Stochastic parrots, cap and gown edition

Chris von Csefalvay

Blog post published December 6, 2023 in *Chris von Csefalvay*
https://doi.org/10.59350/412q0-sbn89

Abstract

It's not every day that you find out you have climbed the exalted heights of another discipline. My work is pretty interdisciplinary, but it shocked me, too, that I'm apparently holding forth on neoliberalism and the epistemic question in African universities (archive link): This, of course, came at some surprise to me, as I have never written anything on the topic.

Keywords

Epidemiology (Subfield), AI, Academia, Writing, LLMs

Recommended citation

von Csefalvay, C. (2023, December 6). Stochastic parrots, cap and gown edition. *Chris Von Csefalvay*. https://doi.org/10.59350/412q0-sbn89

Copyright

Copyright © Chris von Csefalvay 2023. Distributed under the terms of the Creative Commons Attribution 4.0 International License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

It's not every day that you find out you have climbed the exalted heights of another discipline. My work is pretty interdisciplinary, but it shocked me, too, that I'm apparently holding forth on neoliberalism and the epistemic question in African universities (archive link):

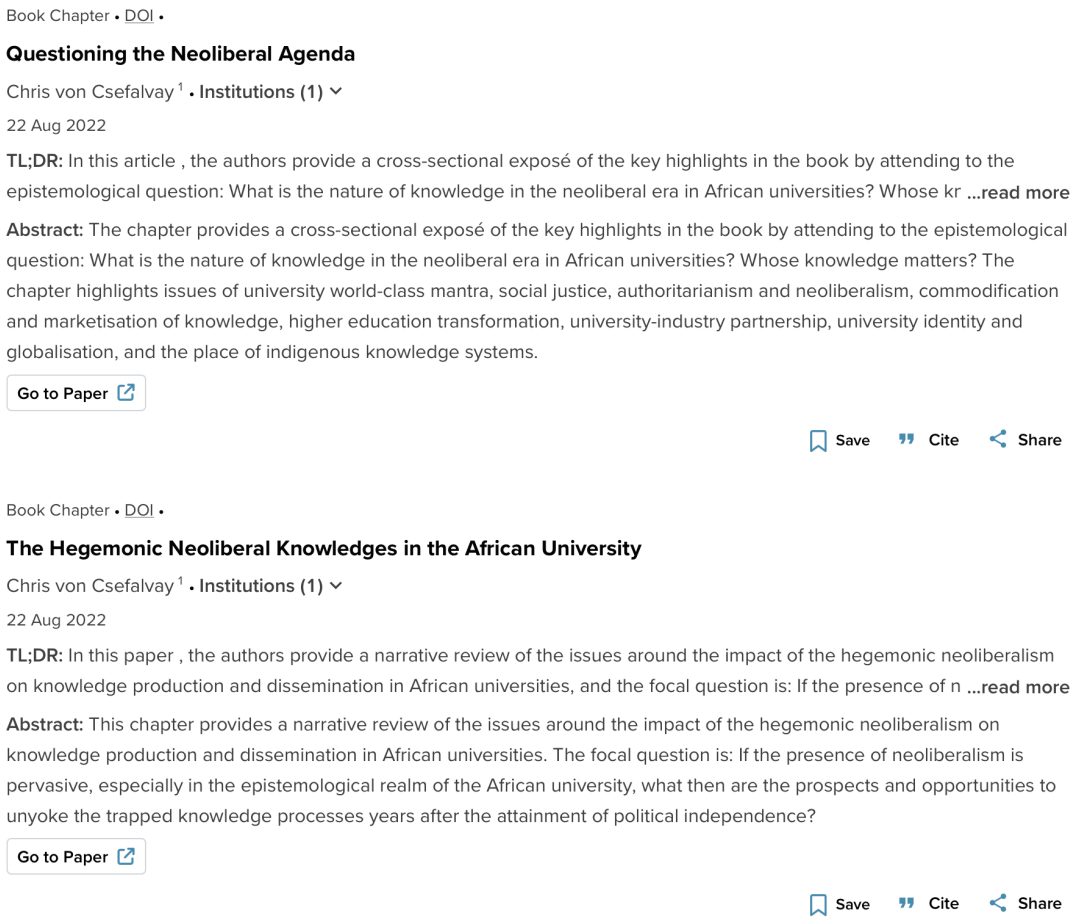

Book Chapter • DOI •

Questioning the Neoliberal Agenda

Chris von Csefalvay [1] • **Institutions (1)**

22 Aug 2022

TL;DR: In this article , the authors provide a cross-sectional exposé of the key highlights in the book by attending to the epistemological question: What is the nature of knowledge in the neoliberal era in African universities? Whose kr **...read more**

Abstract: The chapter provides a cross-sectional exposé of the key highlights in the book by attending to the epistemological question: What is the nature of knowledge in the neoliberal era in African universities? Whose knowledge matters? The chapter highlights issues of university world-class mantra, social justice, authoritarianism and neoliberalism, commodification and marketisation of knowledge, higher education transformation, university-industry partnership, university identity and globalisation, and the place of indigenous knowledge systems.

Go to Paper

Save Cite Share

Book Chapter • DOI •

The Hegemonic Neoliberal Knowledges in the African University

Chris von Csefalvay [1] • **Institutions (1)**

22 Aug 2022

TL;DR: In this paper , the authors provide a narrative review of the issues around the impact of the hegemonic neoliberalism on knowledge production and dissemination in African universities, and the focal question is: If the presence of n **...read more**

Abstract: This chapter provides a narrative review of the issues around the impact of the hegemonic neoliberalism on knowledge production and dissemination in African universities. The focal question is: If the presence of neoliberalism is pervasive, especially in the epistemological realm of the African university, what then are the prospects and opportunities to unyoke the trapped knowledge processes years after the attainment of political independence?

