

The Irish Journal of Technology Enhanced Learning

Volume 9 | Number 1

Article type: Position Paper

The Humanities in the Age of Generative AI

James O'Sullivan

University College Cork

james.osullivan@ucc.ie

<https://orcid.org/0000-0002-4214-9933>

<https://doi.org/10.22554/26yk4973>



This work is licensed under [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/)

Recommended Citation

O'Sullivan, J. (2026) The Humanities in the Age of Generative AI. *Irish Journal of Technology Enhanced Learning*, 9(1). <https://doi.org/10.22554/26yk4973>

Position Paper

The Humanities in the Age of Generative AI

James O'Sullivan*

University College Cork

Ireland

<https://orcid.org/0000-0002-4214-9933>

<https://doi.org/10.22554/26yk4973>

Keywords:

GenAI, Artificial Intelligence, Humanities, Arts, Large Language Models, LLMs, Generative AI

Introduction

The humanities are once again being declared obsolete. With the advent of generative artificial intelligence, particularly large language models, commentators have begun to ask whether humanistic skills are still needed at all. If machines can write fluently and critique persuasively, what remains for disciplines premised on interpretation, argument, and, most notably, the (often intentionally) slow cultivation of insight?

These questions are not new, and the humanities have faced successive waves of scepticism over the past century: the utilitarianism of post-war science policy, the expansion of vocational higher education in the 1980s and 1990s, the rise of data-driven research paradigms associated with the digital turn, and most recently, the claim by Palantir CEO Alex Karp that AI will ‘destroy’ humanities jobs (Munis 2026). Each time, the question has been posed in similar terms: what use are the humanities in a world preoccupied with speed and measurable (typically economic) outcomes?

The emergence of generative AI represents something more than the latest ideological pressure. Since the post-war era, pressures on the humanities have largely been a function of institutional incentives that advantage STEM disciplines (Small 2013), but the threats posed by generative AI feel different, and they are different, because they fundamentally shift the structures of knowledge production. Some subfields, ethics and philosophy of mind among them, have clearly gained some ground as AI has become a matter of public concern. The humanities are not uniformly in retreat, but the cumulative effect of these structural pressures, together with the capabilities of generative AI, produces a threat to humanistic inquiry that exceeds what any single one of them would produce on its own. Earlier pressures operated largely on the humanities’ institutional footprint, on resourcing and prioritisation, while leaving the underlying intellectual activities intact. Large language models operate on

* Address for corresponding author james.osullivan@ucc.ie

The Irish Journal of Technology Enhanced Learning Ireland is the journal of the Irish Learning Technology Association, an Irish-based professional and scholarly society and membership organization. (CRO# 520231) <http://www.ilta.ie/>. This is an Open Access article distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0>).

writing, the medium through which that activity is conducted, which is why it warrants a different kind of attention even from scholars who have lived through several previous declarations of crisis. Large language models significantly reshape the processes by which textual meaning is generated and legitimised, and their implications are epistemic as much as pedagogical or institutional.

The challenge, then, is not how the humanities can survive generative AI and large language models, but how they can respond on their own terms, by interrogating the assumptions and politics of these technologies. Generative AI is a threat to humanistic authority, but the response to that threat lies in re-articulating that very authority in dialogue with the technical systems that increasingly mediate our intellectual life.

Humanities graduates and humanities-trained researchers already contribute to model development, policy formation, and applied ethics work inside technology companies, and that contribution is neither denied nor displaced by what follows; the question addressed here is the narrower one of how universities, as the institutions that still train most humanists and still produce most humanistic research, should think about teaching and learning, research, and governance under conditions in which generative AI has become infrastructural.

