Ethan Mollick's line: these may be the best models most of us ever use. Patrick raises the equity problem this creates, where the people who can afford to plagiarise are the haves. He pushes back on whether a crash is really impending; Mark replies that even without one, we should expect the enshittification of AI products after the IPOs.

They discuss adverse selection, which Noel grants as a risk but doesn't think severe so far, and Noel's worry as a parent about what these tools are doing to children, before grasping for a silver lining.

C O N T E N T S

2:11	What a review journal is for, and why that remit matters here
2:39	A spike of submissions where AI use seemed clear
3:25	Two papers published, readers who felt something was awry, and a corrigendum
4:38	Eight desk rejects, and writing the editorial
5:47	The discomfort of sensing something is wrong and being unable to demonstrate it
6:23	Articulated norms make the problem legible rather than private
7:21	The symmetry with what we already ask of students
9:53	Putting a marker down: will we know if the policy has worked?
12:21	Review journals as a special case worth protecting
14:08	What counts as authorship, and what we add
16:16	Reading the statement as epistemically realistic rather than restrictive
17:44	The hill Mark would die on, and why he is becoming more pessimistic
18:56	Publish-or-perish, and models designed to preempt intention
21:14	Interrogating your own archive: going deeper rather than faster
22:50	C. Wright Mills on intellectual craftsmanship
23:47	Blanket prohibition as a strategy doomed to failure
25:52	Would a thousand prompts against your own reading be permitted?
26:43	Standing up for authorship, and where Noel parts company with Mark
28:41	Tackle the ball or tackle the man: the industry behind the technology
29:51	Professional norms as the middle ground between prohibition and drift
30:58	An antagonist of posthumanism finds its critique of authorship relevant
31:27	Epistemic agency as an alternative to reinscribing authorship
32:27	Podcast ideas that end up in papers, and citation infrastructure
33:44	Models offering contributions that do not feel like your own
37:28	Authorship 1.0: social, but never crediting the books you read
38:21	Authorship 2.0, and why treating a model as co-author is the problem
39:01	Launched into the world: Wild West rollout and forced adaptation
39:41	Cognitive offloading, and the antonym we do not have
41:55	One of two things is locked in: IPO or the bubble bursts
42:47	Blitzscaling without network effects
43:10	The pricing shock, and assumptions upended within eighteen months
47:47	On balance a destructive force, and where to attribute causation
50:57	What does a post-AI university look like?
53:31	The crash that might never come

- 53:39 Doctorow's enshittification as a working model
- 54:49 Inverting Mollick: these may be the best models we ever use
- 55:50 Adverse selection: does a stricter policy crowd out good actors?
- 57:13 Crash or enshittification, and a journal willing to change its policy
- 57:58 Worse than nuclear weapons, because they are ubiquitous
- 58:40 What bothers Noel as a parent
- 59:24 Waiting hopefully for the enshittification to come
- 1:00:25 Low tech, no tech, better tech
- 1:01:05 Integrating with one hand, protecting spaces of judgment with the other
- 1:02:21 Syllabus reform, with thanks to the tech overlords

P A T R I C K A D L E R 0:00

Alright, so we're here for our third conversation. We're raising the stakes a little bit with two people I haven't talked to before, and I'll start by introducing them. It's gonna be a great conversation about institutional responses to AI in academia. So, we've got... Mark and Noel with us. Noel Castree is Professor of Geography at the University of Manchester, and an editor of *Progress in Human Geography*, having previously held posts at Liverpool, Wollongong, and the University of Technology, Sydney. He writes on the political economy of environmental use, theories of people-environment relations, and cultural conceptions. of nature. And then we've got Mark Carrigan, also University of Manchester. He's a senior lecturer in education. working across philosophical, anthropology, digital sociology, platform studies, and psychoanalysis. His research develops conceptual frameworks for understanding large language models, you know, appropriately for this conversation, in social life, and treats higher education as a site where the dynamics are being worked out in practice. A concern reflected in generative AI for academics, and the forthcoming, the AI crisis in higher education. So what we want to do is... you know, it might be an AI crisis, but we don't have to be frightened of it, or anxious that much about it, and, you know, we can just have a nice discussion, jumping off from... *Progress in Human Geography's* really generative, I think, statement on AI use, and so it comes with a set of guidelines, but some really well-reasoned motivation For this new AI policy. So, Noel, do you want to just start out by talking about the policy and maybe the context for the policy? Like, what kind of work was arriving on your desk that maybe necessitated this?

N O E L C A S T R E E 2:05

Yeah, thanks, Patrick. So, yeah, *progress in human geography*, so well established. Journal, but it's a review journal, so it has a very specific role in the wider family of journals in academic geography. So, we publish review articles, that survey fields of knowledge, and debate up to the research frontier. And try to push that frontier forward by critically, but sympathetically engaging with quite substantial bodies of published work. So earlier this year, we started to... we had a sort of spike of submissions to the journal, where we suspected AI was being used, you know, inappropriately, I mean, fairly clearly, inappropriately, And then we had some papers where we weren't sure. We didn't have a policy at the time, so... I vet all the papers as they come in before they go out to editors for, sending out for review. So I led a couple through, I wasn't sure about them, and we put them out for review. They passed through peer review. You know, with the usual requests for changes and modifications. And then they were published. And... Very quickly, readers of two articles in particular. felt something was awry. And so, we then had to investigate, which we've never had to do before. We had to investigate whether there'd been some sort of misuse of AI. And obviously then we had to think about what constitutes misuse for the first time. So, quite laborious investigative processes, quite stressful for the authors involved, of course, because... You know, we didn't really have a clear AI policy before their pieces were published. One has to tread very carefully when you're not quite sure what constitutes misuse. You don't want to start throwing accusations out when you don't really know what the criteria look like. Anyway, so, we've recently published a corrigendum about one of those papers. The other paper, passed through the investigation, successfully. But we still are not entirely sure that it's, been written in the right way, and AI has not been misused in the process, but we can't demonstrate it, so... And then finally, I should say that I did a lot of desk rejects earlier this year, so those were the two that got through, but I did a lot of desk rejects, maybe something like 8. Where, quite aside from whatever weaknesses or Inappropriate comment they had relative to the remit of our journal. there was clear AI, flags for me. So that's the background to the writing of the editorial, which... Was put together about 3 months ago, and came out about 2 weeks ago online, maybe 3 weeks ago.

