Vol.9 No.2 (2026)

Journal of Applied Learning & Teaching

ISSN: 2591-801X

Proudly owned and sponsored by Kaplan Business School, Australia

Content Available at: **https://jalt.open-publishing.org/index.php/jalt/index**

Not anti-technology, but anti-capitalist: A critical theory of AI. An interview with Professor Simon Lindgren

Simon Lindgren [A] — [A] *Umeå University, Sweden*

Jürgen Rudolph [B] — [B] *Murdoch University, Singapore*

Fiona Xiaofei Tang [C] — [C] *Kaplan Business School, Australia*

Fadhil Ismail [D] — [D] *Kaplan Higher Education, Singapore*

Shannon Tan [E] — [E] *Amity Global Institute, Singapore*

Keywords

Artificial intelligence;
critical theory;
discourse analysis;
generative AI;
higher education;
inequality;
solutionism.

Correspondence

tshannon@singapore.amity.edu [E]

Article Info

Received 9 March 2026
Received in revised form 10 April 2026
Accepted 13 April 2026
Available online 25 May 2026

DOI: https://doi.org/10.37074/jalt.2026.9.2.1

Abstract

In this interview, Professor Simon Lindgren – Professor of Sociology at Umeå University, Sweden, and founder of the DIGSUM Centre for Digital Social Research – reflects on the intellectual trajectories, institutional frustrations, and critical theoretical commitments that underpin his recent and forthcoming scholarship on artificial intelligence. Drawing on *Critical Theory of AI* (2024), the *Handbook of Critical Studies of Artificial Intelligence* (2023), and the forthcoming *Discourse Machines* (2026), Lindgren argues that AI functions as an "empty signifier" – a concept sufficiently loose to serve almost any ideological agenda, yet firm enough to reshape policy, power relations, and everyday life.

He critiques responsible AI and AI ethics frameworks as structurally insufficient and introduces his concept of "bias bias" to describe the tendency to reduce systemic inequalities, such as racism and sexism, to technical calibration problems. Lindgren addresses the deepening of inequality through corporate concentration of AI development, the marginalisation of minority languages, and the collision between cheating discourse and efficiency discourse within universities. He advocates for a literacy-centred, de-dramatising pedagogical response to generative AI, stressing that domain knowledge remains essential for meaningful human–AI collaboration. The interview concludes with Lindgren's clarification that critical scholarship is not anti-technology per se but is fundamentally opposed to the capitalist structures and determinist logic that currently govern AI's development and deployment.

‘I realised there was something called sociology’

Jürgen Rudolph (JR): Thank you for agreeing to this conversation, Simon. Your prolific publication record spans from the late 1990s to the present, beginning with work on youth culture and violent crime in Sweden. Could you tell us about your own schooling and what initially drew you to sociology and media studies? Were there formative experiences or mentors who shaped your intellectual trajectory?

Simon Lindgren (SL): My earliest academic work was indeed on youth culture and media, though it extended into historical sociology and, in part, criminology, owing to the research environment at the department at the time. Those publications on youth and violent crime came from a project I worked on as a research assistant early in my career.

My situation is perhaps unusual in that I was born and raised in Umeå, a small university town in northern Sweden, and I have spent my entire academic career there. My path to sociology traces back to the final year of high school – the Swedish gymnasium – when I was around 19. We were required to complete a small thesis project, and at the time, I was an avid record collector with a deep interest in the history of rock music.

I decided to focus my project within the social sciences, writing about youth groups and movements across history – the Beat Generation of the 1950s, the hippie movement of the 1960s, and punk in the 1970s. Writing in the early 1990s, I also tried to make sense of the various youth subcultures of the 1980s. As I worked on it, I found myself less interested in the music itself and more in these cultural movements as social and political phenomena. It was i-

Figure 1. Simon Lindgren as a young student.

-n the school library that I discovered sociology: the books I kept reaching for, on youth culture and youth rebellion, turned out to be sociology books. Until then, I had assumed I would become a lawyer, inspired, I confess, by American TV dramas like L.A. Law – the whole courtroom drama seemed compelling.

Figure 2. L.A. Law (Steven Bochco Productions) logo. Fair use. L.A. Law was an American legal drama television series that ran for eight seasons from 1986 to 1994.

When the time came to apply to university, however, I thought: 'Well, I'm not ready to become a lawyer'. That felt like something for a future version of myself who had properly grown up, which did not yet feel like me. So I enrolled in sociology, drawn by the possibility of turning these interests into a career. The realisation that one could study youth and popular culture at university was genuinely formative: it pulled me away from the pragmatic plan to study law and towards something that, at the time, simply seemed more fun.

After my undergraduate studies, I was fortunate to enter a doctoral programme, working on a project about young people in rural Sweden, focusing on media representations of youth culture and incorporating a historical dimension. The next significant shift came after completing my PhD. I secured a lectureship and, fairly early on, became director of undergraduate studies – primarily an administrative role that I found stimulating for a few years, but eventually I wanted to return to research. By then, I had become fascinated with the emergence of the Internet.

I was drawn to the emerging field of Internet research, which at the time had an exciting, pioneering energy – many of its leading figures were young American women doing genuinely pathbreaking work. In 2006, I had a research proposal on online piracy approved, my first funded project, and from that point I regarded myself as an Internet researcher of sorts, though the field has naturally evolved – through datafication and, most recently, AI.

I hold a Full Professorship in sociology today, and the path through youth culture research introduced me to the cultural studies tradition – the Birmingham School, Stuart Hall, and, more broadly, critical theory. Pierre Bourdieu, Michel Foucault, and critical theory in its wider sense have been lasting influences. That critical, qualitative foundation, combined with a gradual shift towards digital objects of study and then digital methods, eventually led me to write the book, *Data Theory,* in 2020.

It is not a biography or memoir, but I still view it as an attempt to bring together the different strands of the academic identity I had accumulated over those years. I wrote it during a sabbatical semester – a stocktaking, in a sense, of everything that had brought me to that point. I was also fortunate to be part of establishing the DIGSUM Research Centre in 2016, and the work has built from there.

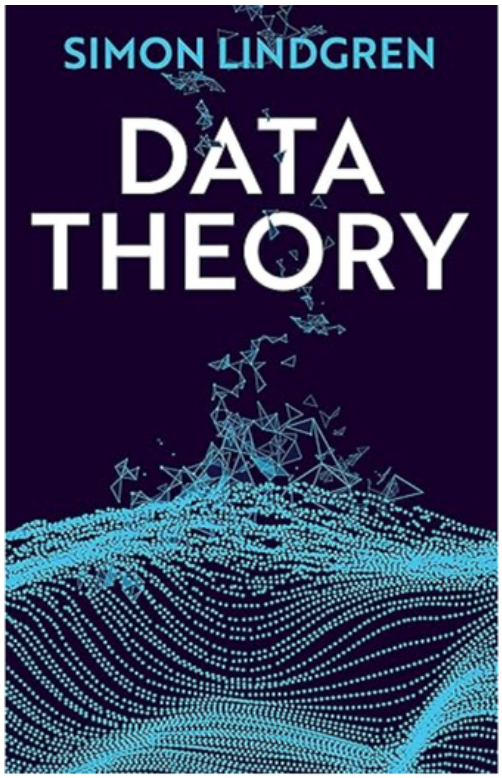

Figure 3. Cover of Lindgren's *Data Theory* (Polity Press, 2020).

The underdog university: Periphery as critical vantage point

Fiona Tang (FT): Your career has been based primarily at Umeå University in northern Sweden, where you have held positions, including as a Professor and Research Centre Director. How has your geographic location influenced your research perspectives?

SL: Historically, Sweden had four universities – Stockholm, Gothenburg, Uppsala, and Lund – the northernmost of which, Uppsala, lies some 550 kilometres south of Umeå. In the early to mid 1960s, roughly a decade before I was born, the decision was taken to establish a regional university in the north.

For many years, Umeå was only the fifth university in Sweden, and it always had a distinctive profile within the Nordic context. In the late 1960s, it was known as the Red University – the student movement was strong, and there has always been a tradition of social engagement and progressive politics. It was younger and more irreverent than the established institutions; the underdog, really. And culturally, it has been somewhat easier here to develop research areas that were not yet institutionally established – digital social research being a clear example, having emerged in response to broader changes in the world.

