

'CO-WRITING' LEGAL RESEARCH WITH AN AI BOT: CRITICAL REFLECTIONS OF AN AI-SCEPTIC RESEARCHER

ALYSIA BLACKHAM*

Universities are increasingly encouraging their staff to make use of artificial intelligence and large language models in their work processes. This article critically explores the process of 'co-writing' research with Microsoft Copilot, to empirically evaluate the usefulness of such tools. It concludes that legal research undertaken with current large language models risks producing poorer quality legal research, and undermining the inherently human research dialogue that is foundational to legal research.

I INTRODUCTION

In 2026, a legal research article was published in the United States ('US') journal *SMU Law Review Forum*.¹ The author acknowledged that the paper was co-written with Anthropic's large language model ('LLM') Claude.² The author described Claude's role 'as an assistant':

I used Anthropic's Claude as an assistant when producing this Article (mostly in early 2026). This Article took me around fifteen hours, a small fraction of the time it would've taken a few years ago ... I personally formed the thesis and penned the title, abstract, introduction, and most of the body. At no point was Claude freestyling arguments; from start to finish it worked to implement its principal's vision. And while Claude astonished me with a first draft that was better than I would have received from the average research assistant or junior associate, I discarded or substantially revised nearly all of it. But some of Claude's research and writing was impeccable, and it would have been silly to omit those contributions merely because of their source.³

That this partly LLM-generated article has been accepted for publication in a university law review has caused considerable consternation online. It raises questions of intellectual property and authorship; research integrity, methods and

* Professor, Melbourne Law School, University of Melbourne, alysia.blackham@unimelb.edu.au.

1 Michael Showalter, "'Mere Machines': Why Originalism Requires Robotic Judging" (2026) 79 *SMU Law Review Forum* 84 <<https://doi.org/10.25172/slrf.79.1.4>> ('Why Originalism Requires Robotic Judging').

2 Cf Andrew M Perlman's earlier paper, which was solely written by ChatGPT: Andrew M Perlman, 'Generative AI and the Future of Legal Scholarship' (Research Paper No 5072765, Social Science Research Network, 26 December 2024) <<https://doi.org/10.2139/ssrn.5072765>>.

3 Showalter, 'Why Originalism Requires Robotic Judging' (n 1) 85.

rigour; job displacement and job loss for human ‘assistants’; and the environmental impact of relying on artificial intelligence (‘AI’) and LLMs to undertake work which would previously be undertaken by humans.

At the same time as this article was causing concern online, I was taking part in an Australian university-based pilot of the use of Microsoft Copilot (‘Copilot’). In taking part in the trial, we were encouraged to integrate Copilot into our daily ‘workflow’. For me, in my role as a teaching and research academic, my daily ‘workflow’ often entails (on a good day) thinking, researching and writing. Integrating Copilot into this workflow potentially raises many of the same issues flagged by the *SMU Law Review Forum* paper. And, as I write these words, I am encouraged by a magical pen icon to ‘[d]raft with Copilot’ in Microsoft Word.

In the interests of full disclosure, though, I perhaps lean more to the side of AI-sceptics than techno-optimists. I have written papers critiquing the discrimination risks of AI and algorithmic management,⁴ and the use of AI and LLMs to ‘shortcut’ established legal research methods.⁵ My scholarship has a strong focus on empirical legal research methods and ensuring legal research of a high calibre.⁶ I am well-versed in the risk of bias and hallucinations that might emerge from LLMs and other AI tools.⁷ I have yet to think of a use case in which LLMs and Copilot will improve my day-to-day ‘workflow’ or tasks. Indeed, I enjoy the act of writing. I personally see writing as a way of developing and refining my thinking. I think by writing.