P A T R I C K A D L E R 5:13

So, this is fascinating, because Mark, in a fairly recent post, I believe, about Substack, you commented on academic cultures around suspicion of AI use, right? And it seems like a lot of this is going on in Progress's case. The readership, Noel is saying, is... a little suspicious that they don't know for sure, Noel doesn't know for sure when he's desk rejecting things, the AI is being used, and this creates this kind of atmosphere of suspicion, right?

M A R K C A R R I G A N 5:47

Yeah, and it's really uncomfortable as a position if you find yourself feeling something is wrong, but being unable to demonstrate it. And I feel like we need to talk more about that discomfort, because if it remains an entirely private, individualized phenomenon. I think it could, curdle into something... very unpleasant, which is why I really like the statement, because I think it does... Make clear the reasoning process behind its generation, but also the need to establish norms through which the assessment can be made. And I think once we have clearly articulated norms. We have the capacity to make this, legible. It becomes something that we can talk about, and something that we can reflect on in terms of publicly available standards. In itself, the problem doesn't go away. But I think, as with our students, we've been in a really weird, liminal grey zone for the last few years. Where we are all working with a sense that is basically normative about what is appropriate and inappropriate. But we've not been articulating it in a systematic, coherent way. And because of that, we're all just kind of trading intuitions about what we should or shouldn't do. And I think we've lost a lot of time as an intellectual community, where we could have been cultivating a much deeper, more reflective. more constructive conversation. Because we didn't start that process, and I think initiatives like this are a really good... Start to it. I mean... The comparison with what we expect from students, I find increasingly instructive. Because we are willing to articulate these standards increasingly when it comes to our students. And I think there is a degree of symmetry, because we're trying to protect a certain sphere of intellectual activity, whether that's teaching and learning. or knowledge production. And in this case, what stood out to me about how Noel framed this in the statement itself, was the very specific intellectual purpose of the journal. So it's not just protecting knowledge production and scholarly communication broadly, it's the journal's commitment to consolidating knowledge. And I think that's the level at which we need to think about this, not in terms of a broad sense of higher education being under threat from the new technology, but specific functions that, in a changing socio-technical environment, need different kinds of What would you call it? Kind of custodianship to operate successfully. And this is a big shift in our environment, but it's one that, you know, it predates generative AI. I think there are lots of issues we're now confronting that began with social media, and the normalization of social media for academic communication. that are now extending into the use of LLMs. And as a whole, we've just not really got to grips with the platformization of knowledge production. And I think now we're at the stage where we're forced to. But we've got all this kind of unacknowledged backlog from recent years. We've got a lot of work to do, but I think this is a really good starting point for that work, and I hope more journals Do this, and, you know, not just the rules, but the norms. This is what we expect of you, and this is why we expect it of you. That feels really important to me.

P A T R I C K A D L E R 9:09

I mean, since this policy was a response to, like. basically, I think, like, my characterization, but, like, a kind of... unsustainable or toxic environment where... things are getting published, and you can't kind of trust their providence. Is it fair to say, Noel, that, like, we'll know fairly quickly if this policy has worked,

because there will be more confidence? Like, you won't be just rejecting as much, and there will... will kind of restore a sense of... of faith in the quality of the... of the work?

N O E L C A S T R E E 9:46

Yeah, I think... I think so. I mean, look, let's not exaggerate. I'm not saying we have a major problem here in the journal. We had... a spike in submissions that seemed a bit, dubious. The submissions seemed a bit dubious. I'm not saying this is going to carry on, But it did feel important to put a marker down now and absolutely clarify as best we can, because there are grey areas, obviously. But to clarify as best we can for authors, yeah, as Mark said, what are the norms and what are the expectations? So, I must say, earlier this year, I did find it quite easy to spot, some sort of rogue papers. had been, you know, written using too much AI assistance. Yeah, it does get harder, though, when the use is a bit more subtle. And so that will be an ever-present... for all editors, I think. But the good... the good thing about being my age, and I think most of the editors of the journal are probably in their late 40s, early 50s. Is we come from a slightly older generation, where, you know, the technology wasn't around to do this kind of stuff for us. So, we feel very comfortable being a little bit strict. We're not sort of apologetic, I suppose you would say, about being quite strict. And I think we do say at one point in the editorial that some authors may feel that we've made a bad call by death-rejecting say, a paper, because it doesn't feel right, but we're willing to do that. We're willing to do that. Not all editors maybe would be willing to do that. But I don't mind being unpopular and having someone grumbling privately about me, that I'm too strict, or I'm being overly suspicious, or whatever. But that just reflects, you know, the fact I've been around a long time, and I'm of a certain age.

P A T R I C K A D L E R 11:42

And, you know, Progress is one of many journals, and it does seem... You know, so the statement goes to great lengths to talk about your specific remit, interpretive reviews, and would you say that, like, a policy that's basically this strict is probably going to be better suited to a progress, and maybe a different kind of journal would, by its nature, require a different kind of AI policy.

