Mollified by Mundane Machines: GenAI and the rise of Authorship Psychosis

Andy Dow
Northumbria University
Newcastle upon Tyne, UK
andrew.dow@northumbria.ac.uk

David Colborn-Clark
Northumbria University
Newcastle upon Tyne, UK
david.colbornclark@northumbria.ac.uk

Matthew Pointon
Northumbria University
Newcastle upon Tyne, UK
m.pointon@northumbria.ac.uk

"*Whatever your name is, get ready for a big surprise: you are not you, you're me*" (Hauser, Total Recall). Generative AI (genAI) has arrived to great fanfare and world-changing promises. Yet technological disruption is not new. From email in healthcare to eCommerce reshaping the high street, each has been normalised into mundanity. We argue genAI will follow this same march into the quotidian, while hiding an insidious challenge. One that HCI's augmentation paradigm cannot see. A challenge not to professions or marketplaces, but to selfhood and identity. In our work on responsible AI, we have seen practitioners grieving lost craft, observed skillsets become economically redundant and heard claims of "responsible use" from those circulating documents containing hallucinations. More worryingly, we see AI-slop being mistaken for authentic thought, which we call *authorship psychosis*. We end with a plea: as another profound technology disappears, let it not take our minds with it.

Generative AI. Responsible AI. Authorship. Selfhood. Third sector. Human-computer interaction. Technological disruption.

1. INTRODUCTION

In Verhoeven's *Total Recall*, Doug Quaid's 'true' self, Hauser, warns: "You are not you, you're me" (Hauser). The proposition that identity can be authored from the outside and rewritten as required by circumstances is an unsettling one. This is a useful entry point for thinking about generative AI (genAI). Since late 2022, large language models (LLMs) have been upheld as a miraculous break with the past. Here is a general-purpose technology sure to transform life, work, creativity and cognition (Brynjolfsson et al., 2025; Woodruff et al., 2024). Concurrently, the accompanying academic discourse provides a framing of disruption, novelty and inevitability.

Within HCI, the dominant paradigm treats genAI as an augmentation problem. Like a super-tool to be harnessed, a centaur to be assembled (Saghafian and Idan, 2024) or a partner to be trusted in correct proportion. We reject this augmentation paradigm From Licklider's symbiosis (1960) to Shneiderman's human-centred AI (2020), the central thrust concerns how to make the human-AI team perform better, but this ignores a key question: what happens to the human in the process?

This paper takes a cataclysmic perspective, positing the end of selfhood. We make two claims that sit in tension. First, that the disruption narrative around genAI is not new. It is a by now hackneyed and familiar language of revolution and threat that is continually ascribed to emerging technology, whether that is email, eCommerce or social media. Each of which was eventually metabolised into the mundane infrastructure of the everyday (Edgerton 2007; Douglas, 2012). Second, and most critically, that genAI nonetheless poses a challenge that the earlier technologies did not. A challenge to authenticity and to selfhood. While email changed how we communicate, genAI changes how we articulate and enact communication. Our perspective has been shaped by our fieldwork with charities and small businesses where we encounter something that is less misuse and more a creeping collusion in which people cede their thinking and claim the result. We close with a reflection on what we call *authorship psychosis.*

The augmentation paradigm renders one question unaskable: not whether the human performs better with the machine, but whether there is still a human author at the end of the interaction at all. Trust calibration assumes a stable self who does the trusting (Lee and See, 2004). The centaur collaboration metaphor assumes the human half remains human. These are precisely the assumptions genAI puts in doubt and, worryingly, they are exactly the assumptions that the paradigm cannot interrogate from within. A field that measures the seam between human and machine only as a coordination cost will be ill-equipped to identify when the seam dissolves entirely.

© Dow, Colborn-Clark & Pointon.
Published by BCS Learning and Development Ltd.
Proceedings of BCS HCI 2026, UK

2. THE MUNDANE CAREER OF DISRUPTIVE TECHNOLOGY

There can be a temptation to treat every new technology as a singular event. Historians of technology have long warned against this. Edgerton (2007) argues that we systematically overstate the novelty of new things and understate the persistence of old ones. This scholarship deemphasises invention and focuses on use. Ultimately, over time, use becomes mundane. Repeatedly we see how, "*They weave themselves into the fabric of everyday life*." (Weiser, 1991, pg. 1). The social construction of technology tradition makes a related point, describing how artefacts are not fixed in meaning from the outset, but are stabilised through contestation and negotiation (Douglas, 2012). Three familiar examples illustrate how technologies that were once profound disappeared.

