The Death of the Ghost in the Machine:

A Dimensional Approach to Artificial Intelligence in Digital Archives

LJ Fulugan

HIST 7372: History of Archiving and Archival Records

6 April 2026

> The Sultan asked Solomon for a Signet motto,
>
> that should be good for Adversity or Prosperity.
>
> Solomon gave him, "THIS ALSO SHALL PASS AWAY."
>
> –Edward Fitzgerald, "Solomon's Seal," *Polonius*.[1]

Artificial intelligence (AI) is a poorly-explained buzzword. This is true in the ambiguity with which "AI technology" is referred to in conversation as well as the opacity of the internal mechanisms of AI models beyond input and output. Regardless of one's stance on the topic, it is not contentious to recognize the technology's prevalence in whatever form it takes and its implications for broad-spectrum change. Scholars in the disciplines of information systems, digital humanities, and archives have begun to grapple with the unexplainability and opacity of AI as it pertains to the structures and values of traditional academic research and society writ large. At the same time, the institutions and organizations that those critiques emerge from are increasingly incorporating generative AI and machine learning (ML) into their administrative management systems. The implementation of AI in corporate organizational suites (such as Microsoft 365's Copilot, Adobe Acrobat's AI Assistant, and Google AI's Gemini) means that users of those applications will frequently encounter AI regardless of whether they choose to engage with it or not. Because of the breadth at which AI technologies are applied in practice, across industries and academic disciplines alike, this essay follows overarching trends in the implementation of AI in archives and recordkeeping processes as the outset for further discussion. AI technologies, I argue, can be understood as historically- and socially-mediated, nuanced, and contextually situated in development and use. Without an awareness of the social,

[1] Line breaks added. While the quote "this too shall pass" was popularized in English by Fitzgerald's 1852 *Polonius: A Collection of Wise Saws and Modern Instances*, "Solomon's Seal" was based on a translation of a Farsi parable attributed to the twelfth century Sufi poet Fariduddin Abu Hamed Mohammad Attar Nishapuri, or Attar of Nishapur. See "This Too Shall Pass," in *Tales from the Land of the Sufis,* trans. Mojdeh Bayat and Mohammad Ali Jamnia (Shambhala Publications, 2014). Similar story traditions appear in Ottoman Turkish, as well as in Hebrew and Judaic practice.

socio-technical, and technological dimensions in which AI becomes implicated alongside the institutional and infrastructural networks that inform records management, the rampant pace of AI's implementation in the present moment appears to run counter to traditional archival principles of accountability and transparency. The perspective on AI illustrated by this paper, on the scholarly front, follows a framework termed digital provenance. The digital provenance framework is rooted in a postmodern conception of archives that foregrounds conceptual fluidity and materiality, articulated in a transitional period of Canadian national record-keeping at the turn of the twentieth century, marked by drastic shifts and obsolescence in information and communication technologies. This makes it a strong corollary for the contemporary moment of technological change of today. Rather than a statistical or empirical review of AI in archival literature, this essay falls more in line with digital humanities scholarship to consider the relationship between the multiple dimensions of digital provenance that scaffold AI systems in archival work.

A Definition of AI: Change, Systems, and Parent Categories

I stress the fleeting nature of the present moment. As of the writing of this paper at the beginning of 2026's second quarter, major annual reports on AI have already been published by leading figures in the industrial and academic arms of the technology's research and development. The phrase "a time of unprecedented change" has entered the common parlance of anywhere with access to broadband internet and, while it no longer bears repeating, the sentiment is true. AI represents a novel technology of greater scale than the personal computer and World Wide Web according to its rate of adoption, however uneven (a motif echoed throughout publications on AI).[2] But have we, as the current living of our collective species, ever lived in a time of relative

[2] Stanford University Institute for Human-Centered AI, AI Index Steering Committee, *AI Index 2026 Annual Report*, (Stanford, April 2026), 3, 199, https://hai.stanford.edu/ai-index/2026-ai-index- report; *International AI Safety Report 2026,* Department of Science, Innovation, & Technology and AI Security Institute, vol. 2 (United

