



AI-ENABLED CULTURAL HERITAGE PRESERVATION AND A BROWSER-BASED VIRTUAL MUSEUM: DESIGN AND IMPLEMENTATION OF A KENYA COASTAL HERITAGE PLATFORM

Anwar Hood Ahmed ^{1,*}, Mathew Egessa ¹, Patrick Oginga ², Collins Becky ² and Titus Nyandoro ²

¹ Technical University of Mombasa, School of Business, Mombasa, Kenya.

² Kenya Coastal Heritage Digital Platform Technical Team, Kenya.

* Corresponding Author

Open Access Research Journal of Science and Technology, 2026, 18(01), 046–058

Publication history: Received on 15 August 2026; revised on 27 September 2026; accepted on 29 September 2026

Article DOI: <https://doi.org/10.53022/oarjst.2026.18.1.0095>

Copyright © 2026 Author(s) retain the copyright of this article. This article is published under the terms of the Creative Commons Attribution License 4.0

ABSTRACT

Coastal Kenya holds a wide-ranging cultural heritage expressed in material artefacts, oral traditions, indigenous and local languages, historical knowledge, and everyday community practices. Yet these resources are increasingly strained by modernisation and global circulation, shifting modes of cultural transmission, and gaps in systematic documentation. This study reports the design and implementation of the Kenya Coastal Heritage Digital Platform produced under the Leveraging Artificial Intelligence for the Preservation and Promotion of Coastal Cultural Heritage in Kenya (LAIPCH-K) project. The study examines how the project's initially AI-focused aims were converted into a connected digital heritage system comprising a public heritage repository, a curator-mediated curation dashboard, and a browser-based virtual museum. Structured metadata, verified artefact records, multilingual resources, storytelling and oral-history materials, interactive 3D artefacts, and community-oriented governance arrangements are integrated into a single workflow. The virtual museum combines Three.js, WebGL, WebXR, Web Audio, and Web Speech, with procedurally generated Swahili-inspired architectural elements, first-person navigation, on-demand loading of validated 3D artefacts, and an animated multilingual cultural guide named Amina. The study follows a digital-heritage design and implementation pathway grounded in the LAIPCH-K proposal, project documentation, engineering specifications, and prototype development. The results demonstrate an operational pipeline linking heritage contribution to curation and validation, metadata management, publication, and immersive presentation. The evidence does not, however, establish visitor learning outcomes, community satisfaction, AI-model performance, or broader socioeconomic effects. The study argues that durable digital heritage preservation depends on aligning technical systems with metadata practice, provenance, cultural validation, multilingual representation, accessibility, governance, and sustained community participation.

Keywords: Digital Heritage; Artificial Intelligence; Virtual Museum; Coastal Kenya; Language Preservation; Community Participation

1 INTRODUCTION

Kenya's coastal region contains culturally diverse heritage expressed not only through material artefacts but also through oral traditions, indigenous and local languages, historical knowledge, artistic production, and community-based cultural practice. Heritage in this setting is sustained through transmission and interpretation as much as through physical custody, meaning that participation by communities is central to what is preserved and how it remains

meaningful. In this respect, the UNESCO Convention for the Safeguarding of the Intangible Cultural Heritage frames safeguarding as a process that should involve the communities, groups, and individuals who create, maintain, and transmit heritage, while supporting their active role in its management [1].

Preservation difficulties become sharper where knowledge circulates orally, or through practices that are not well represented in conventional archives. In the original LAIPCH-K proposal, pressures linked to globalisation, modernisation, urbanisation, and socioeconomic change were described as affecting indigenous knowledge, languages, oral histories, and cultural artefacts, with Bajuni and Pokomo heritage identified as priority areas for documentation and revitalisation. The present work converts that broad aim into a concrete platform that combines an artefact repository and curation dashboard with multilingual resources, storytelling functionality, and a browser-based virtual museum.

Digital technologies increasingly shape how cultural heritage is documented and encountered, and this forms a key context for the platform. UNESCO's 2024–2026 work on culture and digital technologies recognises digital mapping, online museums, 3D modelling, and related tools as expanding routes for documentation, access, and experience. Yet UNESCO also notes that digitisation brings governance, ethical, intellectual-property, and cultural-diversity issues that have to be addressed alongside the technical build. [2]

Museum scholarship reflects a similar shift, moving beyond straightforward digitisation toward connected digital exhibition environments. A systematic review on technologies used in museum exhibition transformation reports growing use of artificial intelligence, immersive systems, cloud computing, and related components for exhibition design, accessibility, education, and visitor experience [3]. Their synthesis also points to a persistent requirement, namely, technologies need to be fitted to specific museum and heritage-use settings rather than treated as an objective in themselves.

