OpenHeritage3D.org: 3D Atlas Archival and Streaming Infrastructure for Open Cultural Heritage 3D Data

Scott McAvoy[1] Aviral Agarwal[1] Dominique Rissolo[1] Filiberto Chiabrando[2]
George Bent[3] Dave Pfaff[3] John Ristevski[4] Falko Kuester[1]

[1]Cultural Heritage Engineering Initiative (CHEI), California Institute of Telecommunications and Technology (Calit2), University of California San Diego
[2]Laboratory of Geomatics for Cultural Heritage (LabG4CH), Department of Architecture and Design (DAD), Polytechnic University of Turin, Torino, Italy
[3]Florence as it Was, Washington and Lee University, Lexington Virginia, USA

[4]CyArk, Oakland USA

smcavoy@ucsd.edu https://openheritage3d.org

Submission to World Internet Conference, Wuzhen Summit — Submitted April 2026

Abstract

An investment in digital heritage is an investment in societal technology literacy. High definition cultural heritage documentation has reached a level of technical maturity at which terrestrial LiDAR, aerial photogrammetry, and multi-sensor survey systems routinely produce datasets of extraordinary fidelity and scientific value. Yet the infrastructure for preserving, sharing, and delivering these datasets at scale has lagged critically behind the pace of data production. The result is a fragmentation crisis: high-value giga-resolution point clouds and photogrammetric models are routinely stranded in institutional storage, locked in proprietary formats, or degraded to low-fidelity surrogates before publication. OpenHeritage3D.org's 3D Atlas systems addresses this gap directly through a purpose-built, open-source archival and streaming platform designed for datasets of billions of points and terabyte-scale volumes. The platform combines octree-based WebGL streaming, CesiumJS georeferenced spatial integration, FAIR-aligned metadata, and mandatory open licensing to deliver original-resolution heritage data to any web browser without specialized client software. Tools developed for the platform have been applied in wildfire monitoring,coral reef conservation, and built infrastructure conservation contexts, demonstrating broad transferability beyond heritage documentation. This paper describes the platform's technical framework, implementation process, and a forward-looking roadmap for a globally distributed federated network of regional nodes coordinated with European partners at Politecnico di Torino in Italy, with active outreach toward Asian partner institutions.

Keywords: cultural heritage, point cloud streaming, LiDAR, photogrammetry, open data, Potree, CesiumJS, FAIR data, federated infrastructure, change detection, digital preservation

1 Introduction

Cultural heritage documentation has undergone a profound transformation in the past decade. Terrestrial LiDAR scanners, drone-mounted sensors, and close-range photogrammetry systems now routinely capture archaeological sites, historic monuments, and built environments at sub-centimeter resolution, generating three-dimensional point clouds, mesh models, and georeferenced datasets of extraordinary fidelity. These technologies promise a new era of heritage stewardship: sites threatened by conflict, climate change, urban development, or natural disaster can be captured before they are altered or lost, and the resulting data can serve as both a scientific record and a public resource.

Yet the realization of this promise has been persistently obstructed by a structural gap between data production and data accessibility. The community has developed sophisticated tools for acquiring and processing 3D data, including LiDAR systems capable of capturing hundreds of millions of points per scan and photogrammetric pipelines that reconstruct geometry at giga-resolution scale. The infrastructure for long-term preservation, open access, and interactive web-based delivery of such datasets has lagged critically behind. The result is a fragmentation crisis: high-value 3D heritage data is routinely stranded in institutional hard drives, locked in proprietary formats, or published only as degraded derivative products (low-polygon meshes, compressed image textures) that strip away the scientific detail embedded in the original capture.

This paper describes OpenHeritage3D.org, a platform developed specifically to address these structural deficits. Section 2 characterizes the domain challenges in detail. Section 3 describes the technical architecture and imple-

Figure 1: Florence on Google Maps

mentation. Section 4 discusses key innovations and cross-domain transferability. Section 5 presents the roadmap for a globally federated network. Section 6 concludes.