What, then, does integrating Copilot into my legal research ‘workflow’ mean? Does it add value (and meet the AI ‘productivity’ agenda being pushed by governments⁸ and universities)? What are the risks and problems that are likely to emerge, for this and other legal research projects? In an Australian legal context, are hallucinations actually rare?⁹ In asking these questions, this article is intended

4 Alysia Blackham, ‘Setting the Framework for Accountability for Algorithmic Discrimination at Work’ (2023) 47(1) *Melbourne University Law Review* 63; Alysia Blackham, ‘Artificial Intelligence and Equality at Work: Evaluating the Adequacy of the UK Legal Framework for Responding to the Discrimination Risks of New Technologies’ (2026) *Industrial Law Journal* dwag013 <<https://doi.org/10.1093/indlaw/dwag013>>; Alysia Blackham, ‘Artificial Intelligence and Equality at Work’ in Marco Biasi (ed), *Artificial Intelligence and Labour Law: A Global Overview* (Routledge, 2026) 103 <<https://doi.org/10.4324/9781042003266-8>>.

5 Alysia Blackham, ‘Interrogating New Methods in Socio-Legal Studies: Content Analysis, Case Law and Artificial Intelligence’ (2025) 50(2) *Alternative Law Journal* 85 <<https://doi.org/10.1177/1037969X251325869>> (‘Interrogating New Methods in Socio-Legal Studies’).

6 Alysia Blackham, ‘When Law and Data Collide: The Methodological Challenge of Conducting Mixed Methods Research in Law’ (2022) 49(Supp 1) *Journal of Law and Society* S87 <<https://doi.org/10.1111/jols.12373>>; Alysia Blackham and Sean Cooney (eds), *Research Methods in Labour Law: A Handbook* (Edward Elgar Publishing, 2024) <<https://doi.org/10.4337/9781803925257>>.

7 Blackham, ‘Interrogating New Methods in Socio-Legal Studies’ (n 5).

8 Andrew Leigh, ‘Fair Work in the Age of AI’ (Speech, Canberra, 3 September 2025) <<https://ministers.treasury.gov.au/ministers/andrew-leigh-2025/speeches/address-seizing-opportunities-ai-while-protecting-fair-go>>.

9 Cf @MikeJShowalter (X, 23 March 2026 AEST) <<https://x.com/MikeJShowalter/status/2035977171437560207>>, archived at <<https://perma.cc/KDA9-Z7LE>> (‘X Post 23 March 2026’), quoted in @orinkerr.bsky.social (Orin Kerr) (BlueSky, 24 March 2026 AEST) <<https://bsky.app/profile/orinkerr.bsky.social/post/3mhr6w2ftgs2m>>, archived at <<https://perma.cc/QG2R-VY9X>> (‘BlueSky Post 24 March 2026’).

to partly act as a check on the techno-optimism that often pervades universities and their management. But I am also keen to critically evaluate the usefulness of these tools, when some of our colleagues have become smitten with them as 'assistants'.

This introduction has been written before engaging with Copilot on this project. It is solely human authored. In the Parts that follow, though, I have attempted to 'co-write' this article with Copilot.¹⁰ Copilot is perhaps not the best LLM tool for this job; using Claude might have been a better fit for this project. Equally, given Microsoft is rolling out Copilot at universities and across the public service, often without notice or clear consent,¹¹ it offers a critical case study of the risks and emerging problems with integrating 'AI' into our 'workflow'. My experience of using Copilot is also unlikely to be replicable, as Copilot is trained and draws on all internal documents to which I have access (such as all my email, SharePoint documents and so on). I have also personalised my Copilot 'experience' and settings to request that the outputs be professional yet accessible, and to not include emojis, for example. This article therefore captures my experience at a particular point in time (March 2026); Copilot and its usefulness will likely change over time. I have also used the more accessible Copilot, not deep reasoning or agentic tools like Researcher, which are more tailored to finding sources. I am writing this article as an experienced academic who finds joy in the process of writing. The benefits for me, personally, of an LLM are perhaps going to be less significant than for others for whom writing is a chore, or where English is their second or third language, for example. The risks of undertaking such a project, however, are likely to be more universal.