For the Kenya Coastal Heritage Digital Platform, this concern is practical rather than rhetorical. Digitisation is not approached as the mere creation of images or 3D objects; instead, the implementation ties together contribution, metadata production, cultural validation, publication, and immersive interpretation. The resulting architecture is organised as three main subsystems, a public heritage website, a back-office curation dashboard, and a browser-based virtual museum. All three rely on a shared, validated heritage database, so that a record curated in the back office can appear in the public repository and, when suitable, within the immersive museum.

In the virtual museum, heritage objects are presented inside a navigable spatial environment. Rather than depending on installed specialist software, the prototype is built on browser-native technologies, including Three.js, WebGL, and WebXR. WebXR offers a standards-based route for interacting with virtual and augmented reality devices through the web, and its ongoing development reflects the continuing consolidation of immersive web capabilities.

The study's contribution is presented in two parts. One, it documents how an AI-oriented cultural-heritage proposal was operationalised into an integrated digital heritage platform. Two, it shows how a browser-based virtual museum can act as a public-facing interface connected to a structured, curator-mediated heritage repository. Critically, the study does not claim that all AI capabilities initially envisaged within LAIPCH-K have already been demonstrated empirically. Implemented platform functions are separated from AI capabilities that still require later development and evaluation.

2 BACKGROUND AND PROJECT RATIONALE

2.1 Coastal Heritage Preservation Challenge

Preserving cultural heritage increasingly calls for approaches that can hold tangible and intangible forms together. Physical artefacts can be photographed, scanned, and modelled, yet their cultural meaning often depends on context, language, oral history, community knowledge, and patterns of use. A system that archives an object while stripping away the surrounding context risks producing a technically sophisticated representation that remains culturally incomplete.

UNESCO's framing of intangible cultural heritage aligns with this point, since it stresses safeguarding as well as continued transmission. Documentation and access are recognised as necessary, but so is respect for customary norms that regulate who can access particular knowledge and under what conditions.

In LAIPCH-K, this challenge is addressed by pairing physical-object representation with structured metadata and intangible cultural content. Project materials point to artefacts such as dugout canoes, fishing baskets, carved Lamu doors, kanga cloth, and ceremonial objects. On this basis, the system treats each artefact as one element within a larger cultural record, not as a standalone 3D model.

The approach also matches broader directions in digital heritage practice. UNESCO's Recommendation concerning the Preservation of, and Access to, Documentary Heritage Including in Digital Form promotes structured, machine-readable, linkable heritage content, and it places emphasis on inclusive online access, suitable infrastructure, and interoperability [4].

European initiatives likewise stress that 3D digitisation should be guided by purpose and user needs. The 2024 Twin it! call to action suggests that the quality and precision of 3D digitisation should be selected in relation to intended use, since research-grade reconstruction and online public access often require different technical targets [5].

That logic supports the architecture used for the Kenya Coastal Heritage Digital Platform. The platform does not assume that every object must first reach research-level geometric fidelity before contributing to public engagement. Instead, web-ready models are supported so that delivery through a browser remains efficient, while structured metadata and provenance details are retained.

2.2 From Preservation to Digital Experience

Beyond storage, the platform links digital objects with multilingual materials, oral histories, educational resources, and immersive presentation. Users can search artefact records in the public repository, while the storytelling component holds folktales, fishing songs, proverbs, and oral histories. Selected objects are then introduced into an explorable environment through the virtual museum.

This movement from archive to experience reflects wider changes in museum technology. Li et al. (2024) note that digital-transformation technologies can strengthen exhibition design, accessibility, and educational potential when they are integrated appropriately into museum settings. Research on immersive virtual heritage also suggests that 3D reconstruction, interactive narrative, and immersive interfaces open new routes for interpretation, although cultural sensitivity and user experience remain decisive design constraints [6].

Related work, such as a 2024 study of the HimmapanVR virtual museum, illustrates how 3D digitisation and immersive environments can increase access to culturally specific material and support perceived authenticity, while also identifying continuing user-experience challenges [7].

The Kenya Coastal Heritage platform follows a comparable general direction, yet it places stronger emphasis on the connection between community-controlled heritage data and the public immersive interface. That emphasis differentiates the system and frames the core contribution presented in this study.

3 ALIGNMENT WITH THE ORIGINAL LAIPCH-K PROPOSAL

The LAIPCH-K proposal articulated five core aims: an AI-enabled digital archive, AI-supported language revitalisation, VR/AR-based cultural engagement, mechanisms for community participation, and ethical guidance for AI use in cultural preservation. The delivered platform maps onto large parts of this agenda, yet the strength of the supporting evidence is uneven across objectives.