1.1 The 3D Atlases

The 3D Atlas systems enable simple spatial referencing on a global basemap using the open worldmap system. We currently have eight 3D Atlas projects, including UNESCO world heritage sites in Chichen Itza in Mexico, Torino in Italy, The Oracle of Delphi in Greece, The Historic Center of San Marino, The Island of Cyprus, along with natural heritage site comprised of coral reef restoration projects, and wildfire datasets in California. Here we will use one example of an Atlas showing the city center of Florence in Italy, the entirety of which is designated as a UNESCO world Heritage Site. Within the same system, the user navigates from a global basemap to sub-milimeter detail on a single painting. The ability to move through time and place, and scale, are well communicated within this case study CesiumJS [McAvoy et al., 2026],

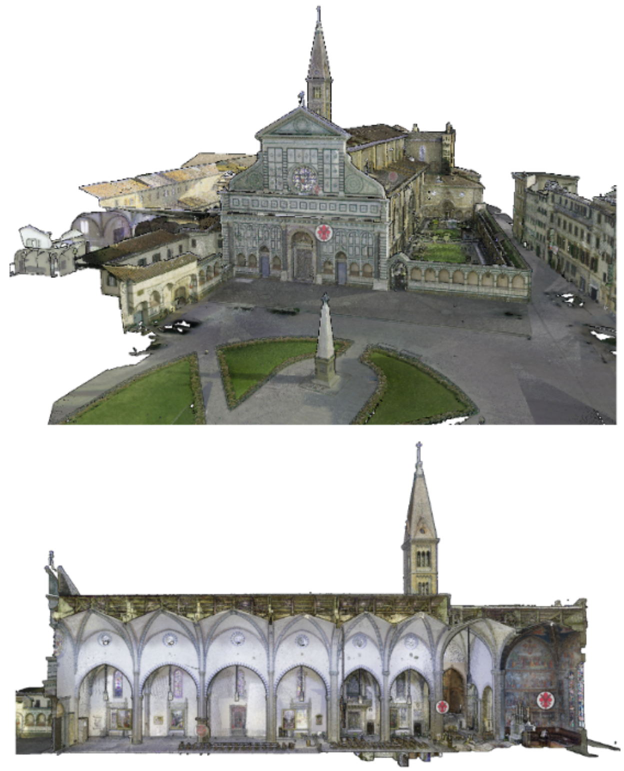

Figure 2: Basilica of Santa Maria Novella

The user navigates to Florence from a global basemap. The basemap can be easily configured with a a wide variety of map layers showing labels, 2D and 3D details. Here we employ the Google Photorealistic 3D tiles to better contextualize the buildings in place. For the City of Florence there are 41 individual sites with data provided, and hundreds of high-resolution individual models of artworks and architectural features plugged in to lower resolution sitemap models. From the global basemap, as the user approaches a beacon indicating the presence of additional data, a 68 billion point terrestrial LiDAR scan is loaded, which acts as a site basemap for a number of individual analysis and storytelling projects (Figure 2). Cross sections can be made on the fly revealing angles and relationships which cannot otherwise be seen. Behind the main altar, the Tornabuoni Chapel features frescoes by the renaissance artist Ghirlandiao, under whom Michelangelo had apprenticed. These frescoes, were painted on top of another set of

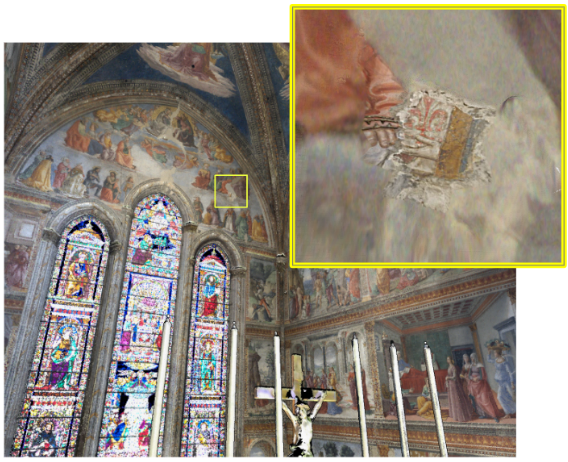

Figure 3: Tornabuoni Chapel in Santa Maria Novella

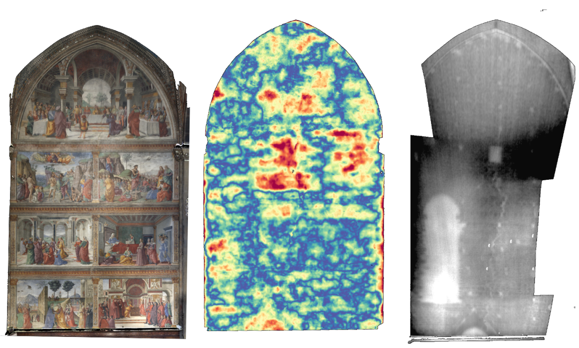

Figure 4: East wall Tornabuoni Chapel, visible, planar deformation map, thermography