To ensure research integrity, the chat log with Copilot which informed this article is available from the author on request.¹² Further, where material produced by Copilot is included in this article, it has been explicitly demarcated using quotation marks or indentation. If writing is not in quotation marks or indented, or expressly attributed to Copilot, it remains my own. In Part II, I explore the process of framing and developing an article with Copilot. Part III turns to questions of authorship and integrity, before considering whether the use of LLMs and AI represents the future of legal research, and the risks that might emerge, in Part IV. Part V concludes.

II FRAMING AND STRUCTURING AN ARTICLE

To start, then, I sought Copilot's advice on how to structure and frame an article of this nature, using the article's title as a prompt. Copilot suggests

10 Though, in practice, I have used little of what was produced by Copilot directly in this article.

11 In relation to consumers: see Australian Competition and Consumer Commission, 'Microsoft in Court for Allegedly Misleading Millions of Australians over Microsoft 365 Subscriptions' (Media Release, 27 October 2025) <<https://www.accc.gov.au/media-release/microsoft-in-court-for-allegedly-misleading-millions-of-australians-over-microsoft-365-subscriptions>>.

12 This chat log has had to be redacted in part given Copilot has referred to sensitive information and documents in places. This privacy risk is discussed further below.

‘foreground[ing] your scepticism as a methodological stance’ (as, indeed, I have above); and that the scope of the article should be a ‘reflective account of *process*, *judgment*, and *control* in legal scholarship’. My initial reaction (even as a sceptical recipient) is that this seems a persuasive and elegant framing of what I am attempting. Frankly, it is annoying to be one-upped by what is self-proclaimed (by Copilot) to be ‘a probabilistic text generator’.

Copilot also prompts me to include a section in this article on methods, and what ‘co-writing’ actually entails. Copilot suggests that I use AI for ‘structuring, re-phrasing, stress-testing arguments, or generating counter-positions’ but not for finding legal sources, stating legal rules, or making evaluative claims without verifying information. Oddly, this seems to be the very opposite of what a research assistant might typically do; that is, find legal sources and references, to help state and articulate legal rules. Equally, if Copilot is not being asked to identify legal rules or sources, the risk of hallucination is arguably far reduced. That said, it is perhaps unlikely that most people using AI tools would abstain from asking about legal rules and sources. Copilot encourages me to articulate that there is human control at all stages of the article, in relation to ‘argument selection’, ‘framing’, ‘tone’ and ‘citations’, and to ensure that AI use is acknowledged.

The article structure that Copilot prompts me to use is arguably underwhelming and lacking in sophistication: methods; what AI is good at; what AI is bad at; authorship, responsibility and academic integrity; and reflections for legal researchers. (And yet, Copilot posits that its strengths as a tool include ‘Structural prompts’ and ‘Offering alternative ways to order an argument.’) The problems with using AI that Copilot flags reflect some of those already noted above: of mis-citing or mis-stating legal rules; of providing sources that cannot be verified or are not authoritative; and of ‘Producing consensus-sounding language where the law is contested.’

As a somewhat sceptical user, my key concern relates to the final risk Copilot posits: ‘False confidence: The risk of polished prose masking weak or incorrect analysis.’ It requires significant mental effort to interrogate and challenge the authoritative-sounding writing produced by LLMs. Interestingly, too, in writing this article, I have found my own writing often defaulting to the language used by Copilot, even if that is not my natural style of writing (for example, trying not to use the word ‘polished’ in the previous sentence to describe Copilot’s language required some mental effort). This might reflect our tendency to trust or defer to automated outputs (what is known as ‘automation bias’).¹³ It also perhaps reflects broader concerns that LLMs are reshaping language and how we communicate as humans.¹⁴ As a researcher, I am not delighted that an unexpected consequence of

13 Thomas Dratsch et al, ‘Automation Bias in Mammography: The Impact of Artificial Intelligence BI-RADS Suggestions on Reader Performance’ (2023) 307(4) *Radiology* e222176:1–9 <<https://doi.org/10.1148/radiol.222176>>.

