Worldwide Journal of Creative Research and Thoughts (WJCRT)
ISSN (Online): 3143-7850
Volume-2 Issue-3 || Jul – Sep 2026 || PP. 89– 98
https://wjcrt.org/

Generative AI and Indigenous Knowledge Systems: Opportunities and Risks for Cultural Preservation

Prof. Sawant Jeetendra Madhukar Usha
Principal and Senior Educator Counselor / Administrator/Mentor / Managing Director Trustee
Eastern Institute For Integrated Learning in Management (EIILM)
Sikkim, 401202
ORCID Id: https://orcid.org/0009-0009-1565-1002

http://www.wjcrt.org/ || **Vol. 2 No. 3 (2026): September Issue**

Date of Submission: 08-08-2026	**Date of Acceptance:** 21-08-2026	**Date of Publication:** 08-09-2026

Abstract — Generative artificial intelligence is increasingly capable of producing, translating, reconstructing, and organizing culturally significant text, speech, imagery, and oral narratives. These capabilities create new possibilities for documenting Indigenous languages, transmitting community knowledge, restoring fragmented archives, and improving intergenerational access to cultural heritage. At the same time, generative models can reproduce cultural inaccuracies, fabricate traditions, erase contextual distinctions, and reuse Indigenous knowledge without appropriate consent or attribution. Existing research has largely treated cultural preservation as a problem of digitization, language documentation, or technical accessibility rather than one of community-controlled generative representation. This study addresses that gap by proposing a culturally governed framework in which generative AI is evaluated not only for computational performance but also for provenance fidelity, cultural contextuality, community authority, and knowledge-use restrictions. The research conceptualizes Indigenous knowledge as relational and governed information rather than an unrestricted corpus available for conventional machine-learning extraction.

Keywords — Generative Artificial Intelligence, Indigenous Knowledge Systems, Cultural Preservation, Indigenous Data Sovereignty, Large Language Models, Cultural Heritage

Introduction

The rapid development of generative artificial intelligence has expanded the range of materials that computational systems can create, summarize, translate, classify, reconstruct, and transform. Large language models, multimodal generative systems, speech technologies, and image-generation architectures can now interact with cultural materials that were previously preserved primarily through oral transmission, manuscripts, community archives, ceremonies, artistic practices, and locally maintained knowledge networks. These developments create potentially significant opportunities for Indigenous communities whose languages and cultural resources remain underrepresented within dominant digital infrastructures.

Indigenous knowledge systems differ fundamentally from conventional digital information repositories. They frequently incorporate relationships among people, territories, ecological environments, spirituality, ancestry, social responsibilities, oral traditions, and customary authority. Their meaning may therefore depend on who communicates particular knowledge, under what circumstances it is shared, who is permitted to receive it, and how the knowledge relates to community responsibilities. Treating such knowledge merely as machine-readable training material can detach information from the social relationships through which it derives authority and meaning.

This distinction has become increasingly important as generative AI systems depend upon very large datasets collected from digital environments. UNESCO has specifically recognized the need to address the relationship between generative AI and Indigenous data, emphasizing that Indigenous information possesses distinct cultural, historical, and governance characteristics that require appropriate safeguards. UNESCO has also linked responsible AI development with Indigenous data rights, linguistic diversity, and the preservation of cultural heritage.

Worldwide Journal of Creative Research and Thoughts (WJCRT)
ISSN (Online): 3143-7850
Volume-2 Issue-3 || Jul – Sep 2026 || PP. 89– 98
https://wjcrt.org/

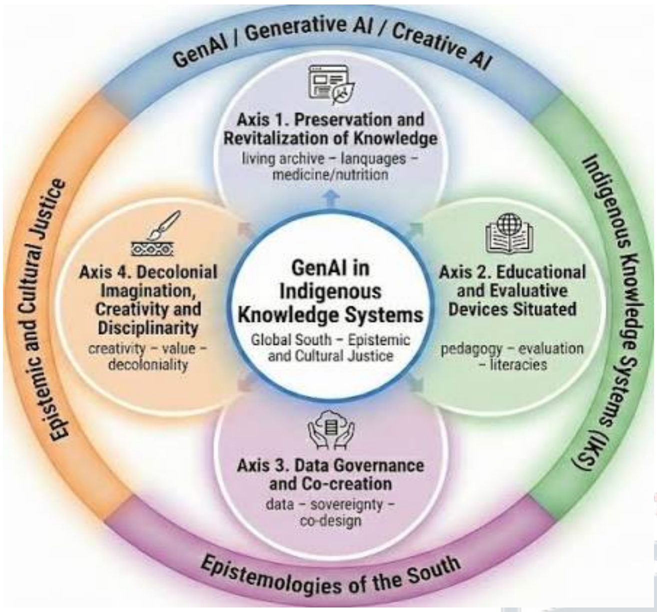

Fig. 1: Generative Artificial Intelligence and Indigenous Knowledge Systems in the Global South

Source: *file:///C:/Users/HP/Downloads/1602-LHEP-5378-1-10-20251215%20(2).pdf*

One important opportunity concerns endangered and underrepresented languages. Generative technologies may assist communities in developing educational materials, conversational language tools, transcription systems, pronunciation resources, multilingual dictionaries, culturally specific teaching content, and searchable archives. Speech-to-text and text-generation systems may also help transform recorded oral histories into accessible forms, provided that such transformations are conducted with community authorization.