stability? While drawing on past reports of other novel technologies in the time of their developments might grant insight into this premise as it applies to the current topic (and in a more thematically suitable manner for this paper), it is not difficult to understand change as a human constant. The impermanence of things as it pertains to the experience of being alive has long been a mainstay in canons of philosophical, cultural, and artistic thought throughout times and spaces. Similar notions of change and impermanence are profuse in discussions on AI. Yet, in recognizing change, the emphasis of AI's transformative effects calls for a more critical evaluation beyond its moral adjudication: what changes, specifically? Where? Why? What does change mean in this moment, this place, this context of use? What can we learn from the adoption of novel technologies like the personal computer and World Wide Web (predecessors, to some respect, to AI) and the dissonance between intention and result? Change is constant, yes, but that is no excuse for wholesale acceptance or generalization where the systems we interact with are deployed in the most literal and lethal sense of the verb.[3] A causal series of discrete events is an easy way of approaching a monolith; this is not how a survey grounded in academic literature is written, nor is it an adequate treatment of something so enflamed in public discourse. Take "AI is the worst it will ever be right now,"[4] for instance, and consider what "worst" refers to: technological capability, human likeness, resource consumption, or societal repercussion?

The relative lack of clarity with which AI is referred to stems, in part, from the technology's nomenclature itself. The definition of "AI," or artificial intelligence, eludes

---

Kingdom Government, February 2026), 85, https://internationalaisafetyreport.org; OECD Legal Instruments, *Recommendation of the Council on Artificial Intelligence,* OECD/LEGAL/0449, 2019, preamble (last amended 3 May 2024), https://legalinstruments.oecd.org/en/instruments/oecd-legal-0449#background Information.

[3] The military applications of AI technology represent a facet of the discussion beyond the one I can offer here, but I utilize it to stress the implications of the interconnectedness of AI's development.

[4] Geoffrey Hinton, Robert and Elizabeth Distinguished Knight Lecture (University of Manitoba, 13 April 2026).

consensus. Where close similarities can be made between the different terms put forward by different organizational and administrative bodies, there is no single widely-accepted definition; the colloquialism of "AI" in regular conversation is thus doubly ambiguous. For this paper, I refer to modified versions of the terms "AI model" and "AI systems" when I use the phrase "AI technologies" in general. Put forward by Sun et al. based on the most recent iteration of the OECD AI Principles outlined by the *Recommendation of the Council on Artificial Intelligence*,[5] "AI systems" refer to "machine-based systems" with the capability to self-infer its generated output (for example, action, prediction, content, decision, in both digital and physical environments) following input data and structured objectives.[6] The parent category of systems, in this definition, is system lifecycle: that is, the entirety of a given system's development, test, use, and retirement process, in concert with the AI actors (notably, in OECD, a subset of "stakeholders") as well as the related knowledge, skills, and resources involved in an actor's participation in a system lifecycle. Sun et al.'s definition of "AI system" is the same as OECD's in essence. Their definition of "AI model" is as follows:[7]

> A computational construct that makes inferences from inputs to produce outputs and forms (or is intended to form) the reasoning or knowledge core of an AI system.

Sun et al.'s definition follows a boundary problem in distinguishing between AI models and systems, tracing the technical ambiguity from application-specific large language models prior to 2020 through a contemporary trend towards general-purpose systems. Sun et al.'s systematic literature review highlights a number of definitions for the terms "AI models" and "AI systems" from 2019 on, put forward in academic/private research literature as well as by the OECD,

[5] OECD Legal Instruments, *Recommendation of the Council on Artificial Intelligence,* §I.

[6] Yuanyuan Sun et al., "Defining AI Models and AI Systems: A Framework to Resolve the Boundary Problem," (preprint, arXiv [Cs.CY], last revised 16 March 2026), 20–21, https://arxiv.org/abs/2603.10023v2.

[7] Sun et al., "Defining," 20.

European Data Protection Supervisor, ISO/IEC, UNESCO, and numerous organizational bodies in the European Union. In doing so, the preprint identifies three conceptual perspectives informing the two definitions: models/systems as units belonging to respective parent categories, models/systems as they relate to one another, and their respective features. "Features" include a model or system's respective productions and functions (and, in a system's case, the mechanism by which it operates).[8] Crucially, these dimensions are not mutually exclusive and allow for a granular perspective on the inter-related facets of AI models and systems. There can be plural parent categories attached to a subcategory. This reading underscores the relational dimension between system and model, squaring the model as the "core" aspect of the system's inference and production of output, and reveals a taxonomic logic to the definition of AI; the system becomes structured as the parent category of the model, which each may act as the parent categories of more granular components of AI technologies or belong to other parent categories.