Precedents & Parallels

Before turning to those precedents, it is worth acknowledging that the disciplinary categories through which the argument proceeds are themselves historical artefacts, shaped by nineteenth- and twentieth-century settlements about what counted as a legitimate object of university study, and maintained by a set of institutional and professional arrangements that could in principle have developed otherwise (Abbott 2001; Graff 1987). Nothing in what follows rests on the claim that ‘the humanities’ possess an inherent or unchanging authority; the claim is that the specific intellectual traditions gathered under that name have built up, over several centuries, methods for interpreting texts and contexts that remain useful for reasoning about the systems now reshaping those same texts and contexts, and that starting from those traditions is a more defensible epistemological move than beginning elsewhere and importing humanistic concerns afterwards.

It is tempting to treat generative artificial intelligence as an unprecedented rupture in the history of knowledge, because the speed, scale, and apparent autonomy of large language models can create the illusion that we are confronting something wholly novel, an alien logic that stands apart from the humanistic traditions of meaning-making. But a longer historical view suggests otherwise, and the humanities have always evolved in relation to the technologies of their time, often through periods of disruption that later came to seem foundational.

The invention of the printing press in the fifteenth century is perhaps the most obvious precedent. The shift from script to movable type transformed the structures of intellectual authority itself, standardising editions and reproducing arguments through a material change in how words were stored and shared. The humanist scholars of the Renaissance did not retreat from this shift, but exploited it, establishing new commentarial practices and textual norms in the process. Erasmus and his contemporaries worked directly with printers such as Aldus Manutius to produce critical editions of classical texts whose apparatus and systems of cross-reference gradually became naturalised as the scholarly form itself, while the standardisation of spelling, the emergence of the title page, and the development of the

alphabetical index each began as printerly innovations before settling into the taken-for-granted infrastructure of humanistic work (Eisenstein 1980; Grafton 2011).

Similar tensions re-emerged with the rise of broadcast technologies in the twentieth century, particularly with radio and television, which disrupted older hierarchies of literary culture, creating new publics and new anxieties. The mass reach of media prompted fierce debates within the humanities: what counted as legitimate cultural production? How should critics respond to formats designed for mass consumption? The emergence of cultural studies and media theory, associated with figures such as Hoggart (1957), Williams (1958, 2003), Hall (1980), and McLuhan (1964), can be traced to this moment of disciplinary adaptation.

More recently, Massive Open Online Courses—the fabled ‘MOOCs’—of the early 2010s were briefly heralded as a transformative threat to traditional universities, but neither the connectivist cMOOCs nor the later xMOOC offered by platforms like EdX and Coursera replaced the lecture theatre, but they did force disciplines to articulate their pedagogical purposes for audiences that could no longer be presumed to be seated in a classroom. The analogy to generative AI is imperfect, since MOOCs just extended existing content to new audiences rather than producing new content, but the episode is a useful reminder that predictions of disciplinary obsolescence have a poor track record, and that the teaching practices that emerged from the response often outlasted the technology that prompted them. The xMOOC variety in particular forced educators to reconsider the question of audience, and although massification remains problematic as an end in itself, working in potentially massified environments pressed disciplines to articulate their purposes more clearly. There are almost certainly cases, however limited, wherein MOOCs opened the door to creative pedagogies that reimaged humanistic teaching for online environments.

The humanities responded to online delivery by revising their methods and their objects of study, but they did so when the technology in question left the basic activity of reading and writing in human hands. Generative AI is the first case in which the technology can itself produce the kind of text that humanistic inquiry has historically treated as its evidence, and the precedents are therefore limited in what they can settle. What they do establish is that the disciplines most visibly under pressure have, historically, shown a capacity for methodological agility and reflective critique that the rhetoric of obsolescence tends to underestimate.

Computational Humanities in Practice

Despite persistent caricatures of the humanities as resistant to technological change, scholars across fields such as literary studies, history, media studies, and philosophy have long engaged with computational tools as provocations to theory and method. The current wave of generative AI may feel disruptive, but it follows decades of work in the digital humanities that anticipated many of today’s epistemological and ethical debates.

