

CONVERSATION ONE

Two Centaurs and Two Human Beings Discuss Progress in Human Geography

Patrick Adler, Ben Gerlofs, Tim Wotherspoon and Giancarlo DeFrancesco on a journal drawing a line, what is lost when the steps are skipped, and whether more research is a good thing at all.

Recorded 24 August 2026 · Runtime 1:10

Speakers Patrick Adler, Geography · Ben Gerlofs, Geography · Tim Wotherspoon, Physics ·

Giancarlo DeFrancesco, Earth Sciences

Transcript generated by Zoom, edited by Claude Opus 5

Learn more about [Ben](#) here.



Generated by ChatGPT Images 2.0 (GPT Image 2)

CONTENTS

01:29	The statement: interpretive reviews, disclosure, and only light use of AI
04:55	Editing the Journal of Urban Affairs: you suspect it, and can rarely prove it
06:54	Why the frustration is sharper as a peer reviewer than as an editor
08:20	Submission agreements, and students bringing generated text to office hours
12:08	AI coming for the work worth doing: still washing his own dishes
16:11	The hard scientist wants liberation; the social scientist wants better science
17:09	Why the review article the journal refuses is the one a scientist would want
19:10	The scientific-paper Mad Lib, and where the ideas are supposed to live
20:55	Data obfuscation: stripping sample locations before handing over a dataset
28:37	Is there a path to teaching discernment? "I don't, but I want there to be"
30:34	ArcMap to ArcPro: the model as friction remover when the tools change
32:32	Cory Doctorow, declining quality, and the cancer-screening argument
34:39	Monmonier and cartographic critique: rules a model can check, maps it cannot make
37:23	Statistical software lock-in, and R-to-Python becoming trivial
41:10	Undergraduates meet the technology first, so how can they check anything?
42:49	An earnestness to knowledge that is lost when the steps are skipped
43:28	Integrity or utility: two readings of what it means to earn a skill
45:32	Training checkers instead of experts, and what that concedes
46:49	When policing fails you just grade quality, and the essay is always poor
48:35	The challenge to the productivity premise: why is more papers good?
49:35	Volume as noise, and the cherry-picked citation as a peer review problem
52:26	More slop, but also more good papers at the ninetieth percentile
53:59	Recent results in mathematics, and the clever mathematician in the loop
54:53	Collaborator rather than checker, and why fieldwork resists the handoff
57:07	The centaur and the reverse centaur
58:09	Why the lab and PI model absorbs AI more comfortably than solo work
59:52	If students write worse, should he be teaching writing?
1:03:01	Defending your sentences to four classmates who read your draft
1:07:37	Four years ago it was bad; now he fears losing the ability to critique

PATRICK ADLER 00:01

Sorry, once again to my legions of followers on Disembedded. Haven't figured out all the, technical aspects of this yet, but we are recording this, we will upload it. And I will just start by welcoming our wonderful panel here.

Today, we've got representatives from 3 departments at HKU, here to unload all of their hot takes about AI. In the classroom, in research, and I think we're gonna start. With the most bombastic and negative, and then we'll continue until we get to the most decent and well-meaning. No, we're gonna start with Ben who I think, and so we have as a, like, conversation piece, and I'll share this.

BEN GERLOFS 01:24

Walter Benjamin's The Work of Art in the Age of Mechanical Reproduction.

PATRICK ADLER 01:29

No. Not that. Don't be ridiculous. We have fresh hot off the press.

Hot off the luddite press, no. Hot off the press. A [statement by Progress in Human Geography](#), which is the review article journal in our field, my field and Ben's field, Human Geography. Talking about what it's been like to edit this journal since generative AI came out, and announcing that they will remain committed to what they call interpretive reviews.

Which I really think of as, like, author-helmed, auteurred reviews. They will reject, kind of, meta reviews. They will require transparency about how AI is used in preparation, writing, and article formatting. And I think most, concretely, they expect for AI writing to be AI assistants, to be light, and I'm just gonna search for that, so the word light.

We seek to publish articles that should reflect only light use of AI in their preparation and writing. So I'm putting that on the table. I'm inviting Mr Ben Gerlofs to speak. I have a few things to say about him, and then we'll continue with Mr Wotherspoon and Mr DeFrancesco.

Ben Gerlofs is my colleague here in geography. He's an urban, cultural, political, and historical geographer, studies neighborhood change, urban social movements. His book, [Monstrous Politics](#). I think is admirable for anticipating the topic of gentrification in Mexico City.

Ben was probably the first to describe to me in, like, social scientific terms, and then thereafter, I just heard a lot of people talk about it, so Ben, don't let this conversation suggest otherwise. Ben does know how to find hot topics.

And here he is. Ben.

BEN GERLOFS 04:14

Hey, that's a very generous introduction, you want me to make a statement on this statement?

PATRICK ADLER 04:22

Why don't I start – I have one more thing to say about you to introduce you, and that is that you are an editor. Of a section of a very good journal that I work with, the book review section, the Journal of Urban Affairs. So, like the editors of Progress in Human Geography, you've probably noticed generative AI? And so, can you sympathize with their editorial?

BEN GERLOFS 04:55

Yeah, most definitely, and I'll say, just on that side of things, that our so our section of the Journal of Urban Affairs doesn't currently have its own unique AI policy. The journal has a general policy, which involves the use of AI, among other things, and the publisher, of course, has its own set of policies around these things that applies across outlets. We've discussed for our own guidelines with book reviews at JUA, we've discussed having some sort of statement on the use of AI, but we currently don't. But we do encounter it a lot, in much the same way that I imagine the editors who wrote this statement for Progress, and basically everyone else who does any kind of editorial work these days with academic publishers, does, which is to say, we mostly suspect that we encounter it.

It's very difficult to really know, right? We have, like, a sense that something is AI-generated, or it feels very artificial, or something like this, but these are incredibly difficult things to know, and they're super dangerous accusations to make, of course, of anyone's work, that they've done you suspect that they've done something they're not supposed to do, especially when there's not a policy that says you're supposed to do this and not supposed to do that. And so, we feel like we see a lot of influence of AI, but it's not something that, outside of conversations between the other book review co-editor and myself, that we really discuss very much with authors or anyone else. So, I do understand where these editors are coming from, and I think that they've made a really bold statement.