At the level of digital archiving, the first objective is realised primarily through the platform's underlying infrastructure. A shared artefact database, defined metadata arrangements, a public repository, and a curation dashboard together form a coherent archival backbone. What is not demonstrated in the available materials, however, is measurable performance for AI functions such as image recognition, automated categorisation, or intelligent search and retrieval. Given this gap, the current system is better described as an AI-oriented digital heritage infrastructure than as a fully tested AI archive.

Language revitalisation, the second objective, is addressed through multilingual interface and translation support for English, Kiswahili, Kibajuni, and Kipokomo. Documentation points to an AI-assisted layer built on top of a vocabulary resource that has been vetted by the community. This design choice aligns with known constraints affecting African low-resource languages. AfricaNLP work continues to report limited annotated data, substantial linguistic variation, compute constraints, and unresolved ethical issues as persistent obstacles for speech technologies across African languages [8].

Work focusing specifically on Swahili automatic speech recognition also indicates that, even for an African language with comparatively stronger support, phonetic and morphological complexity, along with named-entity recognition, remain non-trivial problems [9]. In that context, combining technical language services with community-validated content is methodologically reasonable, while any future language models integrated into the platform still require standalone evaluation.

For the third objective, VR/AR cultural engagement, the platform offers a direct instantiation through its browser-based virtual museum. The implementation relies on Three.js/WebGL, incorporates WebXR, and provides interactive navigation, 3D artefact loading, audio elements, and an animated guide. Because WebXR is intended to support VR and AR access via the Web, the selected stack is consistent with current Web standards.

Community participation, the fourth objective, is implemented through a workflow that couples contributor submissions with curator review and validation. Crucially, contributions do not appear publicly by default, and publication is mediated through curation. This approach tracks UNESCO's emphasis on active community participation in safeguarding intangible cultural heritage.

Ethical governance, the fifth objective, is only partially evidenced at this stage. While no formal empirical assessment of the intended ethical framework is presented, the system embeds practical governance mechanisms, including role-based access control, validation steps, activity logging, and a clear separation between contribution and publication. These controls provide an operational foundation for responsible handling of cultural data. UNESCO's guidance on AI ethics stresses human oversight, transparency, fairness, diversity, and inclusive participation, and its cultural policy work also calls for participatory approaches when AI is applied to endangered and indigenous languages and to cultural heritage contexts [10].

4 MATERIALS AND METHODS

This work is most appropriately read as a digital heritage implementation and design study. The analysis draws on the original LAIPCH-K proposal, alongside project documentation and the engineering specification for the virtual museum. Rather than attempting to measure outcomes directly, the study concentrates on how stated aims were translated into architecture, data workflows, content structures, and user-facing functions.

Assessment proceeded across four connected planes: system architecture, governance of heritage data, public-facing functions, and the immersive museum implementation. A central concern throughout was traceability, specifically, whether and how proposal objectives can be followed through to implemented features.

Such a stance matters because technology-focused heritage research often collapses technical feasibility into claims about empirical effectiveness. Here, the materials substantiate implementation and integration, but they do not provide controlled evidence for visitor satisfaction, learning effects, community acceptance, language-model accuracy, or wider socioeconomic impact.

Accordingly, the study adopts a technology-demonstration orientation. The prototype stands as the principal implementation artefact, and the documented architecture, data movement, and functional descriptions constitute the primary basis for analysis.

This methodological position also aligns with current digital heritage practice, where preservation and access increasingly depend on connected repositories, structured metadata, 3D representations, and online delivery. UNESCO recommends that digital heritage materials be structured in machine-readable form and encourages web publication and interoperability. In parallel, Europeana's work on 3D content has stressed the role of metadata and standards in ensuring that 3D heritage objects remain discoverable and reusable [11]. [4]

5 SYSTEM ARCHITECTURE AND DATA GOVERNANCE

The platform is organised as a three-tier architecture comprising a presentation layer, a curation layer, and an immersive experience layer. Validated heritage content is delivered through the public presentation layer. Contributor, curator, and administrator activity is handled within the curation layer. The immersive layer provides the virtual museum experience.

As summarised in Figure 1, the architecture can be read by following the linkage from the original LAIPCH-K proposal to the curation and ingestion dashboard, the validated shared database, the public heritage site, the JSON API, and the virtual museum. This arrangement matters because it avoids treating the virtual museum as a standalone showcase. Instead, the museum functions as one presentation surface within a broader heritage information system.




Figure 1: Integrated platform architecture linking proposal intent, curation, shared data and the deployed virtual-museum experience.




Figure 2: Artefact ingestion, validation and publication pipeline implemented across the platform.