It has been a productive environment for developing new research areas. We are distant from major policy centres – far from Silicon Valley. In today's globalised world, we are, of course, all connected, but a peripheral position does carry certain advantages. It puts less pressure to adopt the dominant framings, and there is more room to develop alternatives. The broader Nordic context shapes our scholarship in other ways too: there is a long democratic tradition, a strong tradition of labour rights – even if we are now witnessing a turn towards the populist right in Swedish politics.

> ***A peripheral position does carry certain advantages. It puts less pressure to adopt the dominant framings, and there is more room to develop alternatives.***

Figure 4. Professor Simon Lindgren, Umeå University, 2012.

There is something particular about this place that is perhaps more noticeable to people coming from outside than to those of us who grew up here. We are close to the Arctic, and we have the Sámi Indigenous population – both of which shape much of the university's research and have spurred work on questions of whose perspectives count. A very concrete example is the question of minority languages in relation to large language models. More broadly, I think it is a mindset thing: we are in a small but energetic and creative town, the university still feels youthful, and there is a culture genuinely open to critical perspectives. I have held visiting positions in the UK and now at Helsinki University, and I travel as much as any academic, but being rooted in this particular environment has probably shaped my scholarly outlook more than I sometimes appreciate.

FT: I was particularly struck by the mention of the Sámi, who, to the best of my knowledge, are among the oldest and most long-established Indigenous communities in Europe.

SL: Yes, indeed. The university has a dedicated centre for Sámi research and a centre for Arctic research. Those perspectives feed into the university's broader critical sensibility – alongside work on gender, social class, and related issues. It certainly contributes to the university's identity. The Sámi dimension has shaped some of our scholarship in direct ways, and others more indirectly, through a kind of shared sensibility and identity. Many of us from northern Sweden have Sámi heritage not far back in our family trees.

The politics of AI definitions

Fadhil Ismail (FI): In *Critical Theory of AI*, you argue that AI functions as an "empty signifier" (Lindgren, 2024, p. 7). You cite Lanier's provocative claim that "AI is a fantasy, nothing but a story we tell about our code" (2018, p. 135), and Ruha Benjamin's observation that "many things dubbed 'AI' today are, basically, just statistical predictions rebranded in the age of big data" (2019, p. 211). In the *Handbook*, you identify among the key tenets of critical AI studies that "AI is the subject of evolving wars of definitions" (Lindgren, 2023, p. 17). How do you navigate this definitional terrain in your own work, and what is at stake when we accept AI as a settled, coherent concept?

SL: This is really at the heart of what the book tries to do. If you work as a developer, you probably do not lose much by treating AI as a settled concept – in fact, it is more pragmatic that way. But from a sociopolitical perspective, AI's definitional openness is precisely the problem. It is a good example of what discourse theorists call an 'empty signifier' – though from a science studies perspective one might also call it a 'boundary object'. It is a concept loose enough to mean almost anything, depending on the context and the agendas of those deploying it, yet firm enough to carry real discursive power, shaping policies, power relations, and people's everyday lives.

> ***It is a concept that becomes loose enough so that you can mean almost anything depending on what the context is – but it's also firm enough to have quite a strong discursive power.***

The quotes you cited from Jaron Lanier and Ruha Benjamin both illustrate how AI can be filled with different content by different actors, depending on their interests and positions. Lanier's claim that AI is only a fantasy is deliberately provocative, but it pinpoints a real mythological element – we construct narratives about what this technology can do, and those narratives take on a life of their own. Benjamin draws attention to something more concrete: that much of what is currently marketed as AI is the same kind of statistical pattern-matching we saw in, for example, search engines before, only rebranded.

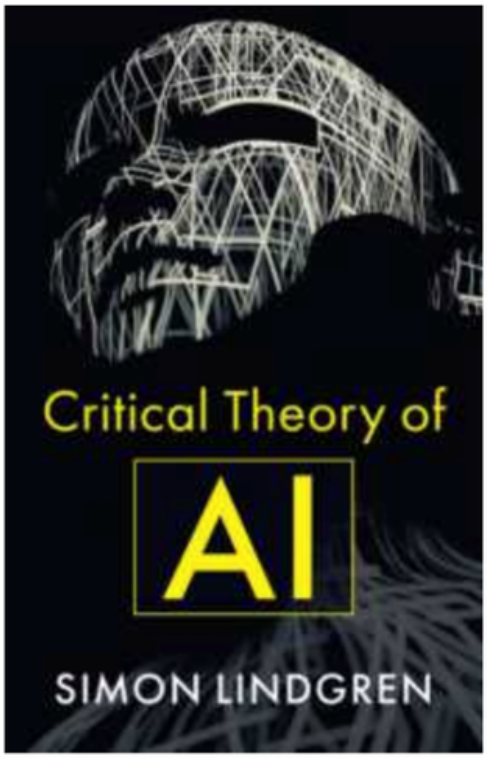

Figure 5. Cover of Lindgren's *Critical Theory of AI* (Polity Press, 2024).

So, how does one navigate this? Not by chasing a true definition, but by treating the definitional instability itself as an object of study – as the site of ongoing struggles over meaning: who gets to define AI, who is excluded from that process, what we use it for, in what direction it develops, which risks are taken seriously, and which communities become marginalised.

Put simply, it becomes a buzzword – deployable in any debate, at any moment, in service of almost any position. Once AI is naturalised as an unstoppable, self-propelling transformative force – the vision promoted by a certain strand of tech-capitalist thinking, associated for instance with figures at OpenAI – we stop asking who built it, who benefits from it, what we miss by calling it 'intelligent', and what alternatives might exist that our current way of talking about it renders unthinkable.

This is what the notion of the 'AI assemblage' in the book gets at: if we do not open the black box of the concept itself, deeper critical analysis becomes very difficult. This is also what is lacking in approaches such as responsible

AI or ethical AI – they tend not to unpack the concept sufficiently, so critique becomes more like auditing, or what people call 'ethics-washing': 'We had social scientists review it, and they said it was fine'.

From a social constructionist and discourse-analytical standpoint, these are always empirical questions that need revisiting, because the same technology can have very different consequences in different settings. The only way to navigate this, in my view, is to accept that we cannot chase a final definition, and that definitional instability is itself an object of study – and not merely a matter of language games. As discourse theorists have long argued, definitions have a very material dimension for people's everyday lives.

JR: To follow up on this, in the light of recent events: the term 'artificial intelligence' traces back to the 1956 Dartmouth Conference, where several competing terms were in use, and it was the one that most captured the Pentagon's imagination – and secured its funding. Given that foundational military connection, and given more recent reports about AI's role in active military operations, do you see a deep structural relationship between AI and the military-industrial complex? And does that sit comfortably alongside the 'empty signifier' argument – or does the military dimension give AI a very particular kind of fixity?

SL: It is worth recalling that the Internet's origin story places it squarely within the American defence establishment – ARPANET, packet switching, and so on – and we know how that chapter ended: commercialisation, and eventually the Internet as we know it today.

With AI, though, the military dimension is frightening in a qualitatively different way. When AI is deployed in warfare, the stakes are no longer whether a model makes it harder for someone to access social support or secure a flat – they are matters of life and death and global security. What I find particularly troubling is how openly some corporations appear willing to engage in these collaborations, according to reports. The image of tech company executives gathered at the same table with Donald Trump at the start of his second presidential term speaks for itself. Whether or not this connects neatly to the empty signifier argument, it exemplifies the logic of: because we have this technology, we should use it. And when it is deployed in these contexts, a great deal can go wrong – and already does. Where AI ends, and human agency begins in all of this, I honestly cannot say.

JR: Absolutely. The proximity of figures like Elon Musk and Sam Altman, alongside many other tech oligarchs, to Donald Trump is striking – not least because Silicon Valley has long cultivated an image of progressive liberalism. Google's motto was famously 'don't be evil'. Things have changed quite dramatically: these companies not only supported Trump's campaign but appear willing to share data with remarkable openness. Figures like Peter Thiel have thrived precisely where democratic oversight is weakest – Palantir being the exemplary case of a firm that treats the surveillance of populations not as a regrettable necessity but as a core and scalable product (Zuboff, 2019).