That work, of course, sits alongside and often in conversation with adjacent literatures that have addressed human relationships with computing for more than three decades. Science and technology studies has examined the co-production of technical systems and social arrangements (Jasanoff 2004; Latour 1987), human-computer interaction has pursued human-centred design as an explicit counterweight to purely technical framings (Bannon 1991), digital sociology and the sociology of algorithms have analysed the organisational and cultural consequences of computational infrastructures (Lupton 2014; Beer 2017), and research in educational technology has built a long record of critique concerning the

assumptions embedded in learning platforms (Selwyn 2010, 2022). The present argument is not offered in ignorance of these traditions, and in many places it depends on them; the argument is that the work of reasoning about what large language models do to textual authority is best begun from within the disciplinary cultures that have spent longest attending to those particular questions.

Stylometry, for instance, once a niche technique used to detect disputed authorship through word frequency patterns, has matured into a rigorous and widely deployed method. Corpus-based studies have demonstrated how statistical modelling can augment, rather than displace, traditional forms of literary interpretation. These methods rely far less on algorithmic magic than on careful human calibration, the selection of most frequent words, the treatment of function versus content vocabulary, the contextual framing of results. Computational tools become a means of surfacing patterns that without the accompanying interpretive judgement would prove quite useless.

Image and object-based disciplines have likewise begun to integrate AI techniques for classification and comparative analysis, and computer vision techniques can assist with recognising recurring visual motifs, architectural structures or stylistic elements across large corpora of visual culture. But these outputs still require the expertise of art historians and theorists to become meaningful—the pattern alone does not constitute knowledge.

What these examples demonstrate is that generative AI in the humanities, like the computational tools that preceded it, becomes productive only when used as an instrument under sustained interpretive direction. The now-common rhetoric of ‘partnering’ with AI should itself be approached with some suspicion. We did not describe our relationships with calculators or virtual learning environments in those terms, and the vocabulary of collaboration risks obscuring the quite ordinary question of who is accountable for the work that a tool helps to produce. While commercial discourse often presents AI as a force of automation, the humanities have approached it as a catalyst for methodological reflection, raising questions about what counts as evidence and interpretation in a hybrid human-machine research environment. Segessenmann et al. (2025) have recently argued for what they call a humanities ‘work program’ for deep learning that goes beyond the narrow remit of AI ethics to address the way such systems reshape human self-understanding and culture, and their distinction between the technical mastery evident in these systems and the interpretive mystery they continue to pose is a useful corrective to the tendency to treat fluency as evidence of understanding.

Generative AI does, however, sit in an unfamiliar position relative to these earlier methods. Stylometry and computer vision analyse material that already exists and defer to the interpreter for meaning, whereas large language models generate material that did not previously exist, in forms designed to be read as though it had been written by a human with intent. The challenges particular to this posture—the black-box nature of the underlying models, their tendency to fabricate sources, and the difficulty of tracing authorship in co-produced texts—cannot be met simply by extending the critical protocols developed for earlier computational methods, though those protocols remain a useful starting point. But what the digital humanities experience does establish is that the incorporation of new technologies into scholarly workflows need not be the end of humanistic inquiry, provided that the terms of incorporation are worked out by scholars rather than imposed on them.

Teaching & Learning with LLMs

Generative AI has entered university classrooms through several channels at once. Students have adopted tools such as ChatGPT, Claude, and Google Gemini—openly or covertly—to summarise readings and draft essays across disciplines. At the same time, university administrations have promoted adoption under the banner of innovation and competitive positioning, while vendors have incorporated generative features into the institutional software stack itself, so that products such as Microsoft Copilot and the AI features embedded in Google Workspace and learning management systems are now encountered by staff and students whether or not they sought them out. This uptake has been uneven, often unregulated, pedagogically ambivalent, and rarely subjected to critical evaluation at the point of use. When a feature is pushed into familiar software, users tend to trust by proxy on the assumption that a reputable vendor would not ship an unreliable tool, even though reputational and commercial incentives offer limited protection against precisely that outcome. But it has made clear that teaching and assessment in the humanities cannot proceed as before.