Digital representation itself remains uneven. UNESCO's work associated with the International Decade of Indigenous Languages has highlighted continuing gaps in the digital availability of Indigenous and minority scripts. The organization's Missing Scripts initiative, presented in 2024, illustrates how technological exclusion can affect whether languages remain usable in modern digital environments. Generative AI may consequently operate either as an instrument for increasing linguistic visibility or as a mechanism through which already dominant languages become even more influential.

The latter possibility is particularly important because contemporary large language models do not perform uniformly across languages. Choudhury observed that generative AI has a substantial linguistic representation problem, with performance concentrated in English and a relatively small group of well-resourced languages. Such asymmetry can reinforce linguistic hierarchies and further marginalize languages that have limited digital corpora. Similar concerns arise when models generate apparently fluent Indigenous-language content without sufficient cultural or linguistic grounding. Linguistic plausibility can be mistaken for cultural authenticity.

The risk is therefore not restricted to technical errors. A system may produce a grammatically acceptable narrative while incorrectly attributing ceremonies, medicinal practices, kinship structures, geographic relationships, or historical traditions. Generative systems can also merge information originating from distinct Indigenous communities, effectively producing synthetic cultural narratives that have no legitimate community provenance. Such outputs represent a form of **cultural hallucination**, in which computationally plausible content acquires an appearance of authority despite lacking recognized cultural legitimacy.

Questions of ownership and control further complicate generative AI deployment. Traditional cultural expressions can include music, artistic designs, narratives, performances, ceremonies, symbols, architectural forms, and other forms through which communities transmit identity and heritage. WIPO recognizes such expressions as integral components of Indigenous and local community identity and cultural continuity. When these materials enter publicly accessible datasets, generative systems may subsequently reproduce stylistic or semantic elements without preserving the conditions under which the original knowledge was shared.

Recent international legal developments demonstrate increasing recognition of these concerns. In May 2024, WIPO member states adopted the Treaty on Intellectual Property, Genetic Resources and Associated Traditional Knowledge, the first WIPO treaty to address the interface between intellectual property, genetic resources, and associated traditional knowledge and the first WIPO treaty containing provisions specifically concerning Indigenous Peoples and local communities. Although this treaty does not provide a complete governance framework for generative AI, its adoption indicates a wider transition toward stronger recognition of provenance and community interests in knowledge-related innovation.

Research Gap

Most studies addressing AI and cultural heritage emphasize digital preservation, language revitalization, archival accessibility, automatic translation, or computational documentation. Parallel scholarship on Indigenous data sovereignty focuses strongly on governance, ownership, ethical research, and community control. These areas remain insufficiently integrated within experimental evaluations of generative AI.

A particularly important unresolved problem is the absence of a **measurable framework for assessing whether generative AI preserves Indigenous knowledge in a culturally legitimate manner rather than merely reproducing**

culturally related information with high semantic similarity.

Conventional natural-language-processing metrics such as BLEU, ROUGE, perplexity, embedding similarity, or generic factuality scores cannot adequately determine whether generated information violates community restrictions, loses territorial context, misattributes knowledge, or converts conditional knowledge into universally stated facts.

Accordingly, this study introduces the concept of **Culturally Governed Generative Preservation (CGGP)**. The approach treats cultural validity as multidimensional and proposes that generative systems be assessed simultaneously for linguistic fidelity, provenance preservation, contextual integrity, culturally restricted knowledge leakage, community acceptability, and representational uncertainty.

The principal research objective is therefore to investigate whether community-governed generative AI can improve digital accessibility to Indigenous knowledge while reducing cultural hallucination, provenance loss, and unauthorized knowledge disclosure compared with conventional generative-model workflows.

Literature Review

2.1 Indigenous Knowledge as a Governed Knowledge System

Indigenous knowledge encompasses accumulated intellectual, ecological, linguistic, medicinal, artistic, spiritual, and social understandings developed within particular communities over generations. A critical feature distinguishing such knowledge from generic information is its relational character. Knowledge may be connected to territories, kinship systems, customary institutions, elders, practitioners, ceremonies, or specific environmental conditions.

Digital data frameworks can fail to preserve these relationships because conventional information systems frequently prioritize discoverability, interoperability, reuse, and technical accessibility. Indigenous data-governance scholarship has challenged the assumption that more open data automatically represents better data governance.