I have lots of questions.

PATRICK ADLER 06:39

So you sympathize with them. But it doesn't sound like you're kind of where they are. I don't sense, like, a lot of frustration in your voice, but then when I'm reading this statement, it's like, I think, boiling over.

BEN GERLOFS 06:54

I'll say it from another angle, as a peer reviewer, for instance. I have a pretty different reaction to this as a peer reviewer than I have had as an editor. I encounter what seem to be much more blatant abuses, let's say, or inappropriate uses, to use the language that these editors use in this statement, of AI in articles that I'm asked to serve as a peer reviewer for. This is more and more common all the time, I feel, where something is, to my eyes, and I would never say this about an individual piece that I was discussing, but we see a lot, I see a lot of them as a reviewer, that to my eyes are really blatantly inappropriate uses of generative AI.

So, in that realm, I feel the frustration more than I have so far as an editor, at least. My experience has been not quite so frustrating on that front as an editor, but increasingly frustrating as a peer reviewer, I would say.

PATRICK ADLER 07:58

Awesome. But then, I guess, and maybe this'll be how we bring in the other two gentlemen. As an instructor. I'm sure you've encountered it, just like everyone else.

And, where's the frustration meter there, and how do you kind of square your different levels of frustration in each of these contexts?

BEN GERLOFS 08:20

A couple things to say about that. On the one hand, I have tried in some of my bigger lecture classes to do things that seem sort of like what these editors are doing here, like requiring a certain kind of disclosure, I've had assignments and submission agreements that I've used in my classes over the past few years. It kind of they're, like, increasingly. What's the word I want here?

They seem increasingly aggressive, these statements that I draft and then ask my students to sign, saying they've not done this, and they've not done this, and they put their name on this piece of paper four different times. I have not found this to be a terribly fruitful way to try and police, the AI use that students, are sending my way. I don't but I don't have a better way. I've talked to lots of instructors about what they do in their own courses, and it feels like instructors across the board are fairly frustrated with how students use AI.

These days. Especially in writing, I think. But in my own courses, I see this in how they study for exams and other stuff, they come to me with a piece of AI-generated text about, like, a key term that I've lectured on. And then they read this to me, and then ask me, like, is that right?

Like, what's wrong? This, I feel, is not a very fruitful way to learn course content, and I find it kind of a frustrating way to interact with students as well, it doesn't seem very earned. Like, this stuff that they've brought to me is not their own idea that needs to be, like, worked on and challenged. Rather, it's a piece of text that they have basically no relationship to, and they're asking me to interact with it, and tell them whether they're interacting with it is correct or not.

This I find super frustrating. So, in my own teaching, I'm really struggling with how to assign writing anymore to, especially, like, lower division undergraduate courses, what use is there in assigning a course essay anymore if the majority of my students will simply use ChatGPT or Claude, or DeepSeek, or something else to produce the ideas, and the initial draft, and any subsequent drafts. And then they'll simply tell me that they didn't do that, or they'll use language that one of those programs generated to describe how they used the programs to arrive at this result. I don't know how to address that problem, really, other than in my most frustrated moments, I suppose I just retreat from assigning things that are AI capturable in so obvious a way.

I don't think that's a very satisfying pedagogical response. So I'm super curious about what others think about that.

PATRICK ADLER 11:09

Well, maybe that's a good opportunity to bring Tim in, and I'll just introduce Tim very quickly. Tim is a bona fide astrophysicist. Knows a thing or two about complex systems. As applied to astrophysics, he mainly at HKU, leads a number of courses and tutorials in the Common Core program, which is a very difficult program to teach in, because you have to teach to every type of student.

We require all students to take a lot of courses in this program, and he's very seasoned. I call him the dean of the Westerners here at HKU, because he's been here longer than the rest of us, and so Tim. I do where's your frustration meter at? And, you know, is it similar to where Ben's is when he's describing his teaching and grading.

TIM WOTHERSPOON 12:08

Oh, it's similar, yeah. I'd say my frustration with it is always increasing. So I'm fortunate enough that I don't usually assign writing or deal with student writing, at least on, like, the 500-word-plus scale, anymore. I was doing that for a while, and it's exactly, like actually, I'm very uncertain why we even have humans in this.

And, actually, I really enjoyed the piece about from the editors and the Progress in Human Geography. I mean, these editors, I don't know, they're you know, the they're screaming at the tide, for sure, right? And hoping it won't come in, but I think they're absolutely right. What I took from their letter was that, they want to do a human activity.

And the humanness of the activity. Like, they talk about doing close reading or deep reading. Coming from science, this is not something I really do. And wanting to interpret what's going on, like, something where the author is bringing their own interpretation.

This is a very human, artistic, creative endeavor. I don't know why we're using AI for this. Right? I don't know why.

I feel like the promise of AI should take me out of the dullness of my work. But this is the work that we're meant to be passionate about and creative about.

And this is the work that AI is coming for. I hate this. I'm so frustrated with this. I'm still doing my own dishes, I'm still doing my own laundry.

Right? But the things that I want to do, like music, or art, or reading, AI's coming for that. I hate that.

For our students. I think what we're learning from our experience with our students is that, the academic pursuit is not an end to itself. For most of our students. Right, the end for our students is a qualification, a grade point average, something along these lines, right?

And I think as long as we are pursuing our lives in terms of metrics, like for music, how many listens we get on Spotify for our academic study, what kind of grades we get in the classroom, AI is going to

be very good at you know, producing content that gets high scores. Because that's how machine learning works, is that it creates content that gets a high score, with some definition of score. And I think us are shouting, or like, but the score is just a incomplete, incorrect way of judging the quality of the thing. There's a intangible quality of the thing that's what we treasure and what we love.

And I think that's correct. But this is not how the world works, right? We want to publish a certain number of papers. We don't want to express something intangible to our bosses.

We want to tell them that we have this H-index, we've published with these high-impact journals. I think AI is gonna win this game. I don't know. Is that a reasonable take?

