Algorithmic Curation and Cultural Memory: Ethical Implications of AI-Driven Archiving in Digital Humanities Projects

Dr. Anu Sharma
Teerthanker Mahaveer University, Moradabad, Uttar Pradesh, India
Email: er.anusharma18@gmail.com

Abstract

Digital humanities projects increasingly rely on machine learning pipelines to select, classify, describe, and surface cultural materials at scales that exceed human archival capacity. While algorithmic curation promises efficiency and discoverability, it also transfers consequential judgments about what a society remembers, and how, from accountable human institutions to opaque computational systems. This paper develops a conceptual framework for analysing the ethical implications of AI-driven archiving along four dimensions: selection (what enters the archive), description (how items are represented), retrieval (what is surfaced to whom), and preservation (what persists over time). Drawing on cultural memory theory and the ethics of information, we argue that algorithmic curation does not merely accelerate archival work but restructures the memory function itself, introducing distinctive risks of representational harm, silent exclusion, contextual flattening, and memory drift. We propose a set of governance principles — provenance transparency, contestability, plural curation, and periodic human re-review — and map them to concrete design interventions for digital humanities infrastructure. The paper contributes a vocabulary and an evaluative checklist that project teams, funders, and cultural institutions can use when deciding where, and under what safeguards, algorithmic curation is appropriate.

Keywords: Algorithmic curation; Cultural memory; Digital humanities; AI ethics; Digital archives; Representational harm; Archival governance

DOI: 10.21590/ijtmh20230904022

1. Introduction

Archives have never been neutral. Every archive is the residue of decisions — about what was worth keeping, whose voices were recorded, and which categories organised the record. Cultural memory scholars have long emphasised that societies remember through institutions, media, and practices rather than through any collective mind, and that the shape of the remembered past is therefore an artefact of curatorial labour. What is genuinely new in the current moment is not selectivity, but the delegation of selectivity to computational systems. Optical character recognition, image classification, entity extraction, automatic metadata generation, recommender systems, and, most recently, large language models are being embedded across the archival lifecycle of digital humanities (DH) projects, from mass digitisation programmes to community memory platforms.

This delegation is usually justified on capacity grounds: national libraries hold collections that could never be catalogued item-by-item by human specialists, and community archives frequently operate with minimal staff. Algorithmic curation is thus positioned as the only realistic route to access. Yet the systems doing this work were rarely designed with archival ethics in mind. They inherit biases from training data, they compress ambiguous cultural material into fixed category schemes, and their errors are systematic rather than random — meaning that the same communities, languages, and formats tend to be misdescribed or excluded again and again. Because archives feed education, journalism, scholarship, and public commemoration, errors in the memory infrastructure propagate outward into how a society understands itself.

This paper asks a focused question: what specifically changes, ethically, when curation becomes algorithmic, and what governance arrangements would allow DH projects to capture the benefits of scale while containing the distinctive risks? Our contribution is threefold. First, we decompose algorithmic curation into four archival functions and identify the ethical failure modes characteristic of each. Second, we synthesise these into a comparative framework contrasting human-institutional and algorithmic curation. Third, we derive governance principles and translate them into implementable practices for DH infrastructure.

2. Background and Related Perspectives

2.1 Cultural Memory as Infrastructure

Cultural memory theory distinguishes between communicative memory — the informal, lived memory of recent generations — and cultural memory proper, the institutionalised memory carried by texts, monuments, rituals, and archives. The archive occupies a privileged position in this account: it is the storage layer from which future acts of remembrance draw. Whoever controls admission to and description within the archive therefore exercises a quiet but durable form of power over the future past. Traditional archival science acknowledged this through professional norms — appraisal criteria, provenance principles, and descriptive standards — that made curatorial judgment inspectable and, in principle, accountable.

2.2 The Turn to Algorithmic Mediation

Scholars of platform governance have shown that when algorithmic systems mediate cultural visibility, they do so according to optimisation objectives (engagement, relevance, throughput) that are orthogonal to, and sometimes in tension with, cultural value. Work on search and recommendation has documented how ranking systems can reproduce and amplify social bias, and how personalisation can fragment shared informational environments. Critical data studies have further shown that classification schemes embedded in software carry political consequences: to classify is to make some things visible and others invisible. Applied to archives, these insights suggest that algorithmic curation should be analysed not as a faster version of cataloguing but as a new memory institution with its own political economy.

2.3 Ethics of Information and the Archive

Information ethics provides the normative vocabulary for this analysis. Concepts such as informational harm, epistemic injustice, and the right to be represented fairly extend naturally to archival contexts. A community can be wronged not only by surveillance or exposure but also by erasure — by being systematically absent from, or misdescribed within, the record that a society treats as authoritative. Because DH projects often work with material concerning marginalised or colonised communities, these harms are not edge cases; they are central design considerations.