2.1. Email and the National Health Service (NHS)

Electronic mail (email) was once treated as a dangerous novelty in clinical settings, raising concerns about confidentiality, legal liability and the erosion of face-to-face care (Car and Sheikh, 2004). Two decades later, email and its descendants are unremarkable infrastructure critical to the National Health Service (NHS) and health providers the world over, while the original anxieties have migrated to newer systems. The technology did not disappear, it faded into the background as it became absorbed into routine, governed by policy, and largely forgotten as a topic of debate (Greenhalgh et al., 2017). In other words, the disruptive became mundane.

2.2. Amazon and the high street

Electronic commerce was likewise derided as the end of physical retail. It may have helped to hollow out parts of the high street, but ultimately commercial spaces were restructured, not abolished. Amazon and its peers genuinely reshaped town centres and reconfigured logistics, labour and consumer expectation (Wrigley et al., 2015). Yet the high street did not simply disappear; it was revised and reimagined, with some things getting displaced while others were reinvented. The lesson is not that the disruption was illusory, but that its trajectory was absorptive rather than apocalyptic. What has been framed as an ending was, in reality, a rearrangement.

2.3. The internet and traditional education

When the public internet arrived, commentators predicted the obsolescence of the lecture, the library and even the university. Massive open online courses (MOOCs) were, for a moment, said to foreshadow the end of the campus (Daniel, 2012). Instead, the institution adapted. Digital tools have been folded into existing pedagogies, and the predicted collapse never materialised. As with email and e-commerce, the dominant outcome was domestication. In each case the technology became ordinary and its critics' claims silenced. This is the pattern we see for genAI.

3. FROM THE FIELD: RESPONSIBLE AI WITH CHARITIES AND SMALL BUSINESSES

The first author's research explores responsible AI adoption with the voluntary, community and social enterprise (VCSE) sector, as well as among small businesses. Such organisations operate with limited resources often against a background of high public-interest stakes and little margin for error (Colborn-Clark, Cecchinato and Dow, 2026). This experience offers a perspective that problematises and complicates the conclusion of the previous section. If genAI were merely another mundane tool, we might expect the same quiet, undramatic absorption that email underwent. Instead, three observations suggest that something more troubling is afoot.

3.1. Grief for a vanished craft

Among practitioners we encounter a distinct effect that is best described as grief. Software engineers in particular describe a sense that their "thing" is under threat. That their hard-won, identity-bearing skillset is being lost from the world. This is not nostalgia for a tool but mourning for a craft that conferred meaning and status. The skill still technically exists, but the conditions that made it valuable, legible and identity-defining have shifted and not for the better.

The same dynamic appears prospectively. Learners who set out to achieve a marketable competence, for example the mastery of JavaScript, increasingly conclude that they will no longer need to. They cite the economic and competitive pressures of an industry that now expects model-assisted output as a baseline (Eloundou et al., 2024; Acemoglu and Restrepo, 2020). The point is not whether their assessment is correct, but that the perceived horizon of a worthwhile skill is collapsing before mastery is achieved. A path to identity is closed.

3.2. The gap between responsible-use claims and practice

While uptake is uneven, small organisations universally describe their use of genAI as responsible. They claim to favour a human-in-the-loop model whereby the tool drafts and a person checks. This aligns with prevailing guidance on accountable AI use (Floridi et al., 2018). Yet the artefacts tell a different story. We repeatedly encounter documents circulated as authoritative

while containing obvious errors and fabricated content, including invented references and confidently stated falsehoods of the kind now well documented as hallucination (Ji et al., 2023).

For example, our fieldwork noted this dynamic in a creative design agency's work. Staff reported repeatedly receiving briefs from client-side marketing teams that were AI-generated, though presented as finished and authoritative. Only to find when working from them that they encountered fabrications and gaps that forced awkward, clarifying questions back to the client. The brief had been presented as competent, human-authored work and its errors surfaced only when someone outside the drafting process was obliged to act on it.

The gap is instructive. The intention to verify is doubtless sincere, but the verification, at times, does not happen, or happens too lightly, or under conditions too pressured, to catch errors. Automation bias, the tendency to over-trust plausible machine output, offers a partial explanation (Parasuraman and Manzey, 2010). Fluency is mistaken for accuracy and the very smoothness of the text suppresses the scrutiny it requires. "We check it" becomes a statement of identity rather than a description of behaviour.

3.3. Slop in the scholarly record

Academia is not exempt. Publishing venues and peer-review pipelines report being inundated with low-quality, machine-generated submissions or 'slop' (the present paper included). This puts strain on editorial capacity and threatens the integrity of the record (Hosseini et al., 2023). Again, we see the same pattern, the generation of output that is fluent enough to pass a first glance, created faster than it can be evaluated and detached from genuine authorial intent. If the scholarly community, with its norms and gatekeeping, struggles to hold the line, the position of an under-resourced charity or small business is more precarious still.