N O E L C A S T R E E 12:14

Yeah, look, it's a great question, so I'll give you two answers, and then maybe Mark will have a view on this, so... On the one hand, I think there is something special about review journals. It does need to be protected, so... When you author a review article, what you are doing particularly with progress reviews, which, as you said, Patrick, we're calling interpretive reviews, so they're not systematic. Sort of semi-quantitative, you know, digitally assisted. sort of meta-surveys. These are much more about close engagement, deep reading. And... As an author, being really and mentally immersed and engaged with material. And sort of navigating through landscapes of published research, in ways that seem authentic and... Respectful of the research, but also stamping an interesting interpretation on the research to say something of value to readers, right? So it's a very immersive experience, and so therefore, you can well imagine that there will only be very limited situations in which AI applications or AI software are appropriate to that exercise that I've just described. Okay, so that would be... and that's what we articulate in the editorial, right? To your question, though, I guess I could say that for regular journals, normal journals that publish empirical articles and, you know, methodology-focused articles, or pure theory articles, or whatever it would be. that we can be much more liberal about when and when not it's okay to use AI, but actually, I'm quite strict about it at the moment, anyway, and I've got an open mind, and I'm willing to be persuaded otherwise. But at the root of it, for me, is this issue of what counts as authorship. What is the value added that we as, you know, living, thinking human beings are adding

when we Produce research papers, or chapters, or books. Having either engaged with prior literature and or you know, generated data, information, and then tried to analyze it. What are we adding to that? And again, there are grey areas, but obviously the advent of large language models and these, very capable AI applications are raising deep questions about what is authorship, and, you know, unsurprisingly, there have been some interesting pieces on this in places like the New York Times, The Atlantic, there's a whole slew of publications where writers and commentators have thought about this, and I think that issue is not one we're going to be able to resolve, Easily, it's a very tricky terrain. So I would be quite strict about it, Patrick, yeah. Yeah, I really would.

P A T R I C K A D L E R 15:16

I mean, so moving back to Mark, like, okay, so your remit is over... a lot of activities, a lot of them research-related at the university, and is there a, like, version and an appropriate setting for a more, we could say, co-piloted form of research? And, you know, I know your work is pretty progressive. in thinking through alternatives to this, like, authorship-driven approach that Noel was just commenting on, are there, like, other models that you were playing around with institutionally at Manchester?

M A R K C A R R I G A N 15:53

Well, it's really interesting to hear you talk about the policy Noel, because I didn't read it as, Yeah, in a sense, it's strict, but I didn't read it as illiberal. And, you know, we had a quick conversation via email about it. But, like, I just found myself flicking back through it to check I hadn't just kind of misread it through my own prejudices. And... I really like this, because I think it's epistemically realistic. So, like, I read this as a statement of context, identifying a series of epistemic risks that are particularly connected to a review journal. Then asking authors to proactively provide a statement which explains, at every stage, how they've negotiated those risks effectively. So, you know, the basic norm here is that In this environment, we need to... Be willing and able to account for how we've used the affordances of these tools or not used them, and to show that we are kind of capable of being epistemic agents in spite of these temptations. And I think there's a lot of... Creative, intellectually enriching use that does fit through this policy as it's identified. But only as I read it, if an author is willing and able to be reflexive about that, and to... account for it. And... it would be really interesting to see how these statements play out in practice. And, you know, kind of a year or two from now, I'd love to kind of collate these statements and do an analysis of them, to see what are the kind of, the tropes, the... the kind of syntax, the, like, the poetics, I guess. Of how people begin to narrate their use or non-use. Because I think there are a lot of opportunities here, and... I mean, the hill I'd die on is that you can use LLMs in creative, intellectually enriching ways. The caveat to that, which is why I'm becoming more pessimistic with time. Or, firstly, the institutional pressures that lead people to do this, because... It took me a while to see this, because, like, I really like writing. I mean, I write quite compulsively. I've written a blog with 7,000 posts over the last 15 years, as well as my core writing. When I worked outside of academia for a year, I carried on writing. I really like writing. And for me, the temptation of AI would be to substitute other activities so that I have more time for writing. And it was... took me a long time to understand that not everyone feels that way, to put it mildly. That writing can be an arduous chore, you know, kind of defined by anxiety. And so something that will write for you becomes a relief. And also, the kind of publish or perish pressure, and the particularly neurotic ways in which that plays out, means that There's a constant acceleration, because even if just 10% of us are accelerating the pace of our outputs, that creates a ratcheting up pressure about what being a productive academic looks like. And so the structural pressures are a problem, and also the LLMs themselves are, over time, being designed to preempt user intention. Because, you know, in the early models, you really had to prompt in a kind of

systematic way to get them to do effective work. now you can kind of, cough out some half-formed thoughts. And they've been trained on enough user data that they're able to more or less infer what you want to do. And so, both of those dynamics mean that It's harder and harder to use them in reflective ways. But I don't think we should give up on that use. I think that unless we see a situation in which the models are largely withdrawn, and I think there is a small but real possibility that the economics of the sector just fundamentally don't work. And you cannot make a profit while offering Frontier models as a mass consumer product. And so they might remain as very expensive subscription products, but, you know, we might see a withdrawal. But as long as this is available. we need to distinguish between appropriate and inappropriate use. And this, statement of policy is about the inappropriate use defined by the epistemic risks, which I think is 100% the way to do this. Not generalized moralized concern. Not engagement with the political issues, which I think are very real, particularly in relation to the environment, but given what we're doing, it's the epistemic risks that matter. But I think the parallel project is defining the epistemic opportunities. I think we need to be realistic about the fact they are shrinking, but the epistemic opportunities are there. And much as I think we need to teach our students to do this, while, again, being realistic about the obstacles, we need to work out how to do this as a professional community. We need to teach each other to do it. I think it's two sides of the same coin. And I think large research-intensive universities have a really crucial role to play in that mission.

P A T R I C K A D L E R 20:53

For the purposes of our discussion, Mark, can you provide just maybe one or two examples of what you see as, like, generative and ethical and actually compatible AI uses with something like the progress policy?