14 Jason Snyder, ‘AI Is Rewriting Reality, One Word at a Time’, *Forbes* (online, 14 May 2025) <<https://www.forbes.com/sites/jasonsnyder/2025/05/14/ai-is-rewriting-reality-one-word-at-a-time/>>; Vanessa Bates Ramirez, ‘ChatGPT Is Changing the Words We Use in Conversation’, *Scientific American* (online, 11 July 2025) <<https://www.scientificamerican.com/article/chatgpt-is-changing-the-words-we-use-in-conversation/>>; Hiromu Yakura et al, ‘Empirical Evidence of Large Language Model’s Influence

writing this article is that I sound more like the bot, and less like my human self. While Copilot offered few genuinely new ideas or helpful suggestions for this article, it did subtly influence my writing style, and perhaps not in a positive way.

III 'AUTHORSHIP, RESPONSIBILITY, AND ACADEMIC INTEGRITY'

My broader concern, in embarking upon this project, was to consider how the use of LLMs and AI might reshape our understanding of authorship. If a paper is drafted, even in part, by an LLM, how can we rightly claim authorship of our outputs? In discussions online relating to the *SMU Law Review Forum* paper,¹⁵ the author and some commentators compared AI-generated text to ghost-written text for judges or law firm partners, where judges' associates or junior lawyers and their contributions would rarely be acknowledged.¹⁶ But academic authorship conventions are different, particularly in Australia. Indeed, where a contributor makes a 'significant' contribution to the writing of a paper, they should be acknowledged as an author. The *Authorship* guide supporting the *Australian Code for the Responsible Conduct of Research* states that 'An author is an individual who: has made a significant intellectual or scholarly contribution to research and its output, and agrees to be listed as an author.'¹⁷ Does this mean, if we draw on text generated by an LLM, authorship should be shared with that LLM? (Copilot's outputs emphasise that authorship rules only apply to individuals, not LLMs, though the *Authorship* guide was drafted before LLMs came to be widespread.)

More broadly, though, if LLMs have been trained on vast amounts of scholarly materials, including those used without attribution and potentially in breach of intellectual property laws,¹⁸ how can we acknowledge the authorship of those thousands of authors who have fed into this 'probabilistic text generator'?¹⁹ In this article, where I have used Copilot's outputs, I have quoted them as I typically

on Human Spoken Communication' (Research Paper, 8 July 2025) <<https://doi.org/10.48550/arXiv.2409.01754>>.

15 Showalter, 'Why Originalism Requires Robotic Judging' (n 1).

16 See @orinkerr.bsky.social (Orin Kerr) (BlueSky, 21 March 2026 AEST) and responses, <<https://bsky.app/profile/orinkerr.bsky.social/post/3mhiy7psqb22v>>, archived at <<https://bsky.app/profile/orinkerr.bsky.social/post/3mhr6w2ftgs2m>>; Showalter, 'X Post 23 March 2026' (n 9), quoted in Kerr, 'BlueSky Post 24 March 2026' (n 9).

17 National Health and Medical Research Council, Australian Research Council and Universities Australia, *Authorship: A Guide Supporting the Australian Code for the Responsible Conduct of Research* (Guide, 2019) 2 <<https://www.arc.gov.au/sites/default/files/2023-05/Authorship-Guide.pdf>>.

18 In the US: see *Bartz v Anthropic PBC*, 787 F Supp 3d 1007 (ND Cal, 2025); Stephanie Schmidt, 'Bartz v. Anthropic: Settlement Reached after Landmark Summary Judgment and Class Certification', *Inside Tech Law* (Blog Post, 5 September 2025) <<https://www.insidetechlaw.com/blog/2025/09/bartz-v-anthropic-settlement-reached-after-landmark-summary-judgment-and-class-certification>>.