PATRICK ADLER 16:11

I mean, there's a lot of really great stuff in there. I think what's really fascinating about your perspective, Tim, is you're representing the, like, hardest and apologies to Giancarlo, I mean, he's almost there, but, like, the hardest of the hard sciences, and, you're the one who's expecting from AI like, liberation, and the ability to be more romantic, think about ideas, and be more human, right? So, like, it seems like.

You want, like, social contract with the AI, where you get to be liberated a little bit more. Like, that's how you'll accept AI, is to allow you to be more human. But, like, I think there's another scientific approach. Which is, like, AI lets us do science better, maybe, you know.

We'll let you try that one on, and then maybe bring Giancarlo in to see what his perspective is.

TIM WOTHERSPOON 17:09

Yeah, also when I read this piece from the Progress in Human Geography, I thought that this review article that they don't want is a really scientific review article. Is kind of the one a scientist might want.

PATRICK ADLER 17:29

I mean, but why would a scientist want that one, Tim?

TIM WOTHERSPOON 17:35

Oh, because we're not!

Well, no, because we want an efficient summary of other work, right? And we don't necessarily need someone to have a take on that other work. We just want to know what's been done, who did it, how it interacted with other work. Right, and so, what's the work?

What are the connections between the work? If AI can do that, it'd be very useful, right? Or even, using AI to summarize a scientific paper, might be quite reasonable compared to having AI, unpacked something that I don't know the word I'm looking for, but the informal word would be woo.

Right, but the scientific work should be pretty concrete. And our target, you know, writing a scientific paper. My PhD advisor basically had me fill out a Mad Lib whenever he wrote a paper with me.

So, like, we study, we find. You know, and just fill in the blanks.

PATRICK ADLER 18:54

It was kind of already AI. I mean, if you kind of, like, boil AI down to its, like, fundamentals, where it's just this, like. They identify structure in the world, and then just kind of reproduces that same exact structure. It was like.

Algorithm, right?

TIM WOTHERSPOON 19:10

Yeah. So this should free up things, right? So in the world of science, maybe this AI use is okay. But, I mean, in a scientific paper, where's the author's ideas?

Right, where do those ideas exist? Giancarlo, I mean.

GIANCARLO DEFRANCESCO 19:36

I'm in the throes of exactly that right now, actually, yeah. So, I mean, for me.

PATRICK ADLER 19:41

Well, let's just introduce. Moderator's prerogative. Let's introduce you, Giancarlo. Okay, so Giancarlo, but just, I think it's important background.

In the, unlikely event that someone else watches this. Giancarlo is, about to get a PhD in geology. He's doing the hard science right now. And he, I mean, we were just talking before the stream started, he's, like, a model AI user for me, as a relatively ambitious AI user.

You know, he was telling me things to do that I, you know, I should have been doing. So, like, two things to know about Giancarlo. A, he really needs to get a PhD this year, and B he knows how to use a lot of AI, and I think those two kind of work together, where, like, he doesn't really have the luxury of being, like, human with this technology, he needs to get shit done, so maybe we can get the get-shit-done perspective from Giancarlo, or deviate if I'm getting you wrong, your perspective wrong.

GIANCARLO DEFRANCESCO 20:55

Well, I think there's definitely a balance that I'm trying to strike, right? And in part, I think that's been driven by using the tools and having at least to me, as someone who's intimately familiar with a lot of the topics that I'm working with, having text, you know, a response, an output, text output from a model, and reading it and being like, this is misleading. This is or even, in some cases, wrong. Right?

So being very careful when you're analyzing data with AI, but also making sure that when you're putting something on a document. That your thoughts are your own. And so, for me at least, with the ways that I'm using it, when Patrick was, mentioning that, you know, I think you described me as a

model user. What I'm talking about in there as well, right, is being careful with the data that I'm giving the model as well.

And so, yeah, as I'm finishing up my work, you know, I'm preparing a lot of data, I'm interpreting a lot of data, and I'm also checking my data against some other people's, reported findings. And so, for me, you know, that also includes data obfuscation. So, like, if I'm giving it a large data set, some of which is data that I is mine and I've not published, I might strip out data that I don't want a model to have. So, in the case of the geologic work that I do, I look at a lot of river systems.

If I'm giving it the exact sample location. Right? And that's fed into a system, and it's kind of more broadly put in a model, and it can see the relationships between my samples, the locations, and the concentrations within the river water at that location, right? That can kind of be fed into other interpretations, and my data can get out if I'm not being careful, with how I'm working with the data and the model.

But also kind of going into how I'm using it, and how I'm getting responses back that are incorrect, right? If I'm struggling with an idea, if I'm struggling with a concept, I'm looking for ways to, reinterpret something that I have that I think that I'm thinking differently about than maybe what the literature might suggest a conclusion might be, you know, I also have to go and check that as well. And so, you know, a model recently told me that I need to look into sulfates and a connection with, an elemental fractionation. And I went to Google Scholar, because I think, right, it's like, okay, I've not read about that whatsoever.

Did I miss something in the literature? Because, you know, I'm reading so many documents per day and trying to incorporate that into my work, it can happen, right? Especially when you're trying to move quickly, and like Tim said, when you're trying to hit certain metrics, right, we are trying to move fast so that we can publish significant works that has a high impact factor, and you're trying to do that, you know, multiple times, repeatedly, and maintain quality, right? And so, if you're trying to do that fast, you might miss something.

And so, for me, I went to Google Scholar to just, and to their credit, I think that this is an example of where AI can actually be very helpful, is I can now use Google Scholar Labs, where I can ask a question in a more native. English, or not native, natural English, a way of phrasing what I'm looking for, and I can be very specific at what I'm looking for. And so I asked it to kind of show me if there's any relationship between this element and a sulfate. And it basically came back with stuff that was unrelated, to what I was looking for in terms of a geologic context.

And I went back to the model and I said, hey, like, you know, I just looked for research papers on this. Where did you pull that suggestion from? Why did you suggest that in particular? And I'm not joking, the response was, oh, well, you know, based on this property, and these kind of aspects of data that we've been talking about, this should be you know, what happens.