frescoes, after the church had been damaged by lightning in the 15th century. Here we can see (Figure 3), towards the top of the 30 meter tall chapel, pieces of the Ghirlandaio fresco crumbling away, reveling vibrant pieces of the work below.

An investigation of the frescoed walls (Figure 4), using terrestrial LiDAR and thermographic imaging, shows hidden features behind the plaster. There had been a door. There had been a window. Most importantly, it highlights a structural crack, going all the way up the wall, introducing moisture and temperature variations from the outside, contributing to the deterioration of the frescoes.

Zooming in to a single panel of the frescoes (Figure 5), a 0.6 millimeter resolution structured light scan of a scene showing the nativity of the Virgin Mary, we can see areas where paint is fading and cracks are forming. By analyzing the geometry we highlight the shape of the fresco. We can see the mounds of plaster built up, where artists laid

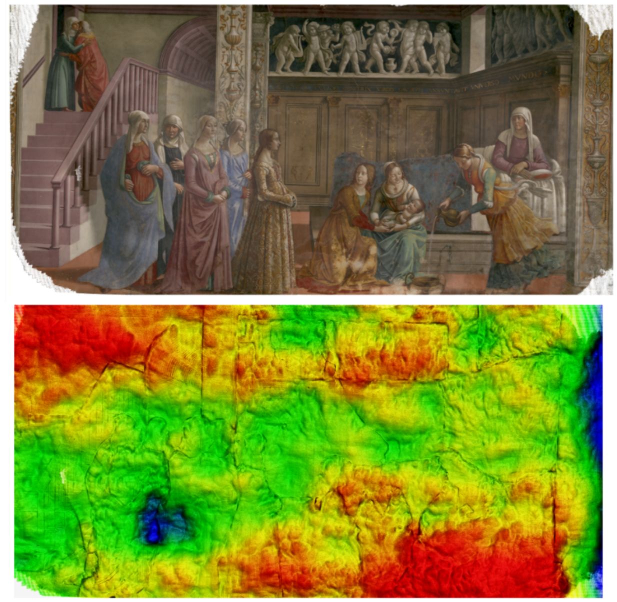

Figure 5: West wall Tornabuoni chapel, Visible, geometry/planar deformation map

new wet plaster every day. Tool marks, hand prints, inscribed guidelines. Clearly the wall, which appears flat, is a complex organic topology.

2 Background and Problem Statement

2.1 Data Scale

A single terrestrial LiDAR scan of a complex heritage site, such as a cathedral interior, an archaeological excavation, or a monumental landscape, may contain billions of points and occupy hundreds of gigabytes in raw form. When sites are documented across multiple campaigns and sensor types, combined datasets routinely reach the terabyte scale. Conventional data repositories, designed around document and tabular data paradigms, are architecturally unsuited to ingesting, indexing, and serving data at this volume and structure.

2.2 Technical Heterogeneity

3D heritage data is captured with diverse instruments including phase-based and time-of-flight terrestrial scanners, mobile and aerial LiDAR systems, structured-light scanners, and multi-camera photogrammetric rigs, each producing data in different formats, coordinate reference systems, and resolution profiles. Integrating these sources into a coherent, queryable archive demands normalization workflows that preserve scientific metadata while enabling interoperability.

2.3 Absence of Open Streaming Infrastructure

While commercial platforms exist for hosting 3D models, they impose file size restrictions, apply lossy compression, obscure underlying data, and preclude scientific reuse. The heritage community has historically lacked access to open-source, scalable infrastructure that allows original point cloud data, not surrogates, to be streamed interactively worldwide.

2.4 Spatial Isolation

Heritage datasets are typically treated as isolated objects with no spatial relationship to one another. Meaningful interpretation often requires relating individual datasets within a shared geographic coordinate framework, a requirement that assembling heterogeneous 3D captures into a georeferenced global environment introduces significant technical and metadata complexity to meet.