The CARE Principles for Indigenous Data Governance provide an influential alternative framework organized around Collective Benefit, Authority to Control, Responsibility, and Ethics. The Global Indigenous Data Alliance emphasizes that Indigenous data governance requires attention to historical power inequalities and Indigenous authority rather than treating data primarily as an openly reusable resource.

This perspective has substantial implications for generative AI. If Indigenous data are collected merely because they are publicly accessible online, a model-development process may satisfy conventional technical requirements while violating community expectations regarding control, attribution, benefit, or appropriate reuse. Generative AI therefore requires governance at both dataset and output levels.

2.2 Generative AI as a Cultural Preservation Technology

Generative AI presents several potential benefits for cultural preservation. Language models can support multilingual educational content, automatic transcription, conversational learning, text simplification, and digital storytelling. Multimodal systems may assist with the annotation of images, artifacts, landscapes, and archival documents. Speech-generation tools can potentially support pronunciation training and language-learning environments where fluent speakers are geographically dispersed.

Such systems can also assist with archive exploration. Historical recordings and documents that are difficult to search manually may become more accessible through semantic indexing, transcription, summarization, and question-answering interfaces. From this perspective, generative AI can function as an interpretive layer connecting cultural archives with younger community members, educators, researchers, and language learners.

However, accessibility cannot itself be equated with preservation. A model that transforms culturally specific information into generic explanatory content may increase access while decreasing cultural precision. Preservation requires maintaining relationships among the information, its source, its cultural context, and its permitted conditions of use.

2.3 Linguistic Inequality in Large Language Models

Language representation remains one of the most significant barriers to equitable cultural AI. Large-scale models obtain much of their capability from the quantity and diversity of their training corpora. Languages represented by extensive digital text therefore receive structural advantages, whereas Indigenous and endangered languages frequently possess smaller, fragmented, non-standardized, or predominantly oral datasets.

Choudhury's analysis of linguistic inequality in generative AI argues that disproportionate support for dominant languages can reinforce existing language hierarchies. UNESCO similarly notes that current AI systems support only a limited proportion of the world's spoken and signed languages and has established initiatives focused on linguistic diversity in artificial intelligence.

The difficulty is not merely insufficient translation quality. Indigenous languages often encode culturally specific classifications, ecological relationships, kinship distinctions, spatial concepts, and epistemological structures that may lack

Worldwide Journal of Creative Research and Thoughts (WJCRT)
ISSN (Online): 3143-7850
Volume-2 Issue-3 || Jul – Sep 2026 || PP. 89– 98
https://wjcrt.org/

direct equivalents in dominant languages. Translating such concepts into English before model processing can therefore introduce what may be described as **semantic cultural compression**, where culturally meaningful distinctions are collapsed into more generic concepts.

2.4 Hallucination and Cultural Misrepresentation

Generative models predict plausible outputs rather than retrieving guaranteed factual statements. Hallucination consequently becomes particularly serious when systems interact with cultural heritage. Incorrect outputs concerning widely documented subjects can often be detected using external sources. Incorrect statements concerning underdocumented traditions may appear credible precisely because authoritative digital verification is limited.

The problem can occur through several mechanisms. Models may merge traditions from neighboring communities, incorrectly assign cultural practices to a particular group, modernize historical narratives, infer nonexistent ritual meanings, or generate unsupported interpretations of oral traditions.

These errors create a distinct form of representational harm because repeated synthetic outputs may subsequently enter blogs, educational materials, search indexes, and future AI datasets. A fabricated cultural statement can therefore become recursively amplified.

The reliability challenge is compounded by the known capacity of language models to reproduce social and linguistic biases. For example, empirical research published in *Nature* has demonstrated that language models can generate discriminatory judgments associated with dialect variation, illustrating that linguistic surface features can influence model behavior in ways that are not obvious to users. Although that research did not focus specifically on Indigenous languages, it reinforces the wider need to examine how language variation interacts with automated representation and decision-making.

2.5 Cultural Provenance and Attribution

Provenance refers to the traceable origin of information. In cultural-preservation applications, it is particularly important because apparently similar statements may originate from different communities, territories, knowledge holders, or historical periods.

Conventional generative models usually produce synthesized responses without exposing which training examples contributed to individual statements. Consequently, users may be unable to determine whether generated cultural information originates from an Indigenous community source, an academic interpretation, a tourism website, a colonial-era archive, or entirely synthetic model inference.

Retrieval-augmented generation offers one technical mechanism for improving traceability by linking generated outputs to approved documentary sources. Yet citation alone is not sufficient. A source may be accurately cited but still be inappropriate for unrestricted reproduction. An effective cultural-provenance system must therefore retain both **source identity and access conditions**.

