TRANSFORMATIONS
A DARIAH Journal

A Reproducible Workflow for the Creation of Digital Twins in the Cultural Heritage Domain

Sebastian Barzaghi[1], Alice Bordignon[2], Federica Collina[1], Francesca Fabbri[2], Bruno Fanini[3], Daniele Ferdani[3], Bianca Gualandi[2,4], Ivan Heibi[2], Nicola Mariniello[5], Arcangelo Massari[2], Marcello Massidda[6], Arianna Moretti[2], Silvio Peroni[2], Sofia Pescarin[6], Maria Felicia Rega[3], Giulia Renda[2], Mattia Sullini[7]

1 Department of Cultural Heritage, Alma Mater Studiorum – Università di Bologna, Via degli Ariani 1, Ravenna, Italy
2 Department of Classical Philology and Italian Studies, Alma Mater Studiorum – Università di Bologna, Via Zamboni 32, Bologna, Italy
3 Digital Heritage Innovation Lab, Institute of Heritage Science, National Research Council, Area della Ricerca di Roma 1, Strada Provinciale 35d, 9, 00010, Montelibretti, Rome, Italy
4 Research Division (ARIC), Research Services and Project Coordination Unit, Alma Mater Studiorum – Università di Bologna, Via Zamboni 33, Bologna, Italy
5 Engineering Ingegneria Informatica S.p.A., Via Giovanni Porzio CDN Torre Saverio Isola C1, 80142, Naples, Italy
6 Digital Heritage Innovation Lab, Institute of Heritage Science, National Research Council, Via Madonna Del Piano, 10, 50019, Sesto Fiorentino, Florence, Italy
7 Department of Architecture, Alma Mater Studiorum – Università di Bologna, Viale Del Risorgimento, 2, 40136, Bologna, Italy

Transformations, A DARIAH Journal
Volume 1, 2025
https://transformations.episciences.org

Dates
Received: 15/11/2024
Accepted: 02/05/2025
Published: 25/07/2025

DOI: 10.46298/transformations.14773

Abstract

This article explores how to create reproducible workflows for the 3D acquisition and digitisation of cultural heritage objects to ensure sustainability and reusability across various institutions. In addressing two main research questions, the paper proposes a workflow that involves the systematic acquisition, processing and digitisation of cultural heritage artefacts. In particular, the workflow focuses on developing digital twins for cultural heritage settings and exhibitions and proposes baseline standards for both technical and interpretative aspects of digitisation. The workflow has been derived from and tested on the pilot case of the temporary exhibition *The Other Renaissance: Ulisse Aldrovandi and the Wonders of the World*, in the context of the CHANGES project. The article reflects on software and hardware equipment to select, the procedures and techniques to use and the formats to adopt to comply with openness, accessibility, transparency, reproducibility, reusability and sustainability in the research workflow, building on previous works on fostering reproducibility in research and improving the interoperability of 3D data across different systems. It highlights the need for transparent documentation of every step of the process, focusing on accountability and practices in the context of cultural heritage research. Finally, the article suggests improvements for enhancing the sustainability of these kinds of workflows and discusses future directions for digitisation efforts and sharing practices.

Keywords: digital twin, reproducibility, 3D models, digitisation

EPIsciences

To cite this article: S. Barzaghi, A. Bordignon, F. Collina, F. Fabbri, B. Fanini, D. Ferdani, B. Gualandi, I. Heibi, N. Mariniello, A. Massari, M. Massidda, A. Moretti, S. Peroni, S. Pescarin, M. Felicia Rega, G. Renda, M. Sullini. (2025). A Reproducible Workflow for the Creation of Digital Twins in the Cultural Heritage Domain. *Transformations? A DARIAH Journal*, vol. 1, p. 1-63. https://doi.org/10.46298/transformations.14773

Un flusso di lavoro riproducibile per la creazione di gemelli digitali nel settore del patrimonio culturale

