

The Shadow of the Machine

Hypertext in the Age of Artificial Intelligence

David E. Millard¹

¹University of Southampton, UK

Abstract

This paper examines the parallels between the transformational impacts of hypertext and those that will come as a result of generative AI, arguing that we are witnessing not incremental change but a paradigm shift in the nature of writing and reading. Generative AI, like hypertext before it, breaks down the boundaries of authorship and textuality, making meaning more fluid, collaborative, and adaptive. I identify four central challenges that arise in this new era: the Thief of Reason (outsourcing of human thought), the Sea of Sludge (proliferation of generic, homogenous content), the Dialogic Web (the dissolution of documents, unsettling authority and intent), and the Death of the Reader (loss of a stable, interpretive audience). Drawing on the legacy of hypertext, I propose that its core principles – non-linearity, transparency, navigability, and reader agency – offer vital conceptual and technical resources for responding to these challenges; but also that success will depend on our ability to interpret these hypertextual values for the AI world and will require active stewardship by both designers and users. As text is becoming fully fluid, collaborative, and adaptive; our greatest task is to ensure that writing remains a space for intentional dialogue, creativity, and human understanding even in the age of intelligent systems.

Keywords

Hypertext, Artificial Intelligence, Large Language Models, Human In The Loop

1. Introduction

The way we write has always been shaped by our tools, and each major technological shift – from scroll to codex, printing press to hypertext – has redefined not just how we write, but what writing is [1].

The scroll facilitated longer, continuous texts, but posed challenges for navigation and search. Texts were experienced communally and recited sequentially, reinforcing a more linear, performative model of authorship. The codex, in contrast, introduced discrete pages, enabling non-linear navigation, referencing, and private study—inviting a more individual and analytical engagement with the written word. The printing press expanded access to knowledge, standardized texts, and gave rise to the modern concept of the author. For the first time, texts became fixed. There was a “definitive edition” – a single, authoritative version of what the author meant. This fixity underpinned intellectual property laws and shaped the foundations of authorship, copyright, and scholarly citation.

Later came the typewriter and word processor, which democratized the mechanics of writing and editing. Writers could now move text, revise freely, and experiment with layout, all of which blurred the once-clear line between drafting and publishing. Writing became faster, more collaborative, and increasingly open to iteration. But it was hypertext that posed the most radical challenge yet to traditional textual paradigms.

Hypertext was born from a vision of augmentation – a digital extension of Engelbart’s foundational idea that computers could amplify human intellect [2, 3]. Hypertext shattered the illusion of the singular, linear text. It introduced non-linearity, reader agency, fluidity, and multiplicity—features that are now central to how we engage with information in the 21st century. In hypertext environments, meaning is constructed not just by the writer, but also by the reader, who chooses their own path through a network of fragments and connections. The text becomes dynamic, customizable, and co-authored in real time.

Narrative and Hypertext Workshop 2025

✉ dem@soton.ac.uk (D. E. Millard)

🌐 <https://davidmillard.org/> (D. E. Millard)

🆔 0000-0002-7512-2710 (D. E. Millard)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

This earlier shift offers crucial lessons for understanding the present turn toward AI. Hypertext forced us to confront new questions: Who controls meaning? Where does the boundary of a text lie? What does authorship mean when readers can restitch the story with every click?

It also introduced concepts that feel strikingly relevant to the world of generative AI, and in particular Large Language Models (LLMs) (as initiated by [4]). The idea of multiple approved versions of a text, customized for different readers, foreshadows today's personalized outputs from such models. Similarly, reader-driven navigation in hypertext is echoed in user-driven prompting, where meaning is not merely conveyed, but collaboratively created. We now live in a world where hypertext is everywhere: in the way we browse, scroll, post, stream, and read. The core attributes of that transformation – non-linearity, reader agency, and fluidity – set the terms for how we have engaged with digital culture. As we now face the rise of AI-mediated writing, hypertext offers both a precedent and a provocation. It teaches us that writing is never just the arrangement of words, but the expression of a system – and when the system changes, so too does the meaning of authorship itself.