Writing *Critical Theory of AI*: From institutional frustration to scholarly intervention

Shannon Tan (ST): What motivated you to write *Critical Theory of AI*, and what are its central arguments? You note in the book that, when it comes to critical theory, there is both a narrow and a broader definition (Lindgren, 2024). The narrow definition points to the Frankfurt School – Horkheimer, Adorno, Marcuse, Habermas – while the broader one encompasses feminism, postcolonialism, queer theory, and other perspectives unified by a focus on "the analysis and questioning of power structures in order to overcome them" (Lindgren, 2024, p. 20). How did you arrive at this broader framing, and what do these perspectives specifically contribute to understanding AI?

SL: There were two levels to what drove me to write the book – a personal or local one and a broader one. At the personal level, from around 2020 onwards, and with a particular boost when generative AI became widely known with the arrival of ChatGPT at the end of 2022, my university – like many academic institutions around the world – began talking a great deal about AI, about major investments in this type of research, and essentially everyone started asking: 'In what ways am I an AI researcher? What can I contribute?' There was also, for understandable reasons, a certain frustration. In fields like computer science, AI had long existed as an established specialisation. But the way we began talking about AI in the early 2020s was about something more: not only the technological advances in machine learning and neural networks, but, at least as much, the impact on society.

As a sociologist sitting on various university committees where AI research strategies were being developed, I repeatedly felt a systemic frustration. The question of what AI looks like as an object of study for critical social science and humanities scholarship was simply not being asked. Instead, the discussions tended to circle around questions like: 'What counts as true AI research? Does it need a technological component? Does it have to involve building systems?' The implicit assumption seemed to be that AI research meant building the systems, and that social scientists or humanities scholars might be brought in alongside – but that studying AI as a social and political phenomenon from the outside was not quite the real thing.

That frustration is what personally drove me. I wanted to say: as a sociologist, this is how I see AI as an object of study, and this is how the conceptual vocabulary I bring – drawn largely from critical theory – can be used to approach it. But this also goes beyond personal experience. At the policy, industry, and media levels, the discussion around AI since the early 2020s has engaged very little with critical social science or humanities perspectives. Even many of those trying to do something valuable – developing ethical AI or responsible AI frameworks – have tended to remain at the level of principles and guidelines, without asking the deeper structural questions.

So what drove me to write the book was, first, to show those who already share these theoretical commitments how we might collectively approach AI as an object of study – they are probably the core audience. But I also hoped to reach those less familiar with this tradition: to make the case that engaging with AI critically at a deeper socio-political level requires bringing in feminist theory, postcolonial theory, queer theory, critical race theory, disability studies, and so on. Without that background, people tend to reinvent the wheel, and what results is something considerably less critical and more narrowly pragmatic. I suppose I was writing for an imaginary audience I wanted to persuade.

ST: How did you arrive at this broader framing, and what do these perspectives specifically contribute to understanding AI? Why does your book rely on the broader definition of critical theory rather than the narrower one identified with the Frankfurt School?

SL: I did receive some criticism for this in one review – that I used too broad a definition of critical theory and that it risks becoming diluted – and I can see the force of that. 'Critical Theory of AI' is certainly a striking title if we take critical theory in a very expansive sense, meaning simply a theoretically rigorous and critical stance. But it is, of course, a concept with a great deal of baggage and history. The narrow definition points to the Frankfurt School – the post-Marxist and Marxist perspectives of Adorno, Horkheimer, Marcuse, Habermas, and others. There are many other traditions that are, in a broad sense, critical – feminism, postcolonialism, queer theory, and so on. These share the same emancipatory agenda that the Frankfurt School articulated: taking sides, standing with marginalised and oppressed groups.

Not everyone would agree that there is a straight line between these traditions, and it is certainly not a simple family tree descending from the Frankfurt School. I do not have a definitive answer to this – this is simply how the book is framed. Interestingly, there is a Swedish translation – which feels slightly odd given that Swedish is my first language and I wrote the original in English – and the Swedish title translates back roughly as 'AI: A Critical Perspective'. The content is the same, but the title is different, presumably because a direct equivalent did not quite work in Swedish. When I teach critical theory, I follow the same structure: start with the Frankfurt School and branch out into these other traditions. I am not committed to arguing they constitute a single family, but they do share a fundamental concern with critiquing power and making visible the imbalances in society.

JR: I think my own position is very close to yours on this. The Frankfurt School thinkers were, after all, white, bourgeois men, however brilliant – and however real their own suffering: Walter Benjamin did not survive the Nazis, and Marcuse and Adorno only found their footing in the United States after considerable struggle. That background meant that race, feminism, anti-colonialism, and ableism were systematically underdeveloped in their work. I remain a great admirer of Marcuse and Adorno, and, to a lesser extent, Habermas, but they need to be updated and extended beyond class analysis to address colonialism, racism, sexism, and ableism. In my view, you-

-r broader framing is not only legitimate but necessary. Critical theory cannot be cast in stone. Indeed, Adorno and Marcuse themselves insisted that critical theory must remain self-critical – it should not be taken as doctrine or treated as a set of final answers.

In your earlier article with Holmström (2020), you outlined building blocks for a research agenda on AI from a social science perspective. How has your thinking evolved since then, particularly in the wake of the generative AI explosion?

SL: We wrote that piece when AI was only just emerging as a field for social scientists – the generative AI explosion had not yet happened. Now, of course, it is all about GenAI. My thinking has certainly evolved, or at least been extended, because the most pressing questions now feel rather different: they are about knowledge production and epistemic authority. The moment you embed LLMs into everything, the relevance of a discourse perspective becomes very sharp. A good example of how my thinking has shifted is the book I have just been finishing – submitting the final edits this month [March], in fact. It is with Polity Press, the same publisher as *Critical Theory of AI,* and it will be called *Discourse Machines*. It is a narrower and more specialised treatment of some of the issues raised in the *Critical Theory of AI* book. It is very much LLM-focused. The subtitle is *Power and Knowledge in the Age of Large Language Models*. My thinking has not fundamentally changed, perhaps, but the discourse analysis expertise I have been developing has become considerably more relevant in the age of LLMs.

In this forthcoming book, I try to develop the argument in detail: not how critical theory, as a broad field, is relevant to AI, but how discourse analysis, as a specific field, is relevant to large language models, as a specific phenomenon within AI. I draw on various discourse theorists and work with notions such as Donna Haraway's concept of situated knowledges – and argue that this is precisely what these models lack: their text is not grounded in any lived experience, however convincing it may sound.

The same applies to Foucault's account of how regimes of truth emerge – these are precisely the kinds of regimes that get locked into language models. So, to answer your question directly: the trajectory from that paper with Holmström to my most recent work has been a movement from a broad, programmatic mapping of what social science can contribute to AI research – covering everything from labour markets to ideas – towards a more focused engagement with generative AI specifically, and in particular with the politics of knowledge production. The broad outlines from that earlier paper still hold, but the object of analysis has transformed, and I have found a sharper specialisation in the political dimensions of large language models and their outputs.

Beyond ethics washing: The handbook, solutionism, and what responsible AI leaves out

FT: In 2023, you published the *Handbook of Critical Studies of Artificial Intelligence*. You also co-authored a chapter with Virginia Dignum on AI solutionism. In that chapter you challenge "the common view that AI, if only constructed in the right way, can be a catch-all solution to a range of social problems" (Lindgren & Dignum, 2023, p. 163), drawing on Morozov's notion of technological solutionism – the belief that technology can function as "catch-all remedies for making society better" (Lindgren & Dignum, 2023, p. 164). You also argue that "attempting to solve bias algorithmically falls short of understanding the complex and multi-faceted different ways in which bias manifests in society" (Lindgren & Dignum, 2023, p. 167). What does this solutionism look like in practice today, and what alternatives do you envision?

SL: I think Morozov wrote this book *To Save Everything. Click Here* – in 2013, if I remember correctly – and it was not focused on AI specifically but on the solutionist mindset around the Internet and new technologies more broadly. The solutionist mindset predates digital technologies, of course, but its core logic is this: you invent the solution before you have fully understood the problem you want to solve.