M A R K C A R R I G A N 21:07

Yeah, I mean, the kind of generic advantage that academics have, at least once you've been publishing for a while. Is, the use of your own archive to have conversations. I mean, environmentally, this is a huge problem, and it needs to be treated very carefully. But the more of your own context you can throw at a model, the more effectively it can tune into your intellectual area. And it's clearly not a substitute for, kind of human interlocutors, collaborators, but it is a supplement to them, and it can be a really powerful supplement, where you have an interlocutor that has read everything you've ever written. in great depth, and can talk within that general domain. And so, it will occasionally make mistakes, but the capacity of the models to tune in to the domain of your work has risen astonishingly over time. They can also make connections between different aspects of your work, and as someone who has a kind of theoretical career, and a much more policy-focused one. Increasingly, talking to the models about things I published is a really effective way of drawing out those connections. And so, I think it works best as a way of interrogating your own archive. I mean, it's the kind of... contemporary equivalent of simply reading your work. I think it's still good to read your past work, but to read your work with the machine and have a conversation about your past work with the machine can be incredibly generative. And these are uses of LLMs that are not about going faster, but about going deeper, about inhabiting your own knowledge production more intimately, more creatively. And I think there's a long tradition of this. One of my favourite authors, C. Wright Mills. In *The Sociological Imagination*, his kind of classic book in the early 1960s, he had an appendix called *On Intellectual Craft* that I read early as a graduate student and has always stayed with me, where he talks about the mechanics of kind of maintaining a notebook, running a filing system, and how to make this kind of what Bourdieu once called the technology of scholarship, to make this a kind of creative and lively process. And I would see it as part of that. It's about working with your own ideation. And it is... Obviously, it seems to me possible to do that without

substituting your ideation for the machine's ideation. It's possible, but it's not always easy, and I think it's becoming more difficult with time. But if we accept that Prohibition... blanket prohibition is a strategy doomed to failure. then we need a positive account of the norms, as well as a negative account of the norms. And that's where statements of use and non-use become interesting, because they will, as a device, surface interesting statements of use. alongside more problematic statements of use. And so it would be really interesting to see how these develop over time. Because I suspect a particular kind of intellectual sensibility is drawn to writing reviews for a journal like this. And I think the reflections and AI use will likely have particular qualities to them as a consequence of that.

P A T R I C K A D L E R 24:27

Thanks, Mark. So, zeroing in a little bit on that use case that you identified, Noel, I think one thing that's interesting about the statement to me is there's this really well contextualized policy And then there is, I... like, I read it as a kind of a side that's like... we don't really... like, I think the term light AI use is used. So just, as a... as a, you know, theoretical idea, imagine a researcher who is, in a very intensive way, doing what Mark is talking about, and having a conversation with themselves, and they're, like, prompting a thousand times, you know, their own thought. So, you know, like, kind of... in any kind of real word, real meaning of the word, like... heavier... words, heavier light, they are heavy users. What do you make of that specific use? And, you know, is heavy use like that something that this policy kind of would allow?

N O E L C A S T R E E 25:32

Yeah, potentially. I mean, the way that Mark just described, using his own archive. You could well imagine somebody, you know, writing, preparing to write a review essay, reading... Articles, chapters, and books in the original, but then also... getting AI to do digests, and then to pose questions about It's understanding, as it were. Of those publications, and that could be quite enriching. I'm not speaking for my co-editors, that seems to me to be acceptable, right? Because... The first-hand reading has been there, as well as the use of AI. So that seems to me to be okay. And it would be more than a light use. But it wouldn't be an inappropriate use, in that case. And I did say before there are grey areas, right? So, obviously, we recognize that. But going back to the comment I made about authorship, yeah, look, I think, and again, I'm speaking for myself, not my co-editors, while there are grey areas, though, I do think we have to really, sort of stand up for something called authorship, Where we are doing the thinking, the work, the... The creative, acts of, you know, intelligence and joining dots and other things. Mark said he loves writing, I love writing, but I also love reading. I love reading, and the last two years in particular, I've been very lucky in that I've had more time than usual. to read, and... it's just... it's a thrill, you know? You're engaging with other human beings' thoughts, and having your own at the same time, so... although it's a sort of hazy thing, what is an author, I do think we have to sort of rough out some sense of that and seek to defend it. And I guess where I'd both agree and disagree with Mark, because Mark said that the opportunities for creative use are shrinking, and I agree. Or appropriate use, rather, of shrinking. So I agree with that, but where I disagree, I think, unless I've misunderstood Mark, is... I think sort of leaving things voluntary. I know we can't prohibit large language models, but I think leaving things voluntary. Unfortunately, there will be a significant minority of people who will abuse that freedom. And they will just try to cut corners, effectively plagiarize. And so, going back to the concept of an author, in those cases, you are not authoring anything, really, you are deceiving the reader. You're presenting machine-generated ideas as your own, and that's fundamentally dishonest. Quite aside from the fact it's just an aberration, it's not what we should be doing. So, I guess that's sort of broadly where I personally stand on this. I'm not anti-technology, I'm not a Luddite, And my closing

comment on this would be, you know, there is that saying, I think it's probably more of a British saying than a North American one, Patrick, that don't... Tackle the ball, don't tackle the man. So, if the ball is large language models. I think there are reasons to critique the way that they've been sort of, rolled out at high speed across the globe, and we're seeing that with students cheating and having AIs write their essays for them. But I also... I want to tackle the man, which is the, you know, the big tech companies that have basically generated these models and thrown them out with very little regulation from. So, I do think we need to look at the whole thing. It's not just the technology, it's the wider industry behind it, and how it's been sort of improperly regulated. So, I think we've sort of got two frontiers in which we have to sort of push back, yeah.

P A T R I C K A D L E R 29:25

Do you want to say something, Mark, in response?