Initial institutional responses have largely focused on academic integrity. Updated plagiarism policies, position papers from public intellectuals and agencies, and hastily drafted faculty guidelines have framed generative AI as a threat to authorship and meaningful assessment. The shape of the discourse has since shifted considerably, moving from an earlier emphasis on prohibition during 2022 to 2024 toward a rival framing in which non-adoption is recast as a kind of educational malpractice, with many staff finding themselves having occupied both positions within the space of two academic years. While necessary, such documents often rely on inherited assumptions about what constitutes ‘cheating’ in an environment where text can be algorithmically generated on demand. The traditional essay, long treated as a proxy for understanding, is no longer unproblematically secure, and it is worth recalling that its security was always partial, since paper mills and other forms of commissioned ghostwriting had for decades offered human-driven means of circumventing the same assessment form.

But these developments, as unwelcome as they may be, create an opportunity to revisit foundational questions. What do we want students to learn in a humanities degree? How should that learning be demonstrated, and by whom? If essays can be convincingly written by machines, perhaps our attention should shift to the quality of human feedback, the process of iterative drafting, and the ethical framing of inquiry. Rather than retreating into surveillance, some educators have begun to design assessments that explicitly incorporate AI tools, asking students to critique or improve upon generated texts, evaluate source fabrication, or reflect metacognitively on their own writing process. Such assessments are considerably more resource-intensive than the essays they displace and therefore sit uneasily with the massified conditions under which much humanities teaching now takes place, which is part of why the present moment is also an opportunity for the profession to articulate what teaching is for, what humanising practices it depends on, and how these connect to the broader question of what it means to hold a degree in a humanistic discipline.

There are also deeper curricular implications, as the humanities can no longer treat digital literacy as the preserve of STEM disciplines. If students are to use large language models responsibly—though it is sometimes difficult to imagine the responsible use of inherently irresponsible tools—then critical AI literacy must become a core competency of higher education, and curricular responses such as the ‘AI digital humanities’ model advanced by Chun and Elkins (2023) offer one concrete template for how such literacy might be

institutionalised rather than bolted onto existing degrees as an afterthought. This includes technical fluency as well as the ability to interrogate bias, representation, and the social implications of algorithmic systems. In this respect, the humanities are not just reactive but essential, as they offer precisely the interpretive frameworks and ethical vocabularies needed to situate AI in its broader cultural and political contexts.

Generative AI is not the end of teaching in the humanities; it is a moment that forces clarity about what good teaching involves and presses the profession to design learning experiences that foreground reasoning. Prohibition at the classroom level is unlikely to hold given how deeply these tools are already embedded in the software environments that students and staff use by default. The reality is that we cannot ban this machine, but we can seek to humanise its use—no small task, admittedly, but one that the humanities have risen to before. The demanding question is whether universities can develop pedagogical practices equal to a tool that produces fluent text on demand, and the answer to that question will not be settled by policy language alone.

The Epistemic & Ethical Challenges of LLMs

If the use of AI in the humanities raises urgent pedagogical and institutional questions, it also poses deeper challenges to the epistemic foundations of humanistic inquiry. These are not only matters of academic practice, but of knowledge itself. What counts as an author? What constitutes interpretation? How do we distinguish between the production and recognition of meaning in a system where machines can mimic fluency without understanding?

Large language models function by predicting word sequences based on statistical patterns derived from vast datasets. They are, by design, indifferent to truth, genre, or rhetorical intention. As such, they produce a kind of synthetic language that bears the surface features of coherence while often lacking referential integrity or argumentative depth. This epistemic shallowness is particularly visible in academic contexts, where ungrounded claims (fabricated sources not least among them) and a lack of disciplinary nuance can undermine the credibility of generated outputs. But the illusion of competence is often compelling, especially for students unfamiliar with the deeper structures of scholarly reasoning.

This raises acute ethical issues, perhaps most notably around authorship and credit. When students or researchers incorporate AI-generated text into their work, questions arise about intellectual ownership and the boundary between authorship and automation. Traditional norms around plagiarism and citation are strained by tools that contribute form without necessarily supplying content, or that interpolate user intent in ways that defy clear attribution. The distinction between human and machine input becomes blurred, but many institutional policies still presume a binary.