2.5 Institutional Sustainability

The heritage documentation community generates data under a wide range of funding structures (national research councils, international development programs, emergency response initiatives, independent expeditions), and the data produced often has no designated custodian once a project concludes. Without persistent, open repositories that enforce standards for metadata, licensing, and long-term preservation, the record of documented heritage is at continuous risk of effective loss.

3 Technical Framework

OpenHeritage3D.org is built on a modular, open-source technical stack designed to handle the full lifecycle of large-scale heritage 3D data: from field acquisition and archival ingest, through server-side processing and format normalization, to interactive web-based delivery and spatial integration. Each layer was selected or developed in response to a specific failure mode observed in prior approaches to heritage data sharing.

3.1 Storage and Archival Layer

At the storage and compute layer, the platform uses cloud object storage to hold archival-quality datasets in their original or losslessly compressed form. A major infrastructure milestone was achieved through an award from Western Digital, resulting in one petabyte of dedicated storage supporting the platform. This contribution has been foundational to the platform's ability to archive full-resolution datasets at scale.

Raw sensor data is stored as images and pointclouds, preserved in appropriate archival formats preserving all key data attributes, in open lossless file standards that retains full precision floating-point coordinates. Photogrammetric inputs (images) outputs, including dense point clouds, mesh geometry, and associated texture assets, are similarly preserved at full resolution. This commitment to archival fidelity distinguishes the platform from commercial 3D hosting services.

3.2 Web Streaming Layer

The streaming layer is built on the Potree ecosystem [Schütz, 2016], an open-source WebGL-based point cloud renderer. Potree uses an octree spatial data structure to enable level-of-detail streaming: as a user navigates a dataset in the browser, only the spatial cells relevant to the current view and zoom level are streamed from the server, making it possible to interactively explore datasets of billions of points without downloading the complete archive. PotreeConverter is run server-side during ingest to generate the streaming-optimized octree tile structure from the archival LAZ source, alongside, never replacing, the preserved original.

3.3 Spatial Integration Layer

For global contextualization, the platform employs CesiumJS [Cesium, 2011], an open-source WebGL-based virtual globe library. Individual site datasets, once processed, are registered to real-world geographic coordinates and embedded within a Cesium-based world viewer. This allows users to discover datasets spatially, navigating from continent to region to site, and situates each dataset within its authentic landscape context. The Cesium layer also supports integration with satellite imagery, terrain models, historical maps, and GIS feature layers, providing the foundation for VR/AR access and full-resolution data querying.

3.4 Metadata and Discovery Layer

The metadata layer implements a schema aligned with FAIR data principles [Wilkinson et al., 2016] (Findable, Accessible, Interoperable, Reusable), capturing dataset provenance, acquisition parameters (instrument type, resolution, coordinate reference system), licensing terms, and persistent identifiers. The schema was developed with reference to CIDOC-CRM and the Europeana Data Model, extended to accommodate the specific attributes of 3D spatial datasets. Datasets are assigned persistent identifiers at ingest, making them citable in academic publications.

3.5 Implementation

The platform was developed iteratively, with each major feature release driven by the requirements of an active documentation campaign. The initial implementation was validated using terrestrial LiDAR campaigns at archaeological sites in the Yucatán Peninsula of Mexico, conducted in collaboration with UC San Diego and partner Mexican institutions [McAvoy et al., 2023]. These multi-scan registrations of monumental architecture, totaling tens of billions of points across multiple sites, stress-tested the ingest pipeline and established baseline performance benchmarks for the streaming layer.

Subsequent campaigns at historic ecclesiastical buildings in Italy, including survey work conducted under formal agreements with Italian heritage authorities and in collaboration with Prof. Grazia Tucci at the GECO Lab (Università degli Studi di Firenze), introduced E57 as an additional ingest format and required the development of coordinate transformation workflows for datasets registered in local site coordinate systems. These campaigns also introduced multi-sensor datasets combining terrestrial LiDAR, close-range photogrammetry, and thermal infrared imaging [McAvoy et al., 2026], demonstrating the platform's capacity for heterogeneous data types.

The georeferenced global environment was implemented following accumulation of sufficient site datasets to validate the Cesium integration at scale. A key challenge was the development of a seamless level-of-detail transition between the global Cesium view and the Potree point cloud viewer, requiring careful management of coordinate reference system conversions, camera state handoff, and progressive loading behavior [McAvoy et al., 2022].