Functioning as the main locus for ingestion and oversight, the dashboard allows contributors to upload artefacts, photographs, 3D models, and associated metadata, after which curators examine, verify, and approve entries prior to release. A clear boundary is therefore established between the act of submitting data and the form in which it is publicly displayed. For cultural heritage collections, this separation matters because questions of meaning, ownership, sensitivity, and suitable modes of display often require deliberation by community representatives or domain specialists.

Alignment is also evident with UNESCO guidance on digital documentary heritage, where attention is placed on accurate cataloguing, machine-readable description, interoperability, and inclusive web access, while still acknowledging that justified constraints can apply in situations involving privacy, security, or rights-holder claims.

Under this framing, the database should not be read as a mere technical repository. Instead, it links each digital object to provenance, cultural identity, historical period, materials, patterns of use, and other contextual descriptors. Platform documentation makes this relationship explicit by tying the metadata design to both validation procedures and decisions about public presentation.

Considered from the perspective of Indigenous and community data governance, the design also carries practical implications. The CARE Principles, focused on Collective Benefit, Authority to Control, Responsibility, and Ethics, argue that openness by itself does not necessarily meet Indigenous rights and interests. Within the Kenya Coastal Heritage platform, curator-mediated publication offers an initial operational route for embedding authority and responsibility in digital heritage practice. Open-data approaches alone may not adequately address indigenous peoples' rights and interests [12]

6 THE DIGITAL HERITAGE PLATFORM

6.1 Public Heritage Repository

Presented to end users as a searchable, filterable gallery, the public repository is complemented by detailed artefact pages. Individual records can contain descriptive text, cultural context, periodisation, materials, usage, provenance, attribution to uploaders, and linked media. Alongside these elements, interactive 3D viewing extends the ways in which digital objects can be examined outside an immersive museum setting.

International recommendations are broadly reflected in this approach, particularly the expectation that digital heritage collections provide accurate catalogues, web access, and structured information that supports interoperability. The emphasis on 3D also follows current dissemination trends in cultural heritage. Europeana's ongoing 3D programme indicates a shift toward more information-rich representations, which in turn depend on suitable metadata schemes, rights administration, and shared technical standards [11].

6.2 Language Preservation and Multilingual Access

A central pillar of the platform is language preservation. Support is provided for Kibajuni, Kipokomo, Kiswahili, and English, while the virtual museum integrates multilingual interface dictionaries as well as spoken narration.

The relevance of this design choice becomes clearer given the uneven coverage of African languages in language technologies. Recent studies point to persistent obstacles for low-resource African speech systems, including limited datasets, linguistic complexity, constrained computational capacity, and ethical risks [8, 13].

On these grounds, the community-vetted language resource is better interpreted as a governance layer, not simply as preparatory material for automated translation. UNESCO's AI ethics framework encourages AI applications for endangered and Indigenous languages, yet it also cautions that automated language systems can shape linguistic diversity and cultural expression.

6.3 Storytelling and Educational Resources

The storytelling hub expands the platform beyond material culture by collecting audio-oriented folktales, fishing songs, proverbs, and oral histories, paired with transcripts and translations. Educational materials are also included, such as lesson plans, presentations, quizzes, and interactive learning resources.

This multimodal composition fits the character of intangible heritage, which is often passed on through performance, speech, music, and intergenerational exchange. In the UNESCO Convention, education, awareness raising, and non-formal transmission are described as key components of safeguarding.

By combining audio, text, images, and 3D objects, the platform offers a representation that exceeds an object-only catalogue. Even so, it remains necessary for later empirical work to test whether these resources measurably improve learning outcomes, rather than presuming that richer media automatically yields educational gains.




Figure 3: Functional map of the deployed prototype interface observed at the public virtual-museum endpoint.

6.4 Community Participation and Curation

Community involvement is built into the contribution and review workflow. Members can submit photographs, narratives, and artefact information; contributors can save drafts or submit records for validation; curators are able to revise and approve entries; administrators oversee system-level functions.

This design feature warrants attention because digital heritage systems can reproduce unequal power dynamics when communities are positioned only as content suppliers, rather than as actors in interpretation and governance. UNESCO's intangible heritage framework calls for communities to participate actively in safeguarding and management, and the CARE Principles similarly stress authority, collective benefit, and responsibility in Indigenous data governance.

For these reasons, the validation process provides a foundational mechanism for community-oriented governance in digital heritage settings. At the same time, the present study should avoid asserting that community acceptance or satisfaction has already been demonstrated, since such claims depend on systematic engagement and evaluation with community participants.