But there was no literature to support a statement like that. And that doesn't mean that the model is wrong, right? That's maybe something that you can look into in the future, and you can go down a

rabbit hole and look at that, but if I'm suggesting that as a conclusion for what I'm currently doing. Without any evidence to support that.

That's a really big claim that I'm making, and that's something that maybe I shouldn't be throwing out there without being able to check that claim, essentially. So yeah, data obfuscation, I think, is really important, but also understanding that right, even today, and I think some of us might remember, right, when AI was first being kind of pushed on us as a community, a lot of the responses were definitely inaccurate, and I think that's less of the case now, but to say that it doesn't happen, or to think that it doesn't happen, especially for new people writing, and up-and-coming scientists, who are using it a lot, like, I see my peers use it quite a bit. I obviously use it quite a bit. Making sure that they're that it's my writing, that I'm putting the words on the page, and, you know, you might want to check it, say, hey, where's my argument weak?

What am I being unclear about in my writing? That's one thing. But you do have to struggle with the text a bit, yourself. Those are important skills that we need to have, because if we're losing that, then if we're losing that as a community.

I think that we could be in for a difficult future as a society. I mean, maybe this is a good opportunity we all had a.

PATRICK ADLER 27:09

Yeah, that was, I mean, that was awesome. I mean, I think your, like, so you talked about there at the end, and previously, kind of this judiciousness you need, this, you know, like, basically. Your skill of, like, maybe as a PhD student, you know, at this stage, you would be regularly exercising critical judgment, like, evaluating some approach. A little bit, like, you'd start to be reviewing articles.

But not a lot. And now, on a kind of like, dozens of times, maybe an hour, you have to use that judgment. So returning to what Ben's saying, and then you talk about obfuscation, these other skills, what Ben's saying, like. Ben, do you think, like there's a sort of social scientific Urban geography, human geography.

Kind of version of where we can, like, train our students to use these models, like, with discernment. Is there kind of a path forward there where, you know, instead of the student that you were describing coming in and just belching back, something they got from a model. We're kind of figuring out how to train them to figure out the best responses and get the most out of these models. Do you see some hope there?

BEN GERLOFS 28:37

I don't, but I want there to be some. And actually, I was gonna say earlier, and I should have said in response to your question about, like, a kind of frustration, where's the frustration meter at with these things? But part of my own frustration with this is my own inability to envision how to better use this technology. In my own work, but also for other people.

I feel super frustrated by not being able to make this thing work in a way that I think is useful, and also by other people's inappropriate or poor or lazy use of these things, which I feel kind of pollutes the whole thing. It gives kind of a negative impression of, like, this whole genre of work and thought and activity that makes it difficult for someone like me to engage with it. I feel frustrated by all of that, by my own kind of old person. Luddite inability to use technology in a way that people that seem to be younger and or smarter than me, of any age, are doing.

I don't I talk to lots of people about this. I feel like I'm here to be impressed with what social science and humanities and other parts of the university or the world can do with AI, and I just have yet to see it. I'm sorry for any colleagues who might view or hear this and hear themselves in conversations they've had with me, because I really I ask a lot of people who are enthusiastic about this, like yourself, like other people who I really respect as scholars, as teachers, as human beings who are enthusiastic about this. I ask, like, what is it that makes it cool?

What can you do with it? What are your students doing with it? How is it helping your own work? I occasionally hear cool things like what Giancarlo was just saying, where this sort of makes sense, that you would use a program to do this kind of a thing.

But in my own areas, I just have yet to find it, and I've genuinely tried. I'm not very good at this sort of thing anyway, but I don't see how I wish I did.

GIANCARLO DEFRANCESCO 30:34

I'm curious, Ben, I'm not super familiar with how much your work overlaps with GIS, but I found, at least in my experience, and this kind of also touches a bit, maybe, upon what Tim was saying, for me, when I'm using tools that I'm not always super familiar with, I just transitioned from ArcMap, I know I'm a dinosaur, to ArcPro, but, you know, with that transition comes a lot of tools that have been tweaked, that have been moved, and so there are times where it's like, I know how to do this in ArcMap, but I can't use ArcMap right now. How do I do this in ArcPro? Right? And it gives me step-by-step instructions.

And, you know, in part, my job is to know how to interpret that data and work with the data. But also, right, it's like, I don't I might not want to spend my entire afternoon troubleshooting why a layer output isn't showing up properly because I didn't click a parameter correctly, or if there is a parameter that I don't know how to use in a tool that I'm familiar with because it's a new feature, hey, what does this tool do? How do I use it properly. And so I think when we're when Tim was kind of mentioning a bit about, kind of, music creation and, other, creative aspects, at least for me and my personal life, using AI to help me you know, save time on the, kind of nitty-gritty of some of the coarseness of learning a new skill and learning a new tool, and helping me be a bit more creative has been really awesome.

I've been using it to build a home lab for myself. But again, right, I'm doing it step by step, and I'm saying, okay, why are we doing it? Why are we using this variable? Can we do this in a different way?

And I have been caught up in times where, okay, maybe that wasn't the best approach, and I should have asked a better question, right? But so it depends.

BEN GERLOFS 32:24

I really like that. I'll be back in 2 seconds. I'm gonna grab something from Michelle.

PATRICK ADLER 32:29

Oh, this is very convenient, Ben.

TIM WOTHERSPOON 32:32

I listened to this conversation with Cory Doctorow recently, and I think he would largely agree with Giancarlo and with Ben. Right, I think he is quite familiar with all the judicious and productive uses but he's not hopeful, the use will be judicious at large, he's not hopeful for this at all. And, his, I mean, he's a smart guy, I'd say. And knowledgeable about how to use these things, both, and he's quite skeptical about it.

And what I heard him saying is that the quality is going to decrease. For whatever it is, the quality is going to decrease. In human geography, if the quality decreases, how does the world change? I mean, the world changes for human geographers, I right?

For the editors of human geography, the world changes. And they don't like the new world. But for the rest of us, how does it change? I don't know.

But, Cory Doctorow's example, though, was using AI to read medical scans to detect cancer. This is a huge app, but the it still requires a human to check the work. Right, but then you're gonna fire all the humans who are doing all this work, and then the people who check are gonna retire, and no one will be in this field, no one will know how to do it. And people gonna die, because the AI is not going to read their medical scan correctly.