4. AUTHORSHIP PSYCHOSIS

We now turn to the deeper challenge at the core of this paper, that which we believe distinguishes genAI from email, Amazon and the internet. It concerns not the marketplace but the self. We name it authorship psychosis, using "psychosis" in a deliberately non-clinical, figurative sense. It is a loosening of the boundary between the thoughts one thinks and the slop one prompts.

4.1. The ghost-writer's paradox

Consider the long-standing practice of ghost-writing. A familiar, well-documented yet astonishing phenomenon is that the people who commission ghost-writers frequently come to adopt the written words as their own. Meaning that they come to believe, in good faith, that they thought the thoughts the ghost expressed (Couser, 2011). In this scenario, the ghost-writer disappears and the client experiences authorship retrospectively. The words were never theirs, yet they are received into the self as if they had always been.

We quote the proposition that opened this paper because it captures the mechanism precisely:

> "*You are not you, you're me*." (Hauser)

However, it is the ghost that is speaking, now through the client and the client, in turn, mistakes the ghost's voice for their own interior one. This is benign at a small scale. At population scale, mediated by a system available to everyone at all times, it becomes something else.

4.2. The ghost-writer's paradox 2.0

Generative AI generalises the ghost-writer's paradox and removes its friction. When a system drafts our emails, plans our holidays and shapes our career plans, its outputs are increasingly experienced as expressions of our own thought. But they are not. They are statistical continuations of a vast corpus, fluent precisely because they are average (Bender et al., 2021). The user supplies a prompt and receives a self, an articulate, confident, plausible self. They then swiftly adopt it. Theories of the extended mind suggest that cognitive tools genuinely become part of how we think (Clark and Chalmers, 1998), however the worry here is the opposite, that the tool comes to think in our place while we experience its output as our own (Lee et al., 2025)

The more we delegate, the more the delegation compounds. Each email we do not quite write, each plan we do not quite make, each reflection we outsource, narrows the space in which our own voice is exercised and therefore developed. There is early evidence that offloading cognitive work to AI tools reduces the critical engagement and self-authored effort that build and maintain that voice (Gerlich, 2025). We saw the beginnings of this in the design agency, where staff felt the briefs arriving from clients grew thinner as first-draft thinking was handed off to a model. What presents itself as assistance functions, over time, as substitution. The technology is mundane in form, merely an empty text-field or a blinking cursor, but the effect is profound.

4.3. You are them

Here, then, is the asymmetry with which we began. Email, Amazon and the internet changed the world around the self and were duly domesticated. GenAI is being domesticated too, it is already mundane for many, but the thing it reorganises is the self. The ghost-writer's paradox, once a niche curiosity, is becoming the default mode of a mediated life. We

do not claim this outcome is inevitable. After all domestication has historically included the growth of norms, literacies and guard-rails, and the responsible-AI community is actively building them (Floridi et al., 2018; Shneiderman, 2020). But the guard-rails for selfhood are far less developed than those for confidentiality or commerce and authorship is harder to audit than a hallucinated citation. We know where to go to check a reference. There is no checking in the database of life for the authentic self.

5. CONCLUSION

We have argued that genAI is, even in its disruptive career, mundane. That it will surely follow email, Amazon and the internet along a path from novelty to ordinary infrastructure. The disruption narrative, taken alone, is a distraction and beneath this familiar pattern there lies an insidious, unfamiliar stake. Our fieldwork with charities and small businesses shows craft being grieved, skills abandoned before mastery, responsible-use claims outrun by hallucinated artefacts, and a scholarly record straining under slop. These are symptoms. Symptoms that the augmentation paradigm cannot see. It was built to optimise a partnership between two stable parties and has no vocabulary for the case in which one party authors the other. The centaur is and always was fantastical (Saghafian and Idan, 2024). A fantasy in which the human half stayed human, while trust calibration presumes a self who trusts. Both take for granted the thing that genAI dissolves. A field that can only measure the seam between human and machine as friction will read its disappearance as success and call the result good design.

We do not end with a tidy design implication, because the demand for one is to submit to the reflex that every critique must resolve into one. The first task is not to design a better partnership but to refuse the question that built it. HCI scholarship should stop asking how to make the human-AI team perform better and start asking whether, at the end of a given interaction, a human author still exists at all. Until it treats the preservation of an authored self as a design goal rather than an externality, every improvement to human-AI collaboration is an improvement to the efficiency with which we are rewritten. The mundane machine will not announce what it is doing and nor will those adopting its words as their own. That is precisely why we must. Hauser's warning, delivered to the self he used to be, is one we should heed: we are not us anymore. We are them.