M A R K C A R R I G A N 29:28

Yeah, this is really interesting, and, like, I fully agree about the need to, not leave it to individuals, because I think if it's a matter of individual decision-making, then the incentives of the workplace will win out, and the logic people feel they have to follow as academic employees and the competitive pressures in the difficult job market will win out. But I think there's a middle ground between that and prohibition, which is, kind of... Professional norms, professional standards. I think it's insufficient, but I think it is necessary. And I think if we can establish that kind of, working consensus, at least within disciplinary communities, then that is a foundation to deal with some of these wider issues. And the same is true of teaching and learning, I would argue. I'm increasingly convinced that the way I don't think these problems can be solved in any meaningful sense, but the way to get traction on them is to protect disciplinary communities as loci of epistemic judgment. That this is how we do things here, and we need to agree on that, make it very articulate, elaborate it, and then try and enforce it. And I agree as well about the need to stand for something, and I think that's why I like the statement so much. But honestly, I'm increasingly unsure that authorship is the way to do that. And I come to this as someone who saw myself as an antagonist of post-humanism for a long time. In 2021, I edited a book which was a long critique of posthumanism. I saw post-humanism as the theoretical enemy. And in a sense, I still do, but... The posthumanist attack on authorship feels eerily relevant to me in the contemporary context. Not because I accept the kind of poststructuralist baggage that comes with it. But because I'm not sure reinscribing authorship can do the work that we want it to do. I don't know what the alternative is. In my, kind of, teaching and learning work, with Peter Kahn, another Manchester colleague, we're doing a lot of work on epistemic agency, which I think is promising. Is this something where we're making an active contribution? And I think you can operationalize epistemic agency in some quite, sophisticated ways. I don't know if that is the right answer, but... I worry that trying to reinscribe authorship will not address the practical questions and problems that we're confronting. Because it does rest on the idea of the author as the sole originator of the ideas. And the post-humanist critique that they're always rendered obscure, the kind of network of relationships through which ideas and influence is developed. I find increasingly plausible. And this is something we saw with social media as well. As someone who studied academic use of social media for a long time. One of the experiences that is very common amongst podcasters, and I mean, Patrick, you might have had this experience, is, ideas that were first discussed in a podcast. Ending up in a paper without attribution. And the number of podcast hosts I've spoken to who've had that experience is really striking. And I think the reason for it is that... Or cultural notions of authorship always rested on a certain infrastructure of citation to solidify and render them concrete, and that infrastructure is not adequate to digital media. Like, as a

blogger for a long time, this has got better, but I think it got better largely because Google Scholar started indexing blogs. And as Google Scholar started indexing blogs, people would discover blog posts and treat them as citational objects in a way that's hard to do for a podcast. I'm... I'm not sure citing language models is the way to go, to put it mildly, but I think there is a difference that on a case-by-case basis, you can usually distinguish. between an LLM helping you surface something latent with your own thought. There's a sense of recognition, right? This is actually something that was implicit in what I've already argued, but the conversation with the LLM pulls out some latent content to it. But increasingly, in the last 6 months with the new generation of models, I find the LLM offering a contribution that does not feel like my own. And it's often doing it in a way that presents it as if it's my own. And that really worries me, and that's making me think I'm gonna stop using them in the way I've been using them, because I don't know how to handle those objects. Sometimes it is obviously something that I already thought, I'd just not kind of framed in this way. Sometimes it's ambiguous, but for the first time in the last 6 months. LLMs are presenting me with ideas that I do not feel any sense of cognitive ownership over. And, like, that's when it becomes complex, but do we then treat the LLM as an author that we cite? Like, I don't know, but my point is that I don't think the... notion... the concept of authorship, I'm not sure it can do the work that we need it to do in this context. But the fact I don't have an alternative makes that a fairly... an impressive critique, I guess.

P A T R I C K A D L E R 34:47

I mean, I think in this conversation, there is, among all three of us, a sense of, like, epistemic humility, so, like, Noel talks about gray areas, you... I think Mark are, like, pretty, you know. it sounds like, you know, within the last few months, you've been variously more enthusiastic or more terrified of these technologies, and, like, I don't think any of us are prepared to say what... where to draw lines right now, and so, like. you know, a big focus of this project, Disembedded, is in... coming at the AI question from, like, a very different but equally honest approach, where you just, like, are transparent about everything, and you use that transparency, and, you know, you use your kind of integrity around how you're doing this to learn more. And, I'm wondering, Noel, you know, because you're, you know, I'm talking to you, we're talking to you not just as editor of this one journal, but as a distinguished geographer, do you think there should be spaces in our discipline to figure this out, and Mark, do you think there should be spaces at universities that are really permissive, but transparent, that, you know, we, you know, kind of sandboxes to figure this stuff out? Or maybe is it too terrifying, or too environmentally taxing, or all these other things, for us to even have that kind of space of experimentation?

N O E L C A S T R E E 36:14

Well, I think what you're doing is very interesting, Patrick, and I think my answer to your question is... so long as the large language models are here, and the sort of big applications, like the Geminis and the Claudes and all the rest. As long as they're here and available, whether through paid subscription within a university or free, Yeah, look, so the answer is yes. What you're doing, I think, is worthwhile, and I like the honesty of it, so you're all in, you're fully immersed, as it were. He's trying to see what comes out of it in a very honest and transparent way, so... I much prefer something like that, and the learnings you would get from that. As opposed to the kind of what you might call the sneaky use of AI, where people are not fully disclosing, but getting a lot of help. They're cognitively offloading, essentially, in many cases. But they're not disclosing, and it's very hard to tell. And of course, you know, many of our students are doing that, and in many cases, it's rather easy to tell. But, you know, you can't always actually prove it, so we're in this sort of quagmire with many of our students. But my other answer to your question is, I would prefer it if the large language models weren't available, so I'm going to go back

to Mark on this notion of authorship, so... when I used the term authorship before, I'm going to call it authorship 1.0, okay? And it's a very social model of authorship, it's the classic model of authorship. It's not asocial, it's a social model of authorship. In which even the greatest writers and thinkers that we've ever had, of course, do not think in a vacuum. They're deriving things from their conversations, from their readings of other people's work, and so on and so forth. So... That's 1.0. But in authorship 1.0, which is, you know, the last four or five hundred years, 1,000, 2,000 years since we started You know, to write things down on parchment, In Authorship 1.0, you would never say that, you know, when you published a book, really, the 6 other books you read are co-authors with you. You would never say that, because that just goes against the conventions of Authorship 1.0. The difference with Authorship 2.0, which is the era we're in now with large language models, which is what Mark was referring to. Is the capability of this digital, These digital models and applications. is such that, as Mark intimated, you could conceivably think of them as co-authors. And I guess my point is that's a problem. That's a problem. Drawing from them isn't a problem, but drawing from them so much that you would be willing to attribute them as authors, that really concerns me. And... I think we're being forced to have these conversations because, going back to my earlier comment. These models have basically been launched out into the world. It's a bit Wild West-ish. And billions of people are being forced to adapt, and I think that's completely wrong. End of story. But it's happened, and we do have to adapt. So that takes me back to my first answer to your question, which was, you know, what you're doing is great, yeah.