3. A Functional Decomposition of Algorithmic Curation

We decompose curation into four functions and examine how algorithmic mediation alters the ethical character of each.

3.1 Selection: What Enters the Archive

In mass digitisation and web archiving, selection is increasingly performed by crawl policies, deduplication heuristics, format filters, and quality-assessment classifiers. Each embeds assumptions: a legibility classifier trained on clean print will silently discard handwritten, vernacular, or degraded material; a spam filter tuned on majority-language content will disproportionately exclude minority-language pages. The ethical shift is from contestable exclusion (a curator can be asked why an item was rejected) to silent exclusion (no human ever sees the rejected item, and no log records a reason intelligible to outsiders).

3.2 Description: How Items Are Represented

Automatic metadata generation — captioning, entity recognition, subject tagging — determines how items can later be found and understood. Models trained on contemporary, Western, majority-language corpora routinely mislabel cultural artefacts, apply anachronistic or offensive category terms, and strip context that human cataloguers would have preserved. Description errors are compounding: a mislabelled item is not merely wrong in itself but becomes unfindable by the communities best placed to correct it. We term this contextual flattening — the reduction of thick cultural meaning to whatever the model's label space can express.

3.3 Retrieval: What Is Surfaced to Whom

Search ranking and recommendation decide which fraction of a vast archive is ever actually seen. Popularity-weighted ranking creates a rich-get-richer dynamic in which already-canonical materials dominate attention, while personalisation risks fragmenting the shared reference points that cultural memory requires. When a national photographic archive is accessed primarily through a relevance-ranked search box, the effective archive — the archive as experienced — is the ranking function's output, not the collection itself.

3.4 Preservation: What Persists Over Time

Storage tiering, deduplication, and automated retention policies increasingly determine what is kept at what fidelity. Model-driven decisions about "near-duplicate" content can collapse meaningful variants (editions, annotations, provenance copies) into a single canonical object. Over successive migrations, small automated losses accumulate into what we call memory drift: the archive of 2040 differs from the archive of 2020 in ways no one decided and no one can reconstruct.

4. Comparative Framework

Table 1 contrasts human-institutional curation and algorithmic curation across the four functions and cross-cutting governance properties. The point of the comparison is not that human curation is unbiased — it demonstrably is not — but that its biases are located in accountable institutions with professional norms, whereas algorithmic biases are located in artefacts (training data, objective functions, thresholds) that are typically undocumented and unowned.

Table 1. Human-institutional versus algorithmic curation across archival functions

Dimension	Human-Institutional Curation	Algorithmic Curation	Distinctive Ethical Risk
Selection	Appraisal by trained archivists under published criteria; exclusions recordable and contestable	Crawl policies, filters, and classifiers applied at scale; exclusions unlogged and invisible	Silent exclusion of minority languages, formats, and communities
Description	Cataloguing under descriptive standards; ambiguity preserved in free-text notes	Automatic tagging within a fixed label space learned from majority corpora	Contextual flattening; offensive or anachronistic labelling
Retrieval	Finding aids, curated exhibitions, reference interviews	Relevance ranking, engagement-weighted recommendation, personalisation	Canon amplification; fragmentation of shared memory
Preservation	Deliberate retention schedules; provenance principle protects variants	Deduplication and tiering driven by similarity models and cost optimisation	Memory drift; unrecorded cumulative loss
Accountability	Professional norms; institutional answerability; decisions attributable to persons	Diffuse responsibility across vendors, models, and pipelines	Responsibility gap when harms surface
Error profile	Idiosyncratic, locally correctable	Systematic, correlated across the collection	Structural misrepresentation of the same groups repeatedly
Scale	Low throughput; backlogs common	Very high throughput; full-collection coverage feasible	Efficiency rationale crowds out deliberation

5. Governance Principles and Design Interventions

5.1 Provenance Transparency

Every algorithmically generated archival act — inclusion, exclusion, label, ranking — should carry machine-readable provenance: which system, which version, which confidence. Provenance converts silent decisions into inspectable ones and is the precondition for every other safeguard. In practice this means treating model outputs as assertions attributed to a named agent, exactly as archival description standards already treat human cataloguers' assertions.

5.2 Contestability

Affected communities need low-friction channels to challenge descriptions and exclusions, and challenges must be able to reach the pipeline, not just the individual record. A correction that fixes one mislabelled photograph while the classifier continues mislabelling ten thousand others is not contestability; it is triage. DH platforms should therefore log challenge patterns and route clusters of similar challenges to model-level review.