Go to Paper

Save Cite Share

Figure 1: Apparently, I'm commenting on neoliberalism in African universities.

This, of course, came at some surprise to me, as I have never written anything on the topic. I have, however, written a lot about AI, and I have written a thing or two about Africa, so I guess it was only a matter of time before I was conflated with someone else. This time, the unwitting victim deprived of his credit was Prof. Amasa P. Ndofirepi, who is an educational studies scholar at the University of Johannesburg. I have no idea how I ended up being credited with his work, but I'm sure it was an honest mistake.

The problem is, with AI, mistakes compound. So if an unwitting student were to ask for a quick literature review of neoliberalism on the subject, they might get something like this from Scispace:

> *The literature on the impact of neoliberalism on knowledge production and dissemination in African universities has been extensively explored by various authors. Qosimova Gulbahor, in her paper "Placing Knowledge at the Centre of an Alternative Public Good Imaginary of African Universities," discusses the alternative public good mission of African universities and the need for them to apply their knowledge infrastructure to community development challenges. Chris von Csefalvay, in his paper "The Hegemonic Neoliberal Knowledges in the African University," examines the pervasive presence of neoliberalism in African universities and explores the*

> *prospects and opportunities to unyoke the trapped knowledge processes. These authors, along with others, highlight concerns about the dominance of Western knowledge, the commodification of knowledge, and the need for African universities to prioritize socially-just knowledges that serve African priorities and challenges.*

I mean, that's flattering, but I'd really rather be credited mostly for my own work. I'm sure Prof. Ndofirepi would agree.

Why I care

This is, of course, not good for academia. We've generally been coasting from one crisis to another. We've got a replication crisis, there's enough dodgy Western Blots to blot out the sun, we've got the Tessier-Lavigne mess, and that's just what I can think of off the top of my head before my first coffee. A predatory publishing industry doesn't help this at all. We need another crisis on top of this like we need a hole in the head, and yet, here we are.

Now, as far as I'm aware, no serious academic is actually using these tools to do their research. On the other hand, non-academics *are*. For journalists, in particular, such tools are a godsend – literature reviews are annoying, and if you can get a computer to do it for you, why not? The problem is, of course, that you're supposed to double-check this stuff and, well, journalists are known for many things, but double-checking stuff properly isn't really one of them.