Riassunto

Questo articolo esplora come creare flussi di lavoro riproducibili per l'acquisizione 3D e la digitalizzazione di oggetti del patrimonio culturale, al fine di garantirne la sostenibilità e la riutilizzabilità in diverse istituzioni. Affrontando due principali quesiti di ricerca, l'articolo propone un flusso di lavoro che prevede l'acquisizione, l'elaborazione e la digitalizzazione sistematica di manufatti del patrimonio culturale. In particolare, il flusso di lavoro si concentra sullo sviluppo di gemelli digitali per allestimenti e mostre che espongono patrimonio culturale e propone standard di base per gli aspetti tecnici e interpretativi della digitalizzazione. Il flusso di lavoro è stato derivato e testato sul caso pilota della mostra temporanea *L'Altro Rinascimento: Ulisse Aldrovandi e le Meraviglie del Mondo* nell'ambito del progetto CHANGES. Nell'articolo introduciamo le attrezzature software e hardware, le procedure e le tecniche da utilizzare e i formati da adottare per garantire apertura, accessibilità, trasparenza, riproducibilità, riutilizzabilità e sostenibilità del flusso di lavoro di ricerca, basando i nostri risultati su lavori precedenti volti a promuovere la riproducibilità nella ricerca e a migliorare l'interoperabilità dei dati 3D tra sistemi diversi. In particolare, il nostro contributo evidenzia la necessità di documentare in modo trasparente ogni fase del processo, concentrandosi sulla responsabilità e sulle pratiche relative alla ricerca sul patrimonio culturale. Infine, l'articolo suggerisce come migliorare la sostenibilità di questi flussi di lavoro e discute possibili direzioni future per far avanzare i processi di digitalizzazione e le pratiche di condivisione.

Parole chiave: gemello digitale, riproducibilità, modelli 3D, digitalizzazione

Acknowledgements

The authors want to thank all the people from different departments of the University of Bologna and of the Institute of Heritage Science involved in this project: Alessandro Iannucci, Anna Forte, Diego Ronchi, Elena Formia, Federica Bonifazi, Federica Giacomini, Filippo Fantini, Francesca Tomasi, Gabriele Bitelli, Giulia Fiorini, Giuliana Benvenuti, Laura Travaglini, Luca Cipriani, Luca Vittuari, Maria Alessandra Tini, Marilena Daquino, Rachele Manganelli Del Fà, Roberto Balzani, Simona Colitti, Valentina Alena Girelli.

This work has been partially funded by Project PE 0000020 CHANGES - CUP B53C22003780006, NRP Mission 4 Component 2 Investment 1.3, Funded by the European Union–NextGenerationEU.

CRediT authorship contribution statement

Authors' contribution according to CRediT (https://credit.niso.org/):

Sebastian Barzaghi: Data curation, Methodology, Writing – original draft, Writing – review & editing
Alice Bordignon: Methodology, Visualization, Writing – original draft, Writing – review & editing
Federica Collina: Methodology, Writing – original draft, Writing – review & editing
Francesca Fabbri: Methodology, Visualization, Writing – original draft, Writing – review & editing
Bruno Fanini: Methodology, Software, Writing – original draft, Writing – review & editing
Daniele Ferdani: Methodology, Writing – original draft, Writing – review & editing
Bianca Gualandi: Methodology, Writing – original draft, Writing – review & editing
Ivan Heibi: Methodology, Writing – original draft, Writing – review & editing
Nicola Mariniello: Methodology, Writing – original draft, Writing – review & editing
Arcangelo Massari: Methodology, Software, Writing – original draft, Writing – review & editing
Marcello Massidda: Methodology, Software, Writing – original draft
Arianna Moretti: Methodology, Software, Writing – original draft, Writing – review & editing
Silvio Peroni: Investigation, Methodology, Supervision, Writing – original draft, Writing – review & editing
Sofia Pescarin: Investigation, Methodology, Supervision, Writing – review & editing
Maria Felicia Rega: Methodology, Writing – original draft, Writing – review & editing
Giulia Renda: Methodology, Software, Writing – original draft, Writing – review & editing
Mattia Sullini: Methodology, Writing – original draft, Writing – review & editing

1. Introduction

Reproducibility has long been a topic of discussion across scientific disciplines—from the medical domain, as introduced by Niven et al. (2018), to the humanities (see, for instance, Peels and Bouter 2018). A number of researchers around the world agree that we are currently facing a reproducibility crisis in research (Peng 2015; Baker 2016; Hutson 2018). However, our understanding of reproducibility and its impact depends to a large extent on the type of research and the scholarly domain in question:

> There seems to be an agreement on the fact that research can be reproducible in varying degrees, from an "ideal" computational reproducibility all the way to fields where multiple interpretations of a certain phenomenon coexist. (Barzaghi et al. 2024a)

Considerations about reproducibility also affect those areas of the humanities that deal with the adoption of computational methods in the valorisation and preservation of cultural heritage (CH). Even though technological tools mediate the acquisition and digitisation of cultural heritage artefacts, the researchers involved in these activities act on subjective interpretations, bringing in different viewpoints, theoretical backgrounds and previous assessments that affect the entire process (Barzaghi et al. 2024a). These activities, situated within the humanities domain, often involve objects "with *meaning* and *value*, objects such as paintings, texts, statues and buildings—in opposition to, say, objects such as atoms and viruses that are studied in the natural sciences" (Peels 2019, 2). As a result, reproducibility in CH cannot be achieved without accounting for all influencing factors, including subjective elements, individual institutional practices and national cultural policies.

The present article aims to answer the following two research questions (RQs):
1. Which materials, methods and tools can be used to implement a reproducible workflow for acquiring and digitising 3D cultural heritage artefacts?
2. How can current practices be improved to enhance the reusability and sustainability of these workflows?

In the work we have done in the context of the project CHANGES–Spoke 4 (https://www.fondazionechanges.org/spoke-4/), we have faced reproducibility issues in defining and implementing specific sub-projects whose main activities focus on acquiring and digitising cultural heritage objects and creating digital twins[1] of museum exhibitions, with the aim of improving the valorisation and preservation of cultural heritage. In this context, we

1. It is important to highlight that, in the context of this article, we are adopting the definition of *digital twin* introduced in Niccolucci et al. (2023), which we discuss extensively in section 2.1.

have considered a set of case studies involving heterogeneous cultural institutions, which include natural history and science museums, art galleries, site museums with (in)tangible heritage and landscapes, historical palaces and anthropological museums, as well as museums with extensive collections and high-tech approaches.

To come up with acquisition and digitisation guidelines that can be used by the various partners (including researchers, professionals and external companies) involved in all the case studies, we ran a pilot study that could serve as a common experimental ground and a baseline to define a precise and reproducible workflow for the acquisition and digitisation of 3D cultural heritage assets: the temporary exhibition entitled *The Other Renaissance: Ulisse Aldrovandi and the Wonders of the World* (https://site.unibo.it/aldrovandi500/en/mostra-l-altro-rinascimento). As described in prior research on this topic (Balzani et al. 2024), the goal of this activity was "to obtain a digital version of the experience at the exhibition, starting from its digital twin, connected to the digital asset of the different items (3D and multimedia) of the collections, organised and accessible online by users, using various devices (from home computers, smartphones, to tablets and VR headsets)."

Building on the above-mentioned work, this article introduces the reproducible workflow we devised for planning acquisition and digitisation campaigns across all nine case studies in the project. It complements the guidelines for FAIR (findable, accessible, interoperable, reusable) management of 3D data in cultural heritage (Barzaghi et al. 2024b) by adding in-depth information about all the phases in the acquisition and digitisation workflow. In addition, it outlines the steps required to create a digital twin of cultural heritage events like exhibitions, placing an emphasis on research interoperability, transparency, openness and accountability.

Finally, the proposed workflow is designed to make its output—the digital twin—reusable and adaptable for a range of future applications that go beyond the original goal of valorisation and preservation. In particular, we showcase a psychological experiment in which the digital twin produced by the workflow was used as a test environment for gathering requirements and identifying features that capture user behaviour in 3D and VR environments.

The rest of the article is organised as follows. In section 2, we introduce the most relevant related works on reproducible workflows for 3D acquisition and digitisation of cultural heritage, interoperable 3D formats for storing and exchanging 3D models and techniques for creating digital twins. Section 3 answers RQ1 by outlining the main steps of our reproducible workflow for acquiring and digitising objects and creating digital twins of events in the cultural heritage domain. RQ2 is addressed in section 4, where we discuss possible further improvement of the workflow to foster better reusability, interoperability, transparency, openness and accountability in the process. In section 5, we conclude by sketching out some future work.