Highlighting the relational aspects of AI concepts, one can expand the scope of analysis to ancillary institutional and organizational discourses that also inform the development and operation of AI. One way these can be read is through AI development indices. Examples of indices that have been published in 2026 by AI oversight/advisory organizations thus far include the OECD.AI Index, a method of measuring participant countries' implementation of OECD's *Recommendation* with results from the 2023-2024 cycle),[9] the ninth annual edition of the *AI Index Report* by Stanford University's Institute for Human-Centered Artificial Intelligence (HAI),[10] and the second *International AI Safety Report* written and administered by an independent group of members from over thirty national governments and intergovernmental

[8] Sun et al., "Defining," 8–9.

[9] OECD, *The OECD.AI Index,* (OECD Publishing, 2026), https://doi.org/10.1787/32c01014-en

[10] Stanford University Institute for Human-Centered AI, *AI Index 2026 Annual Report,* 2.

organizations.[11] This a short-list of selected non/intergovernmental indices; publications continue to be produced by networks of organizations committed to the governance and risk management of AI technologies throughout their system lifecycles. Indices and their publishing organizations supply recommendations for AI policies and principles at the institutional level and reflect ongoing observations of AI development across different environments. For instance, the 2026 Stanford HAI index is organized into nine chapters on Research and Development (1); Technical Performance (2); Responsible AI (3); Economy (4); Science (5); Medicine (6); Education (7); Policy and Governance (8); and Public Opinion (9).[12] The HAI index's chapter organization has been fluid across its nine-year tenure while preserving the general conceptual schema of its scope. Where indices might fail to capture the lay interactions of AI at an everyday-use level (for instance, an individual user's chat prompt with ChatGPT or Microsoft Copilot), societal engagement with AI technologies can be contextualized through the methodological indicators used to refine AI development indices. Empirical evaluations of AI using development indices represent an extension of systemic literature reviews; precursory examples from the last two years can be found examining the independent Tortoise Media *Global AI Index*, the Stanford HAI index, and the Oxford *Government AI Readiness Index.*[13] Ojeda-Castro et al.'s examination of Tortoise's *Global AI Index* identifies global AI development trends from the index's three pillars (Implementation, Innovation, and Investment), finding that talent and technical infrastructure play strong roles in a country's technological innovation.[14] In having highlighted

[11] *International AI Safety Report 2026,* 9.

[12] Stanford University Institute for Human-Centered AI, *AI Index 2026 Annual Report,* 7–8.

[13] Angela Ojeda-Castro et al., "Global drivers of artificial intelligence development: An empirical study using the global AI index across 62 countries," *Issues in Information Systems* 26, no. 2 (2025): 125–126, https://doi.org/10.48009/2_iis_110.

[14] Tortoise GAI, qtd. in Ojeda-Castro et al. 126–127.

the usefulness of AI development indices to understand AI's implementation in society, however, it is equally important to identify the need to navigate bias in the methods of such reports from the point I write. In the 2026 HAI index, only a single reference to records management is made under a section on the use of generative AI with patient's health records.[15] Throughout the most-recent AI development indices, records (administrative records, academic and organizational publications, institutional reporting, for example) are treated as sources for systematic reviews and analyses. Records professionals such as archivists, whose domain includes the theoretical conceptualization of records and their organization, can remediate the treatment of the concept of record in AI development beyond source. I return to AI development indices later on, in accordance with the framework of digital provenance, as well as the fluidity of archival theory in the main. This last point reveals, however, another shade of AI's opacity as it pertains to this discussion: an underemphasized sector of AI as it pertains to records management, archival and otherwise, from the perspective of AI development.

A Digital Provenance Framework

These lines of questioning are not absent from the disciplinary vantage of digital archives and preservation I write from. In 2019, the same year as the first HAI index, Pendergrass et al. published a framework for sustainable digital preservation drawing on comparative definitions of the concept of "sustainability" in cultural heritage spaces (*ie.,* finance, staffing, and environmental).[16] In doing so, the framework centers the technological dimension of digital preservation as enabled by information & communication technology (ICT) and argues for an awareness of environmental sustainability in digital preservation that spans the entirety of a

[15] Stanford University Institute for Human-Centered AI, *AI Index 2026 Annual Report*, 277.

[16] Keith Pendergrass et al., "Toward Environmentally Sustainable Digital Preservation," *The American Archivist* 82, no. 1 (2019): 167–168, https://doi.org/10.17723/0360-9081-82.1.165

given technology's life cycle. I argue that such an approach is necessary for the implementation of AI in archival and records management processes. Where Pendergrass et al.'s article focuses on sustainability and the relationship between digital preservation and ICT system lifecycles, understanding AI systems as constituent parts of system lifecycles allows one to begin tracing the contextual web that enravels the implementation of AI system lifecycles with archival processes and workflows. The anchor points of this web are not bounded; they exist across the conceptual dimensions and taxonomic hierarchies raised. How, then, does one unravel a multidimensional web?