Over time, the paradigm of hypertext has expanded dramatically [5]. What began as a “tool for thought” [6] has grown to encompass interface design, digital literature, and even a world view: a way of understanding and constructing systems that are fundamentally human-centred and responsive [7]. This epistemological stance frames hypertext not merely as a technical artifact, but as a philosophy of system design that privileges human needs, agency, and meaning-making. The ubiquity of hypertext means that how it responds to the rise of AI is not a niche concern, but a question of fundamental importance for the future of digital knowledge, creativity, and human agency.

We have a tendency to think of technological progress as gradual change, demonstrating an inability to understand the importance of paradigm shifts that lead to technological discontinuities [8]. The speculative Sci-fi of the 50s was all jet engines and atomics, because no-one foresaw the silicon chip.

In this paper I want to consider AI and writing, but not as an incremental change, instead I think it would be useful to speculate about the radically different future that seems to be coming.

We will need to set aside some of our technical baggage, but Hypertext may yet save us all.

2. The Thief of Reason

To write is to think [9]. Allowing AI to write for us is therefore, fundamentally, an act of outsourcing our thinking. This means that AI's encroachment into the writing process challenges not only our practices but the very cognitive labour at the heart of our intellectual life. This challenge is often met by optimists with an argument for augmentation – that we should build AI that assists rather than replaces us as writers. Yet the reality is more nuanced. There are different ways to engage with AI as an augmenting force.

AI's most immediate impact within the hypertextual landscape is its capacity to serve as an intelligent companion or guide. Rather than simply presenting static information, AI can offer clarifications, generate summaries, and highlight relevance cues that are dynamically tailored to each user's context, interests, and needs [10]. Writers working in this way, let's call them the **Poets**, engage with AI at arms length. Here, the human writer remains the primary creator, using AI as an assistant for grammar checks, phrasing suggestions, and ideation. Authorial control and deliberate word choice are central. The AI acts as a research assistant offering feedback, an editor proposing revisions, or a first reader highlighting areas for improvement. The authentic voice of the author remains strong.

The Poet may be the way that many of us imagine an augmented person will write in the future. But there is an entirely new way to work that is only possible when the generation of prose is fast and cheap. Writers working in this way, let's call them the **Potters**, mould their text like clay on a wheel, allowing the AI to generate copy after copy that they revise and iterate. The writer gives prompts, direction, and editorial vision, but the AI produces substantive chunks of text – drafting, revising, and reworking as part of a collaborative process. This role echoes the collective and distributed authorship found in film production, game design, or television writing, where the creative director shares ownership with others who are contributing much of the material. The critical difference is that this team now includes

silicon intelligence, yet our established notions of authorship and attribution need not fundamentally change. The author is still the author, the text a product of an augmented human.

These roles are not rigid categories; instead, they reflect a spectrum of practices, and writers may blend or alternate between them within a single project [11]. Crucially, neither role is inherently superior nor limited to a single context. The role of the Poet suits tasks requiring human empathy or authenticity, such as personal letters or reflective journals. While the role of the Potter is ideal for collaborative, functional, or rapid-turnaround texts like reports, manuals, or commercial copy.

Both keep alive the notion of writing as thinking, as a process, that seems so central to the philosophy of hypertext. Yet the sorts of skills needed to be an effective Potter are best developed by first experiencing life as a Poet [12]. Co-authorship requires not only trust in our partner, but confidence in ourselves [13]. In educational terms, much like teaching children handwriting and Shakespeare – even though most will never write a play with a fountain pen – we must teach ourselves how to be Poets in order to develop the critical skills we need to be good Potters. Lest the tools of augmentation become a mechanism for the attrition of our minds.

3. The Sea of Sludge

Pure AI writing, a third mode, emerges when we hand full control to the machine – acting not as Poet, nor as Potter, but as absent witness. This mode will not vanish, even if our educational systems get everything right. The simplicity of production ensures that, in the future, the majority of published text will be composed by AI, with only a minority produced by augmented humans and a vanishingly small quantity by unassisted writers.