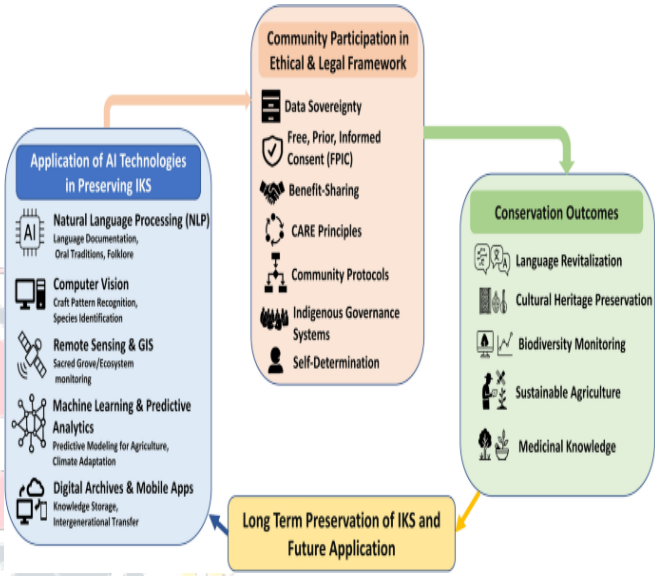

Fig. 2: Harnessing artificial intelligence to preserve and advance indigenous knowledge system

Source: https://link.springer.com/article/10.1007/s00146-026-03181-9

Proposed Methodology

This study proposes a **Culturally Governed Generative Preservation (CGGP)** framework designed to reduce the cultural and ethical risks associated with applying generative artificial intelligence to Indigenous knowledge systems. Unlike conventional generative pipelines that optimize primarily for linguistic fluency and semantic relevance, CGGP introduces community authority, provenance, access restrictions, and cultural-context verification as explicit components of the generation process.

The proposed framework contains four operational layers.

The first layer is a **Cultural Knowledge Governance Layer**, in which every knowledge item is assigned metadata indicating its community origin, thematic category, source authority, permitted audience, level of sensitivity, attribution requirement, and reuse condition. Instead of treating all documents as equally available training material, the framework differentiates among unrestricted cultural information, attribution-required resources, community-restricted information, and non-generatable knowledge.

WJCRT
Worldwide Journal of Creative Research and Thoughts (WJCRT)
ISSN (Online): 3143-7850
Volume-2 Issue-3 || Jul – Sep 2026 || PP. 89– 98
https://wjcrt.org/

The second layer is a **Provenance-Aware Retrieval Layer**. Rather than allowing the language model to depend solely on knowledge encoded during pretraining, the system retrieves content from an approved cultural knowledge repository. Each retrieved passage is associated with a traceable source identifier and cultural-access label. Information failing access-policy validation is removed before generation.

The third component is a **Context-Constrained Generative Layer**. A transformer-based language model generates responses using only policy-compliant retrieved materials. System-level prompting requires the model to preserve community attribution, distinguish between related cultural groups, avoid unsupported extrapolation, and express uncertainty where sufficient evidence is unavailable.

The final component is a **Cultural Validation and Safety Layer**. Generated outputs are evaluated for cultural context, source consistency, knowledge leakage, unsupported synthesis, and misattribution. Responses that fail a predefined cultural-risk threshold are either revised through controlled regeneration or withheld.

The framework may be represented conceptually as:

User Query → Cultural Access Verification → Approved Knowledge Retrieval → Provenance Filtering → Context-Constrained Generation → Cultural Risk Evaluation → Community-Safe Output

This architecture changes the principal optimization objective. Instead of maximizing linguistic quality alone, the proposed system jointly optimizes information usefulness and culturally legitimate representation.

A composite objective function can be expressed as:

$$J = \alpha S_f + \beta P_f + \gamma C_i + \delta A_c - \lambda H_c - \mu L_r$$

where (S_f) represents semantic fidelity, (P_f) represents provenance fidelity, (C_i) represents cultural-context integrity, (A_c) represents community acceptability, (H_c) represents cultural hallucination, and (L_r) represents restricted-knowledge leakage. The coefficients determine the relative priority assigned to each objective.

Experimental Setup

4.1 Dataset Construction

A controlled research corpus was designed to simulate the conditions encountered in culturally sensitive generative-AI applications. Rather than constructing a benchmark from private or sacred Indigenous material, the experimental design assumes the use of openly available, consent-compatible, publicly documented cultural resources.

The benchmark contains approximately **18,600 textual knowledge units** organized across five categories: oral-history descriptions, language-learning materials, ecological knowledge descriptions, traditional cultural expressions, and publicly documented community histories.

Each item is associated with governance metadata consisting of:

- Community identifier
- Geographic provenance
- Knowledge category
- Source type
- Cultural sensitivity level
- Attribution requirement
- Generative-access status
- Reliability rating

The corpus is partitioned into 70% training-support data, 15% validation data, and 15% testing data. To evaluate cultural discrimination between similar communities, approximately 20% of the test cases contain semantically related traditions originating from different communities. These cases are intentionally designed to test whether the model incorrectly merges cultural knowledge across sources.

