Verfassungsblog
ON MATTERS CONSTITUTIONAL

Who Cares Who Wrote It?

Anmol Jain

16.10.2026 doi: 10.59704/a9938c9d038949fd

I remember, as a child, fantasising about having the ability to retain everything I read for future use. It'd work as pure magic. Today, as major publishers strike unprecedented deals to share their entire journal catalogues with AI platforms, we are witnessing something remarkably close to that childhood dream becoming reality. AI systems can now read, process, and synthesise more legal scholarship in seconds than any human could absorb in a lifetime. What should we make of this?

The legal academy, and perhaps many other fields, is at a curious juncture. We have largely accepted AI as a tool for refinement—for polishing language, organising footnotes, checking citations, and formatting documents to meet journal specifications. These uses are increasingly being normalised, particularly because they democratise access to professional presentation standards, level the playing field for non-native speakers, and free scholars from mundane but time-consuming tasks to focus on substantive analysis. AI, in this limited role, functions as an equaliser and an efficiency engine.

Yet we have drawn a bright line. AI must not be used for ideation, argument development, and initial writing. One can test arguments with AI-based platforms, probe for weaknesses, and refine logic through dialogue with these systems. However, the generation of core ideas and legal reasoning must remain human, and writing must remain distinctly personal. This boundary has already become something of a professional orthodoxy, and is often repeated in editorial guidelines, conference discussions, law school policies, and informal conversations.

But is this the right standard to adopt? I want to suggest that it may be time to critically interrogate this boundary to question whether our assumptions about authorship, originality, and value in legal scholarship can withstand the technological forces reshaping our field. The following arguments do not call for abandoning this distinction entirely, but to seek a deeper rethink of the commonly accepted positions.

The Inevitability of Indistinction

Let's start with a practical problem. The line between human-generated and AI-generated content is becoming increasingly difficult to discern, and this difficulty is only going to intensify. As AI platforms start to generate responses following set identifiable patterns, it is simultaneously giving rise to another set of platforms involved in humanising such texts. As these systems become more sophisticated—as they are trained on more comprehensive datasets, develop more nuanced understanding of legal reasoning, and learn to mimic not just style but analytical frameworks—detecting their involvement will become nearly impossible. We may establish disclosure requirements and honour systems, but enforcement will prove elusive. The technology simply will not cooperate with our categorical distinctions.

This is also more than a mere technical problem. It raises a deeper normative question. If we cannot reliably identify AI-generated scholarship, what exactly are we trying to protect? If a piece of legal analysis is novel, methodologically sound, analytically rigorous, and contributes meaningfully to doctrinal understanding, does it matter whether its core arguments emerged from human cognition or computational processes? Does it matter that it was primarily authored by a human? Our instinctive answer is "yes, of course it matters." But why?

There is another dimension to this enforcement problem, i.e., the erosion of trust. When we commit ourselves to policing the boundary between human and AI authorship, we inevitably create a culture of suspicion where every innovative argument becomes suspect, and every elegant formulation invites scru-

tiny. Every piece of scholarship that demonstrates unusual comprehensiveness or analytical sophistication could trigger the question: Did a human really write this? We are training ourselves, in other words, to ask "who wrote this?" before we ask "is this argument sound?" I, of course, note that AI detection may be relatively easy in 2026, largely because its responses tend to be exactly opposite to what I claim here (as people tend to call it, "AI slop"), making their hollowness easy to identify. However, as systems evolve and develop further, it is not hard to imagine a near future where AI responses shed their hollowness and become very sophisticated and increasingly harder to identify.

This culture of suspicion has corrosive effects. It diverts our attention from substantive evaluation to a kind of authentication exercise. In other words, scholarly discourse risks devolving into procedural policing rather than rigorous engagement with ideas. We become less trusting of our colleagues, more paranoid about academic integrity, more focused on detecting violations than on advancing knowledge. The academy, which depends on a baseline of trust and good faith engagement, then becomes characterised by doubt and verification exercises.

Moreover, if we cannot reliably enforce this distinction anyway, as sustained technological advancement suggests, then this erosion of trust serves no constructive purpose. We are paying a high social cost to maintain an increasingly unidentifiable boundary. This suggests we may be asking the wrong question entirely. Instead of obsessing over "who wrote this?", perhaps we should be asking "what is legal scholarship for?"