Three consequences follow. First, AI's drive toward consensus, coherence, and generality risks flattening the vibrancy of human expression [14]. The danger is not merely technical but cultural: a slow drift toward uniformity, where the distinctive voices and idiosyncratic perspectives that animate our literature and scholarship are eroded by a tide of algorithmic sameness [15]. This sea of sludge – prose that is grammatically correct but devoid of passion or originality – threatens not only literary creativity but the diversity of thought that hypertext, at its best, was designed to nurture.

Second, the texts necessary to train future machines will become ever rarer. If AI models are fed on their own dead, they risk inbreeding and intellectual sterility. Worse for us, relying solely on the past risks freezing the cultural canon. Hamlet may be on the holodeck, but it's telling that in the 24th century there are no new famous playwrights.

Third, text will lose its role as a visible marker of understanding and a foundation stone for trust. In many domains – such as education, journalism, and research – writing has long served as a proxy for deeper competence. The arrival of AI unsettles this tradition, enabling surface-level fluency without genuine understanding, challenging long-held assumptions about the relationship between writing and expertise [16].

While this will be challenging it also means that the sort of textual make-work that we commonly encounter in our lives will lose its value. If text has no communicative purpose other than to signal our competence, then in the future it will have no purpose at all. The grounds for credibility must shift. The need for writing as a filter evaporates when anyone can produce superficially flawless text, pressing us to invent new, more meaningful markers of expertise and sincerity.

4. The Dialogic Web

Early hypertext systems were designed within a document-centric paradigm, where knowledge was stored in discrete, authored units such as books, memos, or scholarly articles. These systems reflected a linear and static conception of information, emphasizing fixed boundaries and curated content. Hypertext began to break these boundaries, creating an explicit inter-textuality with profound impacts on meaning-making [17] and knowledge representation [18]. However, AI disrupts this model further

by enabling just-in-time synthesis, generating responses that draw seamlessly from multiple sources and often dissolving the boundaries between original documents completely.

Hypertext was once celebrated as the realization of Barthes' notion of the "Death of the Author" [19] enabling text to be co-created by networks of nodes and links, by the author's design, and by the reader's navigation. Even so, the reader's experience remained curated, their epiphanies planned, and agency limited [20]. In the world of AI, the line moves further still. Readers now possess real power: the capacity to reshape, synthesize, and integrate ideas from myriad sources, including their own notes and creations. In this emerging dialogic web, meaning is not merely transmitted but co-constructed in real time through the interplay between human and machine [21].

We thus move from a world of contiguous, fixed texts to one of continuous text, always in motion – personalized, permanently evolving. Writing and reading as posthuman cognitive assemblage [22]. AI may let us realize the dream of the adaptive, branching garden of forking paths [23], but as we approach that infinity, the very notion of boundaries must dissolve. The far future belongs not to static documents, but to a living, endlessly changing conversation with the Permascroll.

5. The Death of the Reader

This leads to a final conundrum. If text is endlessly malleable and lacks fixity, for whom do we write? Writing has always involved an implied reader, an imagined participant whose knowledge and sensitivities shape our choices of metaphor, explanation, and style. This process serves both to clarify and also to gatekeep, demarcating communities and excluding outsiders.

AI upends this practice. Documents lose their sacredness, becoming endlessly deconstructed and re-combinable. No longer do we write for a single reader; we write for all readers – or, perhaps, for the machine itself. From SEO-driven content to writing optimized for algorithmic sorting and retrieval, we already live in a world where much of what's produced is engineered to be parsed and processed by artificial readers [24].

In such a world, the traditional reader is not the end recipient of a crafted message, but a secondary participant. Readers increasingly rely on AI tools to reshape texts according to their personal preferences – summarising, translating, adjusting voice or tone, or re-framing content to meet different levels of expertise. What was once fixed becomes flexible. The notion of a stable, singular audience is eroding, replaced by reader-driven customization, where the same text can appear differently to different readers based on how their AI agents remix it.