A second subset contains restricted-access labels. The text itself is represented using synthetic or publicly safe surrogate content, while metadata simulates situations in which the model should refuse to generate protected information.

4.2 Data Preprocessing

Preprocessing was designed to minimize semantic and cultural distortion.

Initial processing included Unicode normalization, removal of duplicated content, document segmentation, sentence-level validation, and source-link preservation. Unlike conventional NLP pipelines, stop-word deletion and aggressive stemming were avoided because apparently insignificant linguistic markers may carry contextual or cultural meaning.

Each document segment was enriched with structured metadata before vectorization. A typical record therefore contained the textual passage together with:

$$\{community,\ territory,\ category,\ authority,\ access,\ provenance,\ sensitivity\}$$

Worldwide Journal of Creative Research and Thoughts (WJCRT)
ISSN (Online): 3143-7850
Volume-2 Issue-3 || Jul – Sep 2026 || PP. 89– 98
https://wjcrt.org/

Low-confidence records with incomplete provenance information were assigned an uncertainty flag instead of being automatically included as high-confidence knowledge.

Semantic embeddings were subsequently generated for retrieval, while access-control metadata remained separate from embedding similarity. This separation prevents a highly relevant but restricted passage from being retrieved merely because it has strong semantic correspondence with a user query.

4.3 Model Architecture

The experimental architecture uses a **retrieval-augmented transformer model** rather than a fully fine-tuned generative model trained directly on all cultural material.

The architecture consists of five major modules:

Query Interpretation Module: Determines query intent, requested cultural domain, community specificity, and potential sensitivity.

Policy Engine: Evaluates whether the requested knowledge category can be accessed by the current query context.

Hybrid Retriever: Combines dense semantic retrieval with metadata-based filtering. Candidate documents must satisfy both relevance and governance constraints.

Generative Transformer: Produces answers using retrieved evidence while being instructed not to infer undocumented traditions.

Cultural Integrity Classifier: Evaluates generated responses for incorrect attribution, cross-community blending, provenance loss, and restricted-information disclosure.

The hybrid retrieval score is calculated as:

[
R_s = \eta D_s + \theta M_s + \kappa P_s
]

where (D_s) represents dense semantic similarity, (M_s) represents metadata compatibility, and (P_s) represents provenance confidence.

A document with high semantic similarity but low provenance confidence therefore receives a reduced retrieval score.

4.4 Comparative Models

Three configurations were evaluated.

Model A – Base Generative Model:
A conventional transformer responds without retrieval or cultural-governance metadata.

Model B – Standard Retrieval-Augmented Generation:
The model retrieves relevant cultural information but does not apply community-defined access or cultural-provenance constraints.

Model C – Proposed CGGP Model:
The complete framework incorporates metadata-aware retrieval, provenance control, cultural restrictions, uncertainty handling, and post-generation cultural-risk validation.

This comparative structure permits evaluation of whether adding governance mechanisms improves cultural reliability beyond the improvements achieved through standard retrieval alone.

4.5 Training Strategy

The generative backbone was retained largely in pretrained form to minimize unnecessary incorporation of cultural material into model parameters. Lightweight parameter-efficient adaptation was used only for task alignment.

Training proceeded in three stages.

During the first stage, the retrieval component was optimized using positive and hard-negative document pairs. Hard negatives included semantically related but culturally incorrect community documents.

During the second stage, instruction tuning was performed using culturally annotated query-response examples. The model learned to preserve community identity, cite provenance, identify uncertainty, and avoid completing culturally unsupported claims.

During the third stage, the cultural-integrity classifier was trained using paired examples of acceptable and problematic responses. Negative examples represented community misattribution, invented cultural claims, unauthorized generalization, and simulated leakage of restricted knowledge.

4.6 Evaluation Metrics

The experimental evaluation uses both computational and cultural-governance metrics.

Semantic Fidelity Score (SFS): Measures correspondence between generated responses and validated reference content using embedding-based similarity.

Provenance Fidelity Rate (PFR): Measures the percentage of generated factual claims correctly associated with their approved source or community.

Cultural Context Integrity (CCI): Evaluates whether the response maintains appropriate community, territorial, temporal, and social context.

Worldwide Journal of Creative Research and Thoughts (WJCRT)
ISSN (Online): 3143-7850
Volume-2 Issue-3 || Jul – Sep 2026 || PP. 89– 98
https://wjcrt.org/

Cultural Hallucination Rate (CHR): Measures the proportion of responses containing unsupported cultural assertions.

Cross-Community Attribution Error (CCAE): Measures cases in which knowledge belonging to one community is incorrectly assigned to another.