19 'AI Best Practices for Authors', *The Authors Guild* (Web Page, 11 May 2026) <<https://authorsguild.org/resource/ai-best-practices-for-authors/>>. Indeed, some tools can be prompted to reproduce other authors' work, nearly in full: Xinyue Liu et al, *Alignment Whack-a-Mole: Finetuning Activates Verbatim Recall of Copyrighted Books in Large Language Models* (Columbia Public Law Research Paper No 6449179, 20 March 2026) <<https://papers.ssrn.com/abstract=6449179>>.

would quote from any other source. But if authors use these outputs directly, can (sole) authorship really be claimed with integrity?

In Copilot's guidance for writing this Part of the article (this Part's title is Copilot's own), its primary point is that '[a]uthorship remains human'; when said with such confidence, it is perhaps unsurprising that some authors are willing to take full credit for an LLM's outputs. But can authorship really remain human in this context, when LLMs are reshaping our thinking and our language? And how can we effectively acknowledge the thousands of authors whose original work has informed these models? On one thing, though, Copilot and I agree: 'Responsibility for claims cannot be delegated to an AI.' As authors, outputs published under our names must remain our responsibility. Potentially, though, this leads us to a scenario where we cannot claim full credit for an output, but must take full responsibility. As an author, this makes me distinctly uncomfortable.

After drafting the paragraph above, I followed Copilot's suggested further prompt and asked it to '[d]raft the authorship section [of the article] in my voice', to compare what I produced as a human author with what Copilot produced. What was produced, though, was more akin to a methods and transparency section, not a reflection on the importance of human authorship. The only exception to this was one paragraph on authorship and its relationship to accountability:

Authorship in legal scholarship has always been tied to accountability: for claims made, authorities relied upon, and values defended. That accountability cannot meaningfully be assigned to a system that lacks understanding, intention or professional obligation. For this reason, the AI tool used in preparing this article is not listed as a co-author, nor could it be. To do so would risk confusing assistance with authorship and undermining established norms of scholarly responsibility.

I would question, though, whether authorship is just about accountability. While accountability seems critical, authorship is surely about more than just being accountable for our ideas and outputs. Ideally, legal research should be a scholarly dialogue: across and between authors, drawing and building on others' ideas and contributions, and extending the field through our own research, ideas and scholarship. A view of authorship as just accountability seems an impoverished view of academic scholarship. Rather, authorship represents a form of voice, participation and engagement with a legal and scholarly community. Arguably, where words are not our own, but produced by an LLM, our own voice is missing from this dialogue (however much the LLM seeks to approximate our 'voice'). LLM 'assistance' potentially undermines and dilutes our scholarly voice, and might serve to constrain and limit our best or most innovative ideas.

Even if LLM-produced text is edited, amended and refined by a human author, it remains unclear to me how we, as scholars, might meaningfully claim authorship over such text. Further, given the challenges of looking past the veneer of LLMs and their overconfidence, it seems unlikely that we would edit or amend such text with sufficient diligence. In practice, it is likely to be easier to defer to the LLM's language and articulation of a point, with all its limits, overconfidence and lack of nuance, further undermining any claim to human authorship.

IV THE FUTURE OF LEGAL RESEARCH?

This study implies, then, that while LLM-produced research might be easier or quicker,²⁰ it is not necessarily better, and might undermine the critical 'dialogue' that gives legal research its richness and vibrancy. LLMs appear to often produce material that is 'good enough' or potentially passable, but not necessarily good or outstanding. As legal scholars, surely we are seeking something more than work that is 'good enough'. This has implications, too, for education and research supervision. Our role as educators and supervisors is increasingly likely to involve convincing students and researchers that the hard work of independent thinking and writing is important and necessary. During a PhD, our brains should hurt (at least, on occasion) from deep thinking. Without these formative research processes and experiences, we are all likely to be poorer scholars and thinkers.