7 VIRTUAL MUSEUM: DESIGN AND TECHNICAL IMPLEMENTATION

7.1 Design Philosophy

Four design commitments structure the virtual museum, namely zero installation, progressive enhancement, cultural grounding, and disciplined performance. Ordinary web browsers are treated as the default runtime, with compatible immersive headsets positioned as an extra route rather than a requirement.

What follows from this browser-first stance is a deliberate reduction in reliance on specialised, on-site museum setups. Under the current WebXR specification, standardised web interfaces can address VR and AR devices, yet the same web application structure can still run across varied hardware configurations.




Figure 4: Browser-native virtual museum implementation stack documented in the engineering study.

7.2 Rendering Engine

Implementation is delivered as a browser application built with Three.js on top of WebGL. The reported stack combines GLTFLoader, DRACOLoader, MeshoptDecoder, RoomEnvironment, and PointerLockControls, and it pairs these with practical rendering adjustments such as a capped device-pixel ratio, soft shadows, tone mapping, and fog.

For browser-based 3D distribution, glTF/GLB fits the requirements well. Three.js supplies GLTFLoader for glTF 2.0, and it accommodates Draco and Meshopt compression extensions [14, 15]. While Draco can markedly shrink geometry payloads, it shifts cost onto the client via decoding, so file-size gains must be weighed against runtime decoding speed.

7.3 Procedural Swahili-Inspired Environment

Instead of depending only on a large, pre-authored scene, the museum generates key environmental components procedurally. In the documented build, marble-like flooring, walls, openings, Swahili-inspired arches, and Islamic geometric patterning are produced alongside procedural textures that serve decorative surfaces and artwork panels.

Two practical benefits follow. One is a smaller dependence on heavy, static asset bundles. The other is the ability to tie the visual language of the space to the cultural frame being presented. Even so, the wording “Swahili-inspired” should remain, so that the environment is not read as a complete or definitive reconstruction of historical Swahili architecture.

7.4 Navigation and Interaction

PointerLockControls anchor first-person navigation, supported by keyboard movement and constraints on camera pitch. Artefact selection is handled through raycasting, with selected objects highlighted and accompanied by contextual information. A quick-selection panel provides direct exhibit access and also acts as an accessibility-oriented table of contents.

These interface choices align with recent work that points to interaction design and user experience as central concerns in immersive heritage contexts [16]. Still, usability claims require empirical evaluation with representative visitors. Evidence from immersive-heritage studies also indicates that higher technological immersion can raise engagement, yet it does not remove usability problems.

7.5 Dynamic Artefact Pipeline

A key technical contribution is the linkage between the validated database and the immersive presentation layer. Through a JSON endpoint, the museum can request validated artefacts and load models prepared for viewing at runtime. Compressed 3D assets are supported through GLTFLoader, DRACOLoader, and MeshoptDecoder, and bounding-box normalisation helps ensure that objects produced under different modelling practices are placed consistently.

This division creates a clear separation between content administration and application engineering. Once an artefact is validated and packaged in an appropriate format, it can enter the museum without developers manually rebuilding

the scene. Such a structure supports gradual collection expansion and matches the broader shift toward interoperable and reusable 3D cultural heritage outputs.

7.6 Avatar Guide and Multilingual Narration

The platform includes Amina, an animated cultural guide. Greeting, pointing, speaking, walking, and idle states are coordinated with artefact interaction, and when an artefact is chosen, Amina can turn toward it and deliver contextual narration. English, Kiswahili, and Bajuni are supported, and spoken delivery is produced via browser speech synthesis.

Language is treated as part of the exhibition experience rather than an external translation layer. This matters in heritage interpretation because linguistic forms convey cultural meaning that is not always preserved through direct, word-for-word translation. The direction also matches UNESCO's emphasis on linguistic diversity and the need for human oversight in AI-assisted language technologies.

7.7 Audio, Spatial Sound and WebXR

Background music, positional ambience, and proximity-aware audio are provided via the Web Audio API and Three.js audio components. The system lowers background music while Amina speaks, and it starts audio only after user input to satisfy browser autoplay constraints. WebXR is offered as an optional immersive route, while desktop browsing remains the primary access mode.

By combining visual presentation with spatial audio and interaction, the museum follows the ongoing shift toward multimodal digital heritage experiences. Current WebXR work also shows that web-delivered VR/AR remains an active, changing technical area. The chosen architecture therefore matches contemporary web-based immersive practice, while keeping standard browser access for visitors who do not have VR devices.

8 RESULTS: IMPLEMENTED PLATFORM CAPABILITIES

The implemented system demonstrates an operational chain connecting heritage-data ingestion, curation, validation, metadata management, publication, and immersive display. The main outcome is not only the existence of a virtual museum scene, but an integrated workflow in which preservation records are converted into public, immersive cultural resources. As documented, the platform brings together a searchable digital repository, a curator-mediated dashboard, a 3D exhibition space, multilingual material, storytelling resources, an animated guide, and a WebXR option.