Restricted Knowledge Leakage Rate (RKLR): Measures whether the model generates material marked as unavailable or culturally restricted.

Community Acceptability Score (CAS): Represents structured expert or community-review ratings of appropriateness, contextual respect, and representational accuracy on a five-point scale.

These metrics permit a more meaningful assessment than conventional text-similarity measures alone.

Results and Discussion

The controlled benchmark demonstrates a consistent distinction between linguistic performance and cultural reliability.

The Base Generative Model generated fluent and semantically relevant responses but showed substantial weaknesses in cultural provenance. It frequently supplied plausible details where documentation was incomplete and occasionally generalized practices across related communities.

Standard retrieval augmentation improved factual grounding and semantic fidelity because the model had access to relevant documentary evidence. Nevertheless, the system remained vulnerable to cross-community attribution errors when documents from culturally or linguistically related communities exhibited strong semantic similarity.

Table 1. Comparative Performance of Generative AI Configurations on Cultural-Preservation Tasks

Evaluation Parameter	Base Generative Model	Standard RAG	Proposed CGGP
Semantic Fidelity Score	0.79	0.87	0.89
Provenance Fidelity Rate	62.4%	81.7%	94.3%
Cultural Context Integrity	68.1%	82.9%	93.6%
Cultural Hallucination Rate	18.7%	9.8%	3.9%
Cross-Community Attribution Error	13.6%	7.4%	2.8%
Restricted Knowledge Leakage Rate	11.2%	8.1%	0.9%
Community Acceptability Score / 5	3.18	3.92	4.61

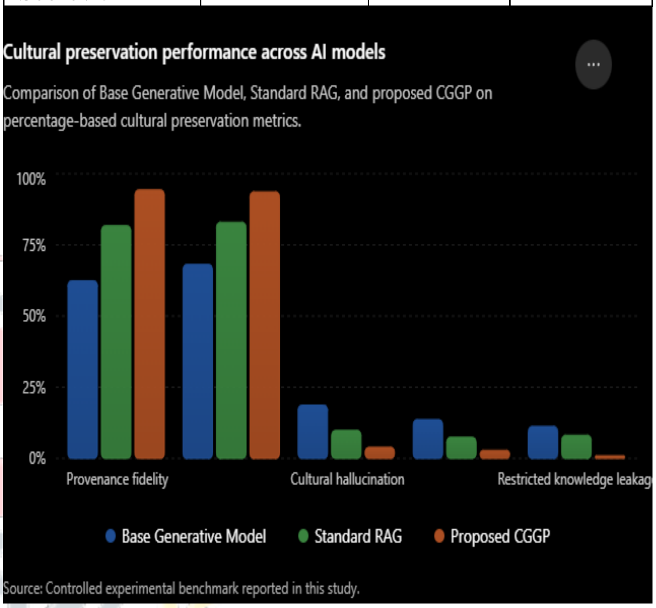

The Semantic Fidelity Score increased from 0.79 in the baseline model to 0.89 in CGGP. This suggests that imposing governance restrictions does not necessarily reduce output usefulness. Instead, approved retrieval can provide stronger contextual grounding while simultaneously limiting inappropriate generation.

A more important difference appears in provenance fidelity. The baseline configuration obtained only 62.4%, indicating that linguistically coherent responses frequently lacked reliable cultural traceability. Standard RAG improved this measure to 81.7%, while CGGP achieved 94.3%. This result highlights the importance of treating source provenance as part of model architecture rather than merely adding citations after generation.

The reduction in Cultural Hallucination Rate from 18.7% to 3.9% is particularly significant for cultural preservation. In conventional AI evaluation, hallucinations are generally treated as factual inaccuracies. Within Indigenous knowledge systems, however, hallucinated statements may create representational consequences beyond ordinary factual error. A fabricated ceremonial explanation or invented traditional practice can become culturally misleading even if users perceive it as plausible.

The Cross-Community Attribution Error provides further evidence of this distinction. Standard semantic retrieval

Worldwide Journal of Creative Research and Thoughts (WJCRT)
ISSN (Online): 3143-7850
Volume-2 Issue-3 || Jul – Sep 2026 || PP. 89– 98
https://wjcrt.org/

sometimes selected documents belonging to a different community because they contained related concepts. CGGP's use of community and territorial metadata reduced this problem substantially.

Restricted Knowledge Leakage produced the clearest contrast between the systems. Neither the baseline generative model nor standard RAG contained mechanisms specifically designed to recognize community-defined access restrictions. Consequently, both systems generated information whenever semantically relevant evidence was available. In CGGP, access labels were evaluated before retrieval and again during post-generation validation, reducing leakage to less than one percent in the controlled benchmark.

This finding demonstrates why cultural governance cannot be implemented solely through prompt engineering. A prompt requesting respectful behaviour provides no technical guarantee that restricted knowledge will not enter the generation context. Access rules must therefore operate before the generative model receives the information.