So, the quality is gonna go down, and in some applications, people are gonna die. So yeah, I agree.

BEN GERLOFS 34:39

Pat, can I just ask you to stop sharing? For a second, I was, distracted by my own inability to use Zoom. How do I make the people bigger? Oh, right, it's because Pat's sharing, okay.

There we go, now we got a nice, four panel there. I went to get this book, [Mapping It Out](#), because I worked for the person who wrote it for a couple of years in graduate school. When I first started grad school, and I was chronically underprepared, this is Mark Monmonier, one of the world's preeminent cartographers and historians of cartography, digital and otherwise. And Mark teaches, or taught, these courses at Syracuse for undergraduate and graduate students in map design and interpretation.

They are foundational courses. Very few geography programs, I think, anymore teach classes like that, and they're super valuable. We have one here that's sort of like this. But the way that Mark teaches cartographic design is super rules-based.

There are visual variables, they can be measured, there are things that you sort of have to get right, and we would do these critiques, where students would bring in their work, some of it good, some of it bad. And just go through, like, a series of things. Are the lines crossing where they shouldn't? Is this, like, the appropriate size?

Have they oriented this correctly? Those kinds of critiques of a map to return to Giancarlo's question about GIS, it strikes me it could be very easily produced by even a low-quality AI model, right? It could tell you no, you've messed up the lines, right? These things are too close together, because you just give it the rules for how to put visual variables together in a compelling presentation, and it can just tell you've violated the rules, or no, you haven't violated the rules.

It can't really tell you how to make a good map.

PATRICK ADLER 36:21

Sorry, Ben — it sounds a little like cancer to me. I don't know all about cancer detection, but isn't this just, like lines? Anyway, I don't know, okay.

Sorry, Ben, go on.

BEN GERLOFS 36:36

Cancer, not cancer? Yeah.

PATRICK ADLER 36:40

Isn't it kind of like that? Isn't it a geometric anyway.

BEN GERLOFS 36:43

But if making maps is just detecting cancer or not cancer, yes. But if making maps has anything more to it than that, I don't know what an AI model does for you. As a student, as a researcher, as anyone else, when you try to make a map. I don't know.

GIANCARLO DEFRANCESCO 37:14

We might have just lost Patrick for a moment there.

BEN GERLOFS 37:17

That ran out of tokens.

PATRICK ADLER 37:23

Boy, when Giancarlo talks about, switching GIS models on the fly, I'm reminded of, like, people who would train me in statistics, who would be using these, like, very antiquated statistical software packages that did have the same capabilities as the things that I was trained on. Like, I was trained on R, and they would be trained on something older. And, then just would, like, live their whole careers. On those old software packages, and not only that, they would have to they would teach their students on those packages.

And this you would just, like, lock in for generations on this, like, suboptimal tool. And, you know like giancarlo, like, I found the same thing, like, you know, so I was trained on R, but then, like, there's all this functionality in Python, and I know, Tim, you use Python? And, like, AI just makes this translation R to Python just completely trivial.

Though I know Python.

TIM WOTHERSPOON 38:38

Yeah, any more of these are the same tool.

PATRICK ADLER 38:43

So, okay, so, like I do feel like, and I'm sure you can confirm, Giancarlo, you're kind of like the swing member of this jury.

You're doing so much more stuff than you used to, or than you would be able to do. You're just, like, way more productive, way more powerful, if only because there's, like, 10% of your workflow that you're doing way better. I mean, right?

Or I don't know, maybe I'm wrong.

GIANCARLO DEFRANCESCO 39:18

Absolutely, and I, you know, kind of going back to Ben's point, you know, there's still artistic choices that you're making, in that creative process, and in that intellectual process, and in that analysis process, right? How do you it's how do you know, use it as a tool to help you transition between maybe software or tool sets that you're not as familiar with, but at the end of the day, right, you need to be the one making the decision, for, this is how I want to represent the data, or these are the comparisons that I want to make. And not letting the AI Choose that for you, but letting it help you use the tool to make that, presentable in a way that your audience understands it, or is, you know, so that it's not, unclear as to how you want how you're trying to make your point. And again, right, like, you still have sometimes, right, there are times where, you know, I've had it, where I feel pretty good about something, and I've kind of checked it with AI, and then I go and I show it to somebody else, and they're like, I don't get it.

I don't get this, or I don't get that, right? So you do you still need other people in that process, right? Like, it can only get you so far because they're it's very sycophantic. And so, yeah, just being careful about how you're using it, and in the context of, say, Python or R, I use it to make plots.

For my, data, and, you know, making sure that it's plotting the data correctly, because sometimes it'll strip out data that, if you don't have it formatted right, and just check you have to check everything. But it is very useful in transitioning it.

That's how I'm still staying in limbo. It's I'm still not giving you.

TIM WOTHERSPOON 41:10

Repeatedly, you've talked about the need to check things, and then we meet undergraduates who. Their interaction with this technology is, like, before they learned any skills. So, how are they even able to check?

GIANCARLO DEFRANCESCO 41:30

Yeah, and part of the issue, right, is when you haven't done it yourself yet, that's part of kind of why I made a point where it's like, you do have to wrestle with the text and how you are phrasing things when you're actually sitting down to write. But, you know, if I didn't have the foundational knowledge that I had, would I have spotted the fact that what the AI just told me was wrong? Or if I didn't have the you know, the knowledge that I need to go and double-check, okay, here's something that I might have missed in the literature. The only way that I knew that I might have missed it in the literature was because I had read the literature, right?

And what I felt to be complete. And so you yeah, you just I keep saying that you have to check it, but yeah, again, you know, how do you get students to that point when they are not used to doing that whatsoever at this point anymore? I think that's I know for.

Well, I know that for some of my classes, I'm a little bit insulated from it, because we do stuff that AI has not been able to automate yet, like field mapping. But when I ask them for a report on what they saw, they might, you know, write down a few notes and then have that expanded upon AI. If they are doing it, I haven't caught them explicitly yet, but I'm sure it's happened and I've missed it. Sorry, Ben, what were you going to say?