2. Related work

2.1. Reproducible workflows for 3D digitisation of cultural heritage

The 3D digitisation of cultural heritage assets has become a common practice, aiding the understanding, preservation and promotion of culturally and historically significant sites and artefacts. As museums seek to innovate traditional exhibitions and enhance the appeal of their permanent and temporary collections, the demand for 3D digitisation of artefacts is steadily increasing (Raimo et al. 2022; Farella et al. 2022a).

In recent years, among existing and novel approaches to 3D digitisation, *digital twins* have been particularly successful—especially with respect to 3D modelling and simulation—for purposes like assessment and prediction due to their role as virtual counterparts of physical entities. The term *digital twin* was first defined by Grieves and Vickers (2017) and later expanded in "The Gemini Principles" (Bolton et al. 2018), which emphasised the importance of accurate digital representations for informed decision-making. In the literature produced during the period in which the term was defined and conceptualised, digital twins are often associated with a tight pairing between the physical world and its virtual representation and display strong capabilities in simulating physical phenomena in a digital environment. As such, a degree of real-time, synchronised and bidirectional interaction between physical objects and digital replicas is expected. On the basis of this degree of interaction, these digital copies can be classified into three categories (Kritzinger et al. 2018): *digital models* (no interaction whatsoever); *digital shadows* (unidirectional interaction, from physical to digital); and actual *digital twins* (bidirectional interaction between physical and digital). Others have argued that the requirements for digital twins may be too rigid and proposed a more flexible interpretation, especially for cultural heritage due to its unique, often intangible nature, which can change over time or even cease to exist physically (Gabellone 2022). In particular, Niccolucci et al. (2023) propose a more liberal definition of *digital twin* based on the separation of data representation from data exchange, arguing that, in the context of cultural heritage, features such as continuous bidirectional interaction, while feasible, should not be mandatory. Thus, a digital twin can also be interpreted as a flexible "model of knowledge" that can and should evolve to gradually meet specific needs and eventually become a digital twin in the traditional sense.

Creating digital twins of museum assets is often challenging and labour-intensive, even in their less restrictive form as 3D digital replicas. The process widely adopted in the literature (Raham et al. 2019; Raham and Champion 2019; Demetrescu et al. 2020; Apollonio et al. 2021; Benítez et al. 2022; Peinado-Santana et al. 2021; Sebar et al. 2021; Farella et al. 2022a; Farella et al. 2022b; Bolognesi and Manfredi 2024) generally includes four phases: data acquisition, data processing, mesh and texture generation and optimisation for web visualisation.

A literature review reveals that the same challenges arise in many digitisation processes within the GLAM (galleries, libraries, archives and museums) sector (Farella et al. 2022a). First, virtualisation projects often encompass extensive collections (Guidi et al. 2013; Malik and Guidi 2018) and the integration of objects with their context. Time constraints introduce further issues, especially for temporary exhibitions or restricted monuments typically closed to the public for preservation reasons. An interest in prototyping exhibitions emerges as a way to improve the setup of physical exhibitions (Pescarin et al. 2018). For example, Benítez et al. (2022) describe the task of digitising the Torre de la Cautiva, where the team had only a ten-hour window to complete the data collection.

Furthermore, environmental factors such as inflexible lighting conditions, limited space and the presence of visitors can impact the acquisition phase (Malik and Guidi 2018; Benítez et al. 2022; Vukovic et al. 2022). Adding to these intricacies, artefacts can vary significantly in shape, size and material, requiring highly adaptable digitisation equipment to manage data acquisition across diverse conditions. In addition to selecting methodologies that optimise acquisition times, adequate time must be planned for processing and delivering 3D results, especially in large-scale digitisation projects. Finally, a significant challenge remains in ensuring the publication and accessibility of 3D model derivatives for reuse, including effective management of all the data generated during the process.