8.1 Prototype Integration and Live Deployment Verification

The work moves past a purely conceptual proposal because the documentation is tied to a publicly reachable prototype of the virtual museum. The deployed interface exposes the Kenya Coastal Heritage Museum environment, artefact display, the Amina guide, VR access, music controls, view reset, conversational interaction, language selection, and introductory exhibit information. Each of these visible functions corresponds to the engineering description of browser navigation, avatar-led guidance, multilingual interaction, audio management, and WebXR access.

This development path is consistent with wider efforts in web-based immersive heritage. The literature suggests that virtual museums can broaden access across distance, though the overall experience depends on usability, engagement, accessibility, and interaction quality. WebXR adoption is also aligned with the present web platform, where the W3C WebXR Device API defines mechanisms for reaching VR and AR hardware through the browser [17]. As of June 2026, WebXR is still a Candidate Standard, so it is more accurate to treat it as an evolving web capability than as a fully settled standard.

9 DISCUSSION

9.1 From an AI Proposal to a Socio-Technical Heritage System

The implementation indicates that LAIPCH-K's AI orientation is better treated as a single element within a broader socio-technical heritage system, rather than as the system's defining core. AI-driven language processing, image analysis, classification, and intelligent retrieval do not operate in isolation; they depend on structured datasets, dependable metadata, clear provenance, workable governance, and validation practices that align with local cultural expectations. In the implemented architecture, these prerequisites are provided through an arrangement that links community contribution and curatorial work to a shared heritage repository, then carries that material forward into public-facing and immersive interfaces.

This matters because contemporary digital heritage infrastructures increasingly braid together databases, 3D reconstruction pipelines, immersive environments, and AI capabilities. Recent work, for example, characterises digital museums as moving past basic repository functions toward interactive spaces that bring generative technologies, immersive experiences, and digital collections into a single environment [18]. From this perspective, LAIPCH-K contributes not only through visualisation, but also by showing how a structured heritage corpus can reliably feed multiple downstream uses.




Figure 5: Traceability from selected LAIPCH-K proposal intentions to documented implementation and prototype evidence.

9.2 Accessibility as a Design Principle

By relying on a browser-native delivery model, the platform reduces reliance on specialist software and fixed VR installations. Standard browsers serve as the main entry point, while compatible immersive devices remain a secondary route for those who have them. The implementation also includes fallback pathways and performance-oriented rendering choices, such as pixel-ratio controls, selective shadowing, shared material strategies, procedural texturing, and compressed geometry.

This design choice aligns closely with the realities of university-based and community heritage projects, where budgets, equipment access, and technical support can be limited. A recent scoping review of virtual-museum research points to usability, engagement, accessibility, and physical comfort as recurring dimensions that require explicit evaluation [19]. Related empirical studies have also examined how immersive VR characteristics connect to visitor experience within virtual museum contexts [20, 21]. Accessibility, then, should not be framed only as the removal of technical obstacles; it also concerns whether different user groups can engage with heritage content in a meaningful and sustained way.

9.3 Cultural Grounding and Community Control

Cultural grounding is expressed through several layers of the platform’s design. The virtual environment draws on Swahili coastal architectural forms, artefacts are presented with contextual metadata, Bajuni is integrated into the multilingual interface and narration, and community-submitted content is channelled through curator validation. Together, these decisions separate the platform from the aesthetics and assumptions of a generic virtual gallery.

Even so, cultural authenticity is not something that digital representation can establish on its own. The UNESCO Convention for the Safeguarding of the Intangible Cultural Heritage stresses the widest possible community participation and active involvement in heritage management, a requirement that becomes especially salient when digital systems represent indigenous or minority languages, oral histories, and culturally sensitive knowledge. UNESCO’s Recommendation on the Ethics of AI similarly promotes participatory approaches when AI is applied to cultural heritage and endangered or indigenous languages, while also cautioning that automated language technologies can reshape linguistic nuance and cultural expression. [1] [10]

9.4 What the Current Evidence Does Not Establish

The documentation supports claims about design choices, integration, and operational feasibility, yet it does not demonstrate downstream effectiveness. No controlled evidence is presented on visitor experience, educational learning gains, community satisfaction, 3D geometric accuracy, or language-model performance metrics such as precision, recall,

or F1. For this reason, socio-economic and tourism effects remain intended or possible outcomes, not empirically demonstrated impacts.

This distinction is consequential, given that recent work on virtual museums increasingly uses user-experience methods rather than relying only on technical demonstration. The present study is therefore best read as an implementation and feasibility contribution.