BEN GERLOFS 42:49

No, I'm sorry I interrupted. I was just gonna say, I think this is part of what those editors are after in that statement, is that there's an earnestness to knowledge that is lost if you skip that step. That's I think that's true, especially for folks who have not acquired these skills, of course, that's obvious, so in the case of many of our students, for instance. But I think that's true regardless, right?

It doesn't matter how old or experienced you are, like, if you haven't earned this knowledge, there's something missing in that step. If you haven't taken the time that it would take to do this thing without the assistance of something that can just do it for you, I think there's something real that's lost in that.

PATRICK ADLER 43:28

Aren't there, like it seems like there's two, kind of, complexions to that, Ben. And the first is, like, a sort of integrity one. Like, earned as in do you deserve something? Yeah.

And this is a tone that I get in this editorial. And then the other is, like, what Giancarlo is talking about, where it's like, are you able to responsibly use this technology because you can judge it, you can restrain it. Judge good and bad outputs. Do you mean it in both senses?

Like, is this kind of like, professional cycling, where it's only really it only really counts if you do it yourself. Or is it something different?

BEN GERLOFS 44:22

The answer is yes, like, I mean it in both senses. I hope that I lean much more heavily on the utility one, and not the deservingness, not the merit side of that. I think that the merit side is an important conversation. Like, in professional cycling, there are rules, and when you violate them, and you're caught doing so, you lose your medals.

And I don't have a huge problem with that. I think it's important to revisit the rules a lot, and ask what the rules are there for, and if the rules should be changed, I think that's super important, much more so in this case than in cycling, per se. But the other side is really important. I think the thing that's lost that I care more about when you skip steps, because AI is doing stuff for you, is that there's a utility that's lost.

You're not able to do the checking, like Giancarlo was talking about. You lose the ability, I think, to or you don't ever acquire the ability, I mean, both, I think, make sense, to come up with new ideas, to create these kind of minor variations on things that are not, just, like, something generated by a predictive model, but something that a human being would genuinely ask or be curious about.

PATRICK ADLER 45:32

And this is the people dying of cancer thing that Tim's talking about. I wonder, though, like. I can see a paradigm, like, a completely transformed educational paradigm, where instead of training, like, independent experts in something, we're training people who can check AI. But even saying that, to return to what Tim was talking about at the beginning.

Seems to, like, seed a lot of, like. The thing, the, like, autonomy, the whatever that we seem to think we that we seem to think it's all about, to begin with. Like, if we're just checkers then is it even worth it anymore to, like, be in this field or something? I feel like there is a path forward, but, like.

What that implies about our role is just so intimidating to us that, like, we don't even entertain it, even though, like, pedagogically, I think that's a way you can do this, where it's like, let's become really good checkers so no one dies of cancer.

TIM WOTHERSPOON 46:49

I think, when it comes to the classroom, what we end up doing, and what me and my colleagues and what I heard Ben saying is so frustrating is it's impossible to police the work, so we just judge its

quality. So, a lot of times, just giving the essay prompt to a language model. And then giving me the essay. That essay is poor.

Right, even with the best models, that essay is always poor. So then we read this poor essay, and we say, this essay is poor for reason A, B, and C. It's very frustrating for us, because it's clear that we're not helping anyone become a better writer, or a better thinker.

All we're doing is saying, oh, however you're approaching acquiring these skills is wrong. But we can't even say, oh, you had a good idea here, it would be better if you did that. Oh, this is a nice idea, doesn't really come across, what were you really trying to say? Those skills that we really want our students to have that they're not gonna have.

So I'm very frustrated. With it. I re- I really think we need to rethink, what are we doing here? Why are we using this tool?

Why is it so why is the incentive to use this tool so strong for a student?

BEN GERLOFS 48:35

Actually, I think that can be translated back into the question that pat was asking Giancarlo earlier, or talking about Giancarlo being so much more productive, as so many folks are with this technology. But I think there's an assumption in there that more is good, that more stuff is better, right? More papers, more essays, more things that AI can help us generate more quickly? Why is that a good thing, I mean, the editors of this statement are partly responding to, like, this massive flood of papers across all disciplines that are suddenly pouring in as if there's just, like, a tenfold increase in the number of academics overnight.

Why is that productivity a good thing? I'm not confident that it is.

TIM WOTHERSPOON 49:20

It's not. I mean, I wish we were I wish it would do my laundry for me. Where are the robots?

Where are the robots, right? So.

GIANCARLO DEFRANCESCO 49:35

As well, right, it makes it challenging for us to then sift through the noise as well, and it slows that process down, because, you know, there and in some sense, right, depending on your perspective, right, it also can stifle progress if your reviewer finds an obscure article that says something that you know, was maybe published, that, you know, is maybe true for the data set that they analyzed, but maybe one in a broader context is not necessarily broadly applicable in your field, and they use that as a way to maybe disagree with your argument for reasons that are not entirely justified, right? But they and it could be an honest mistake, right? It could be that was something that they read, and it doesn't mean that data in that context isn't accurate, but when again, when you take it into a broader picture, how do you compare that? And especially in a peer review process, how do you then kind of make sure

that by cherry-picking those examples and using them in your review process, make sure that's still a fair review.

PATRICK ADLER 50:49

This is, like, a red herring problem and not, like, a hallucination. Like, it's.

It's subtly different.

I mean, but then sometimes, like, don't things that seem like red herrings, like, because of just the nature of, like, radical innovations, like, some things that seem kind of weird can lead to, like, bigger breakthroughs, and.

GIANCARLO DEFRANCESCO 51:12

Absolutely, right? And in some of the data that I even work with, right, it's like, I have stuff that looks weird, and I'm struggling to explain it in some cases. But as well, right, it's like, yeah, how do I bring that into a broader context and make sure in my writing as well, that I also communicate that, okay, this is the problem with this data, and we're still trying to understand how this relationship between these systems work. And emphasize that so that doesn't maybe get misinterpreted and that broad more broadly applied in a way that is maybe inappropriate.

But with right, if you have an increasing volume of noise, like these like this, these editors are talking about in their statement, maybe an increasing volume of noise, then and then that creates an increasing amount of material that we have to sift through, to try to understand how our fields are changing, and how we're supposed to be interpreting, maybe. The data that we have, and the things that we're thinking about. It creates a new complication as well.