Digital provenance is a framework for approaching born-digital records based on literature in archival theory, digital preservation, and media archaeology.[17] The framework connects theoretical discourses on the processing and management of digital records that can be used to situate the implementation of AI within a pre-existing body of recordkeeping literature. This section first examines digital provenance through two precedent archival concepts prior to relating it to the discussion at hand. The first concept, as referenced in Bak's work on digital provenance, is the postmodern archives approach to provenance per Tom Nesmith's 2002 "Seeing Archives."[18] The second is Laura Millar's extensive 1998 treatise on total archives, a Canadian records management system first defined by Dominion Archivist Wilfred Smith in 1972.[19] Nesmith and Millar's respective works were written from a shared recordkeeping context inundated by a rise in electronic and digital media that required the re-negotiation of

---

[17] Greg Bak, "Digital provenance," *Archival Science* 24 (2024): 850, 852, https://doi.org/10.1007/s10502-024-09462-w.

[18] Tom Nesmith, "Seeing Archives: Postmodernism and the Changing Intellectual Place of Archives," *The American Archivist* 65, no. 1 (Spring/Summer 2002): 38–39, quoted in Bak, "Digital provenance," 849.

[19] Laura Millar, "Discharging our Debt: The Evolution of the Total Archives Concept in English Canada," *Archivaria* 46 (Fall 1998): 117, https://archivaria.ca/index.php/archivaria/article/view/12677.

longstanding notions of materiality/digital materiality, originality and authenticity, and the crucial agency of archivists and records managers within the transigent landscapes of preservation. Digital provenance contributes to the re-negotiation of these issues in the current moment; I utilize the framework towards an examination of the implementation of AI in digital archives, and the "spectre of automation"[20] that haunts archival work today.

A Genealogy: Total and Postmodern Archives in Canada

A genealogical approach to digital provenance (*à la* Foucault) finds its roots in the works of thinkers like Tom Nesmith and reflects a theoretical reorientation through the tail end of the twentieth century generally referred to as the postmodernist turn in archival studies and a broader archival turn in academia as evident in cultural studies, literary studies, and other disciplines in the arts & humanities. Central to Nesmith's thesis is a close reading of Jacques Derrida's *Archive Fever*, from which a postmodern understanding of archives follows to complicate traditional understandings of truth and representation by underscoring the interpretative judgements of archivists that structure a record's meaning and reveal their contextuality.[21] Rather than a consistent and organized line drawn straight from creator to archives, the provenance of a record might be entangled by the activities of contributors (archival and non-archival) across time and space, both before and after its acquisition into a repository. Though not a formally organized school of thought, the scholars attributed to postmodernism can be compared according to a shared critique against normative conceptions of meaning and truth and the disruption of ontological, epistemological, and methodological bases of knowledge (see western post-secondary institutes and research) as shaped by and shaping archival processes.[22] To that

[20] Millar, "Discharging," 127–128.

[21] Nesmith, "Seeing," 25, 30–31.

[22] Nesmith, "Seeing," 32–33, 37.

effect, a postmodern appreciation of provenance is identified in contrast to a modern or classical tradition of provenance encapsulated by the concept of *respect des fonds* (itself, defined by the archival context from which the principle was written). Concepts and perceptions of what those concepts represent are contextually subjective and recursively influenced by the contexts of communication within which they are beheld. A postmodern archival understanding of provenance contends with the meanings of any one record or act of preservation by the wider contexts of communication (intellectual, historical, social, legal, for a few examples) from which those meanings are made. Meanings, concepts, and perceptions of what those concepts are intended to mean are contextually subjective and recursively influenced by the contexts of communication within which they are beheld. These have always been changing.

Though predominantly seen as an intellectual shift towards contextualism, the postmodern turn in archival theory ran concurrent with emerging concerns surrounding the management of records—particularly electronic and/or digital records—and was, itself, framed by a particular recordkeeping context. The final section of Nesmith's "Seeing Archives," entitled "Implications for the Computerized Record," traces the crucial relationship between postmodern theories of archival provenance to the advent of electronic records through the latter half of the twentieth century and the natal years of the twenty-first.[23] The increasingly plural digital and analog forms of multimedia acquired by archival institutions in the time of "Seeing Archives" eluded pre-existing practices of the management and preservation of paper/hardcopy documents and presented problems for the management of archives and collections distinct from their non-electronic counterparts. Namely, these problems included a mass increase in the volume of records being created and preserved, the need for a standardized system of describing and