The documentation and accessibility of research outcomes and methodologies are essential for ensuring the reproducibility of workflows in any field. Concurrently, it is equally crucial to define an efficient methodology that meets certain basic requirements (Farella et al. 2022b) in terms of data quality, including a) an accurate, comprehensive and precise reconstruction of the object's shape and geometry, minimising occlusions and preventing any loss of detail; b) high-resolution, realistic textures for optimised geometries, suitable for close inspection in web-based and/or augmented/virtual reality (AR/VR) applications; and c) efficient acquisition and processing times per object to support large-scale digitisation efforts.

Additionally, the literature emphasises a growing need within the GLAM sector to develop accessible pipelines that can be replicated by less experienced staff and within limited budgets, providing a digitisation process that is more sustainable and scalable (Rahaman et al. 2019; Rahaman and Champion 2019).

Concerning the acquisition phase, the cultural heritage digitisation pipelines reveal that scanning technologies and photogrammetry are the most common acquisition methodologies (Apollonio et al. 2021). Photogrammetry, which offers the advantages of lower costs, adaptability and reduced skill requirements, can achieve the same level of geometric accuracy as scanning when supported by metrology kits (Milosz et al. 2020; Sebar et al. 2021). This can be done, for instance, by placing markers that can be identified by the chosen processing software, enhancing accuracy in the capture process. In general, the most suitable acquisition technique is influenced by various constraints, including the object's surface type, size, location and geometric complexity and the space available for acquisition. Scanners, for instance, have strict requirements regarding the acquisition distance range, whereas photogrammetry is more flexible and adaptable. Structured-light projection scanners, although costly, enable the rapid and accurate capture of large volumes of data in a very short time, proving advantageous when time constraints are tight. To optimise photogrammetry acquisition times while maintaining data accuracy, Menna et al. (2016) and Sebar et al. (2021) suggest automated techniques, such as turntables, for digitising small to medium-sized objects, supported by open platform software to control light and shutter release. In this context, using a coded target grid significantly improves camera orientation and model scaling (Sapirstein 2018; Luhmann et al. 2023). To obtain a dense and complete dataset (Luhmann et al. 2023), it is essential to rotate the turntable in angular intervals of at least 30°, ensuring overlap between adjacent shots (Lo Brutto and Spera 2011; Menna et al. 2017). Ideally, the rotation should be repeated at a minimum of two different elevations to simulate a shooting geometry with multiple intersections. The use of acquisition schemas can further speed up the process by allowing for uniform and well-diffused illumination (Webb et al. 2020). Another key aspect is to adjust sensor settings to achieve adequate image quality and sharpness, in alignment with the desired resolution for the intended 3D model. To obtain accurate colour data and exposure settings, colour checkers are generally used, such as the popular X-Rite ColorChecker Classic, which features 24 standardised patches with known reflectance values (Apollonio et al. 2021).

For structured-light scanner acquisition and processing, open-source software options like Open Scan (https://encr.pw/openscaneu) and 3DUNDERWORLD-SLS (https://3dunderworld.org/) are available. However, compared to proprietary systems, these options generally provide lower levels of detail and require technical adjustments that demand specialised skills (Gu et al. 2016).

In photogrammetry, standard workflows for data processing predominantly rely on proprietary structure-from-motion (SfM) software (Jones and Church 2020), with Metashape (https://www.agisoft.com/) being the most commonly used in academic settings (Vukovic et al. 2022; Chapinal-Heras et al. 2023). While proprietary software (as of 14 April 2025) offers higher data accuracy and more advanced functions, open-source alternatives like Regard3D (https://www.regard3d.org/) and Meshroom (https://alicevision.org/) are gaining popularity (Rahaman and Champion 2019). Additionally, Reality Capture (https://www.capturingreality.com/), known for its processing speed, has recently provided a free version for students, educators, hobbyists and companies with annual revenues under $1 million USD, expanding the range of accessible SfM processing options.

In the studies reviewed (Medina et al. 2020; Chai and Li 2022), mesh retopology is mainly done using tools that facilitate automatic procedures, such as Instant Meshes (https://github.com/wjakob/instant-meshes) (Jakob et al. 2015). Modelling and optimisation are often done together, with optimisation being crucial for publishing 3D models online. This typically involves low-poly geometries with high-quality PBR textures, depending on the project's goals. While high-poly models are used for research and documentation, optimised models are preferred for museum displays and AR/VR applications for dissemination and valorisation purposes (Peinado-Santana et al. 2021). In this context, Apollonio et al. (2021) propose creating two versions of a model: a *master model* and a *derived model*, with the latter undergoing remeshing and optimisation for web-based visualisation.