PATRICK ADLER 52:26

I mean, it's fascinating, because, like, I feel like there's no, like, there's no question that there's more shit papers being submitted to journals because of AI, more AI slop. I think that's without question. But I think it's almost equally without question that, like, at the 90th percentile of quality, there's just more great papers. And this is because it's too early for the, like you know, armageddon that Tim was describing to have set in.

Don't know how to exercise judgment. Like, we have a generation, or a couple generations of scientists, or social scientists, who are trained to be discerning, and then also get the power of this technology. So, like.

There have to be more good papers, too. And so, like, going back to, like, Ben.

TIM WOTHERSPOON 53:17

I'm sure the good papers are growing. I'm sure you're right.

Like, the AI in the hands of the best researchers. Is increasing the quality and quantity of their research.

PATRICK ADLER 53:35

But going back to, like, Ben's question, like, what do we make of the fact that there's just more really good papers out there? Probably, like, just taking Giancarlo again as a kind of, like. Case in point, like, top researcher, top young researcher, he's just doing more, he's writing more papers, and he's making those papers better, because he's got access to these tools.

GIANCARLO DEFRANCESCO 53:59

I think in the field of mathematics right now, that's kind of definitely being rocked a bit by AI. There have been quite a few seemingly really large steps forward, in that field as a result of using AI to revisit stuff that was published a while ago and kind of look at something from a first principles perspective, that I think a lot of the people in that field are finding useful. I say this as someone who is not a mathematician, so, and Tim might be able to speak more to that.

TIM WOTHERSPOON 54:37

No, I mean, there's all these results that they're saying that AI found these results, but.

I know that there was a really clever mathematician involved.

GIANCARLO DEFRANCESCO 54:53

And so this is kind of maybe something that I wanted to build on, kind of what Patrick was maybe talking about, or alluded to earlier, and I think that Tim, you saying that is definitely a good asterisk, to that, is that, you know, in some sense, we almost kind of might become more of a instead of, like, a facilitator or a checker, maybe more of, like, a collaborator, where, you know, I might say, okay, you know, we got this paper, and I have questions, about, how new processes work based on the findings that I have from this paper. I can identify new places and work with the AI to help me optimize a way to sample a region if I'm looking at something, and help with the logistics side of it. AI can't help me in the field, right? Like Tim said, we don't have the robots.

And in a sense, I think, specifically in geology, where I do field work, and I have to go out, you know, my understanding of the broader systems is something that I can't translate well into AI when I'm actually standing there in the field. It might be able to aggregate a lot of data that I might not be able to in my head in a different way. But when it comes down to actually, like, subtle things that I'm noticing in the field, it's not necessarily at that point yet. But so, then it becomes, like, my knowledge, in a way that I can't be supported by AI is then going to help me collaborate with it, as I plan, my research as I collect data as I sift through and format stuff.

To just kind of help it along the process. But that's not applicable to all fields. Not all fields require that fieldwork. And to an extent, I think biology as well is another field where that's going to be a much more collaborative process.

Will that be the case everywhere? In mathematics, right, it's always with somebody else. It's always somebody, some very smart individual who is talking with the AI, who is discussing those ideas, but they're also somebody who has a good foundation in the knowledge.

TIM WOTHERSPOON 57:07

Yeah, very clever. The other metaphor that Cory Doctorow made was about centaurs. So, a centaur is a metaphor of a human that is assisted by a machine. But he talks about reverse centaur.

Is this a machine? Is this about a human? Like, the it's a horse's head and a human's legs, right? It's backwards, and I think the reverse centaur is a bad idea.

It's bad for humans, and it's bad for the quality of the work. And so, I think mostly Giancarlo and Patrick are describing the centaurs of the world, like.

What they can do. But then the kids you know, the kids these days, they're not centaurs.

The brain is the horse. Anyway, go on.

BEN GERLOFS 58:09

Just a quick note about these different metaphors. I think they work differently in different parts of the academic exercise and academic life. So in fields where, like, the PI and lab model is dominance, this idea of, like, collaborating with AI agents or whatever, makes way more sense, because this is how labs kind of work already, right? Whether someone's got kind of an idea, and then they assign that to a postdoc or to a junior researcher, please go figure out, like, how to acquire this data, or how to make these two things commensurable so we can do this analysis.

I'll go you know, schmooze with people and get the money, or I'll make sure we get the permissions, or whatever, and then you do the actual work. This sort of makes sense. In fields where that's not the case, where it's, like, one person, in the dark, stumbling around, trying to find some truth somewhere, this seems like a much more dangerous thing. It doesn't seem to work the same way at all.

Like, that kind of collaboration only produces your reverse centaurs, I think, Tim.

TIM WOTHERSPOON 59:08

Hey guys, I think I need to bounce. I'm so tired.

PATRICK ADLER 59:12

No, we've had a good hour of conversation here, but maybe

We can all, just end by saying some. New idea they have from this conversation, or new question? I don't know, I just came up with that. I just stumbled around for that one.

BEN GERLOFS 59:38

I see you typing with a hand there, you got a little AI.

PATRICK ADLER 59:52

I mean, so I'll start. One question I have of you guys, but it's really a question I have of myself, is like you know, like, I probably have the same complaints as all of you, that I'm not or that students aren't writing as well because of AI. But then, like, that should suggest that I should be teaching writing more, now.

That I have. Like, it's up to me to fill in that. And somehow. I don't teach riding very much now.

I didn't teach riding very much before. And it's like, who is responsible for people losing these skills? It's literally fucking me!

So I'm gonna be thinking about how I can integrate and it probably is going to be more of this, like, reverse centaur mode, because I think that's going to be the most accessible and relevant to these students, but how can I teach writing in some form so that these skills don't get lost? I like this reaction from.

BEN GERLOFS 1:01:06

I don't know, man, I don't know. Disgusting. You're talking about, like, a really high-order skill, like the kind of writing that you should be teaching, that you and I and the others in this conversation should be teaching, to university students, is not, like, how to construct a sentence. It's how to reach people, how to reach an audience, how to, like, draw out an idea and fully explain it, how to give something color and flavor.