10 ETHICAL, GOVERNANCE AND SUSTAINABILITY CONSIDERATIONS

The original project identifies informed consent, intellectual property, community participation, privacy, and cultural sensitivity as core ethical requirements. Within the implementation, these concerns are addressed through role-based access controls, curator validation, activity logging, and a separation between contributor submission and public publication.

Such mechanisms reflect the principle that digitisation does not automatically entail unrestricted public release. UNESCO's intangible cultural heritage framework recognises that access conditions should respect customary practices governing particular heritage elements. On this basis, later expansion should differentiate between materials intended for open access and information constrained by cultural norms, privacy needs, intellectual-property concerns, or community decisions. [1]

Long-term sustainability also depends on institutional stewardship. Integration with TUM and relevant cultural institutions provides a plausible route to continuity, while the use of a conventional hosting approach reduces dependence on highly specialised infrastructure. As the repository expands, persistent identifiers, version control, provenance recording, access controls, and documented community review processes should become standard governance features. Where AI-generated translations, descriptions, or classifications are introduced, they should also retain provenance markers and an explicit review status.

Limitations

Several constraints shape the study's evidentiary scope. First, the work is an implementation study, not a controlled comparison among virtual-museum technologies. Second, the available documentation does not include independent evaluation of usability, accessibility, or learning outcomes. Third, while the broader project anticipates AI-assisted functions, quantitative evaluation of downstream AI models is not reported. Fourth, browser behaviour, device performance, and speech capabilities can vary across operating systems and hardware, especially for WebXR and browser-based speech functions. Finally, the current architecture remains an evolving proof of concept, and its scalability, maintenance load, and institutional sustainability require additional testing.

These limitations do not negate the technical contribution; instead, they locate it at an appropriate evidentiary level. What is established is operational feasibility and architectural integration, not causal evidence of cultural, educational, tourism, or economic outcomes.

11 FUTURE RESEARCH AND DEVELOPMENT

The implementation suggests several directions for further development. One major opportunity involves adding a Retrieval-Augmented Generation (RAG) layer that links the Amina avatar to the validated cultural corpus while maintaining source grounding and curator review. Such an approach would support conversational interaction without positioning generative AI as an independent authority over cultural knowledge.

Expansion should also prioritise oral-history recordings, speaker profiles, and community-generated materials. More photogrammetric capture and mobile recording workflows could broaden both the diversity and the geographic coverage of the collection. The engineering study also identifies lip-sync and emotion modelling, VR hand tracking, voice commands, gaze-based selection, guided tours, visitor analytics, level-of-detail management, and progressive mesh streaming as plausible enhancements.

These technical additions should be paired with systematic empirical assessment. Virtual-museum research increasingly shows that usability, immersion, engagement, accessibility, cultural context, and user trust need to be measured, rather than assumed to improve with technical complexity alone.

On that basis, subsequent studies should recruit students, researchers, museum professionals, visitors, and relevant coastal communities. Evaluation could include measures of usability, engagement, learning outcomes, accessibility,

cultural appropriateness, trust, and perceived authenticity. Where AI models are introduced, performance should be reported with suitable quantitative metrics, and assessment should be disaggregated by language and cultural domain.

12 CONCLUSION

The Kenya Coastal Heritage Digital Platform shows how an AI-oriented cultural heritage proposal can be translated into an integrated environment for digital preservation and public engagement. A shared architecture ties together a structured heritage repository, curator-mediated governance, multilingual resources, and storytelling content with a browser-based virtual museum.

Using Three.js, WebGL, WebXR, Web Audio, and Web Speech, the virtual museum extends the archive into an interactive spatial setting. Dynamic artefact loading, procedural Swahili-inspired architecture, multilingual narration, and the Amina avatar indicate that immersive cultural interpretation can be delivered through standard Web technologies. WebXR's continuing development as a W3C Candidate Standard also points to the suitability of this pathway, while signalling that the technical foundation remains in active evolution.

A central implication follows: digital heritage preservation cannot be reduced to AI, 3D modelling, or visualisation alone. Sustainable preservation requires digitisation practices joined to metadata, provenance, cultural validation, multilingual representation, accessibility, governance, and community participation. The platform presented here offers an operational base on which more advanced AI functions and more rigorous empirical evaluation can later be built.

STATEMENTS ON COMPLIANCE WITH ETHICAL STANDARDS

Acknowledgments

The authors would like to acknowledge the financial support provided by the National Research Fund (NRF) Kenya for this project. We would also like to acknowledge the NRF-Kenya Coastal Heritage Digital Platform Technical Team and institutional and community partners for their support and contributions towards the development, documentation and implementation of the platform. Their collaboration and technical expertise was vital in ensuring the successful completion of the project.