[23] Nesmith, "Seeing," 37–38, 39, 41.

appraising electronic records, and the change and obsolescence of technical infrastructures, which accompanied the postmodern critique and re-capitulation of key disciplinary concepts in archival studies. Laura Millar's treatise on the Canadian concept of total archives was written from within the same temporal and national context as Nesmith's work and raises similar issues as those posed by postmodern archives. The total archives concept was first coined in 1972 by then-Dominion Archivist Wilfred I. Smith to refer to a Canadian records management program practiced since the establishment of a national archive under Douglas Brymner in the nineteenth century, wherein an institutional archival mandate encompassed the preservation of both private and public records. This mandate would continue to be held through the proliferation of communication technologies (like wireless telegraphy and telephone systems, microphotography and photographic reproduction, and increasingly complex audiovisual technologies) from the mid-twentieth century onwards.[24] Millar heralds "The Spectre of Automation" in reference to these advancing computer and informational technologies and the issues they posed. As it related to the then-contemporary total archives concept, these included a distinction between records and information, dependency on changing media and ongoing technological interventions, and the comparative ease with which electronic records can be created and manipulated in relation to paper documentation. Millar's review of the history of total archives called for more proactive records management programs within institutional archives in the current moment in order to contend with the management and preservation of digital records.

The 1980s shift from a centralized total archives concept to a more integrated and segmented records management system was deeply entrenched in the socio-economic and administrative realities of Canadian budget cuts of the current century, just as postmodern

[24] Millar, "Discharging," 125–126, 127–128.

archival thought effloresced. The history of total archives grants insight into the social, technological, and institutional realities that informed the Canadian archival and recordkeeping context from which both Nesmith's and Millar's works came from. Both total and postmodern archives concepts were negotiated amidst shared concerns at the end of the twentieth century which unsettled traditional notions of provenance and materiality in a recordkeeping era of profound theoretical and technological change.

Dimensions of Digital Provenance

Millar's "Spectre of Automation" and the distinction drawn between physical records and digital information, while an example of the archival context of the time, forewarns a common conception of digital materiality as opposed to a more "physical" materiality of analog records. Analog and digital records each appear, at first, as if they might be cleanly delineated by a "physical" world inhabited by the fleshy bodies one lives in and an abstract world of digitality unbounded by physicality. Bak's concept of digital provenance allows one to question this perception of the materiality of digital records. As stated earlier, the concept of digital provenance was first written for use with born-digital records.[25] However, just as the materiality of a digital record might appear transient, like a ghost, the technical dimension of digital provenance involves the physical hardware and infrastructure that support digital systems as well as the records created therein. Recalling the definition of AI systems offered by Sun et al., digital systems can be read as a parent category of AI systems. While the concept of digital provenance specifically applies to digital records, rather than digital recordkeeping functions and processes, the social, socio-technical, and technological dimensions that digital provenance invokes are important contexts for understanding AI technologies and AI-modified tools as they are used

[25] Bak, "Digital Provenance," 850.

within archival settings. The theoretical approach to linking AI and archives outlined thus far is taxonomic, in its organization of definitions in parent categories, but flexible, in its application across distinct system lifecycles. Just as digital provenance is a type of provenance, re-scoping the digital provenance framework reflects a taxonomic re-ordering of its original scale of analysis from records to recordkeeping process. It involves joining different systems of knowledge and technology through the dimensions outlined above and, harkening to Pendergrass et al.'s sustainable digital preservation framework, opens the way for a concrete inquiry into AI's implementation and use in archival processes. Granted the dual-pronged ambiguity of AI technology, methods of conceptual organization are in order—methods that digital archivists and recordkeepers are well-situated to supply.

A digital provenance framework involves examining the technical, socio-technical, and social dimensions of digital records and the systems used to create them. As stated above, Bak's concept follows Nesmith's postmodern understanding of provenance. Emphasizing the materiality of born-digital records, digital provenance is a particular type of provenance that focuses on the systems and system users throughout a record's custodial history through three related dimensions.[26] The technical dimension of the framework examines the relationship between technological systems and the socio-cultural contexts they emerge from and foment, as structured by born-digital records and the systems involved in their creation, use, and preservation. The socio-technical dimension captures the systems of knowledge involved in user engagement with digital technologies and records. Finally, the social dimension of digital provenance considers the meanings and symbolisms ascribed to digital systems and born-digital records. The breadth of technologies implicated in the production and maintenance of born-digital records bears grave importance for contemporary archival practitioners, who now