The metadata and licensing framework was implemented in consultation with data governance specialists at the Qualcomm Institute, UC San Diego. Open licensing is required for all archived datasets: contributors must assign a Creative Commons or equivalent open license as a condition of archival.

4 Innovations and Cross-Domain Applications

4.1 Full-Resolution Open Streaming

The most significant contribution is the achievement of full-resolution, open-access streaming of giga-scale point cloud datasets through a standard web browser, without specialized client software or institutional authentication. Prior to platforms of this type, the practical options for sharing large heritage point clouds were limited to physical media transfer, credentialed FTP servers, or commercial platforms imposing size limits and proprietary access controls. OpenHeritage3D.org demonstrates that multi-billion point registrations of complex heritage subjects can be made genuinely publicly accessible on the open web, with interactive query and visualization capabilities supporting both casual exploration and rigorous scientific analysis.

4.2 Georeferenced Multi-Site Spatial Environment

Heritage 3D datasets have historically been published as isolated objects with no spatial relationship to one another. The Cesium-based global viewer establishes a framework in which datasets from different projects, institutions, and world regions coexist in a shared geographic coordinate space, enabling spatial query and cross-site comparative analysis at a scale previously impractical. A researcher studying architectural typology across a regional tradition, or a conservator tracking material change across multiple campaigns, can access and compare datasets through a single interface.

4.3 Archival Fidelity as a Design Principle

Many existing heritage data portals publish 3D data as display artifacts without preserving the original acquisition data or metadata necessary to assess scientific quality. OpenHeritage3D.org treats archival fidelity as a primary design requirement: original point clouds with full acquisition metadata are preserved and available for download alongside the streaming representation. This makes the archive genuinely useful for quantitative research applications including change detection, volumetric modeling, and GIS integration.

4.4 Cross-Domain Transferability

The spatial data management, change detection, and streaming infrastructure developed for OpenHeritage3D.org is actively deployed in allied domains. The platform's tooling is used in ALERTCalifornia, the University of California's statewide sensor network for wildfire and natural hazard monitoring. The same infrastructure supports a growing number of conservation and built-infrastructure management projects, where repeated LiDAR surveys of bridges, retaining walls, historic structures, and coastal features must be managed, compared, and delivered to engineering and planning teams. This cross-sector applicability demonstrates that the platform's contributions are broadly applicable spatial data infrastructure, with direct relevance to public safety, environmental monitoring, and infrastructure asset management.

4.5 Post-Disaster Change Detection

Work connected to the 2025 Southern California wildfires demonstrated the platform's value in high-stakes applied contexts. Pre-fire LiDAR data for affected landscape and built heritage areas, archived and georeferenced through the platform's infrastructure, provided a scientific baseline for post-fire damage assessment and geomorphic change analysis that would not have been possible without the combination of archival fidelity and spatial integration the platform provides.

5 Future Prospects and Global Network Roadmap

5.1 Federated Architecture

The next phase of development moves from a single-institution platform toward a globally distributed network of interoperable regional nodes, unified by a common data standard, shared discovery infrastructure, and a centrally coordinated services layer. Regional nodes hosted by partner universities, national heritage agencies, and research institutes will contribute datasets to a federated global system. This architecture removes centralized ingest bottlenecks, reduces latency for regional users, respects evolving data sovereignty regulations, and builds the distributed institutional ownership that makes infrastructure sustainable across decades.

Regional nodes will operate independent Potree and object storage instances while exposing a standardized API layer that allows the global discovery environment to index and query across all nodes simultaneously. Achieving this requires the development and adoption of a common metadata exchange format and streaming API specification, to be advanced through formal engagement with the Open Geospatial Consortium (OGC) and relevant ISO technical committees.

5.2 Advanced Visualization and AI Applications

The next generation of the platform will integrate WebXR-based virtual and augmented reality access, allowing users with consumer VR headsets to inhabit heritage environments at true scale and full resolution and enabling on-site AR overlay of archival data onto existing physical fabric. Alongside immersive access, the visualization layer will incorporate in-browser quantitative analysis tools: measurement, cross-section extraction, volumetric comparison, and temporal change visualization for sites with multi-epoch datasets.