We see this in user-led innovation: you build a rudimentary app, release it, watch how people use it, and develop the functions accordingly. But in more serious terms, the solutionist logic is now everywhere. The question is no longer whether we need AI, but how we will use it. Institutions – schools, corporations, universities – implement AI because they believe it must always be the right solution, and the problem gets defined after the fact. This becomes particularly stark at the macro level: if our goal is not merely to optimise a recruitment process or mark s-

-tudent assignments but to fix structural inequality in society, the solutionist response is to build a 'fair' model and declare the problem solved.

This is the logic of the technological fix. A model trained on image data reflecting historical patterns of dominance will struggle with facial recognition for people of certain skin tones. The solutionist response is to tweak parameters so the model pays less attention to skin colour – and declare the structural problem solved. But of course it has not been. The critique of solutionism insists on a prior question: before we ask what AI can do here, we should ask whether we need AI here at all – and that question can only be answered in relation to the kind of society we want. Which returns us to the core critical theory questions: who benefits, and what do we lose?

> ***Before we ask what AI can do here, we should ask whether we need AI here – based on what kind of society we want.***

Morozov makes the point that some things in society do not need solutions – they need to be approached in more complex ways. We cannot solve racism, sexism, or ableism; we can only tackle them carefully and collaboratively. That also becomes an argument for meaningfully involving social scientists, humanities scholars, affected communities, and democratic institutions in the development of AI systems – not bringing them in after the fact to tick a box confirming that the system is fair.

> ***We can't solve racism or sexism or ableism, but we can tackle them in more complex ways.***

The chapter was a collaboration with Virginia Dignum, who leads an AI policy lab at Umeå University and runs a research group on responsible AI. What I contributed was the critical theory perspective.

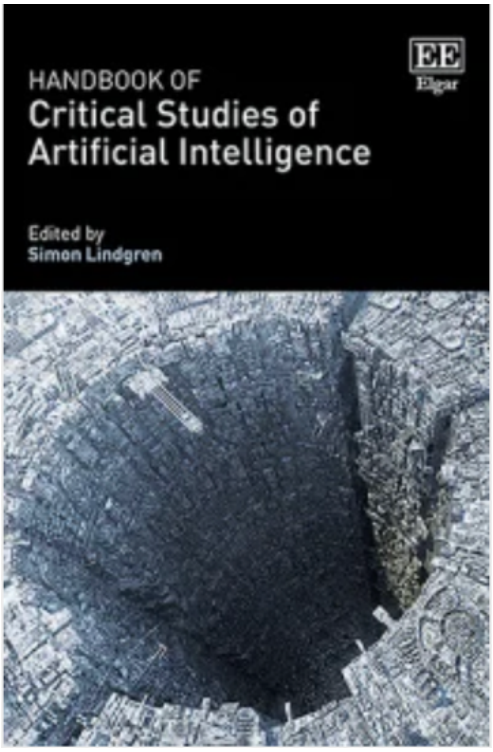

Figure 6. Cover of Lindgren's *Handbook of Critical Studies of Artificial Intelligence* (Edward Elgar, 2023).

JR: We have our own experience with Edward Elgar (the publisher of the Handbook), which was very good, though the hardback pricing does rather limit the readership. The good news is that our *Handbook of AI in Higher Education* (Popenici et al., 2026) is on Google Play. Seeing your book now appear in a more affordable paperback a couple of years on gives us some hope that they will eventually do the same with ours. I am very much looking forward to your new book. Polity Press is an excellent publisher, and the pricing is genuinely accessible.

I wanted to add something on technological solutionism: my sense of it was reinforced by reading Acemoglu and Johnson's *Power and Progress* (2023), which surveys a thousand years of technological promise – and the recurring pattern that those promises tend to benefit the few rather than the many. If you look at our lives today, email arrives incessantly, devices beep around the clock, and expectations and stress levels have risen rather than fallen. The techno-optimists and solutionists have been making the same arguments for centuries, and a critical perspective remains as necessary as ever.

FI: You edited the aforementioned *Handbook of Critical Studies of Artificial Intelligence* for Edward Elgar (2023) – an admirably ambitious undertaking with 76 chapters. In your introduction you write that "we need to push analyses of AI much further into critical territory than what is the case today" (Lindgren, 2023, p. 4), and observe that responsible AI research often plays "second fiddle to the heavy economic and technological drives to push forward no matter the consequences" (Lindgren, 2023, p. 1). You also cite Bassett and Roberts's formulation that "critical studies of artificial intelligence pick up where normative models of 'responsible AI' end" (Bassett & Roberts, 2023, as cited in Lindgren, 2023, p. 2). What was your vision for the collection, and how do you distinguish "critical studies" of AI from AI ethics?

SL: The agenda is essentially the same as in the *Critical Theory of AI* book, though with a 76-chapter collection, you inevitably get contributions pulling in somewhat different directions – it is not entirely coherent, and not everyone in the book would sign off on exactly my vision. There was also a genuinely pragmatic dimension to how it came about. I was approached by Edward Elgar to edit a handbook on AI and social science, just as I was developing the ideas that would become the *Critical Theory of AI* book. I said I would do it, but only if it could be on critical studies of AI – and they were open to that. I then threw myself into finding contributors, emailing people I knew and others I had simply looked up online, and a remarkably large number both accepted and actually delivered their chapters.

That is how the book reached the size it did, which was wonderful. As for how it differs from AI ethics, I think it is all there in the Bassett and Roberts quote. The fields of AI ethics and responsible AI are important and do valuable work, but they operate within a particular framework – one that asks how we can make AI systems fair, transparent, and accountable. These are necessary questions, but such research risks getting stuck at the level of guidelines, accepting the basic premise that AI will be developed and deployed and that our task is to ensure it is done responsibly. Someone has to do that work, and perhaps my critique is somewhat unfair in that sense.

But just as sociology and the humanities have a tradition of theoretically oriented research that is not directly applied, we also need to ask the deeper questions: who controls the development, who profits from it, what are the ideological stakes, what are the implications for labour, exploitation, capitalism? These questions may be present somewhere in the responsible AI framework, but they are seldom sufficiently developed – and not necessarily by scholars with the intellectual traditions to pursue them rigorously, traditions from which we can learn a great deal about language, power, and labour. What I am trying to say is that these perspectives risk becoming lightweight in that context. A lighter version may well be needed, but it is not what I, as a sociologist, am here to do.

> ***Who controls the development? Who profits from the development? Questions about ideology, about labour, exploitation, capitalism – these may not be developed enough in the framework of responsible AI.***

JR: Absolutely. I do not think it is unfair at all – pushing the critical envelope is important. At the same time, it is worth noting that AI ethicists themselves have come under considerable pressure. The structural marginalisation of AI ethics within the industry has been well-documented: Google's forced removal of the leaders of its own Ethical AI team demonstrated that critical internal scrutiny carries career risk, while Grok's explicit design philosophy – privileging permissiveness over harm mitigation – represents a rival commercial logic in which safety guardrails are reframed as ideological liability rather than responsible engineering (Bender et al., 2021; Hammond, 2024). So even the more moderate voices are under fire, and perhaps they deserve a degree of solidarity.

Power, inequality, and the possibility of public AI

FI: Your work returns repeatedly to the argument that AI systems both manifest and exacerbate existing inequities

- something I find personally compelling, given, for example, that the paid versions of tools like ChatGPT are substantially more capable than the free ones, yet many people simply cannot afford them. In the *Handbook*, you write that "AI contributes to conserving capitalism, as it offers it a new device by which to underpin the illusion of innovation, expansion, and effectivization as the one way forward" (Lindgren, 2023, p. 18). Is the so-called AI revolution bound to deepen inequality, or are there conditions under which it might not?

SL: It is worth trying to be optimistic. Right now, AI development is concentrated through a structure of corporate power, with quite exploitative working conditions in places, and still relatively weak regulation in many jurisdictions. But from a social constructionist perspective, things do not have to be this way. We can imagine AI contributing to a more just and equitable world. What might that look like?