Replacing research assistants with LLMs also poses risks. First, it is unlikely that an LLM can replace a trained research assistant, at least where that research assistant is tasked with finding authoritative and reliable legal sources and rules. But some scholars seem to disregard the limits of what an LLM can reliably contribute to legal research.²¹ This likely poses substantial risks to the integrity of legal research, including due to the risk of hallucinations.²² Second, and more broadly, engaging junior scholars as research assistants, including on large research projects (where funding might – exceptionally – be available for legal research), forms an important way of training and mentoring future generations of legal scholars. While LLMs might appear 'cheaper' – at least in terms of the direct costs, which have not been outsourced – the consequences for legal training and researcher development are potentially significant.

Another concern, too, is that LLMs might ignore or miss key issues, yet the overconfidence of LLM outputs hides these substantive gaps. For example, one, potentially critical, point that did not appear in Copilot's initial outputs in this study related to the environmental cost of using LLMs to generate legal research. A legitimate critique of this experiment is that I have used only a handful of prompts, and that Copilot's outputs could be refined, finessed and improved with additional engagement and further prompts. But I am conscious that each prompt draws on energy and water. Is the potential output worth the environmental cost?²³ When explicitly prompted, Copilot provided an estimate that the energy cost of my initial prompt (which generated a 'long, structured academic response') was around 0.3–

20 Showalter, 'Why Originalism Requires Robotic Judging' (n 1).

21 For example, Showalter argues that hallucinations were rare in his study, implying that LLMs were used to find sources and references: see, eg, Showalter, 'X Post 23 March 2026' (n 9), quoted in Kerr, 'BlueSky Post 24 March 2026' (n 9).

22 Blackham, 'Interrogating New Methods in Socio-Legal Studies' (n 5) 89; Matthew Dahl et al, 'Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models' (2024) 16(1) *Journal of Legal Analysis* 64 <<https://doi.org/10.1093/jla/laec003>>.

23 Senate Select Committee on Adopting Artificial Intelligence, Parliament of Australia, *Select Committee on Adopting Artificial Intelligence* (Final Report, November 2024) 143–9 [6.1]–[6.34] <[https://parlinfo.aph.gov.au/parlInfo/download/committees/reportsen/RB000470/toc_pdf/SelectCommitteeonAdoptingArtificialIntelligence\(AI\).pdf](https://parlinfo.aph.gov.au/parlInfo/download/committees/reportsen/RB000470/toc_pdf/SelectCommitteeonAdoptingArtificialIntelligence(AI).pdf)>.

0.5 watt-hours;²⁴ and the water usage was about 0.3–0.6 litres. This excludes the initial training of the model. Further prompts and responses would require further energy and further water. While AI companies are not required to explicitly quantify the environmental cost of using their tools, can I, in good conscience, justify the potential environmental cost of further finessing this (rather average) material?

Truly unique or original ideas are also unlikely to emerge from a ‘predictive text generator’, though LLMs might bring existing ideas together in unexpected or new ways. Research that solely relies on ideas from an LLM is therefore unlikely to be the most original or compelling, though it is at least possible that LLMs might prompt a human user to think in new or different ways, generating innovative ideas.

There is a broader question, too, of what AI and LLMs might mean for legal research methods. Copilot describes legal research as ‘interpretive and normative’, and posits that I should ‘[s]tress the centrality of judgment, context, and contestation’ to legal research. Given Copilot is a ‘probabilistic text generator’, AI cannot “reason” or “understand law”. On these points, Copilot is not wrong.²⁵ But we might further emphasise the critical human and normative aspects of legal research, even legal doctrinal research.²⁶ Law is innately human in its construction and development; judges and law makers are only human, after all.²⁷ Arguably, the interpretation, systematisation and normative development of law should also be an innately human process, not one devolved to LLMs.

LLMs also raise broader concerns of privacy and data security. Data shared might be used for LLM training,²⁸ or become accessible to other users or developers.²⁹ While Copilot is meant to offer more data security than other models like ChatGPT,³⁰ I still found that it was linking to sensitive documents when

24 This may undersell the energy usage for this specific, complex query: see the similar estimate for a simple query in Josh You, ‘How Much Energy Does ChatGPT Use?’ (7 February 2025) *Epoch AI* <<https://epoch.ai/gradient-updates/how-much-energy-does-chatgpt-use>>.