A complementary research analytics API will allow programmatic query of datasets at the point cloud level without downloading entire archives. Spatial subset extraction, statistical sampling, attribute filtering, and co-registration operations exposed as API endpoints will enable a new class of large-scale quantitative heritage research.

The combination of georeferenced, full-resolution, richly documented 3D datasets the network will accumulate constitutes an exceptionally valuable resource for machine learning applied to heritage, environmental, and built environment problems. Semantic segmentation of architectural elements, automated damage classification, structural anomaly detection, and generative reconstruction of damaged fabric are all active research areas constrained by the absence of large, high-quality, openly licensed training datasets. The network's analytics infrastructure will include explicit support for AI training data export, with standardized dataset packaging, integrated annotation tooling, and open licensing for derived training data.

5.3 Phased Implementation Plan

Phase 1 (Years 1–2): Network Foundation. Establish the federated architecture and recruit the initial cohort of regional node partners. Develop the interoperability API specification, deploy reference node implementations at two to three partner institutions, and validate cross-node discovery and streaming performance. Advance metadata standardization through formal submission to OGC's Heritage and Cultural Resources Domain Working Group. Funding will draw on NSF and NEH digital infrastructure programs, supplemented by European Research Infrastructure mechanisms available to consortium partners.

Phase 2 (Years 2–3): Analytics and Visualization Layer. Develop the research analytics API, WebXR visualization layer, and initial AI training data infrastructure in close collaboration with computer science and visualization research groups at partner institutions. Coordinate a dedicated cohort of documentation campaigns across node partners to populate the network with the volume and variety of data needed to validate the analytics infrastructure and seed the AI training data repository. Current point-based visualization systems will be overhauled to include scalable gaussian splat-based visualization systems, merging these new capabilities with existing metrological systems.

Phase 3 (Years 3–5): Global Expansion. Pursue broad network expansion, targeting node deployments in underrepresented regions including Sub-Saharan Africa, South and Southeast Asia, and Latin America. Provide structured on-boarding support comprising technical deployment assistance, metadata training, and data management capacity building. Advance international standards toward formal publication and develop commercial and governmental licensing arrangements for AI training data and analytics API access. In these endeavors, we hope to develop close working relationships with governmental and extra-governmental agencies worldwide.

5.4 International Cooperation

The project participates actively in CIPA (the International Committee for Architectural Photogrammetry), the leading professional organization for heritage documentation affiliated with both ICOMOS (International Council on Monuments and Sites) and ISPRS (International Society for Photogrammetry and Remote Sensing). CIPA membership provides a peer network of institutional partners and a direct channel into international standards development for heritage recording.

Formal agreements are in place with the Politecnico di Torino (Italy) and Cyprus University of Technology to expand the platform's infrastructure into Europe. The Politecnico di Torino brings world-leading expertise in heritage photogrammetry through its DIATI department; the collaboration includes joint dataset contribution, shared curriculum development, and co-authorship on standards submissions [McAvoy et al., 2022]. Cyprus University of

Technology provides a strategically positioned node for Eastern Mediterranean heritage data, a region of exceptional archaeological and architectural significance underserved by existing open data infrastructure.

The project is actively seeking partner institutions in Asia to extend this vision worldwide. Asia encompasses some of the world's most significant and most endangered cultural heritage, from the archaeological landscapes of Central Asia and the monumental architecture of South and Southeast Asia to the rapidly urbanizing historic districts of East Asian cities. Regional node partners in Asia would bring irreplaceable heritage datasets and the institutional context necessary to ensure the network serves local research communities and public audiences. The project welcomes dialogue with universities, national heritage agencies, and research institutes across the region who share a commitment to open science and digital heritage preservation.

5.5 Resource Allocation

Infrastructure requirements will be met through a combination of institutional research computing commitments from node partners and negotiated arrangements with cloud providers offering research program pricing, anchored by the Western Digital one-petabyte storage contribution at the primary node. Core engineering personnel (three to four software engineers and one DevOps specialist at the primary node, supplemented by partner institution contributions) and a dedicated program manager will coordinate consortium governance and partner onboarding across the five-year implementation period.

5.6 Risk Mitigation

The primary technical risk, achieving reliable performant cross-node streaming at global scale, is mitigated through phased validation and investment in robust monitoring, caching, and graceful degradation design ensuring a node outage does not compromise access to data held elsewhere.