[26] Bak, "Digital Provenance," sec. 1, "Dimensions of digital provenance."

face AI systems as another digital system to adjudicate as part of a record's own lifecycle.[27] Digital provenance offers a perspective by which AI systems can be understood in the recordkeeping processes of digital material, throughout and across core functions of archival work (. Where earlier archival literature like Millar and Nesmith's warn of ghosts of futures past in the form of computerization and automation, digital provenance offers a solution to the present future of dynamic technological and social change wrought by the development of AI in the current decade of the twenty-first century. Thus, a similarity can be drawn between the recordkeeping context in which postmodern provenance and total archives were written and the contemporary moment for which digital provenance was; a context of conceptual fluidity facing an increasingly vast corpus of digital records and the rapidly changing networks of systems (including AI systems) that structure their creation, use, and archivy.

The background of digital provenance offered by Bak situates the importance of provenance as it has grown beyond archival studies in the advent of machine learning and large language AI models, citing a 2023 executive order from former American president Joseph Biden that stressed safety, security, and trust.[28] Returning to the topic of AI development indices, one can identify how these sorts of values and principles capture the social and socio-technical dimensions of the implementation of AI technology over time and in our present moment; the conceptual definitions of AI models, systems, and system lifecycles allows one to bridge from the socio-technical dimension of AI research to the technical dimension of infrastructure necessary to support the implementation and use of the technology across different system lifecycles. With respect to different system lifecycles, one can think of how the same AI application programming interface (API), like those offered by OpenAI, Anthropic Claude, or

[27] Bak, "Digital Provenance," 852.

[28] Bak, "Digital Provenance," 860–861.

Google Gemini, can be operated towards distinct functions. The AI development indices capture some of these in their respective analyses. While the discussion on AI models and systems and the conceptual framework used to define them are admittedly generalized, the dimensional bracketing of digital provenance specifies the overarching social, socio-technical, and technical aspects of AI technologies to their specific implementation in archival recordkeeping. Digital provenance represents a terminal block by which a discussion of AI systems and models can be directly understood in relation to archival processes and functions, and the framework's dimensions constitute the crux of the overlap between the two systems of knowledge.

AI Systems in Archives and Records Management

While a consideration of records management is lacking in AI development, the reverse is not true. Archival professionals have been cognizant of the need to negotiate automation and digital materiality since the time of Nesmith's writing, and archival scholars of the present moment have closely followed the AI turn since the second decade of the twenty-first century. One such perspective comes from Dighton, Abel, and Bikmohammadi, who interview archivists from galleries, libraries, archives, and museum (GLAM) institutions from a host of different administrative and organizational structures in order to illustrate current perceptions of generative AI technologies within local archival knowledge systems.[29] From a digital provenance perspective, the article foregrounds a socio-technical dimension of generative AI systems as used by archivists. The method by which the authors synthesize the field at play (via participant interviews) is distinct from typical reviews of AI in archives, which otherwise take empirical statistical approaches in the form of systematic literature reviews. In doing so, Dighton, Abel,

[29] Desiree Dighton, Benjamin Abel, and Mina Bikmohammadi, "Localizing with GAI in the Archives: Exploring Practitioner Attitudes on Challenges and Opportunities," *Technical Communication Quarterly* 34 no. 3 (2025): 323–324.

and Bikmohammadi's review stresses a human-to-human approach and outlines a social dimension of digital provenance through both the plural definitions of AI given in the archivists interviews and the authors' argument for a participatory localization-based approach to integrating archival systems of knowledge.[30] The participatory localization approach calls forth the need to consider the specific ethical, legal, and technical dimensions of archival functions, from appraisal and description to preservation and further, when implementing GAI as part of them. In simple terms, AI systems must be implemented in archival and records management processes in a localized manner, with the participation of the users involved in a given system's lifecycle. Such an approach is sensitive to obligations of access while maintaining the evidential safety and trust ascribed to the archival institution's preservation of documentary heritage.[31] At the same time, the technical capability of generative AI systems to produce output from large volumes of input data is a frequent point raised and potentiates an opportunity to address longstanding issues of access in archival institutions. The interviews conducted highlight key points for our discussion that connect to broader discourses of AI as implemented in the specific system lifecycles of archival and records management processes. Framed in terms of digital provenance, technical applications to archival and recordkeeping processes will never be perfect out of the box, and must be articulated in relation to the local social and socio-technical dimensions in the same manner as digital records.