Perhaps some form of public AI infrastructure, not controlled by these corporations. The field of Internet research has had analogous discussions: could a public service Internet be possible, and what would it look like? Perhaps it could address some of the problems of platform capitalism. One could also envision more democratic mechanisms for governing AI – giving those affected by these systems greater power over how they are used. In the labour market, that might involve a new role for trade unions in relation to AI. More broadly, one could imagine open-source, open-access models distributing the benefits more evenly. It is easy to be pessimistic about all of this, given current trends, but these are political choices, not inevitabilities.

These are political questions. Technological determinism would have us believe that AI 'demands' to be part of US warfare – but it is people who make those choices, and choices can be made differently. I also think we are still at a stage where our understanding of AI is heavily coloured by its novelty: we feel it will change everything. Looking at the history of technological change, there is a recurring pattern of media panics and tech panics – predictions of total transformation that, in retrospect, turn out to have been exaggerated. But when you are in the middle of it, you can never be sure this is not the one time that the transformation is real. AI does have characteristics that make it feel genuinely different. When the Internet arrived, or when Wikipedia appeared, people said students would never learn to research independently again – and yet that is no longer a pressing concern; we simply absorbed it into our understanding of how information works, much as plagiarism has always been a challenge. Generative AI may follow a similar trajectory, once we understand what it actually does and move beyond the mythology.

A small example: I grew up in the 1980s, and I remember when we got our first microwave oven. It came with recipe books suggesting you would do all your cooking with it from now on. Today, the microwave has found its place – reheating leftovers, defrosting things quickly. AI is obviously more consequential than that, but there is something of the same dynamic: we are in a phase where it feels ubiquitous and all-transforming, and perhaps it will settle into a more bounded role. To answer your question directly: it is ultimately an idealistic one, but one worth taking seriously. The hope would be that people gain more power over AI. The history of the Internet is not encouraging on this front – it began as a relatively open, subcultural space, and has since become increasingly centralised in its ownership, losing much of the variety of its early days.

Perhaps I am romanticising somewhat, but Jaron Lanier has written about how the early web – personal homepages before social media – was a patchwork of individual expression: people chose their own layouts, their own images, their own content. Everyone's Facebook profile looks exactly the same. Perhaps with AI, something similar will happen. At any rate, one wants to believe there is an activist dimension – that AI can also be turned against the very developments it is currently accelerating.

FT: This is just my personal view, but I think the issue is not really about AI technology specifically – or any new technology – so much as about how we engage with questions of equality. It is ultimately about the people behind the technology: how they train these large language models, and how they choose to deploy them. The same logic applies to nuclear power: the technology itself was not inherently destructive; that depended on the choices people made about weaponising it.

JR: Inequality here is genuinely complex. There is a fundamental 'garbage in, garbage out' problem: these systems hoover up the entire Internet without compensating anyone – unless you happen to be a large publisher with the

Figure 7. Simon Lindgren (Photo by Markus Naarttijärvi).

leverage to negotiate, as at least one has done recently for a reported $25 million, apparently without consulting the authors whose work was involved. But the point we are making is already very serious: the free and premium versions of these tools are simply different worlds. When I spoke recently with people in South Africa, a BRICS country, supposedly fast-developing, data plans are already very expensive, and everyone is using the free version. And then there is the language dimension you mentioned earlier regarding Sámi: if you do not speak English or Chinese, you are at a significant disadvantage. AI is supposed to be the great equaliser – every student gets a personal AI tutor, every person gets an AI assistant. The reality looks rather different. The so-called 1% are benefiting disproportionately.

SL: There is a paper by two of my colleagues on minority languages in northern Sweden and how they become marginalised through the way these models are trained (Eriksson Krutrök & Poromaa Isling, 2025).

JR: Access to the best models is one dimension of inequality; language is another. English and Chinese work extremely well. I was curious how Swedish fares.

SL: It works quite well – and I think my assessment goes beyond the fact that Swedish is my first language. The models can be a little clunky in Swedish at times, but it is still quite close; perhaps 95% as capable as in English. There is also a task-dependency: if you prompt in English and ask for a translation into Swedish, the result is often poor. But if you conduct the whole conversation in Swedish, the output is generally fluent. I do notice a difference, but it is not large – though it is noticeable enough to confirm that language differences are real.

JR: That makes sense – Swedish is a highly literate language with an enormous written corpus, which will certainly help. Compare that with South African minority languages, many African languages, or certain Asian languages with relatively sparse written corpora, and the drop-off in quality would be dramatic. I have not tested German recently, but I was pleasantly surprised the last time I did.

The vertigo moment: GenAI, scholarly identity, and the collision inside the university

ST: Your *Critical Theory of AI* book opens with a striking scene: you ask ChatGPT to write the introduction, and experience "a momentary sense of vertigo: Will there even be a need for scholars who write books in the future?", before realising the model was "just mimicking us" (Lindgren, 2024, p. 2). Since ChatGPT's release in late 2022, generative AI has transformed discussions in education and beyond. As someone who has studied digital media and society for decades, how do you assess this moment now, with more distance? Is the hype warranted?

SL: To be honest, the sense of vertigo I describe was perhaps somewhat theatrical. I had written most of the book before ChatGPT arrived, and its release made me want to ensure the book felt current – so I added that introduction, staging the interaction with the chatbot, towards the end of the writing process. I was using dramatic language to draw the reader in, and the vertigo question felt more rhetorical at the time. The genuine doubt came later, as I have come to know these tools better. Take, for example, this *Discourse Machines* book: for whatever reason, the publisher felt it should be relatively short, so the contract was for around 50,000 words. It then went out for peer review, and one of the anonymous reviewers commented that the book was on the shorter side – but perhaps, they added, that was fitting given the topic and the possibility that we are entering a post-reading or post-writing era.

That struck me a bit – the implication being that the books being written now may be among the last we ever write, as we move into a post-writing era. I would not go that far. But the remark touches something real. So much of academic work is about producing coherent text, and – if I may say so – the ability to write fluently is, in many ways, what brought me to where I am today. Many of us who ended up in academia did so partly because we were good at writing; that was how we were evaluated throughout our education. These tools make this skill less distinguishable. And there is a related absurdity: I have been writing in a particular academic style for years, and now the long dashes I favour have become known as an AI tell, which makes one inclined to use them less, to avoid the suspicion of having relied on AI.

I ran some experiments with early AI detectors, feeding in things I had written before GenAI models existed – and found that standard academic prose could be flagged, precisely because it already exhibited the characteristics those models had learned from. That is one reason why writing long-form work feels somewhat less inspiring at the moment. If I put my heart and soul into a book and release it into a market flooded with AI-generated titles that many readers cannot distinguish from the real thing – what is that effort worth? It is a sad feeling, and perhaps a passing one, but the question does press itself: is writing losing some of its value?

> ***If I write a book and put my heart and soul into it and put it out in a market where there will be lots of AI-generated books and lots of people won't be able to tell the difference – what's that effort really worth?***

So yes, I have rather more vertigo now than when I wrote that introduction. In the forthcoming book, I try to address this directly: we must understand LLMs as discourse machines – they do not think, but they do intervene in the discursive field, in Foucault's sense.

It is a moment of real insecurity. You are told not to generate papers with AI – yet journals openly use AI as a reviewing mechanism, and funding bodies announce that first-round assessments will be AI-assisted. If you generate your abstract with AI, the AI reviewer is more likely to rate it highly. I hope we emerge from this phase with a renewed appreciation for human writing. But it does expose something about academic practice: it makes visible the fact that merely producing coherent text has never really been enough.

JR: I am genuinely put out about the em dashes – I was using them long before the GenAI era.

SL: I wrote a rather long post about exactly this.