P A T R I C K A D L E R 39:23

And I wasn't... I mean, it seemed like I was, like, kind of, taking advantage of your, of your generosity by being like, oh, you know, do you think this is a good idea, or whatever? But, Mark, I just needed to try to get that out as unconvincing as it might be.

M A R K C A R R I G A N 39:41

Yeah, I think cognitive offloading is a really useful concept for this, because if we look at the, the kind of minutiae of use on a kind of task-by-task basis. In principle, at least conceptually, it doesn't seem difficult to me to distinguish between those tasks in which someone is using the LLM to support a thinking process, to scaffold it, to go deeper, to lead them through something in a way that they still exercise control over. The challenge is how... complex activities like writing a paper, come from an assemblage of lots of different tasks with lots of different relationships, and then how we assess that, talk about it, adjudicate the claims that are... that are made. And I think if we start from, a concept of what it is to, you know, as I suggested, have epistemic agency, that can be helpful in thinking these things through. It occurred to me recently that we don't have an antonym, or not an obvious antonym, for cognitive offloading, or cognitive outsourcing. And, in a sense, I think I'm offering epistemic agency as an antonym for cognitive offloading. And also, like, cognitive offloading isn't always a problem. Like, some cognitive offloading, we give something over to the machine, because giving it to an... externalizing the cognition helps us complete a more sophisticated task more effectively. So sometimes I think it is okay. The problem is not conceptually distinguishing between our contribution and the machine's a lot of the time. It's how to distinguish that when these tasks come bundle together, because we end up with very complex sets of Activities. And I really like what you were saying, Noel, about, kind of, the challenge of authorship under those conditions, and I think I largely agree. I think we're... I think I'm emphasizing the epistemic opportunities more than you are, but I think... we're almost all completely in agreement. And the thing I'm preoccupied by at the moment is what happens in the next year, because as far as I can see, political... I'm not trained in political economy, I've been trying to teach

it to myself over the last couple of years. But as far as I can see, one of two things is pretty much locked in now. Either Anthropic and or OpenAI will go to IPO, or the bubble will burst before they do. The current pricing model seems obviously unsustainable, that investor capital is being used to offer the services at an effective loss. It's not just the cost of training and the competitive economics of training, it's the simple fact of inference. Like, most users are using These models at a rate that what their payment isn't I'm not... sorry, try that again. It's being offered at a loss, is what I'm trying to say. And that can't be sustained, and at some point, investors are going to cease to be willing to subsidize that, as far as I can see what's driving it at the moment. Is a belief that whoever wins this competition will control the world. But, my colleague Richard Whittle at Salford talks about this in terms of blitzscaling. So the strategy of going for rapid growth in social media. It worked in social media because there are network effects. There are not network effects when it comes to LLMs. LLMs get more expensive rather than less, the more people use them. This does not seem to me like it's a commercial model that... can be sustained in the medium term, let alone the long term. And we don't know what's going to happen when the pricing shock hits, and I think a lot of the assumptions that we're making about use in academia are going to be completely upended within the next 18 months. I don't know what to do about that, but I've been trying to start a conversation about that in higher education in increasingly, strident forms. for months now. But this seems to me the crucial feature of our current landscape. It's not the technology, it's the economy. And I do think there's a small chance that LLMs as they currently exist will go away. I think it's gonna crash the global economy in the process, but I think there is a small chance they will go away. And institutionally, we're starting to see signs of this for enterprise models. So at the moment, most enterprise AI is offered on a per-seat basis, and that's moving to being offered by amount of tokens being used. And once organizations have to pay per per use, rather than per user, the entire economics of it internally change, and institutions will have an incentive to distinguish between appropriate and inappropriate use, because inappropriate use is using a lot of money on something we don't want people to do. That in itself will not drive the production of standards, but at the moment, we have institutional incentives driving individuals to behave irresponsibly. But I think in the near future, institutions will have a really strong financial incentive to try and encourage individuals towards appropriate use, and to stamp out inappropriate use. That doesn't stop individuals using their own models, but I think the costs of that use for individuals are going to rapidly increase. I stopped using Claude Fable because I spent £140 in a week on a coding task that actually I really didn't need to do. And that's the first time at which I experienced this, like, I like using this model, but I don't like spending £140 quid a week on nothing. And I think more and more, people are going to start thinking about the actual cost of what they're asking the models to do. And once you're valuing, placing a value on different types of use, you distinguish, why does this use matter to me? Why does this use not matter to me? And so, the kind of landscape of economic choice, I guess I'm trying to say, is going to shift, and that's going to encourage more normative thinking. And I think it's been hard to establish norms over these last 3 years. I think it will still be hard, but I think it will be a more receptive organisational and economic climate. So, with the huge caveat that I'm pretty confident the global economy will be crashed in the process, I think the problems of AI are going to become a bit more tractable. when this process happens. At least I hope so.

P A T R I C K A D L E R 45:55

Yeah, I mean, I think for the purposes of our, you know, like, the thorny issues we're trying to engage with here, this, this, you know, case that you're describing, Mark. makes Noel's job a little bit easier, right? So, you know, suddenly. this is just a little too expensive, and people aren't using it anymore. I think I personally, I don't know if I'm trained in political economy or whatever, but, like, as a more

economics-inflected person, it does seem like maybe there's not network effects here, but there are economies of scale, famously, with things like energy. The biggest input into these models is energy, and... So, I don't know, I... And, you know, I'm kind of grateful as someone whose retirement fund is probably leveraged to... to these things. At that level, I'm grateful that probably keep continuing. I guess there is, like, an access issue, right? So as soon as there's rationing. At some individual margin. there can be, like, these access wedges, right? Where the... the people who get to plagiarize are, you know, the haves, and the people who can't afford to are the have-nots, and that seems really... perverse as an outcome, right? But, probably, probably, getting away from the main point. I mean, I... I guess, like, Mmm... I mean... So... In this world where... like, or just assuming that we still have something like the current level of AI use, and maybe it stops growing so much. Mark, it sounds like you're... You are concerned about... the types of things it's doing, right? Like, that's where your head is. Like, without being bailed out by the collapse of the sector or something.