5.3 Plural Curation

Where material concerns identifiable communities, curation authority should be shared: community-nominated reviewers, parallel community-controlled metadata layers, and the ability to present multiple descriptive perspectives side by side rather than forcing a single canonical label. Technically, this argues for archives that store descriptions as attributed, layered claims rather than as single authoritative fields.

5.4 Periodic Human Re-Review and Drift Audits

Because algorithmic errors are systematic, sampling-based human audits should be scheduled as recurring archival practice, stratified to over-sample the languages, formats, and communities most at risk of misdescription. Preservation pipelines should additionally maintain drift ledgers: append-only records of what automated processes discarded, merged, or downgraded, sufficient to reconstruct — or at least to know — what was lost.

5.5 Appropriateness Boundaries

Finally, some curatorial judgments should not be delegated at all. Decisions about culturally sensitive material, contested historical events, and human remains or sacred objects call for deliberative, accountable human judgment. The framework in Table 1 can be used as a screening instrument: where the distinctive risk column describes harms that are severe and irreversible for the material in question, algorithmic mediation should be limited to assistive (human-in-the-loop) rather than autonomous roles.

6. Discussion

Two objections deserve engagement. The first holds that human curation is equally biased, so the comparison is unfair. We agree about the bias but not about the equivalence: the ethically relevant difference is the location and inspectability of judgment. Professional archival bias can be named, debated, and reformed because it resides in accountable institutions; model bias resides in artefacts whose provenance is often proprietary. The remedy is not to romanticise human archivists but to force algorithmic curation up to the accountability standard archives already accept for humans. The second objection is economic: safeguards cost money that under-funded cultural institutions do not have. This is precisely why the governance question belongs with funders and infrastructure providers, not individual projects. Shared open-source tooling for provenance, contestation, and audit would amortise costs across the sector, and funding calls can make such safeguards eligible — and eventually required — infrastructure expenses.

The framework also clarifies what is at stake in the current enthusiasm for generative models in archives. Systems that summarise, translate, or "chat with" collections add a fifth function — interpretation — that compounds description and retrieval risks: the user no longer sees the record, but a model's paraphrase of it. Extending the present framework to interpretive mediation is a natural next step.

7. Conclusion

Algorithmic curation is becoming the operating system of cultural memory. This paper has argued that its ethical significance lies not in any single error but in the structural relocation of memory judgments from accountable institutions into undocumented computational artefacts, producing characteristic risks of silent exclusion, contextual flattening, canon amplification, and memory drift. We proposed four governance principles — provenance transparency, contestability, plural curation, and periodic re-review — and identified boundaries beyond which delegation is inappropriate. Future work should operationalise the framework as an audit instrument and validate it across live DH projects of varying scale, and should extend the analysis to generative, interpretive mediation of archives. The question facing the digital humanities is no longer whether algorithms will curate cultural memory — they already do — but whether they will do so under archival ethics or under default engineering objectives.

References

1. Jaiswal, I. A. (2023). Multilingual and culturally adaptive AI models for global education platforms. International Journal for Research in Education (IJRE), 12(9), 17–27. https://doi.org/10.63345/ijre.v12.i9.1
2. M. S. Prashanth, R. Aluvalu, and M. V. V. Kantipudi, “Enhancing Health Product Traceability on the Blockchain: A Novel Approach for Supply Chain Management Inspection to AI,” EAI Endorsed Transactions on Pervasive Health & Technology, vol. 10, no. 1, 2024.
3. H. K. Jani, M. V. V. Prasad Kantipudi, G. Nagababu, D. Prajapati, and S. S. Kachhwaha, “Simultaneity of Wind and Solar Energy: A Spatio-Temporal Analysis to Delineate the Plausible Regions to Harness,” Sustainable Energy Technologies and Assessments, vol. 53, Art. no. 102665, 2022
4. Jaiswal, I. A. (2022). Natural language processing for security policy and log analysis. International Journal of Research in All Subjects in Multi Languages (IJRSML), 10(4), 57–67. https://doi.org/10.63345/ijrsml.v10.i4.1
5. Khan, S. (2019). Sales force security compliance: An in-depth study of GDPR, HIPAA, and PCI-DSS enforcement in cloud-based CRM systems and their implications for global enterprises. International Journal of Advanced Research in Electrical, Electronics and Instrumentation Engineering, 8(9), 2211–2219. https://doi.org/10.15662/IJAREEIE.2019.0809010
6. S. Choudhary, C. H. Kandikattu, S. Kumar, M. V. V. P. Kantipudi, and M. Kumar, “Enhancing cybersecurity through combined convolutional neural network-gated recurrent unit approach for distributed denial of service attack detection,” in Proceedings of the 2024 1st International Conference on Innovative Engineering Sciences and Technological Research (ICIESTR), pp. 1–6, 2024.