Another issue is bias and representation, as large language models inherit the assumptions and exclusions embedded in their training data, much of which is scraped indiscriminately from the internet. This raises deep concerns about the reinforcement of dominant perspectives. As scholars like Safiya Noble have shown, algorithmic outputs often encode structural biases under the guise of neutrality, replicating colonial, racialised, and gendered logics while claiming universality (2018).

Humanistic disciplines are built on the premise that meaning is contextual, contested, and historically situated. Scale is part of the problem here, not incidental to it. Berry and Stockman (2024), drawing on Schumacher's critique of industrial gigantism, describe the

current configuration of large-model development as productive of pathologies of meaning in which technical capacity outruns the human interpretive frameworks that would otherwise hold it accountable. Large language models, by contrast, tend to universalise and flatten, offering plausible-sounding summaries and generalisations that obscure complexity and dissent—they excel at simulation, not interpretation. The humanities must therefore guard against the creeping displacement of judgment by statistical association, what Byung-Chul Han has called the ‘expulsion of the other’ (2018), in algorithmic culture.

The humanities are uniquely positioned to expose and interrogate the logics underpinning AI systems, to ask how and why these models produce what they do, and with what consequences. In doing so, they move from being users of technology to shapers of its ethical and epistemological horizons, reclaiming, as Johanna Drucker once advocated (2012), cultural authority from engineers.

Institutional & Policy Implications

While much of the conversation around generative AI in the humanities has focused on classroom practice and intellectual integrity, its institutional consequences are no less significant. The integration of AI into research, teaching, and administration compels universities to revisit long-standing assumptions about academic labour, quality assurance, and the allocation of resources, particularly in a sector already under financial and structural strain.

One of the most immediate impacts concerns workload and labour models, and as AI tools become embedded in curriculum design, questions emerge about what constitutes academic labour. If staff are expected to evaluate AI-generated student work for accuracy or originality, does this introduce new cognitive and ethical burdens that are currently unacknowledged in institutional planning and resource allocation?

There are significant policy-level concerns about how the value of humanities research is assessed in an AI-saturated landscape. The UK’s upcoming REF 2029 exercise, for example, includes increased emphasis on digital outputs and open-access dissemination, but provides limited guidance on how hybrid or AI-assisted research artefacts should be evaluated. If a scholar uses a large language model to generate large-scale comparative textual analysis or to prototype new critical frameworks, how should that contribution be credited, and under what criteria? The question of credit is only the first of several, since even standard disclosures of the form familiar from the sciences (statistical analyses completed with SPSS version x.x) leave open the harder questions of how such work is peer reviewed, how its claims are verified, and whether outputs produced against a proprietary model at a particular moment can be meaningfully replicated by another scholar at another time. The humanities will need to define these standards themselves, or risk having them imposed externally by metrics better suited to the sciences.

At the funding level, there are signs of recognition, but also of constraint, and the UKRI and Horizon Europe have both issued calls that encourage interdisciplinary AI research, with an explicit role for ethics and public engagement. But these opportunities often presume that humanities scholars serve as adjuncts to technical teams, brought in after development to provide ‘reflection’ or ‘context’. A genuinely collaborative model would embed humanistic expertise at every stage of the research pipeline, from framing questions to designing interfaces and evaluating impact.

AI challenges the governance structures of universities. The procurement of proprietary platforms, the outsourcing of assessment tools, and the dependence on opaque algorithmic systems all raise questions about institutional autonomy, data ethics, and accountability. In the absence of strong internal AI governance frameworks, universities risk aligning themselves uncritically with corporate logics, accepting systems that may not reflect educational values or equity commitments. Liu (2025) frames the broader institutional shift underway as a move from a liberal arts centred on contemplative forms of knowledge to a ‘post-liberal arts’ university organised around applied knowledge, and the humanities’ position within that reconfiguration depends heavily on whether humanists are present when infrastructural and curricular decisions are made or are consulted afterwards to supply legitimating language.