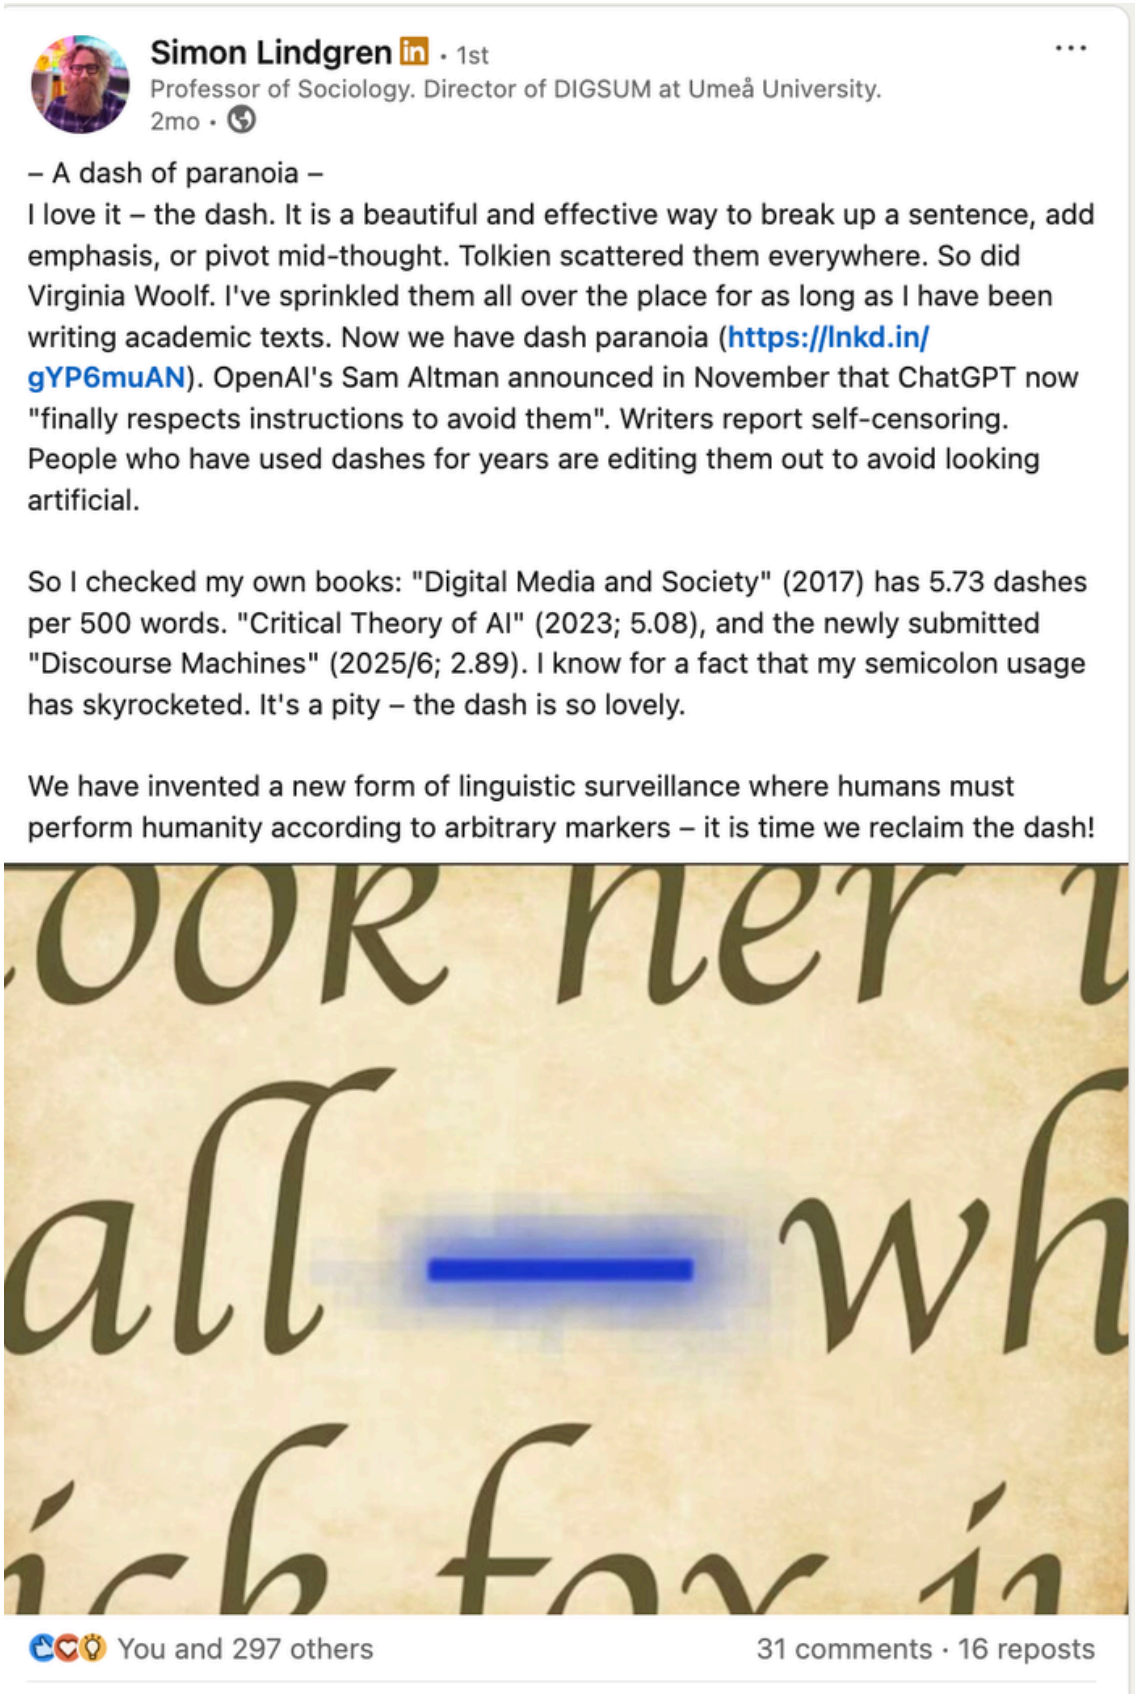

Figure 7. Simon Lindgren's post on em dashes.

JR: I saw it, and others have raised the same complaint. I have actually decided to abandon em dashes, partly because I do not want readers to dismiss the text as AI-generated.

SL: But have you considered the reverse? If you have no em dashes, people may assume you generated the text with AI and then removed them.

JR: Well, I guess [All laughing].

SL: Best to use a moderate number. I have moved to the shorter variant – what are they called again?

JR: The en dashes – yes. I use those too, or sometimes commas in place of em dashes. It is genuinely irritating. Everyone has the same complaint: words they have always favoured are now so copiously deployed by AI that they feel unusable – 'delve' being the canonical example.

FT: In Australia and many other countries, there are sector-wide discussions about how universities should respond to GenAI. A recent article in *The Australian* (Thomas, 2026) reported that over 80% of both students and staff had used GenAI without authorisation in their research or assessments – a finding that has raised considerable alarm in the Australian academic community. What are your views on how universities should respond to generative AI – both as a subject of critical study and as a tool students and researchers are actively using? Do you see a danger of a fully mediated loop: AI-generated assignments, AI-assisted marking, and nobody actually learning anything? And how has GenAI altered the production of research?

SL: As a start, I suspect many people share this situation. As I touched on earlier, in my department at my university, the ongoing discussions about students' use of GenAI tools are still entirely framed around cheating. We handle it exactly as we handled plagiarism – it is about communicating to students that it is prohibited. 'I think this student has AI-generated this answer' has replaced 'I think they have plagiarised this'.

And then, very recently, a completely different dynamic has emerged at my institution. We are locked into the Microsoft 365 ecosystem – Teams, Outlook, all of it – and the university has been offered full access to Microsoft Copilot, embedded across all our communication tools. I am perhaps being slightly conspiratorial, but Microsoft's plan seems clear: everyone will discover how useful it is for routine tasks, and the university will then agree to continue paying for the premium tier. Alongside this, we are receiving a stream of invitations to workshops on how to become more efficient in our teaching with Copilot. The library is running courses on how to produce literature reviews using AI tools.

In all of that, the discourse is one of efficiency – not cheating. That collision – whether at the level of my own institution or more broadly – is clearly bound up with a capitalistic process. To use a blunt phrase: as academic workers, we are having this technology shoved down our throats, while simultaneously being asked to police its use by students. We will probably reach a more settled position eventually, but right now we are still in a somewhat panicked, utopian/dystopian phase where nothing is clear. I do not claim to see it all clearly myself – but I do at least claim to see that we do not see it clearly.