25 Simon Willison, ‘Imitation Intelligence, My Keynote for PyCon US 2024’, *Simon Willison’s Weblog* (Blog Post, 14 July 2024) <<https://simonwillison.net/2024/Jul/14/pycon/>>; Emily M Bender et al, ‘On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?’ (Conference Paper, ACM Conference on Fairness, Accountability, and Transparency, 1 March 2021) 610, 615–17 <<https://dl.acm.org/doi/10.1145/3442188.3445922>>.

26 See, eg, Mark van Hoecke, ‘Legal Doctrine: Which Method(s) for What Kind of Discipline?’ in Mark van Hoecke (ed), *Methodologies of Legal Research: What Kind of Method for What Kind of Discipline?* (Hart, 2011) 1; Alan Bogg, ‘Doctrinal Method in Labour Law in Common Law Systems: Potential, Problems, Prospects’ in Alysia Blackham and Sean Cooney (eds), *Research Methods in Labour Law: A Handbook* (Edward Elgar Publishing, 2024) 186 <<https://doi.org/10.4337/9781803925257.00023>>.

27 See, eg, Alysia Blackham, ‘Reconceiving Judicial Office through a Labour Law Lens’ (2019) 47(2) *Federal Law Review* 203 <<https://doi.org/10.1177/0067205X19831826>>.

28 ‘How Your Data Is Used to Improve Model Performance’, *OpenAI* (Web Page, 16 June 2026) <<https://help.openai.com/en/articles/5722486-how-your-data-is-used-to-improve-model-performance>> (‘How Your Data Is Used’).

29 Chris Stokel-Walker, ‘Exclusive: Google Is Indexing ChatGPT Conversations, Potentially Exposing Sensitive User Data’, *Fast Company* (online, 30 July 2025) <<https://www.fastcompany.com/91376687/google-indexing-chatgpt-conversations>>.

30 Cf ‘Privacy FAQ for Microsoft Copilot’, *Microsoft Support* (Web Page) <<https://support.microsoft.com/en-us/microsoft-copilot/privacy-faq-for-microsoft-copilot>>; ‘How Your Data Is Used’ (n 28).

producing responses to fairly generic prompts (such as 'When should someone be included as an author?'). Given Copilot has access to everything I have access to, and no sense of what should or should not be referred to, there are potentially substantial risks of the wrong information being used, or information being used in inappropriate ways.

V 'CONCLUSION: SCEPTICISM AS AN ETHICAL STANCE, NOT RESISTANCE TO CHANGE'

Close by reinforcing:

- Your scepticism is grounded in method, responsibility, and care for legal reasoning.
- Co-writing with AI is possible only if:
 - The human author remains intellectually dominant.
 - AI use is limited, transparent, and critically examined.
- The experiment strengthened, rather than weakened, your commitment to slow, careful legal scholarship.

For me, personally, I have no desire to engage with an LLM in my writing process in the future. I dislike how my writing style has been influenced, even in this exploratory project. I will revert to my own, independent ways of thinking and structuring legal research; then any errors or gaps in reasoning can truly be my own. Copilot describes a traditional research process as 'slow'; and perhaps it is. Certainly, writing 'with' Copilot took less time than my typical writing process; perhaps, in part, because this article primarily critiques one source: Copilot itself. But my experience in this case study is that writing with the aid of an LLM will likely lead to less innovative and poorer quality legal research, that undermines the inherently human undertaking of legal research, and fails to meaningfully contribute to the research dialogue that is so critical to legal scholarship.

As a community of scholars, we need universities, law journals and referees to build a unified stance regarding how, if at all, AI should be used in legal research. Achieving any form of consensus will require us all to grapple with fundamental questions regarding the nature of legal research and authorship. These are questions that legal scholars cannot shy away from as LLMs become more widespread and are increasingly utilised in legal research. Even once some sort of consensus is developed, though, my words will remain my own.