This is where the humanities have an important role to play, in helping to shape the institutional policies that govern its use, contributing to university-wide AI task forces, drafting ethical guidelines, and participating in procurement decisions. It also means ensuring that critical digital literacy is not siloed in computer science departments or, even when broadly offered, too technical in its focus.

Responding effectively to the AI turn will require the humanities to assert their relevance at the level of policy, governance, and infrastructural design (see Underwood 2018, 2025). Otherwise, agency will continue to shift away from academic communities—from humans.

Reimagining Humanistic Authority in the Age of Gen AI

Whether the humanities are diminished or revitalised by generative AI will depend on how they respond to a rapidly shifting epistemic and institutional landscape. What is at stake is the very authority the humanities claim in shaping knowledge, as well as public values and discourse.

Three broad trajectories are possible, though in practice any given institution is likely to exhibit features of all three at once, and the question is one of emphasis rather than exclusive choice.

In the first, the humanities are excluded from meaningful participation in AI development, with technical design outsourced to industry and ethical oversight remaining nominal as humanistic expertise is consulted only post hoc. This is the path of marginalisation, in which the humanities become increasingly irrelevant to institutional decision-making and public policy.

In the second, the humanities survive by translating their purposes into the language of utility, leaning heavily on the lingo of ‘digital upskilling’, ‘job readiness’, and ‘transferable skills’. This instrumentalisation may secure short-term investment but risks eroding the distinctive methods and critical orientations that define the humanities, a scenario in which AI is accommodated more than it is interrogated.

A third, somewhat more hopeful, trajectory, remains open, one of co-production and critical engagement. This requires reasserting the humanities as guardians of a fading cultural past, but more importantly, as active participants in the construction of technological futures, which means recognising that questions of bias, representation, authority, and meaning are not peripheral to AI. This means embedding the humanities into the core of how universities think about AI, teach with it, and how policymakers govern it.

To realise this future, concrete steps must follow: humanities departments should integrate AI literacy as a field of inquiry in their own right, universities should support interdisciplinary fellowships that embed humanists in data science environments, and vice versa, and ethics boards with real authority should guide procurement and implementation decisions. Assessments should shift from output to process, valuing reasoning and intellectual ownership over surface fluency. And scholars should be supported to intervene publicly in debates around generative AI, particularly in relation to governance.

These reforms ask universities to take seriously the idea that large language models are cultural and political artefacts before they are technical ones, and that the humanities have a unique role in revealing how such systems function and what values they encode.

The question, then, is not whether the humanities can survive AI, but whether we will seize the chance to reshape the terms on which they operate. Humanistic authority in the algorithmic age will depend on the willingness to engage deeply, critically—and maybe even constructively—with the systems that increasingly mediate intellectual life, and to assert, through that engagement, the continuing necessity of the human in this inherently uncertain time of machines.