The Purpose Question

Consider what we understand legal scholarship to be *for*. At its core, legal research exists to help us organise our collective social life better. It is focussed on, among other things, clarifying doctrinal confusion, resolving interpretive disputes, identifying normative tensions, proposing regulatory frameworks, critiquing power structures, and illuminating path dependencies in legal development. The purpose is, therefore, instrumental. We produce scholarship not as an end in itself, but as a means to improve legal systems and, ultimately, human flourishing within those systems.

If this instrumental understanding is correct, and if AI-assisted scholarship can help us accomplish these purposes, why should we resist? (To reemphasise, I am concerned with the positive possibilities of AI systems, and not the *slop* that it often generates). In other words, why should we value a less insightful human-generated analysis over a more illuminating AI-generated one? The analogy to earlier technological transitions may be instructive here.

When legal research migrated from physical libraries to online databases, we did not mourn the loss of manual card catalogues (at least not today) or view digital search as somehow less authentic. Instead, we recognised that accessibility to sources freed researchers to focus on analytical depth, interpretive sophistication, and normative judgment. The shift from books to databases changed research methodology, certainly, but it did not diminish the value of legal scholarship. If anything, it raised expectations as researchers could no longer claim ignorance of relevant materials that comprehensive database searches would have revealed. The transition to AI-assisted research might follow a similar trajectory. Just as databases ensured that scholarship would be judged more by analytical quality than by exhaustive manual source-gathering, AI might shift expectations again.

The Publisher-AI Partnership and Its Implications

The scale of what is happening behind the scenes deserves greater attention. In 2024 and 2025, publishers have signed licensing deals worth millions of dollars with AI companies, granting access to vast catalogues of academic journals, books, and specialised content.[1] The Financial Times, Reuters, Springer Nature, The Associated Press, and countless academic publishers are all making their archives available to train large language models. Some estimates suggest the global AI in publishing market will

[1] Sara Guaglione, *A Timeline of the Major Deals Between Publishers and AI Tech Companies in 2025*, 01.01.2026, Digiday, https://digiday.com/media/a-timeline-of-the-major-deals-between-publishers-and-ai-tech-companies-in-2025/, last accessed 20.08.2026.

grow from $2.8 billion in 2023 to $41.2 billion by 2033. [2]

What this means, practically, is that AI systems will soon have access to effectively the entire corpus of published legal scholarship, across jurisdictions and centuries. No human legal scholar, no matter how brilliant or industrious, could possibly read and retain this volume of material. The comprehensive knowledge that AI systems will possess is fundamentally different in kind, not just degree, from human expertise. Moreover, the possibility for legal research emerging from such databases is unimaginable. It promises to free researchers from constraints emerging from several access restraints, including that of ignorance of knowledge beyond what is immediately available, the need for financial commitments to travel to multiple jurisdictions, and that of language. The same applies equally to the communication of research outcomes.

Setting aside copyright concerns, questions of publisher ethics, and equitable access to such resources for all for a moment —though these obviously matter greatly and need our immediate focus— we must acknowledge the sheer analytical potential this represents. An AI system with access, for instance, to the complete archive of constitutional law scholarship could identify doctrinal patterns, trace conceptual evolution, map argumentative genealogies, and spot contradictions across hundreds of thousands of texts in ways that no individual scholar ever could. It would be wrong to say that it results in replacing human judgment. Rather, it enables forms of analysis that were previously impossible.

The Cultural and Contextual Caveat

I do not mean to suggest that AI can do every kind of legal research. Much legal scholarship requires deep understanding of specific cultural contexts, historical contingencies, social practices, and power dynamics that AI systems cannot yet fully grasp. Constitutional interpretation, for instance, often turns on culturally embedded understandings of dignity, equality, or liberty that resist reduction to textual patterns. Legal anthropology, socio-legal studies, critical race theory and many other approaches depend on contextual knowledge and lived experience that current AI cannot replicate.

Law, being an enterprise to organise and regulate society, is fundamentally shaped by human experience. As academics and practitioners of constitutional design have long recognised, no single institutional or structural design can guarantee similar results across different countries. *Context matters. Culture matters. Experience matters. Original thinking to solve our problems matters. Thinking differently matters.* Many of the guiding legal theories and doctrines that we use today globally were the result of human ingenuity and an attempt to solve human problems while being part of society and deeply understanding its people. AI is just not there yet.