The accessibility of 3D models for data reuse remains challenging. Institutional repositories often limit upload formats, while commercial repositories may lack reliable data provenance or may not guarantee data integrity (Rahaman et al. 2019). One commonly used platform in museum contexts is Sketchfab (https://sketchfab.com/), used by institutions like the British Museum and the Hellenic Museum of Melbourne. Papadopoulos et al. (2025) discuss the transition of Sketchfab to Epic Games' Fab marketplace, highlighting the risks this change poses to cultural heritage, as it may compromise years of 3D modelling work. The study exposes the vulnerability of relying on commercial solutions for data preservation, reflects growing concerns within the community and reinforces the call for developing open and sustainable infrastructures that ensure the accessibility and long-term relevance of 3D cultural heritage in a constantly evolving technological landscape.

In contrast, several academic and open-source solutions have emerged to address these concerns and to develop customised solutions tailored to the specific needs of individual research domains (Ekengren et al. 2024). Notable examples include the Smithsonian's Voyager platform (https://smithsonian.github.io/dpo-voyager/), which offers public access to 3D collections from multiple museums; 3DHOP (https://github.com/cnr-isti-vclab/3DHOP), a web-based 3D viewer developed by the Visual Computing Lab at CNR-ISTI; and the ATON framework (https://github.com/phoenixbf/aton), which was used in our pilot context.

2.2. Overview of 3D formats usage in the international context

Despite the acknowledged importance of 3D in the digitisation of cultural heritage, at present there are still no consolidated standards allowing for the effective design, monitoring and reporting of the process, as highlighted by a recent survey promoted by the European Commission (2022). This survey confirms the persistence of shortcomings that numerous previous studies have already pointed out, including specific aspects such as the absence of suitable repositories (Champion and Rahaman 2020) or the lack of scientific rigour (Statham 2019) in the 3D assets published in light of the recommendations contained in the London Charter (2009) and the Seville Principles (Bendicho 2013). This lack of standardisation even extends to the formats themselves (Hernández-Muñoz 2023), where format selection varies due to contingent boundary conditions such as the prior expertise of the working group or the technological constraints imposed by the chosen publication method, rather than being guided by a systematic evaluation of existing formats according to the specificities of cultural heritage.

Although the present article does not specifically focus on the systematic comparative analysis of existing formats, defining the workflow presented here nevertheless required an evaluation of the existing formats that align with the overall approach and remain efficient throughout the entire processing pipeline—from raw data to publication (table 1).

Table 1: A short overview of select 3D formats, including support for specific properties.

	Type	Mesh	Point cloud	Textures / maps	Animations	Target	Open format
OBJ	ASCII	Y		Y		exchange	Y
FBX	ASCII or binary	Y		Y	Y	exchange	N
PLY	ASCII or binary	Y	Y	Y		raw data	Y
LAS	binary		Y			raw data	Y
E57	binary		Y			raw data	Y
XYZ	ASCII		Y			raw data	Y
glTF	ASCII (JSON) and/or binary	Y	Y	Y	Y	delivery	Y
X3D	ASCII (XML)	Y		Y	Y	delivery	Y
USD	ASCII and/ or binary	Y	Y	Y	Y	delivery	Y

The second crucial factor influencing the selection of a 3D format is its interoperability (Moore et al. 2022), which can be assessed at the technical level by the degree of compatibility with various software tools and at the methodological level by determining whether it is proprietary or neutral. Concerning the latter, to facilitate collaboration between different institutions and scholars and to ensure the long-term preservation of 3D data, the use of open and program-neutral file formats is recommended (Moore et al. 2022). Nonetheless, this general recommendation may pose challenges, as program-neutral formats typically result in the loss of metadata when converted from proprietary formats. On a practical level, however, this issue has a relatively minor impact within the proposed workflow because the injection of metadata occurs during the final stages, using a specific approach and ontology.