References

- Abbott, Andrew. (2001). *Chaos of Disciplines*. University of Chicago Press.
- Bannon, Liam J. (1991). 'From Human Factors to Human Actors: The Role of Psychology and Human-Computer Interaction Studies in System Design.' In *Design at Work: Cooperative Design of Computer Systems*, edited by Joan Greenbaum and Morten Kyng, 25–44. Lawrence Erlbaum Associates.
- Beer, David. (2017). 'The Social Power of Algorithms.' *Information, Communication & Society* 20 (1): 1–13. <https://doi.org/10.1080/1369118X.2016.1216147>
- Chun, Jon, and Katherine Elkins. (2023). 'The Crisis of Artificial Intelligence: A New Digital Humanities Curriculum for Human-Centred AI.' *International Journal of Humanities and Arts Computing* 17 (2): 147–167. <https://doi.org/10.3366/ijhac.2023.0310>
- Berry, David M., and James Stockman. (2024). 'Schumacher in the Age of Generative AI: Towards a New Critique of Technology.' *Thesis Eleven*. <https://doi.org/10.1177/13684310241234028>
- Drucker, Johanna. (2012). 'Humanistic Theory and Digital Scholarship.' In *Debates in the Digital Humanities*, edited by Matthew K. Gold, 35–45. University of Minnesota Press. <http://dhdebates.gc.cuny.edu/debates/text/34>
- Eisenstein, Elizabeth L. (1980). *The Printing Press as an Agent of Change: Communications and Cultural Transformations in Early Modern Europe*. Cambridge University Press.
- Graff, Gerald. (1987). *Professing Literature: An Institutional History*. University of Chicago Press.
- Grafton, Anthony. (2011). *The Culture of Correction in Renaissance Europe*. Yale University Press.
- Hall, Stuart. (1980). 'Encoding/Decoding.' In *Culture, Media, Language: Working Papers in Cultural Studies, 1972–79*, edited by Stuart Hall, Dorothy Hobson, Andrew Lowe, and Paul Willis, 128–38. Routledge. <https://doi.org/10.4324/9780203381182>
- Han, Byung-Chul. (2018). *Trans. Wieland Hoban. The Expulsion of the Other: Society, Perception and Communication Today*. Polity.
- Hoggart, Richard. (1957). *The Uses of Literacy: Aspects of Working-Class Life*. Chatto and Windus.
- Jasanoff, Sheila, ed. (2004). *States of Knowledge: The Co-Production of Science and Social Order*. Routledge. <https://doi.org/10.4324/9780203413845>
- Latour, Bruno. (1987). *Science in Action: How to Follow Scientists and Engineers Through Society*. Harvard University Press.

- Liu, Alan. (2025). 'Data Science and the Post-Liberal Arts University.' *Critical Inquiry* 51 (4): 597–618. <https://doi.org/10.1086/735621>
- Lupton, Deborah. (2014). *Digital Sociology*. Routledge. <https://doi.org/10.4324/9781315776880>
- McLuhan, Marshall. (1964). *Understanding Media: The Extensions of Man*. McGraw-Hill.
- Munis, Jacqueline. (2026). 'Palantir CEO Says AI "Will Destroy" Humanities Jobs, but There Will Be "More than Enough Jobs" for People with Vocational Training.' *Fortune*, April 11. <https://fortune.com/article/palantir-ceo-alex-karp-ai-humanities-jobs-vocational-training>
- Noble, Safiya Umoja. (2018). *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York University Press.
- Segessenmann, Jan, Thilo Stadelmann, Andrew Davison, and Oliver Dürr. (2025). 'Assessing Deep Learning: A Work Program for the Humanities in the Age of Artificial Intelligence.' *AI and Ethics* 5 (1): 1–32. <https://doi.org/10.1007/s43681-023-00408-z>
- Selwyn, Neil. (2010). 'Looking Beyond Learning: Notes Towards the Critical Study of Educational Technology.' *Journal of Computer Assisted Learning* 26 (1): 65–73. <https://doi.org/10.1111/j.1365-2729.2009.00338.x>
- Selwyn, Neil. (2022). *Education and Technology: Key Issues and Debates*. 3rd ed. Bloomsbury.
- Small, Helen. (2013). *The Value of the Humanities*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199683864.001.0001>
- Underwood, Ted. (2018). 'Why an Age of Machine Learning Needs the Humanities.' *Public Books*. <https://www.publicbooks.org/why-an-age-of-machine-learning-needs-the-humanities>
- Underwood, Ted. (2025). 'The Impact of Language Models on the Humanities and Vice Versa.' *Nature Computational Science*, 5: 695-697. <https://doi.org/10.1038/s43588-025-00819-4>
- Williams, Raymond. (1958). *Culture and Society, 1780–1950*. Chatto and Windus.
- Williams, Raymond. (2003). *Television: Technology and Cultural Form*. 3rd edition. Routledge. <https://doi.org/10.4324/9780203426647>