This shift raises essential questions: If no one reads the same text in the same way, is there still one text? If authorship is distributed and readership is mediated, where does meaning happen, and who is responsible for it?

While this could usher in a golden age of accessibility and personalization, it also carries risks. Writing designed for machines may prioritize structure, clarity, and simplicity over subtlety, empathy, and nuance. The richness of human-centred expression – the kind that comes from imagining another's perspective – may erode in favour of machine-legible prose. It is a disincentive to write as Poets, the core role we need in order to develop our skills.

Still, all is not lost. The crafted, expressive, human voice need not disappear. In fact, it may become more valuable. In a world flooded with machine-shaped content, the distinctly human touch may command a cultural premium, much as handmade goods do in an industrialized economy [25].

In this new paradigm, the roles of both author and reader are being remade. Authorship is increasingly curatorial and directional rather than expressive. Reading is becoming interactive, mediated, and transformative, rather than passive or fixed. Instead of lamenting these deaths, perhaps we should see them as opportunities to rethink what writing means, who it's for, and how meaning is made in an age of intelligent systems.

6. Conclusion: What Stars to Steer By?

The future I have outlined here may be alarming. It is an undiscovered country where documents dissolve, intellectual skills wither, and our shared canon deteriorates. ‘You can and must understand computers now!’ Nelson ordered fifty years ago – and perhaps we have failed to heed the full implications of that warning; failed to understand the machines before they started to understand us.

Others have pointed out the ethical role that hypertext might have to play, arguing that hypertext provides an architecture that supports transparency, intentional navigation, and pluralistic sense-making [26]. Within this model, rather than surrendering judgment to automated outputs, users are invited into an interactive process where paths, provenance, and rationale are traceable. Hypertext’s non-linearity helps counteract the flood of undifferentiated AI text, giving readers meaningful tools for discernment and slowing the slide into homogeneity. These systems also create space to re-establish credibility through transparency, and they offer new forms of dialogic engagement by empowering both authors and readers to participate actively in knowledge creation, rather than being swept along by the momentum of generic, automated prose.

All this is true, and we want to believe. However, realizing this ethical potential presents significant challenges. Hypertext systems, especially in their most advanced or open-ended forms, may become too complex or demanding for many users, risking further exclusion if digital literacy does not keep pace. Designing transparent, interactive platforms requires thoughtful ethical and technical decisions that may conflict with the commercial imperatives driving rapid AI adoption. There is no guarantee that readers will consistently choose active, critical engagement over convenient, surface-level consumption – particularly when speed and simplicity are prized. Moreover, the effectiveness of these solutions depends not just on technical design, but on robust institutional, educational, and regulatory support – areas that often lag behind innovation. Most fundamentally, hypertext remains a tool whose value depends on the wider culture in which it is embedded; without a renewed commitment to responsibility and intentionality, even the most carefully constructed systems will not help us escape the shadow of the machine.

The challenge we face is twofold: to safeguard the uniquely human aspects of writing, and to evolve alongside the technological transformation rapidly unfolding around us.

It remains to be seen whether this era of AI capability marks a revolutionary inflection point or is merely the beginning of a longer but no less profound change. Regardless, the conception of text as a stable, authored artifact is giving way to a fluid, generative, and collaborative landscape. We must all learn anew what it means to be writers within this world.

Even as our tools evolve, we remain connected to a longstanding human impulse: the desire to reach across distance and difference to communicate meaning. The text that once served as a direct bridge between two minds now passes through algorithms and collaborative processes. While this changes the nature of that connection, it does not erase it. If we approach these new technologies thoughtfully, and bring to bear our hypertextual tools, we can continue to preserve the essential human moment in writing: the effort to share understanding, perspective, and insight. In this evolving landscape, our challenge is not to mourn what changes, but to shape how those connections are made, ensuring that writing remains a meaningful act of communication.