To ascertain which formats are used in practice, a useful starting point would be the studies by Champion and Rahaman (2020) and Rahaman et al. (2019), where the authors reported on the main formats used in publicly accessible repositories, both commercial and institutional, for the publication of 3D assets. Among the wide range of formats identified, groupings can be established based on both their primary field of use and current patterns of dissemination. The latter makes it possible to identify formats that have been superseded by others, such as X3D; that are now limited to specific uses, such as DXF and PLY; or that persist solely to maintain compatibility with legacy assets, such as 3DS, LWO, IV and WRL. There are also formats linked to specific software tools, such as BLEND to Blender (https://www.blender.org/), and others used for specific purposes, such as STL for 3D printing and IGES for CAD applications.

The remainder can be categorised according to their preferred field of use, which also influences their suitability for different stages of the workflow. For the sake of brevity, a systematic analysis of each format is not reported in the present paper. The comparison was performed at a qualitative synthetic level by compiling the technical specifications of each format when available (table 1).

Concerning raw data, in scenarios where point cloud data may be required, the PLY, LAS and E57 formats have been identified as the most suitable and compatible options. No further assessment has been made, since all three formats support scalar values and allow the incorporation of customised metrics in addition to X, Y, Z coordinates and R, G, B values for each vertex of the point cloud. Secondly, refining the raw data requires adopting specific features that can ease the task, i.e. polygon/smoothing groups, layers, groups, weights, vertex maps and others, across the whole process. A suitable format would support these features, allowing data to be reprocessed starting from any major phase of the workflow and tapping into all such features. Compatible formats, such as ABC, DAE, USD, FBX and OBJ, are the most suitable options in line with these requirements. ABC is mainly used in the film industry whenever very complex scenes and VFX are involved. DAE presents issues concerning compatibility and coherence amongst different software tools. USD, in contrast, is emerging as a highly promising format and experiencing increasingly widespread adoption. However, it is less commonly seen in smaller-scale repositories and pipelines, primarily due to factors like a lack of familiarity and training. Additionally, gaps in available conversion plugins and tools can hinder its seamless integration into existing workflows. A more conservative approach could be to not adopt it in the current iteration of the workflow so that work-in-progress assets can be stored either in FBX or OBJ format.

OBJ is a geometry definition file format written in ASCII and representing 3D geometry along with a limited set of basic features. However, such simplicity comes at the cost of reduced support for more complex features, thus making it a suboptimal choice in complex pipelines that require advanced capabilities that are better supported by the FBX format.

FBX was initially developed in 2001 as a platform-independent 3D data interchange format to enable high-fidelity data transfer between various graphics software platforms. Unlike OBJ, it cannot be considered an open format, as the format itself remains protected by intellectual property rights despite the release of a software development kit by its owner, Autodesk. Nonetheless, its widespread adoption and efficient data compression rates make it a reliable alternative when working with complex, work-in-progress assets that rely on features not supported by OBJ, as in our case.

Finally, the delivery format must be optimal for Web3D applications, both in terms of supported features and format efficiency for real-time interaction and presentation. The latter is determined mainly by the overall compactness in terms of size (network transmission) and memory footprint (textures, etc.), targeting different devices. It must be noted that many of the other aforementioned formats can be successfully used to upload 3D assets to online viewers and repositories. Still, since none is built explicitly for this purpose, they meet the requirements of efficacy only partially. Currently, X3D is becoming increasingly obsolete, while glTF has emerged as the primary dedicated format for Web3D implementation. Additionally, glTF effectively meets the requirements of an ideal standard format for digital cultural heritage, as proposed by the European Commission (2022): it can be used by a vast number of software applications without conversion, it is interoperable and it is designed to be upgradable while maintaining backward compatibility.