A systematic literature review and analysis of AI in archives can be found in an article by Mannheimer et al.[32] The objectives of the review hinge on two questions: how is AI used in the

[30] Dighton, Abel, and Bikmohammadi, "Localizing," 330.

[31] Dighton, Abel, and Bikmohammadi, "Localizing," 322–323.

[32] Sara Mannheimer et al., "Responsible AI Practice in Libraries and Archives: A Review of the Literature," *Information Technology and Libraries* 43, no. 3 (2024): 2–3, 4–6. https://doi.org/10.5860/ital.v43i3.17245.

practices of libraries and archives? What ethical concerns are raised and addressed by the implementation of AI in archival and library practices? Its sample is based on American and Canadian AI case studies in archival and library studies from 2017 to 2023 and distinguishes the projects on the bases of institution type, size, region, disciplinary collaboration, and types of material processed. The number of case studies including AI increased each year over that period (barring 2023, wherein cases were counted from January to July rather than the full year). Mannheimer at al. survey specific AI technologies and ethical concerns noted in each case study project via keyword tag analyses. Out of 51 keyword standardized keyword tags, metadata extraction was the most common function of AI tools in the cases examined, followed by natural language processing and image recognition. Standardizing ethics and values-related keyword tags found that accuracy, human involvement, and bias were the three most frequent ethical values highlighted in case studies: in descending frequency, the other ethics/values keyword tags identified by the review were social justice, transparency, privacy, policy & guidelines, user communities, cost considerations, reliability, accessibility, access, and content.[33] The systemic literature review's associated analysis gives three implications for the implementation of AI in archives and libraries.[34] First, that the three standardized ethical/value principles call for a level of transparency that makes known the implicit biases of AI system users. Second, the technical dimension of AI systems involves a greater degree of interdisciplinarity with fields and units of relevant technological expertise that may reveal disjunctions in how different professionals conceptualize particular terms. Third and final, the majority of published case studies of AI's implementation in archives and libraries do not include extensive discussions on ethical/value considerations. To this last point, Mannheimer et al. indicate that the two most frequently-cited

[33] Mannheimer et al., "Responsible AI Practice," 7–9.

[34] Mannheimer et al., "Responsible AI Practice," 23–24.

ethical/value keyword tags, accuracy and human-in-the-loop, are technical in consideration rather than explicitly ethical.[35] While ethics are inherently values (as ethical values), not all values are based on ethics. The digital provenance framework can delineate the ethics/value keyword tag issue in three dimensions; in other words, the technical and socio-technical dimensions of AI systems are over-emphasized in archival and related literature to the detriment of its consideration in the social dimension of use. "Use" becomes the central action of AI's implementation. In the evaluation of AI's implementation in archives, a flexible understanding of value that can contend with its plurality is vital. Here, Nesmith's concept of postmodern provenance can be fully appreciated; we must be aware of the contexts that inform values across inter-connected systems and conceptual dimensions with balanced nuance, cognizant of *pro* and *con,* rather than polarizing a body of system users that are already connected by the digital network with or without the presence of the ghost.

While I hesitate to identify this paper as specifically postmodern, many of the premises I put forward would likely be called so for my focus on multiple dimensions and contextuality. A common misconception is that postmodern thought operates only in a world of abstract thought and concept relationships. This is not the case; just as the genealogical forebears of digital provenance were theorized in economic and technological realities to offer practical considerations to contemporary practice, the same can be done for AI's implementation in archives and other recordkeeping institutions. In consideration of the ethical/value keyword trends offered by Mannheimer et al. and the overall push for greater transparency and accountability brought forth by AI systems, paradata becomes increasingly important to document.[36] Paradata is a concept rooted in sociology statistics, attributed to the work of Mick P.

---

[35] Mannheimer et al., "Responsible AI Practice," 19.

[36] Jeremy E. Davet et al., "Tracking the Functions of AI as Paradata & Pursuing Archival Accountability," *Society for Imaging Science and Technology (IST) Archiving Conference* 19 (2022): 83–84.

Couper on survey data collection methods, and refers to the functions and data produced as part of the creation of metadata.[37] In the most direct sense, paradata is vital where generative AI is used to produce metadata for records by tracking and documenting the adjacent workflows that support an AI system and its users in a given lifecycle. At a more conceptual level, paradata is meant to address the black box problem of AI, or the general unexplainability of how AI models function in AI systems. Public calls for "responsible AI" run parallel with calls for "explainable AI;" the occluded internal mechanisms of AI models represent a problem for the values of transparency held high in the archival discipline. Paradata does not directly elucidate the contents of the black box. However, taken in tandem with framework approaches like digital provenance, it represents a valuable tool for archivists and records professionals working with digital records and recordkeeping processes that implement AI systems.