Yet the responsibility does not rest solely with designers and developers. Users, too, must approach these systems with discernment and care. Making the most of AI-augmented hypertext requires education, critical awareness, and a willingness to cultivate new habits of engagement. It demands discipline: resisting the allure of effortless automation in favour of practices that foster understanding, originality, and agency.

In this landscape, both human and artificial intelligences will have a role – not as passive wanderers, but as active stewards of the garden of forking paths. The central challenge is to design and to use AI-hypertext systems that amplify human individuality, preserve agency, and foster trust; systems that are not merely intelligent, but wise, and that remain fundamentally in service of the human mind.

References

- [1] K. Houston, *The book: a cover-to-cover exploration of the most powerful object of our time*, WW Norton & Company, 2016.
- [2] V. Bush, *As we may think*, *The Atlantic Monthly* 176 (1945) 101–108. Place: Boston, MA, USA.
- [3] D. C. Engelbart, *Augmenting human intellect: A conceptual framework*, SRI Summary Report AFOSR-3223 1 (1962) 2007. Place: Menlo Park, California, USA.
- [4] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, I. Polosukhin, *Attention is All you Need*, in: *Advances in Neural Information Processing Systems*, volume 30, Curran Associates, Inc., 2017.
- [5] M. W. R. Anderson, D. E. Millard, *Seven Hypertexts*, in: *Proceedings of the 34th ACM Conference on Hypertext and Social Media, HT '23*, Association for Computing Machinery, New York, NY, USA, 2023, pp. 1–15. doi:10.1145/3603163.3609048.
- [6] H. Rheingold, *Tools for Thought: The History and Future of Mind-Expanding Technology*, MIT Press, Cambridge, Mass. London, 2000.
- [7] C. Atzenbeck, P. J. Nürnberg, *Hypertext as method*, in: *Proceedings of the 30th ACM conference on hypertext and social media*, ACM, New York, NY, 2019, pp. 29–38. doi:10.1145/3342220.3343669.
- [8] P. Anderson, M. L. Tushman, *Technological Discontinuities and Dominant Designs: A Cyclical Model of Technological Change*, in: *Organizational Innovation*, Routledge, 1998. Num Pages: 30.
- [9] L. Flower, J. R. Hayes, *A Cognitive Process Theory of Writing*, *College Composition & Communication* 32 (1981) 365–387. URL: <https://publicationsncte.org/content/journals/10.58680/ccc198115885>. doi:10.58680/ccc198115885, publisher: NCTE.
- [10] R. Vansh, D. Rank, S. Dasgupta, T. Chakraborty, *Accuracy is not enough: Evaluating Personalization in Summarizers*, in: H. Bouamor, J. Pino, K. Bali (Eds.), *Findings of the Association for Computational Linguistics: EMNLP 2023*, Association for Computational Linguistics, Singapore, 2023, pp. 2582–2595. doi:10.18653/v1/2023.findings-emnlp.169.
- [11] J. Hutson, *Human-AI Collaboration in Writing: A Multidimensional Framework for Creative and Intellectual Authorship*, *International Journal of Changes in Education* (2025). URL: <https://digitalcommons.lindenwood.edu/faculty-research-papers/720>. doi:<https://doi.org/10.47852/bonviewIJCE52024908>.
- [12] N. F. Bauer, *Does ChatGPT Increase Language Homogenization?*, in: C. Vaih-Baur, V. Mathauer, E.-I. von Gamm, D. Pietzcker (Eds.), *KI in Medien, Kommunikation und Marketing: Wirtschaftliche, gesellschaftliche und rechtliche Perspektiven*, Springer Fachmedien, Wiesbaden, 2025, pp. 11–31. URL: https://doi.org/10.1007/978-3-658-46344-1_2. doi:10.1007/978-3-658-46344-1_2.
- [13] H. Ryan, D. Abramov, S. Acker, S. Elkins, *Can AI Be a Co-Author?: How Generative AI Challenges the Boundaries of Authorship in a General Education Writing Class*, *Thresholds in Education* 48 (2025) 40–56. URL: <https://eric.ed.gov/?id=EJ1468037>, publisher: Academy for Educational Studies ERIC Number: EJ1468037.
- [14] A. R. Doshi, O. P. Hauser, *Generative AI enhances individual creativity but reduces the collective diversity of novel content*, *Science Advances* 10 (2024) eadn5290. URL: <https://www.science.org/doi/10.1126/sciadv.adn5290>. doi:10.1126/sciadv.adn5290, publisher: American Association for the Advancement of Science.
- [15] D. Agarwal, M. Naaman, A. Vashistha, *AI Suggestions Homogenize Writing Toward Western Styles and Diminish Cultural Nuances*, in: *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems, CHI '25*, Association for Computing Machinery, New York, NY, USA, 2025, pp. 1–21. URL: <https://dl.acm.org/doi/10.1145/3706598.3713564>. doi:10.1145/3706598.3713564.
- [16] M. Perkins, J. Roe, D. Postma, J. McGaughran, D. Hickerson, *Detection of GPT-4 Generated Text in Higher Education: Combining Academic Judgement and Software to Identify Generative AI Tool Misuse*, *Journal of Academic Ethics* 22 (2024) 89–113. URL: <https://doi.org/10.1007/s10805-023-09492-6>. doi:10.1007/s10805-023-09492-6.
- [17] G. P. Landow, *HyperText: The Convergence of Contemporary Critical Theory and Technology (Parallax: Re-visions of Culture and Society Series)*, Johns Hopkins University Press, USA, 1991.