M A R K C A R R I G A N 47:47

Yeah, I'm really, I'm really concerned in general. Like, I think on balance, these have been a destructive force in relation to... Across most sectors, like, I had a couple of years where I felt very alienated from AI critics, including some of my friends. And the reason for that alienation was just because they, in my perception, would not accept the obvious empirical claim that you can use these models in creative ways. And I struggled to engage with much AI criticism, because I couldn't get past the fact that they were just denying that, and I thought that was a real problem. But now that... I think that's less of an issue, and I think a lot of AI critics do accept that there are situated uses that are not intrinsically negative. Like, I'm much more... in tune with the direction of AI criticism than I was, and I realized that I agree with the vast majority of it. I just placed the emphasis differently about the causation. A line that I came out with quite a bit in the first couple of years of AI hitting higher education was we shouldn't blame the technology for a pathology of the bureaucracy. And I do stand by that. I think we have to... be careful about where we attribute causation. And I think we're seeing lots of problems that are multifactorial. And there is a tendency that I think is diminishing to imagine that AI caused this problem, and if we just get rid of AI, then everything will go back to being okay. And, like, you know, sociologically, I just find that implausible. You know, we have to see the problems caused by the technology in terms of the current state of organizations and sectors, which are shaping how the technology's being taken up and used. But, like, ultimately, those are, kind of, fairly trivial disagreements, I do think, on balance, this has been destructive. It will continue to be destructive. I just think that there was no plausible strategy, particularly within higher education, of prohibiting use. I mean, this isn't like, you know, VR, where multi-billion dollar companies push a technology no one wants. These are the most powerful companies in the world, with the biggest cash reserves in human history, who are putting their entire future on the line of pushing a technology that is actually pretty useful. And given that balance of forces, it seemed clear to me that diffusion would happen very rapidly, and that we should try and... Deal with the consequences of that diffusion, rather than imagine that we could hold diffusion at bay through force of critique. And now we're in a situation where diffusion has happened, and... I do feel slightly frustrated, because I think a lot of... AI criticism could have been slightly re-gear'd towards managing the consequences and responding to them within institutions. But I think now there is an awareness that actually this is part of the landscape, and the practical problems it's caused are so multitudinous. That we have to think strategically about how to respond to them, and we need to link up these responses, and we need to have a sense of where we're going. what does a post-AI university look like? You know, I don't think any of us have a good answer to that question, but we need to get out of responding reactively, we need to begin to think about what we're arguing for, what

we want to move. towards. And I think that can't just be about the technology. That has to be connected to existing struggles. and new struggles. Because the thing I'm plagued by, I mean, is that if I'm right, that the crash is coming, and I suspect it will be a very large crash. I mean, think about the rise of fascism at the moment. Fascism is flourishing around the world, while the world economy is still technically growing. And so, if we're going to hit a depression, what.

P A T R I C K A D L E R 51:36

Well, in part because of AI.

P A T R I C K A D L E R 51:38

Yeah, props to the whole thing.

M A R K C A R R I G A N 51:41

And so, like, the book series that I'm working on at the moment with my friend Milan Stürmer is about historicizing LLMs in terms of platform capitalism. And, like, just to... I'll stop rambling now, but this, in a nutshell, is kind of how I see it. that there was a tendency to see the LLM as the new thing starting a new phase, whereas from the start, my intuition was, this is the culmination of an existing phase. This is the point at which the existing phase of capitalist development reaches its crescendo, and we begin to pivot into something new. This is the end of a story, not the start of a new one. And I think as we head towards the crash, that's when the new story starts. And as worrying as I find where we are now, I find the prospect of where we might be in 2-3 years' time absolutely terrifying. And... I don't know how to link all this together, but I feel like we just need to... Change the direction of the conversation a little bit to... See this more historically than many of us have been there.

P A T R I C K A D L E R 52:42

It's a fascinating discussion, right, because, you know, you did also, you know, as you're talking about this potential crash, or even impending crash, you mentioned it's this technology that, unlike VR, people really want to use, it's very well capitalized, capabilities just keep coming out, so I guess you would allow, Mark, that maybe there's a possibility that the crash never comes, right? And then we need to start having more of these conversations, and Noel's... Noel probably just quits being editor because it's so annoying, or something like... I don't... you know, something like that, right? Like, that's a possible future world, maybe.

M A R K C A R R I G A N 53:21

I don't know if I've reasoned myself into too much of a kind of tightly knitted framework here, but yeah, I think the crash might never come. A colleague at Manchester, James Sumner, made the suggestion that the military-industrial complex could prop up AI development in the US indefinitely. And I was a bit skeptical when he said that to me, but it's the thought that lodged itself into my brain and realized that, actually, that is a possibility. But the other inflection point is the IPO. I mean, we've seen this process. Like, I think Cory Doctorow's idea of enshittification, should not be seen as a master theory of platform capitalism. I think it's a kind of working model, but I think it's quite an effective model in some ways. And, you know, we've seen this story. Like, I think if these firms go to IPO, we're going to see a kind of deliberate degradation of the product, in order to increase profitability. And so the product and its conditions of access are going to change even if there isn't a crash. They're going to change differently. But, you know, compare Google in 2026 to Google 15 years ago. Like, it's a product

that has obviously got much worse. And it's not got much worse because technology has gone backwards, it's got much worse, because a rapacious international corporation needs to extract as much as possible from it. And I think if the crash doesn't happen, we'll see a rapid enshittification process in the models, which will possibly have some similar outcomes. And this is why, I mean, Ethan Mollick has the famous phrase that gets repeated a lot, these are the worst models we'll ever use. And like, from my point of view, like, these models might be the best, or at least for most of us, these models might be the best. Not because the technology is going to decline, but because, for most of us, the conditions of access are going to deteriorate over time, whether there's a crash or whether enshittification kicks in post-IPO.