On the one hand, we're getting GenAI shoved down our throats as academic workers, and at the same time, we're supposed to be policing the students' use of it.

On the pedagogical question, I lean towards what I would call a de-dramatising approach. There is endless discussion about what policies to adopt, and my own position is fairly simple: in the courses I teach, I tell students they may not use generative AI tools to *produce the text* they submit. I do not say much more than that – it is the same principle as not copying and pasting from the Internet. Sometimes I add that by submitting text, students certify that they have written it themselves.

I came across something interesting recently on a Swedish research funding body's FAQ page. One question asks whether applicants may use AI tools to write their proposal, and whether they must declare this if they do. The answer is that there is no prohibition on using AI tools to write the application. The applicant is responsible for the accuracy of the contents for ensuring the research can be carried out as described, and for ensuring the proposal contains no plagiarised, false, manipulated, or otherwise incorrect content.

The second question asks whether applications will be assessed using AI tools. The answer: AI is not used to assess the applications themselves, but may be used to polish the language of the reviewers' written assessments. This is rather different from what I tell my students – that by submitting, they certify having written it themselves. This funder's position is: you may use AI, but you are responsible for the result.

The principle is: use AI if you wish, but own what you submit. You can delegate parts of the work to AI and then verify – or simply trust it. I am not sure what the right answer is, but one thing is certain: these tools exist, and we will have to live with them. In teaching, reactions vary considerably. Some colleagues say take-home assignments are dead – the only viable option is the return of traditional invigilated examinations.

But students still have to write their dissertations, and they cannot do a master's or PhD thesis in a supervised hall. My general philosophy – ideologically speaking – has always been that students may use new tools, as long as they do so in an informed and literate way. But I find this one genuinely harder than previous tech panics. I have always been relatively confident that the alarm would pass and things would settle – but this one, I think, really does change quite a lot.

And yet – there was a time when students could not Google things, then they could. We thought that would transform everything. Now we have tools that can generate text. Perhaps it is just another tool: you generate some text, you still have to review it, you rationalise part of the writing process. But then you see students who arrive at university for the first time and immediately fall into the practice of generating everything they submit. That is different.

It also depends on where you are in your career when you encounter these technologies. I have spent roughly 15 years teaching myself to code in Python for my research. Now that AI tools can accelerate that process, I feel that I still have enough of a foundation to inspect the generated code and verify it – which is genuinely useful.

My two sons – one is 18, the other 21 – were actually more sceptical than me when these tools first appeared. They would say, 'ChatGPT cannot do that – it will get it wrong'. And I would explain how to craft a prompt: specify this, exclude that. I gave them quite an elaborate example. Their response was the aha moment: 'Yes, but that is because you already have the knowledge to formulate that prompt'.

If I know the relevant social theory, the methodology, the field – if I have written 30 or 40 research applications over my career – and I ask the tool to draft a methods section based on specific parameters, it is like having a research assistant. But if a brand-new student arrives in an entirely new field and does the same thing, it is something else entirely. Where this leads, I genuinely do not know. But my philosophical instinct is that we should not be too frightened of the technology – that this is, at root, a question of literacy.

Perhaps a larger case of literacy than some previous ones – but still about literacy: understanding what the tool does, where it starts and ends, and what the output actually is. In a field as text-centred as ours, it is easy to slide i-

-nto generating your abstract with AI so that an AI reviewer rates it highly, presenting AI-generated slides to an audience whose talk gets AI-transcribed and AI-summarised. Or perhaps, in a few years, we will look back and think: do you remember when we were all so exercised about this, before it settled into just another technology?

In some lines of work, the disruption is less immediate – if text production is not your central activity, an AI navigation tool for truck drivers is useful but not identity-threatening. For academics, though, it cuts rather close to home.

JR: The truck driver, of course, may also eventually be replaced by a self-driving vehicle and reduced to supervising the system, intervening only in exceptional circumstances. My own sense is that the current AI moment is somewhat more significant than previous technological disruptions – the calculator, Wikipedia, Google, spell-check, Grammarly – where the 'everything will change' predictions were eventually absorbed without the catastrophe feared. The qualitative difference is that GenAI generates. Calculators do not calculate by themselves; Internet search tools do not produce text. The self-generative capacity is the new thing.

SL: Yes – and it is already being embedded everywhere. Grammarly, for example, started as something like an enhanced version of Microsoft's spell-checker and has since become much more substantial. I have seen PhD students use it extensively: they write a paragraph, Grammarly suggests improvements, which are now powered by AI under the hood. That is functionally the same as pasting the paragraph into a chatbot and asking it to improve the text. Which means that if you write entirely in your own words but run everything through Grammarly, the resulting text may be flagged by AI detectors as substantially AI-generated.

Some people try to avoid ChatGPT for environmental reasons – the frequently cited claim that every prompt consumes roughly a glass of water – and switch to Google instead. But Google now routes queries through AI models. There is no real opt-out anymore.

Some people try to avoid ChatGPT for environmental reasons – the frequently cited claim that every prompt consumes roughly a glass of water – and switch to Google Search instead. But Google now routes queries through AI models. There is no real opt-out anymore.

JR: What's your view about AI detection?

SL: I do not trust it, largely because of the false positive rate.

JR: As you mentioned earlier, some of your own pre-GenAI writing gets flagged as AI-generated.

SL: And if we accept that text-improving tools – modern grammar-checkers, style suggestions, Grammarly – should not be prohibited, then there will be many cases where surface-level AI detection flags the text without any knowledge of the actual underlying work that has been put in. That is why the focus on detection feels misplaced to me. The question of whether a sentence was produced by my own keystrokes or by a chatbot fine-tuned on my writing style, makes the philosophical issue of authorship genuinely murky.

Whether a sentence was produced by my own keystrokes or by a chatbot fine-tuned on my writing style, the philosophical question of authorship becomes genuinely murky.

AI detection also creates a panopticon effect: knowing that text may be run through a detector, no one wants to be stigmatised for using AI, so everyone self-polices. As the LinkedIn post shows (see Figure 7), I have tracked how my

own use of the dash has declined since chatbots became widespread – which means, by that particular measure, I was writing more like an AI before. These are genuinely complex issues. But in the end, I believe it must come back to literacy: understanding what the tool is doing, and ensuring that you remain in the driver's seat.

Teaching against the algorithm: Defamiliarisation, 'bias bias', and how LLMs reshape what counts as knowledge

FI: You have written course textbooks on digital media. How do you approach teaching students to think critically about technologies in which they are deeply embedded?

SL: This is something I was discussing with colleagues just yesterday – January and February are among my most intensive teaching periods, and I have been teaching partly from my own textbook (Lindgren, 2025). The concept of digital immigrants and digital natives – Prensky's (2001) notion of a generational divide, with those born around 1985 as the rough dividing line – is by now somewhat antiquated and arguably flawed, but it still captures something. I was in my early twenties when I encountered the Internet; my students today are around twenty and have grown up with these technologies entirely.

For students, these technologies are simply given – they take them for granted. Teaching on this involves a balancing act: they have a great deal to say about these platforms, and one has to be careful not to come across as the classic boomer stating the obvious. The key pedagogical strategy, I think, is to work through theory and concepts – a kind of defamiliarisation. They may have a functional everyday understanding of what an algorithm does, but giving them conceptual tools developed by scholars who were thinking about these issues before the platforms even existed can suddenly reframe the familiar. It works both ways. I sometimes have, for example, a module on datafication and targeted advertising, and the whole Shoshana Zuboff perspective on surveillance capitalism. I sometimes overestimate how much of this students will already have absorbed – I worry standing at the lectern that it will feel like old news to them. But in seminars, you discover that seeing it framed this way is still genuinely new, sometimes even an eye-opener: 'Is it really like that?'

There can even be a kind of political awakening, though students obviously hold different views on what to make of it. My approach is to take a specific phenomenon, then introduce a scholar who has a concept for it, and show how the concept illuminates the phenomenon. That is the general format. Years ago, teaching these courses was largely a matter of explaining what the Internet was; now, theory has become more important precisely because students are so embedded in these technologies. Being embedded provides a lot of everyday knowledge, but it can also make it harder to achieve what one might call clinical distance from the object of study. Prensky – or at least some of the debates following his work – suggested that digital immigrants are hopelessly lagging behind. But there is a real advantage in coming to these technologies later: a sociology professor who encounters LLMs at fifty-plus brings a multi-layered understanding, because they have lived through previous media systems and thought carefully about their precursors.