This is a kind of writing that, like, human geography requires. But it's not something you can teach to people who don't know how to construct a sentence. The kind of writing that is being lost is the lower-order stuff. That's the steps that are being skipped when students.

PATRICK ADLER 1:01:51

It's a high school teacher's fault.

This is a very cynical view, Ben, and I'm sure you know. Okay. So there's not there's nothing okay, there's nothing.

BEN GERLOFS 1:02:04

There's nothing you and I can do to teach these writing skills. No, I don't think so. At the university level.

TIM WOTHERSPOON 1:02:11

Not at the scale. Right, and a Common Core course, you have 120 students. Right, maybe if I had 5.

BEN GERLOFS 1:02:24

And, like, an advanced course in my own field, as well. Like, if you had people in a seminar for 3 hours once a week, and they read a book together, and they all write, like, a critical essay about this book, and then you read each other's work, that's still possible, I think you're right, at that scale. But you're also talking about a different kind of writing at that point. I want that to be true, Pat, but I think that's a horrifyingly losing battle.

PATRICK ADLER 1:02:48

You could train an AI to write a software program to prompt students with questions, and they can assess the better question the better sentences.

GIANCARLO DEFRANCESCO 1:03:01

Something that I remember doing, in high school, secondary school, was that, in part of my writing development for English class, we had to write an essay, a draft, right? This was not our final submission. We would submit that to two copies, one to our, teacher, and the other to the other, another student in the class, or even sometimes a group of students, right?

So I might print out five copies. One for my teacher, and four for my peers. And I would receive four essays. I would have a day or two to read those, that would be my homework, and then I would have to come in and give people specific feedback and ask them why they made that choice.

Why did you say this here? if you're sitting around your peers, and you can't explain your own work or why you made that decision, I think that's a way, to kind of almost socially engineer a situation where somebody can't escape that, discussion, right? And can't, sidestep that responsibility. You submitted that work as yours, and someone is asking you about why you said something.

That's a very different environment, and they might not even be confrontational, right? It's not necessarily that they're, you know, criticizing you, but they're asking, oh, like, I really thought that you described that in an interesting way, why did you say that? The first time it happens, I'm sure that some of the students will be like, I don't know. Right?

If their peers are reviewing them, then that's not a good response, and their grade might reflect that, and they might then learn so that for the second time or third time, okay, I actually do have to take this more seriously, I have to put my own thoughts into it. I will say, though, I think that, you know, in this environment, right? Research happens in our native language, and so I do think that for a lot of people who do not speak English natively, since research is so often communicated in English, it is really helpful for minorities in our field who are, trying to, communicate their ideas, but maybe don't have all the tools that we have. And so, you know, how do you balance all those things?

Keeping that creative aspect of it, and that responsibility and accountability, while also, like, supporting people who might still need to use these tools to, you know, work on their own, and again, not theirs, not their fault, but, you know, their disadvantages.

PATRICK ADLER 1:05:55

I think it was great. I mean, tongue-in-cheek again, it's like, you know, you let the we let the AIs grade for us, and then we make the other students shame and regulate. And then we go walk all the way to the bank. Anyway, no, there's some really good, points there, and then, like, the language thing is obviously I mean, because language is the thing that AI is the best at, it's a language model, and our students are just, like, doing so much better with it.

But then, I guess, to, like, Ben's concerns and some other concerns, like maybe they're writing with it, but then they're speaking even more poorly. Because they don't know it Ben, any last words? I mean. You didn't like mine.

You sure didn't like mine.

BEN GERLOFS 1:06:44

No, I loved it. I also hated it.

My one thing was gonna be about your essay, which we did not talk about, which I found to be much more interesting and exciting than this statement by these editors, which, although I completely understand it and largely agree with it, is sort of sad. It has kind of a frustrating tone.

Yeah, your essay is quite hopeful. I maybe another night on this show hear the voice of the new territories, we'll get to discuss, that piece. I think it has a lot of exciting stuff in it. I have lots of questions and critiques to make.

PATRICK ADLER 1:07:26

Well, I look forward to that. Tim, anything and then thank you so much, Ben, so generous with your time, so generous with that, Tim, you got anything else? dean of the Westerners, sending us off.

TIM WOTHERSPOON 1:07:37

Oh, yeah, it's very interesting to talk to you guys about this. I think, maybe 4 years ago or so, like, I wasn't so worried about AI because, it was bad. And, it's better now, and I'm way more worried about it.

So, I think what I've been learning is really that the quality of the AI to do the thing that we thought only a human could do. It's the most, terrifying, because I think we will lose the ability to critique work. And this is really what I'm scared of. So, it's been nice to talk to you guys about it.

But I'm up way past my bedtime.

PATRICK ADLER 1:08:40

No, you get over jet lag, it's a good opportunity to release everyone, I'm really appreciative. This was this is a better conversation than I anticipated.

The highest quality conversation the four of us have ever had, and we've had.

Numerous. I mean, no, just saying, like, we.

GIANCARLO DEFRANCESCO 1:09:02

I was about to be hurt, but then I, you know.

PATRICK ADLER 1:09:05

No, we went high with this, so yeah, I'm gonna.

TIM WOTHERSPOON 1:09:10

I didn't use any AI to prepare my remarks. Let's close that now.

BEN GERLOFS 1:09:16

Skills were harmed in the making of this.

PATRICK ADLER 1:09:19

I'm gonna post and transcribe this, and, you know, I think. People are gonna get something out of it, so thank you guys so much.

Sorry that the stream itself didn't work.

BEN GERLOFS 1:09:32

Thanks for hosting, Pat, thanks for inviting us.

PATRICK ADLER 1:09:34

This was great, thank you guys.

BEN GERLOFS 1:09:36

Thanks for being brave and doing this kind of wild stuff, also.

PATRICK ADLER 1:09:42

Alright. See you guys soon, I think.

BEN GERLOFS 1:09:46

Hey, everybody.

Transcript generated by Zoom and edited by Claude Opus 5: backchannel interjections and broken fragments removed,
a small number of misrecognised proper nouns 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