7. S. Rani, D. Ghai, and S. Kumar, "Reconstruction of wire frame model of complex images using syntactic pattern recognition," in IET Conference Proceedings CP791, vol. 2021, no. 11, pp. 8–13, 2021.
8. Swathi, S. Kumar, S. Rani, A. Jain, and R. K. M. V. N. M., "Emotion classification using feature extraction of facial expression," in Proceedings of the 2022 2nd International Conference on Technological Advancements in Computational Sciences (ICTACS), pp. 283–288, 2022.
9. S. Choudhary, S. Kumar, S. Gowroju, M. Gulhane, and R. Sri Lakshmi, Eds., Genomics at the Nexus of AI, Computer Vision, and Machine Learning. Hoboken, NJ, USA: John Wiley & Sons, 2024.
10. S. Kumar, S. Choudhary, S. Gowroju, and A. Bhola, "Convolutional neural network approach for multimodal biometric recognition system for banking sector on fusion of face and finger," in Multimodal Biometric and Machine Learning Technologies: Applications for Computer Vision, pp. 251–267, 2023.
11. S. Choudhary, S. Kumar, M. Gulhane, and M. Kumar, "Secured automated certificate creation based on multimodal biometric verification," in Multimodal Biometric and Machine Learning Technologies: Applications for Computer Vision, pp. 269–281, 2023.
12. Jaiswal, I. A. (2024). AI-powered observability and incident prediction in distributed enterprise platforms. Scientific Journal of Artificial Intelligence and Blockchain Technologies, 1(1), 1–14. https://doi.org/10.63345/sjaibt.v1.i1.201
13. S. Choudhary, C. H. Kandikattu, S. Kumar, M. V. V. P. Kantipudi, and M. Kumar, "Enhancing cybersecurity through combined convolutional neural network-gated recurrent unit approach for distributed denial of service attack detection," in Proceedings of the 2024 1st International Conference on Innovative Engineering Sciences and Technological Research (ICIESTR), pp. 1–6, 2024.
14. Jaiswal, I. A. (2023). High-Performance AI-Augmented Content Management Systems for Distributed Clouds. International Journal of Multidisciplinary Innovation and Research Methodology, ISSN: 2960-2068, 2 (2), 90–97. Retrieved from https://ijmirm. com/index. php/ijmirm/article/view, 243.
15. Khan, S. (2018). The role of zero-trust models in database security: Eliminating implicit trust and enforcing continuous verification in enterprise data access systems. International Journal of Research in Electronics and Computer Engineering, 6(1), 1622–1628.
16. Jaiswal, I. A. (2024). Self-healing REST services using artificial intelligence in multi-cloud environments. Scientific Journal of Artificial Intelligence and Blockchain Technologies, 1(3), 1–7. https://doi.org/10.63345/sjaibt.v1.i3.201
17. V. Saini, A. Jain, Anurag, and M. V. V. Prasad Kantipudi, "An Efficient Approach for Improving the Performance of Autonomous Vehicle Using Advanced Computer Vision," in International Conference on Soft Computing and Pattern Recognition, Cham, Switzerland: Springer Nature Switzerland, 2023, pp. 164–174.
18. Khan, S. (2017). The role of privacy-preserving techniques in database security: Secure multi-party computation, differential privacy, and homomorphic encryption. The Research Journal, 3(4), 29–35.
19. Khan, S. (2016). A study of data provenance and integrity in database security: Ensuring authenticity, non-repudiation, and accountability in data lifecycle management. International Journal of Research in Electronics and Computer Engineering, 4(3), 180–186.
20. Rallabandi, B. R. (2020). MEC-native 5G systems: Orchestration algorithms for ultra-low latency cloud-edge integration. International Journal of Intelligent Systems and Applications in Engineering, 8(4), 398–408. https://ijisae.org/index.php/IJISAE/article/view/7889
21. Rallabandi, B. R. (2018). Joint deployment and operational energy optimization in heterogeneous cellular networks under traffic variability. International Journal of Communication Networks and Information Security, 10(3), 638–647. https://ijcnis.org/index.php/ijcnis/article/view/8604
22. Cherladine, K., Rallabandi, B. R., Goel, A., Kaushik, K., Dhillon, A., & Soni, M. (2025). Generative adversarial defense models for simulated cyberattack scenarios in virtual networks. In 2025 International Conference on Digital Innovations for Sustainable Solutions (ICDISS) (pp. 1–6). IEEE. https://doi.org/10.1109/ICDISS68238.2025.11320686
23. C. H. Neetha and M. P. Kantipudi, "Adaptive Edge-Based Bilinear Interpolation for Smart Healthcare," in IET Conference Proceedings CP824, vol. 2022, no. 26, Stevenage, U.K.: The Institution of Engineering and Technology (IET), 2022, pp. 7–13.
24. R. Aluvalu, R. Venkatachalam, V. Uma Maheswari, R. J. Anandhi, M. V. V. Kantipudi, and S. Prashanth Mallellu, "IWHO-Based Cluster Head Selection for Vehicle-to-Vehicle Communication in Intelligent Transport System," International Journal of Transport Development and Integration, vol. 8, no. 2, pp. 154–166, 2024