And so, after months of academics fretting about ChatGPT eating their lunch, we're confronted with the actual problem. AI is not better at producing decent science, but it is vastly faster and more efficient at producing *bad* science.

Which we weren't short on to begin with.

What's the problem?

Language is a tool that works on the basis of some conventions of meaning. Language models encapsulate these conventions, but they cannot encapsulate *all* of them – there are compromises to be made if a system with limited resources has to contend with nearly unlimited human imagination. When language models' limits come to blows with domain-specific language, we get into trouble.

Language models are really weak at one thing: reasoned judgment. As a scientist, you are trained to exercise this kind of reasoned judgment in determining what is, and what isn't, worth considering as an authority. That's why we make our master's students (and hopefully most undergraduates) write literature reviews until the cows come home. It teaches them to develop that judgment, and also to know how to explore the fringes of their research question. I have looked at a few 'academic AI' tools that claim to be doing some of this, and they're not very good at it. Scite is so far one of the better ones, and the literature reviews it produces are still pretty bad: results are heavily weighted towards recent publications, towards the specific in preference

to the foundational and often towards meandering misinterpretations of the research question as long as sources for that could be found, in preference to actually identifying a gap.

To be quite fair towards these models, they have to deal with academic literature, which is an abundance of noise with a flicker of signal. There is, not to put too fine a point on it, a ton of crap out there, and it's not always easy to tell the difference between the crap and the good stuff. That's why we have peer review, and that's why we have literature reviews. The problem is, of course, that these models are not trained on the literature, but on the internet. And the internet is a very different place from the academic literature indeed.

Academic writing, especially domain specific writing, has a language of its own. It's not fair to expect a language model trained on English to also master uses of English that might as well be a different language. To give a favourite example of mine: in magnetic resonance imaging of the brain, there's often talk of something called 'flow voids'. Now, normal human reasoning would interpret a 'flow void' to be the absence of flow, or something along those lines. In MRI, a flow void is actually the opposite: it is a 'void' of signal created in a vessel through which something (usually blood, sometimes CSF) flows.[1] This is a very specific term that has a very specific meaning in a very specific context. A language model, however, would not know that. It would assume that a 'flow void' is (de)void of flow.

[1] Flow voids happen in the context of spin-echo imaging. These modalities involve two pulses – an excitation pulse and a refocusing pulse. Blood that moves perpendicular to the image plane will be hit by the excitation pulse but not the refocusing pulse. Therefore, it will not create a signal, which gives us the 'void' appearance of signal hypointense vessels.

Just about *all* of science is like that. We have a language of our own, and it's not always easy to understand. What definitely doesn't make it easier to work with, however, is when the source material is also wrong. Which is what we're dealing with all too often, viz. Figure 1.

What can we do about it?

Most 'academic AI' applications are riding on the crest of a wave of high expectations that surround everything AI-related right now. They offer to be useful aides-de-camp to beleaguered academics who have to contend with exponentially growing literature, but in reality fall far short of that promise. And the inherent 'black box' nature of such models means that it's not always easy to tell when they're wrong.

At this point, perhaps the best we can do is to hold off on using generative AI tools for academic research until they're better. We're not there yet. We're not even close. For highly domain specific applications, retrieval-augmented generation (RAG) approaches utilising a curated knowledge base of publications in that realm has proven to be a very useful approach indeed, but those are specialised tools that are at the present primarily in the purview of private industry. I've seen some great applications in this field, and I see this as another proof point for my assertion that the future belongs to ecosystems of small, specialised language models rather than one big model that does everything.

For the time being, we're going to have to do things the old-fashioned way: reading the literature and writing our own reviews. I know, it's not very fun. But it's the only way to do it right, and we've managed with that for the last few hundred years. Giving it another few years won't hurt.

https://doi.org/10.59350/412q0-sbn89