But consider fields like insolvency law, contract law, tax law, international finance regulation, environmental law, and competition law. These areas certainly have cultural dimensions, but they also operate according to commonly shared goals, technical frameworks, and analytical methods that transcend particular cultural contexts. In such fields, where legal reasoning follows more standardised patterns and draws on more universal principles, AI's potential is considerable. Similarly, AI can be of great assistance for those engaged in cataloguing archival work, drawing statistical results, and tracing patterns from empirical findings.

Hence, while it is true that we cannot yet write off the value of human authorship, it should not automatically translate into an argument that AI-generated works must be treated with less value. It would be a worthwhile research study to examine areas where AI presents the potential to transform research and knowledge. It would also be interesting to understand how effective AI is at innovation. If one were to clearly describe the status quo and a societal problem, could AI develop or suggest novel solutions? Could AI not just reiterate and repackage the information fed to it,

[2] *Top Academic Publishing Trends in 2025,* Amnet, 17.12.2024, https://amnet.com/top-academic-publishing-trends-in-2025/, last accessed 20.08.2026.

but build upon it? If those solutions are workable, how should we treat that knowledge?

Reframing the Standard: From Authorship to Accountability and Method

If we are to move beyond our current resistance to AI-generated research, we need a different framework for evaluation. One way could be to shift our focus from authorship to two alternative criteria: accountability and methods.

Accountability: Despite AI-led generation of content, the human author must remain accountable for the work that is produced and published. Just as we do not grant academic freedom without expecting scholars to meet professional standards and submit to peer review, we should not accept AI-generated scholarship without clear human accountability for its quality, accuracy, and contribution. The scholar using the tool bears responsibility for the final product.

Methods: Legal scholarship has established methodological standards, including evidence-based reasoning, logical coherence, engagement with counterarguments, citation of sources, and transparency about normative commitments. Any scholarship, whether human-generated, AI-assisted or AI-generated, must meet these methodological requirements.

If we focus on accountability and methods rather than on the percentage or stage of human involvement, we can develop evaluation criteria that are technologically neutral but substantively rigorous. A piece of AI-generated legal analysis that is methodologically sound, acknowledges its limitations, engages seriously with existing doctrine, and contributes new insights should, under this framework, be considered acceptable. A piece of human-generated analysis that fails these standards should not be.

An Invitation to Dialogue

My comments must not be taken as arguments for immediately embracing AI-generated legal scholarship without reservations or safeguards. I am not claiming that all concerns about authenticity, originality, or human judgment are misguided. What I am suggesting is that our current consensus—AI for refinement, yes; AI for ideation and argument, never—may need to be interrogated more rigorously than it has been. We should be having difficult conversations about what we value in legal scholarship and why. Is it the human labour of thinking? The originality of insight? The contribution to collective understanding? The methodological rigour? The accountability structure? Different answers will lead to different frameworks for evaluating AI's role.

Perhaps most importantly, we need to recognise that this debate is not only about AI. It is also about what legal scholarship is for, what makes it valuable, and how we ensure that it continues to serve its core purpose of helping organise collective life better, illuminating legal questions that matter, and contributing to human flourishing within systems of law. If AI can help accomplish those purposes, then we need better arguments than "it's not human" to justify its exclusion.

Irrespective of where one stands on this debate, one cannot resist adjustments to AI. The adoption may be incremental, but it is undeniable. Some might argue for careful, staged integration by limiting AI use to specific functions. This approach has merit and may prove prudent. But whichever path we choose, it is important that we remain comfortable setting aside our conventional positions and thinking differently about both possibilities and threats.

ABOUT THE AUTHOR

Anmol Jain

Anmol Jain is an Honorary Senior Fellow at Melbourne Law School and an Associate Editor at Verfassungsblog. He holds law degrees from Yale Law School and National Law University, Jodhpur.Anmol Jain is an Honorary Senior Fellow at Melbourne Law School and an Associate Editor at Verfassungsblog. He holds law degrees from Yale Law School and National Law University, Jodhpur.

SUGGESTED CITATION

Jain, Anmol: Who Cares Who Wrote It?, Verfassungsblog, 16.10.2026, doi: 10.59704/a9938c9d038949fd.

Licensed under CC BY-SA 4.0.

Langzeitarchiviert bei IntRechtDok unter: doi.org/10.17176/20261016-143330-0.