Disclosure of conflict of interest

All authors declared that no conflict of interest to be disclosed.

REFERENCES

- [1] UNESCO. Convention for the Safeguarding of the Intangible Cultural Heritage. Paris: UNESCO; 2003.
- [2] UNESCO. Culture and Digital Technologies [Internet]. Paris: UNESCO; 2024–2026 [cited 2026 Sep 29]. Available from: <https://www.unesco.org/en/culture-and-digital-technologies>.
- [3] Li J, Zheng X, Watanabe I, Ochiai Y. A systematic review of digital transformation technologies in museum exhibition. *Comput Hum Behav*. 2024;161:108407.
- [4] UNESCO. Recommendation concerning the Preservation of, and Access to, Documentary Heritage Including in Digital Form. Paris: UNESCO; 2015.
- [5] Europeana. Twin it! Call to action: Boosting 3D for Europe's culture by 2030 [Internet]. Europeana PRO; 2024 [cited 2026 Sep 29]. Available from: <https://pro.europeana.eu/index.php/page/twin-it-3d-for-europe-s-culture>.
- [6] Wang H, Gao Z, Zhang X, Du J, Xu Y, Wang Z. Gamifying cultural heritage: Exploring the potential of immersive virtual exhibitions. *Telemat Inform Rep*. 2024;15:100150.
- [7] Chernbumroong S, Worrugin P, Wongwan N, Intawong K, Homla P, Puritat K. Digitization of myth: The HimmapanVR Project's role in cultural preservation. *Heliyon*. 2024;10(9):e30052.
- [8] Imam SH, Sani B, Gete DK, Ahmed BY, Ahmad IS, Abdulmumin I, Yimam SM, Bello MY, Muhammad SH. Automatic speech recognition for African low-resource languages: Challenges and future directions. In: *Proceedings of the Sixth Workshop on African Natural Language Processing*; 2025. p. 89-94.
- [9] Sharma AK, Pandya M, Shukla A. Fine-tuning Whisper Tiny for Swahili ASR: Challenges and recommendations for low-resource speech recognition. In: *Proceedings of the Sixth Workshop on African Natural Language Processing*; 2025. p. 74-81.

- [10] UNESCO. Recommendation on the Ethics of Artificial Intelligence. Paris: UNESCO; 2021.
- [11] Europeana. Review of the Europeana Data Model to support 3D [Internet]. Europeana PRO; 2024 [cited 2026 Sep 29]. Available from: <https://pro.europeana.eu/post/review-of-the-europeana-data-model-to-support-3d>.
- [12] Carroll SR, Garba I, Figueroa-Rodríguez OL, Holbrook J, Lovett R, Materechera S, et al. The CARE principles for Indigenous data governance. *Data Sci J*. 2020;19(1):43.
- [13] Moteu Ngoli T, Mbuh C, Yopa N. Challenges and limitations in gathering resources for low-resource languages: The case of Medumba. In: *Proceedings of the Sixth Workshop on African Natural Language Processing*; 2025. p. 136-142.
- [14] Three.js. DRACOLoader [Internet]. Three.js documentation; 2026 [cited 2026 Sep 29]. Available from: <https://threejs.org/docs/pages/DRACOLoader.html>.
- [15] Three.js. GLTFLoader [Internet]. Three.js documentation; 2026 [cited 2026 Sep 29]. Available from: <https://threejs.org/docs/pages/GLTFLoader.html>.
- [16] Gong Q, Zou N, Yang W, Zheng Q, et al. User experience model and design strategies for virtual reality-based cultural heritage exhibition. *Virtual Reality*. 2024;28:69.
- [17] W3C. WebXR Device API [Internet]. W3C Candidate Recommendation Draft; 9 Jun 2026 [cited 2026 Sep 29]. Available from: <https://www.w3.org/TR/webxr/>.
- [18] Xu J, Yan L, Zhang R, et al. A review of the development and application of generative technology in digital museums. *npj Heritage Science*. 2025;13:589.
- [19] Araújo CC, Nunes FLS, Teixeira JMXN. User Experience Evaluation in Virtual Museum Tours: A Scoping Review (2014–2024). *J Braz Comput Soc*. 2025. doi:10.5753/jbcs.2025.5715.
- [20] Jangra S, Singh G, Mantri A, et al. Exploring the impact of virtual reality on museum experiences: visitor immersion and experience consequences. *Virtual Reality*. 2025;29:84.
- [21] Wang N, Jia C, Wang J, et al. Identifying key factors influencing immersive experiences in virtual reality enhanced museums. *Sci Rep*. 2025;15:31990.