Conclusion

Artificial intelligence is not a singular or stable concept, but is instead shaped by shifting definitions, systems, and contexts of use. The ambiguity surrounding AI is not merely a matter of unclear terminology, but reflects the layered relationships between models, systems, and the broader institutional and technological infrastructures in which they are embedded. Approaching AI through system lifecycles and relational categories reveals the extent to which its development and use are contextually-informed beyond isolated changes in discrete areas. The discussion of archival literature on AI underscores implementation's uneven emphasis across the digital provenance framework's social, socio-technical, and technical dimensions. Technical capabilities and user integration are frequently foregrounded, particularly in applications such as metadata extraction and natural language processing, while the social dimension, including

[37] Davet et al., "Tracking," 84, 86–87.

questions of value, interpretation, and ethical responsibility, is less consistently articulated. This imbalance reflects a broader tension between the drive for efficiency and the archival commitments to transparency, accountability, and contextual understanding. Concepts such as digital provenance and paradata emerge as supportive responses to this tension, offering ways for archivists and records professionals to document the processes surrounding AI use in their local institutions, even as the internal mechanisms of such systems remain opaque. The intersection of AI and archives cannot be separated from the wider contexts in which they meet. AI systems, in this sense, do not function as universal solutions, but must be negotiated within the constraints and responsibilities that define archival work. I return to the question of change. While AI is frequently cast as an unprecedented rupture, the archival field has never been a stranger to upheaval. The transition to electronic records, the proliferation of digital media, and the obsolescence of technological infrastructures each unsettled and reconfigured established archival principles. To treat the present as exceptional risks obscuring these continuities and, with them, the insights of prior transformations. The death of the ghost is an understanding of its lifecycle.

Bibliography

Bak, Greg. “Digital Provenance.” *Archival Science* 24 (2024): 849–864. https://doi.org/10.1007/s10502-024-09462-w.

Davet, Jeremy E., et al. “Tracking the Functions of AI as Paradata & Pursuing Archival Accountability.” *Society for Imaging Science and Technology (IST) Archiving Conference Proceedings* 19 (2022): 83–87.

Dighton, Desiree, Benjamin Abel, and Mina Bikmohammadi. “Localizing with GAI in the Archives: Exploring Practitioner Attitudes on Challenges and Opportunities.” *Technical Communication Quarterly* 34, no. 3 (2025): 321–334.

Hinton, Geoffrey. Robert and Elizabeth Distinguished Knight Lecture. University of Manitoba, April 13, 2026.

International AI Safety Report 2026. Vol. 2. United Kingdom Government, Department of Science, Innovation & Technology and AI Security Institute, February 2026. https://internationalaisafetyreport.org.

Millar, Laura. “Discharging Our Debt: The Evolution of the Total Archives Concept in English Canada.” *Archivaria* 46 (Fall 1998): 103–146.

Nesmith, Tom. “Seeing Archives: Postmodernism and the Changing Intellectual Place of Archives.” *The American Archivist* 65, no. 1 (Spring/Summer 2002): 24–41.

OECD. *Recommendation of the Council on Artificial Intelligence*. OECD/LEGAL/0449. 2019. Last amended May 3, 2024. https://legalinstruments.oecd.org/en/instruments/oecd-legal-0449.

OECD. *The OECD.AI Index*. OECD Publishing, 2026. https://doi.org/10.1787/32c01014-en.

Ojeda-Castro, Angela, et al. “Global Drivers of Artificial Intelligence Development: An Empirical Study Using the Global AI Index across 62 Countries.” *Issues in Information Systems* 26, no. 2 (2025): 125–132. https://doi.org/10.48009/2_iis_110.

Pendergrass, Keith, et al. “Toward Environmentally Sustainable Digital Preservation.” *The American Archivist* 82, no. 1 (2019): 165–206. https://doi.org/10.17723/0360-9081-82.1.165.

Stanford University Institute for Human-Centered Artificial Intelligence. *AI Index 2026 Annual Report*. Stanford University, April 2026. https://hai.stanford.edu/ai-index/2026-ai-index-report.

Sun, Yuanyuan, et al. “Defining AI Models and AI Systems: A Framework to Resolve the Boundary Problem.” arXiv preprint, last revised March 16, 2026. https://arxiv.org/abs/2603.10023v2