P A T R I C K A D L E R 55:14

Which, again, seems like, if that happens, I'm strangely, maybe... and, you know, bringing Noel here, we're talking about this moment as, like, a Cuban Missile Crisis-type moment. We were on the verge of disaster, and then just, like. Like, all these kind of hard, difficult, thorny questions we were just sorting through around authorship, what that means. Just by virtue of the technology, we don't have to deal with them anymore. But, assuming we do have to keep asking ourselves those questions, Noel, I just, you know, and you guys have been so generous with your time, Do you... do you... do the editors think through this, like, adverse selection problem a little bit, where, you know, presumably the... the corrected or withdrawn articles, did not conform to existing policies before the new policy? And so this kind of shows that, you know, there will be bad actors who circumvent policy. Are you afraid that by drawing a more, like, restrictionist approach, you crowd out good actors who want to use the policy honestly, and include only, you know, bad actors?

N O E L C A S T R E E 56:29

Yeah, there is a risk. So far, I don't think the risk has been particularly severe for the journal. I think we're pretty confident that... Everything, really, that gets through pretty much is, you know, has been authored in the traditional way, shall we say. So I'm not... no, I'm not overly concerned about it. I mean, clearly it's possible that, you know, particularly devious individuals, should they exist out there, and of course they do, should want to find very clever ways to conceal. Inappropriate use. So... but I think we're going to be okay long-term, and of course, we're willing to change our policy as circumstances change in. I agree with a lot of what Mark's just said. about either a crash, or if not a crash, then a certain enshittification of the technology. I think that's one of those scenarios, or probably both, in fact, are going to apply. So, there will be changes, and the journal will have to respond, accordingly. I don't want to end on a downer, Patrick, but one thing we've not really talked about, which is kind of presumed behind this whole conversation. you know, these products, we're adults, right? So, you know, and people like myself, because I've got a lot of life experience, I'm kind of old enough to compare what they're doing against what we used to have, and I can kind of make a judgment around that. I think one of the issues is that these technologies are, in a sense, worse than nuclear... Weapons, just to go back to your reference to nuclear weapons, because Nuclear weapons are actually very safe until they go off. The problem with AI is they're, ubiquitous. Children. young people get to get used to them, and have got used to them since late 2022, when ChatGPT came on stream, so we can't... the problem is we can't partition the kind of discussion we're having as adults about how we will judiciously, hopefully, use AI. from, and you know, Mark's right, very destructive technologies, because they're rippling through the whole social fabric of people much younger than we are, so... Again, this is one of the things that bothers me as a parent very greatly. I think it... unfortunately, the technologies are undermining. Some basic capabilities that we should be cultivating in ourselves. So, yeah, sorry to end on a... on a low... I say

end, because I'm going to have to go in a minute, but sorry to end on a low, but I'm just sort of expressing my personal feeling, and I'm clearly not alone. I think a lot of people feel this way. And it is very disturbing. Mark's right, it's not the technology as such, it's the institutional conditions that have allowed the technologies to be rolled out in the way they have, which is why we need to tackle the man, the proverbial man, not just the ball.

P A T R I C K A D L E R 59:24

I mean, we can end on a downer, because there's a kind of hilarious irony that I think you English people will appreciate, where, like, we're kind of waiting for enshittification to save us, I guess? Like, you guys are just here, just, like, you can't wait for this technology to get worse. So, you know, that... you know, I think we can all have a little laugh about that, but Mark, I don't know if you have any final... I mean, Noel creates this opportunity, right, because at the university, you guys bring in people who are... Not to... not to, like, offend, you know, your youngest students, but you bring in people who are like. Pretty young. And you're creating an atmosphere, and, you know, there's a kind of responsibility around, like, promoting technology and developing skills to these people, so I don't know if you have any comments about that.

M A R K C A R R I G A N 1:00:14

Yeah, I agree with the diagnosis. I just retain some optimism that we can design education in ways that mitigate some of these challenges. Or even in my more pessimistic moments, we have a responsibility to try. And so, you know, kind of capabilities-wise, my friend Helen Beetham has a phrase, low tech, no tech, better tech. And increasingly, this feels really important to me, so... We have lots of ways of designing teaching and learning that don't involve screens and don't involve technology. And I think those become structurally hugely significant under these circumstances. And it's not just going back to old methods, or older methods, it's... Redesigning them purposefully and integrating them with a clear sense that we're doing it to protect certain competencies. And so, from my perspective, responsible AI integration in higher education is about integrating with one hand, because I think we do need to do that for all sorts of practical and intellectual reasons, but with the other, kind of designing in a way that keeps it out, that protects spaces of judgment. That creates spaces in which students can sit with difficulty and, kind of carefully work through the challenges of reasoning when you're unsure about the outcome. And I think we can do that. I mean... The political economy, I mean, I'm, like, a very active member of the UK Green Party. You know, kind of from an activist perspective, I think there's lots we can do. As academics, I'm not persuaded we can do that much about the political economy other than critique it. But we can design our work and our institutions in ways that respond to that and express those concerns. And I think we have to do both. I am worried about the future, but I do retain a degree of optimism about our capacity to respond creatively to that future, and even if it fails, I'd much rather, you know, kind of go into the techno-dystopia knowing that I tried.

P A T R I C K A D L E R 1:02:12

I mean, to your point, Mark, I think there's... there hasn't been a moment where more people have dusted off their syllabuses and changed them. Like, this is the great moment for, like, syllabus reform, just because.

P A T R I C K A D L E R 1:02:27

So... so maybe that's... and, you know, and all thanks to the great tech overlords, right? Just as they planned, right? They... but I don't know. Anyway, so I don't know if there's any last... I mean, maybe that's as optimistic as we're gonna get here, and, I really appreciate your time, Mark, Noel, and your open-mindedness for this experiment.

N O E L C A S T R E E 1:03:00

Yeah, thanks, Patrick. It was, fun, very interesting, and it's so important, on these issues. So, yeah, I look forward to, staying in touch with you.

M A R K C A R R I G A N 1:03:08

Yeah, that was a brilliant conversation. Thank you both, and thank you, Patrick, for organizing.

M A R K C A R R I G A N 1:03:13

Yeah. Cheers.

Transcript generated by Zoom and edited by Claude Opus 5: backchannel interjections and broken fragments removed, misrecognised proper nouns and terms corrected. Wording is otherwise the speakers' own.

Part of Disembedded – an economic geography paper written live, with AI tools and no human collaborators. disembedded.com