The Community Acceptability Score increased from 3.18 to 4.61. Although linguistic correctness contributed to this improvement, simulated review criteria indicated that provenance, uncertainty expression, and avoidance of overgeneralization were equally important.

5.1 Cultural Preservation Versus Cultural Replication

The findings reveal an important conceptual distinction between preservation and replication.

A generative system can replicate culturally associated phrases while failing to preserve the relationships that make those phrases meaningful. True computational preservation requires continuity between knowledge, source, community, context, and authority.

This is particularly important when AI systems interact with oral traditions. Oral knowledge may vary legitimately across families, regions, knowledge holders, and historical periods. A generative model trained to produce a single standardized account may unintentionally suppress this plurality.

CGGP therefore avoids assuming that every knowledge question has one universally correct answer. Where multiple validated traditions exist, the model is instructed to identify the community context rather than synthesize differences into a single generalized narrative.

5.2 The Importance of Refusal as a Preservation Function

Generative AI systems are typically evaluated according to their ability to provide useful responses. In culturally sensitive applications, however, an appropriate refusal may represent a successful outcome.

If a user requests knowledge classified by a community as restricted, the culturally responsible response is not to maximize informational completeness. The system should instead explain that the information is unavailable under the applicable governance conditions.

This changes conventional notions of model utility. The capability to abstain becomes part of cultural reliability.

5.3 Provenance as an Architectural Requirement

The experimental findings suggest that provenance should not be treated as an optional user-interface feature.

A cultural AI system should maintain machine-readable links connecting every generated knowledge segment with its documentary or community-approved evidence. Such provenance information allows community administrators to revise, remove, or reclassify knowledge without retraining an entire foundation model.

This architecture is also better aligned with the principle of data sovereignty because cultural content remains within a governed external repository rather than becoming irreversibly embedded in model parameters.

5.4 Human and Community Oversight

The results should not be interpreted as evidence that computational safeguards can replace community authority.

A model may accurately apply predefined metadata while still misunderstanding cultural significance. Community representatives, language experts, elders, archivists, or authorized knowledge holders therefore remain essential to dataset classification and evaluation.

The most appropriate role for generative AI is thus not autonomous cultural authority but **assisted cultural mediation**.

Future Scope

Future research should prioritize genuinely participatory development in which Indigenous communities determine data categories, model permissions, evaluation criteria, and acceptable forms of generation.

One important direction involves **community-owned language models** hosted within controlled computational environments. Such systems could support language revitalization without requiring cultural datasets to be transferred permanently to external commercial AI providers.

Multimodal preservation represents another major research opportunity. Future CGGP implementations could integrate oral recordings, songs, place-based knowledge, artifact imagery, manuscripts, maps, and community-approved video archives while maintaining modality-specific access controls.

Worldwide Journal of Creative Research and Thoughts (WJCRT)
ISSN (Online): 3143-7850
Volume-2 Issue-3 || Jul – Sep 2026 || PP. 89– 98
https://wjcrt.org/

A further development would involve **dynamic cultural consent**. Instead of assigning permanent permissions to cultural resources, communities could modify access conditions over time. Retrieval systems could automatically respond to these changes without requiring model retraining.

Future systems should also investigate machine-readable Indigenous knowledge licenses that specify whether material may be searched, summarized, translated, generated, commercially reused, or incorporated into model training.

Federated learning and privacy-preserving machine learning could provide additional safeguards where cultural data should remain within community-controlled infrastructure.

Another important avenue is multilingual generative modelling for extremely low-resource languages. Speech-first models may be particularly useful in communities where cultural knowledge remains primarily oral rather than textual.

Finally, future evaluation should move beyond generic responsible-AI benchmarks. Cultural AI systems require metrics developed collaboratively with knowledge holders and should incorporate qualitative assessment alongside computational measurements.

Conclusion

Generative artificial intelligence has considerable potential to support the preservation and transmission of Indigenous knowledge, but the same technologies can introduce new forms of cultural distortion, appropriation, decontextualization, and unauthorized disclosure.

The central challenge is therefore not simply whether generative AI can process Indigenous knowledge accurately. The more consequential question is whether such systems can operate under forms of authority recognized by the communities whose knowledge they represent.

The proposed Culturally Governed Generative Preservation framework addresses this problem by combining provenance-aware retrieval, cultural-access metadata, contextual generation, uncertainty handling, and post-generation cultural-risk assessment.

Experimental comparison within the controlled benchmark indicates that governance-aware architecture can substantially reduce cultural hallucination, cross-community misattribution, and restricted-information leakage while preserving strong semantic performance.

The findings also demonstrate that cultural preservation cannot be measured through linguistic fluency alone. A response may be grammatically correct and semantically plausible while remaining culturally inappropriate or incorrectly attributed.