25. S. Velamuri, S. K. Sudabattula, M. V. V. Prasad Kantipudi, and N. Prabaharan, “Q-Learning Based Commercial Electric Vehicles Scheduling in a Renewable Energy Dominant Distribution Systems,” Electric Power Components and Systems, pp. 1–14, 2023
26. Singh, M., Rallabandi, B., Gattupalli, K. K., & Sai Medarametla, M. (2025). Autonomous private 5G field area networks: Achieving secure, scalable, and self-healing smart grid communications. In 2025 2nd International Conference on Recent Trends in Electrical, Electronics and Computing Technologies (ICRTEECT) (pp. 1–6). IEEE. https://doi.org/10.1109/ICRTEECT67512.2025.11448712
27. Desai, A. B., Rallabandi, B., Gattupalli, K. K., & Medarametla, M. S. (2025). Agentic AI for rural connectivity and spectrum-aware networks to power smart agriculture ecosystems. In 2025 2nd International Conference on Recent Trends in Electrical, Electronics and Computing Technologies (ICRTEECT) (pp. 1–6). IEEE. https://doi.org/10.1109/ICRTEECT67512.2025.11448879
28. Goyal, S., Rallabandi, B., Gattupalli, K. K., & Medarametla, M. S. (2025). Digital twin-enabled context-aware networks: Enhancing smart grid resilience through predictive intelligence. In 2025 2nd International Conference on Recent Trends in Electrical, Electronics and Computing Technologies (ICRTEECT) (pp. 1–6). IEEE. https://doi.org/10.1109/ICRTEECT67512.2025.11448880
29. Tripathi, N., Gattupalli, K. K., Rallabandi, B., & Medarametla, M. S. (2025). AI-native Cloud-RAN orchestration for enterprise private 5G using digital twin models. In 2025 1st IEEE Uttar Pradesh Section Women in Engineering International Conference on Electrical Electronics and Computer Engineering (UPWIECON) (pp. 851–857). IEEE. https://doi.org/10.1109/UPWIECON67212.2025.11390348
30. Goyal, S., Gattupalli, K. K., Rallabandi, B., & Medarametla, M. S. (2025). Integrating terrestrial and non-terrestrial networks for reliable rural coverage in agriculture and smart grids. In 2025 1st IEEE Uttar Pradesh Section Women in Engineering International Conference on Electrical Electronics and Computer Engineering (UPWIECON) (pp. 846–850). IEEE. https://doi.org/10.1109/UPWIECON67212.2025.11390331
31. N. Pachauri, V. Suresh, M. V. V. Prasad Kantipudi, R. Alkanhel, and H. A. Abdallah, “Multi-Drug Scheduling for Chemotherapy Using Fractional Order Internal Model Controller,” Mathematics, vol. 11, no. 8, Art. no. 1779, 2023,
32. M. V. V. Prasad Kantipudi, S. Vemuri, N. S. Pradeep Kumar, S. Sreenath Kashyap, and S. Eslamian, “Proposing Model for Water Quality Analysis Based on Hyperspectral Remote Sensor Data,” in Handbook of Hydroinformatics, Elsevier, 2023, pp. 317–324.
33. Rana, M., Srinivas, S., Jamili, L. K., Jaiswal, I. A., Nakka, S., & Kasetti, S. (2025, May). Real-Time Monitoring and Prediction of Blood Sugar Levels in Diabetic Patients with Functional Models. In 2025 International Conference on Engineering, Technology & Management (ICETM) (pp. 1-6). IEEE.
34. Jaiswal, I. A. (2023). Intelligent Cybersecurity Framework for Large-Scale RESTful Service Architectures. International Journal of Research Radicals in Multidisciplinary Fields, ISSN: 2960-043X, 2 (1), 178–184. Retrieved from https://www. researchradicals. com/index. php/rr/article/view, 252.