Much as someone who grew up with only radio may have a sharper critical eye on what is truly new about television than someone who has always had it. The risk, of course, is being patronising – walking into the room and announcing that all these TikTok users are being exploited. It is a genuine challenge: embedding makes the subject immediate and relatable, but some of the critical frameworks can come across as moralising.

ST: What research projects are you currently working on? In the *Handbook*, you argue that "AI is driven by ideology" – that "certain ideas, connected to power relations in society, become constructed as given truths and objectives of AI, even though many other truths and objectives might be possible" (Lindgren, 2023, p. 17). You also coined the term "bias bias" to describe "a disproportional inclination towards the notion of bias" in AI discourse (Lindgren, 2024, p. 173). What are the most pressing questions about AI and society – and AI and higher education – for you right now?

SL: It is a hard one, especially the last part. On what I am currently working on: I have already mentioned the forthcoming *Discourse Machines* book, which I am genuinely excited about. It is planned for release in 2026, probably not before the summer. It is very much grounded in discourse theory and post-structuralism – in some w-

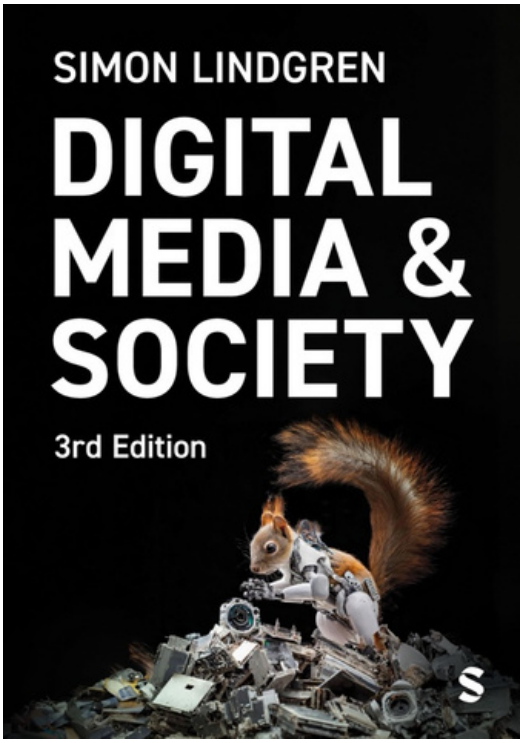

Figure 9. Cover of Lindgren's teaching textbook *Digital Media & Society* (Sage, 3rd Ed., 2025).

-ays a passion project, revisiting theories that shaped me when I first started out, and showing how they bear on large language models.

I also run the DIGSUM Centre for Digital Social Research at Umeå University, which is expanding considerably. We have recently secured substantial funding for a research programme on AI, power and politics, with several sub-strands. The group is growing – in PhD students, postdocs, and a broader ecosystem of projects. Empirically, much of it is about information and opinion dynamics online – disinformation, misinformation, and some emerging work on knowledge regimes in a post-LLM world, which grows out of the book project and is something we are also trying to secure further funding for. On the 'bias bias' term: it connects back to the broader argument about providing a fuller critique of the structural imbalances that LLMs and generative AI models can entrench through their training.

I had grown tired of how the word 'bias' was used in this context. When you say there is bias in an AI model – that it is more likely to assign a male pronoun to a doctor – you simultaneously signal that this is a flaw that crept into the model and can be corrected by reweighting parameters. 'Bias bias' was my attempt to name this risk: the tendency to reduce structural inequalities – racism, sexism – to a technical calibration problem. It connects to everything we have been discussing about moving beyond the technological fix, because 'bias' is a word that almost irresistibly pulls in that direction.

> ***'Bias bias' is a bias towards speaking of bias, where we should speak of racism and sexism.***

The language of bias reframes structural injustice as something more like a measurement error – a mathematical problem, not a social one. As for what is most pressing: it is hard to single out one thing, but the direction my scholarship is taking points towards a central question: how do LLMs reshape what counts as knowledge? That is maybe the key question.

JR: Is there anything you think we should have discussed but did not, or something we only touched on briefly that deserves more attention?

SL: I did think about this when I read through the questions, and I believe we touched on it in an earlier exchange –

but it is worth stating clearly. Critical scholarship on AI is often misread as purely negative, as though one is only finding fault. But the goal of a critical perspective is ultimately constructive: not necessarily to produce hands-on tools, but to keep the conversation about alternatives open, and to resist the determinism – technological determinism, solutionist thinking – that presents one particular development path as inevitable. Silicon Valley's economic logic is not the only logic possible. This is not an anti-technology position. I like technology.

> ***The critical perspective is not an anti-technology perspective. But it's anti other things – such as anti-capitalist.***

And I think that is the right note to end on: the critical perspective is not anti-technology. It is anti-other things – anti-capitalist, among them.

References

Acemoglu, D., & Johnson, S. (2023). *Power and progress: Our thousand-year struggle over technology and prosperity.* Basic Books.

Bassett, C., & Roberts, B. (2023). Automation anxiety: A critical history – the apparently odd recurrence of debates about computation, AI and labour. In S. Lindgren (Ed.), *Handbook of critical studies of artificial intelligence* (pp. 79–93). Edward Elgar.

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*, 610–623. https://doi.org/10.1145/3442188.3445922

Benjamin, R. (2019). *Race after technology: Abolitionist tools for the New Jim Code.* Polity Press.

Eriksson Krutrök, M., & Poromaa Isling, P. (2025). Technolinguistic suspension and the digital futures of minority languages: AI and Meänkieli in Sweden. *New Media & Society*, 14614448251385920.

Hammond, K. (2024). *Misinformation at scale: Elon Musk's Grok and the battle for truth*. Center for Advancing Safety of Machine Intelligence, Northwestern University. https://casmi.northwestern.edu/news/articles/2024/misnformation-at-scale-elon-musks-grok-and-the-battle-for-truth.html

Lanier, J. (2018). *Ten arguments for deleting your social media accounts right now.* The Bodley Head.

Lindgren, S. (2020). *Data theory: Interpretive sociology and computational methods.* John Wiley & Sons.

Lindgren, S. (2023). Introducing critical studies of artificial intelligence. In S. Lindgren (Ed.), *Handbook of critical studies of artificial intelligence* (pp. 1-19). Edward Elgar.

Lindgren, S. (Ed.). (2023). *Handbook of critical studies of artificial intelligence.* Edward Elgar.

Lindgren, S. (2024). *Critical theory of AI.* Polity Press.

Lindgren, S. (2025). *Digital Media and Society* (3rd ed.). SAGE Publications.

Lindgren, S. (in press). *Discourse machines: Knowledge and power in the age of large language models.* Polity Press.

Lindgren, S., & Dignum, V. (2023). Beyond AI solutionism: Toward a multi-disciplinary approach to artificial intelligence in society. In S. Lindgren (Ed.), *Handbook of critical studies of artificial intelligence* (pp. 163–172). Edward Elgar.

Lindgren, S., & Holmström, J. (2020). A social science perspective on artificial intelligence: Building blocks for a research agenda. *Journal of Digital Social Research*, *2*(3), 1–15. https://doi.org/10.33621/jdsr.v2i3.65

Morozov, E. (2013). *To save everything, click here: The folly of technological solutionism.* Public Affairs.

Popenici, S., Rudolph, J., Ismail, F., & Tan, S. (Eds.). (2026). *Handbook of AI in higher education*. Edward Elgar.

Prensky, M. (2001). Digital natives, digital immigrants. *On the Horizon*, *9*(5), 1–6.

Thomas, R. (2026, February 6). How Australia's university students are using AI to cheat their way to a degree. *Weekend Australian Magazine*.https://www.theaustralian.com.au/weekend-australian-magazine/how-australias-university-students-are-using-to-ai-to-cheat-their-way-to-a-degree/news-story/2bd02fe5c5dce5c74914fc01bf883df0

Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power.* Profile Books.