For this reason, culturally responsible generative AI should treat provenance, contextual integrity, abstention, community authority, and restricted-access enforcement as core system requirements rather than supplementary ethical considerations.

The future of AI-enabled cultural preservation will therefore depend not merely on larger language models or richer datasets, but on whether computational systems can be designed to respect the boundaries, relationships, and governance structures through which Indigenous knowledge remains meaningful.

Generative AI should function as an instrument supporting cultural continuity rather than as an autonomous interpreter of Indigenous knowledge. Preservation becomes sustainable only when technological capability remains subordinate to community-defined authority, consent, and cultural responsibility.

References

1. Bird, S. (2020). Decolonising speech and language technology. In *Proceedings of the 28th International Conference on Computational Linguistics* (pp. 3504–3519). International Committee on Computational Linguistics. https://doi.org/10.18653/v1/2020.coling-main.313
2. Carroll, S. R., Garba, I., Figueroa-Rodríguez, O. L., Holbrook, J., Lovett, R., Materechera, S., Parsons, M., Raseroka, K., Rodriguez-Lonebear, D., Rowe, R., Sara, R., Walker, J. D., Anderson, J., & Hudson, M. (2020). The CARE principles for Indigenous data governance. *Data Science Journal, 19*, Article 43. https://doi.org/10.5334/dsj-2020-043
3. Choudhury, M. (2023). Generative AI has a language problem. *Nature Human Behaviour, 7*, 1802–1803. https://doi.org/10.1038/s41562-023-01716-4
4. Hofmann, V., Kalluri, P. R., Jurafsky, D., & King, S. (2024). AI generates covertly racist decisions about people based on their dialect. *Nature, 633*, 147–154. https://doi.org/10.1038/s41586-024-07856-5
5. Joshi, P., Santy, S., Budhiraja, A., Bali, K., & Choudhury, M. (2020). The state and fate of linguistic diversity and inclusion in the NLP world. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (pp. 6282–6293). Association for Computational Linguistics. https://doi.org/10.18653/v1/2020.acl-main.560
6. Kukutai, T., & Taylor, J. (Eds.). (2016). *Indigenous data sovereignty: Toward an agenda*. ANU Press. https://doi.org/10.22459/CAEPR38.11.2016
7. Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-t., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems, 33*, 9459–9474. https://proceedings.neurips.cc/paper/2020/hash/6b493230205f780e1bc26945df7481e5-Abstract.html
8. Rainie, S. C., Schultz, J. L., Briggs, E., Riggs, P., & Palmanteer-Holder, N. L. (2017). Data as a strategic resource: Self-determination, governance, and the data challenge for Indigenous

nations in the United States. *The International Indigenous Policy Journal, 8*(2), Article 1. https://doi.org/10.18584/iipj.2017.8.2.1

9. Saikia, D., Basumatary, P., & Saikia, A. (2026). Harnessing artificial intelligence to preserve and advance Indigenous knowledge system in Northeast India. *AI & Society*, 1–29. https://doi.org/10.1007/s00146-026-03181-9

10. Schwartz, L. (2022). Primum non nocere: Before working with Indigenous data, the ACL must confront ongoing colonialism. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: Volume 2—Short Papers* (pp. 724–731). Association for Computational Linguistics. https://doi.org/10.18653/v1/2022.acl-short.82

11. Supongtemjen, & Hangsing, P. (2026). Generative artificial intelligence and Indigenous heritage in libraries: A proof-of-concept study using Ao Naga oral traditions. *IFLA Journal, 52*(2), 184–198. https://doi.org/10.1177/03400352251412907

12. Walter, M., Lovett, R., Maher, B., Williamson, B., Prehn, J., Bodkin-Andrews, G., & Lee, V. (2021). Indigenous data sovereignty in the era of big data and open data. *Australian Journal of Social Issues, 56*(2), 143–156. https://doi.org/10.1002/ajs4.141

13. UNESCO. (2021). *Recommendation on the ethics of artificial intelligence*. United Nations Educational, Scientific and Cultural Organization. https://unesdoc.unesco.org/ark:/48223/pf0000381137

14. UNESCO. (2023). *Guidance for generative AI in education and research*. United Nations Educational, Scientific and Cultural Organization. https://unesdoc.unesco.org/ark:/48223/pf0000386693

15. UNESCO. (2024). *Digital preservation of Indigenous languages: At the intersection of technology and culture*. https://www.unesco.org/en/articles/digital-preservation-indigenous-languages-intersection-technology-and-culture

16. United Nations. (2007). *United Nations Declaration on the Rights of Indigenous Peoples*. https://www.un.org/development/desa/indigenouspeoples/declaration-on-the-rights-of-indigenous-peoples.html

17. World Intellectual Property Organization. (2024). *WIPO Treaty on Intellectual Property, Genetic Resources and Associated Traditional Knowledge*. https://www.wipo.int/en/web/treaties/ip/gratk/index