- [18] M. Snaprud, H. Kaindl, Knowledge acquisition using hypertext, *Expert Systems with Applications* 5 (1992) 369–375. doi:10.1016/0957-4174(92)90020-S.
- [19] R. Barthes, S. Heath, *Death of the author*, in: *Image, music, text*, Fontana Press, London, UK, 1977, pp. 142–148.
- [20] S. Brooker, Man proposes, god disposes: Re-assessing correspondences in hypertext and anti-authorist literary theory, in: *Proceedings of the 30th ACM conference on hypertext and social media*, ACM, New York, NY, 2019, pp. 39–48. URL: <https://doi.org/10.1145/3342220.3343649>. doi:10.1145/3342220.3343649.
- [21] S. Brooker, Computer, Enhance! Augmentation, Ideation, Hypertext, in: *Proceedings of the 35th ACM Conference on Hypertext and Social Media, HT '24*, Association for Computing Machinery, New York, NY, USA, 2024, pp. 193–196. doi:10.1145/3648188.3678216.
- [22] K. H. N., *Writing//Posthuman: The Literary Text as Cognitive Assemblage*, *Theoretical Studies in Literature and Art* 38 (2018) 6–21. URL: <https://tsla.researchcommons.org/journal/vol38/iss3/20>.
- [23] J. L. Borges, *The garden of forking paths*, Penguin Classics, London, 1941. Pages: 17.
- [24] J. R. Gallagher, Writing for Algorithmic Audiences, *Computers and Composition* 45 (2017) 25–35. URL: <https://www.sciencedirect.com/science/article/pii/S8755461516300652>. doi:10.1016/j.compcom.2017.06.002.
- [25] M. Rubin, J. Z. Li, F. Zimmerman, D. C. Ong, A. Goldenberg, A. Perry, Comparing the value of perceived human versus AI-generated empathy, *Nature Human Behaviour* (2025) 1–15. URL: <https://www.nature.com/articles/s41562-025-02247-w>. doi:10.1038/s41562-025-02247-w, publisher: Nature Publishing Group.
- [26] C. Atzenbeck, Unwinding AI's Moral Maze: Hypertext's Ethical Potential, in: *Proceedings of the 35th ACM Conference on Hypertext and Social Media, HT '24*, Association for Computing Machinery, New York, NY, USA, 2024, pp. 23–28. doi:10.1145/3648188.3678213.