The primary institutional risk, partner attrition due to funding interruptions or shifting priorities, is mitigated by designing node operations to be low-maintenance by default, building data redundancy across multiple nodes for high-value datasets, and structuring governance to allow nodes to transition between active contribution and read-only participation without disrupting the network.

The primary policy risk, evolving data sovereignty and cross-border data transfer regulations, is addressed by the federated architecture itself: data physically remains on regional servers rather than being centrally replicated, aligning the network's design with the direction of travel in international data governance.

Obsolescence risk in a rapidly moving technical landscape is mitigated by the platform's modular architecture, which allows analytics and visualization layers to be updated independently of the archival and streaming infrastructure, and by the commitment to open standards ensuring that archived data retains value regardless of changes in access tooling.

6 Conclusion

OpenHeritage3D.org demonstrates that the gap between data production capacity and data sharing infrastructure is a solvable engineering and organizational problem, not an inherent consequence of data scale or complexity. The technical components required to make big spatial data openly accessible on the web exist, are mature, and are free to use. What has been missing, in heritage as in other data-intensive research domains, is their deliberate integration into purpose-built archival infrastructure governed by open science principles.

The platform provides a working proof of concept that this integration is achievable at professional and scientific scale, and a documented model for others to follow. Its cross-domain deployment in wildfire monitoring, infrastructure management, and post-disaster assessment demonstrates that open heritage data infrastructure is not a specialist niche but a component of broader societal spatial data literacy, reinforcing the principle with which this paper opened: an investment in digital heritage is an investment in societal technology literacy.

The roadmap toward a globally federated network of regional nodes, coordinated through active partnerships with CIPA, ICOMOS, ISPRS, the Politecnico di Torino, Cyprus University of Technology, and prospective Asian partners, represents a concrete path toward ensuring that the 3D records we create of a fragile and threatened global heritage are not just captured, but preserved, shared, and kept alive for the communities and researchers who need them most.

Acknowledgements

The authors thank Western Digital Corporation for the award of one petabyte of dedicated storage infrastructure supporting the OpenHeritage3D.org platform. We gratefully acknowledge the contributions of field campaign partners at Universidad Autónoma de Yucatán, and the CIPA Heritage Documentation community affiliated with ICOMOS and ISPRS. We thank all data contributors whose open-licensed datasets make the platform's scientific mission possible.

7 Use of Artifical Intelligence Disclosure

The AI tool Claude.ai using the Sonnet 4.5 Large Language Model [?] was employed to suggest grammatical corrections to original text, to improve clarity and flow. The authors of this paper have reviewed and approved the final manuscript, and take full responsibility for its content, including the content which has been generated or modified by AI tools.

References

Schütz, M. (2016). Potree: Rendering Large Point Clouds in Web Browsers. *Master's Thesis*, Vienna University of Technology.

Analytical Graphics, Inc. (2011). CesiumJS: An Open-Source JavaScript Library for World-Class 3D Globes and Maps. https://cesium.com/cesiumjs/

Wilkinson, M. D., et al. (2016). The FAIR Guiding Principles for Scientific Data Management and Stewardship. *Scientific Data*, 3, 160018.

McAvoy, S., Tanduo, B., Spreafico, A., Chiabrando, F., Rissolo, D., Ristevski, J., & Kuester, F. (2022). An Archival Framework for Sharing of Cultural Heritage 3D Survey Data: OpenHeritage3D.org. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVI-2/W1-2022, 517–524.

McAvoy, S., Pérez Rivas, M. E., Gallegos Flores, J. M., Dorshow, W., Stanton, T., et al. (2023). Towards a Mexican National Archaeological Atlas: Scalable 3D Web GIS and Archival Systems for Big Data. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*.

McAvoy, S., Ristevski, J., Rissolo, D., & Kuester, F. (2021). OpenHeritage3D: Building an Open Visual Archive for Site Scale Giga-Resolution LiDAR and Photogrammetry Data. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*.

McAvoy, S., et al. (2026). 3D Radiometric Thermography Mosaics with Low-Cost Mobile Sensor Stack *Remote Sensing* (Accepted, awaiting publication).

Anthropic. Claude (Sonnet 4.5) Large Language Model