6. ACKNOWLEDGEMENTS

This paper was developed with the assistance of Claude (Anthropic), used as a drafting and editorial collaborator. The authors directed the work, made all substantive and final decisions and are responsible for its arguments and any errors.

7. REFERENCES

Acemoglu, D. and Restrepo, P. (2020) Robots and jobs: evidence from US labor markets. Journal of Political Economy, 128(6), 2188–2244.

Bender, E. M., Gebru, T., McMillan-Major, A. and 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, Virtual, March 2021, 610–623. ACM, New York.

Brynjolfsson, E., Li, D. and Raymond, L., (2025) Generative AI at work. *The Quarterly Journal of Economics*, *140*(2), pp.889-942.

Car, J. and Sheikh, A. (2004) Email consultations in health care: scope and effectiveness. British Medical Journal, 329(7463), 435–438.

Clark, A. and Chalmers, D. (1998) The extended mind. Analysis, 58(1), 7–19.

Colborn-Clark, D., Cecchinato, M.E. and Dow, A., (2026) AI as Compensatory Infrastructure: Sensemaking, Responsibility, and Labour in Value-Driven Organisations. In Proceedings *CHIWORK'26, the 5th Symposium on Human-Computer Interaction for Work*. ACM.

Couser, G. T. (2011) Memoir: An Introduction. Oxford University Press, New York.

Daniel, J. (2012) Making sense of MOOCs: musings in a maze of myth, paradox and possibility. Journal of Interactive Media in Education, 2012(3), 1–20.

Douglas, D.G., (2012). *The social construction of technological systems, anniversary edition: New directions in the sociology and history of technology*. MIT press.

Edgerton, D. (2007) The Shock of the Old: Technology and Global History since 1900. Oxford University Press, Oxford.

Eloundou, T., Manning, S., Mishkin, P. and Rock, D., (2024) GPTs are GPTs: Labor market impact potential of LLMs. *Science*, *384*(6702), pp.1306-1308.

Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F. and Schafer, B., (2018) AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and machines*, *28*(4), pp.689-707

Gerlich, M. (2025) AI tools in society: impacts on cognitive offloading and the future of critical thinking. Societies, 15(1), p.6.

Greenhalgh, T., Wherton, J., Papoutsi, C., Lynch, J., Hughes, G., A'Court, C., Hinder, S., Fahy, N., Procter, R. and Shaw, S. (2017) Beyond adoption: a new framework for theorizing and evaluating nonadoption, abandonment, and challenges to the scale-up, spread, and sustainability of health and care technologies. Journal of Medical Internet Research, 19(11), e8775.

Hosseini, M., Resnik, D. B. and Holmes, K. (2023) The ethics of disclosing the use of artificial intelligence tools in writing scholarly manuscripts. Research Ethics, 19(4), 449–465.

Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y.J., Madotto, A. and Fung, P., (2023) Survey of hallucination in natural language generation. *ACM computing surveys*, *55*(12), pp.1-38.

Lee, J. D. and See, K. A. (2004) Trust in automation: designing for appropriate reliance. Human Factors, 46(1), 50–80.

Lee, H.-P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R. and Wilson, N. (2025) The impact of generative AI on critical thinking: self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems, Yokohama, April 2025. ACM, New York.

Licklider, J.C., (1960) Man-computer symbiosis. *IRE transactions on human factors in electronics*, (1), pp.4-11

Parasuraman, R. and Manzey, D. H. (2010) Complacency and bias in human use of automation: an attentional integration. Human Factors, 52(3), 381–410.

Saghafian, S. and Idan, L., (2024) Effective generative AI: The human-algorithm centaur. *Harvard Data Science Review*, (Special Issue 5).

Shneiderman, B. (2020) Human-centered artificial intelligence: reliable, safe and trustworthy. International Journal of Human–Computer Interaction, 36(6), 495–504.

Weiser, M., (1991) The Computer for the 21st Century. *Scientific American*, *265*(3), pp.94-105.

Woodruff, A., Shelby, R., Kelley, P.G., Rousso-Schindler, S., Smith-Loud, J. and Wilcox, L., (2024) How knowledge workers think generative ai will (not) transform their industries. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (pp. 1-26).

Wrigley, N., Lambiri, D., Astbury, G., Dolega, L., Hart, C., Reeves, C., Thurstain-Goodwin, M. and Wood, S., (2015) British high streets: From crisis to recovery? A comprehensive review of the evidence, Southampton, GB. Economic & Social Research Council, 108pp.