More specifically, glTF (https://www.khronos.org/Gltf) is an open, royalty-free specification for the efficient transmission of 3D scenes and models by engines and applications (Robinete et al. 2018; Lentz et al. 2021). The format, developed by Khronos (https://www.khronos.org/), is designed for compactness and efficiency in real-time processing and rendering while supporting PBR textures for state-of-the-art photorealistic rendition of physically correct behaviour of materials. glTF 2.0 is designed to be vendor-neutral and runtime-neutral and can be employed by a wide range of native and Web3D applications, regardless of the underlying 3D graphics platforms and APIs. Furthermore, glTF is extensible (as a JSON-based format) and retro-compatible by design, making it a suitable solution for universal delivery and in strong alignment with FAIR principles.

The format has evolved beyond its origins as a standalone 3D format. It is now widely adopted as an international ISO standard (https://www.iso.org/standard/83990.html) and has become the cornerstone of several rapidly growing ecosystems of software tools, standards and extensions. glTF's focus on efficiency is a design goal that sets it apart from typical 3D "authoring" formats, which are typically more verbose and incur higher processing costs to transport data that is no longer needed for the final, interactive application. glTF is complementary to authoring formats, providing a common, interoperable distillation target for publishing 3D resources to a broad audience of end users.

2.3 Accessible digital twin creation with 360° immersive photos and 3DGS

In this section, we explore two complementary approaches to digitisation that stand out for their accessibility and low technological difficulty: 360° immersive photography and the 3D modelling technique known as Gaussian splatting. Both methods offer effective solutions for producing digital twins of cultural objects or environments, especially in cases where resources are limited or a quick and simple piece of digital content is required.

Panoramic 360° photos are a powerful tool for creating digital twins that are lightweight, cost-effective and easily accessible. Unlike traditional 3D scanning technologies, which require complex equipment and significant computational resources, 360° images allow for faithful capture and reproduction of cultural and natural environments with remarkable simplicity. Using 360° cameras or even smartphones, it is possible to produce immersive reproductions of spaces and exhibits without the need for expensive, advanced equipment.

These images can be viewed directly in web browsers, eliminating the need to install specialised software or use devices with high computing capacity. This accessibility makes them a sustainable and scalable solution for cultural heritage promotion, allowing a global audience to be reached with minimal technological effort. In a context where sustainability and accessibility are increasingly central, 360° photos function as interoperable, flexible and ready-to-use digital twins, ideal for museums, cultural institutions and educational projects that aim to offer high-quality digital experiences with minimal economic and environmental impact.

Considering immersive photos as a "lightweight" form of digital twins makes providing access to objects and exhibition spaces easier while enhancing the user experience. Furthermore, this approach aligns with FAIR principles, making data findable and reusable through well-structured online repositories and promoting interoperability with other digital resources. To fully adhere to these principles, additional metadata structure is needed for management. Standard formats, such as JPEG and PNG for images and JSON or XML for metadata, ensure seamless integration with other systems.

Finally, in terms of interoperability, 360° JPEG photos can be converted into other formats, such as 3D glTF, when needed. Open-source tools like Blender make this conversion possible. The process typically involves importing the JPEG as a background or environment texture in Blender, building a simple 3D structure to align with the image, and then exporting the model in glTF format. This method adds flexibility for various applications, making 360° immersive photos an effective tool for documentation, accessibility and the enhancement of the cultural heritage experience.

Developing new techniques for creating accessible and reasonably accurate digital twins remains one of the main challenges in the field. In this context, innovative methods such as Gaussian splatting (Kerbl et al. 2023) are revolutionising the field by offering a more streamlined and accessible solution.

Gaussian splatting—unlike traditional methods such as photogrammetry or laser scanning, which require specialised tools and long processing times—reconstructs a scene by projecting millions of coloured points into a 3D space, each carrying its own information, including position, colour, transparency and size. These points are called Gaussians because their shape is inspired by the Gaussian distribution, a mathematical function which defines a soft and nuanced probability distribution around a central point. The term splatting, on the other hand, refers to the action of projecting or "splatting" these Gaussians onto the screen, as if they were brush strokes floating in space, and which, when superimposed on each other, visually reconstruct the appearance of the object or an entire scene. The final effect is a continuous and detailed 3D representation that is obtained without generating a real three-dimensional mesh. The whole creation process starts from a dataset of sequentially obtained images that have a high level of overlap (about 70%) and represent the object or environment in all its parts. The major advantage of this technique is its ability to generate realistic 3D models using photos taken with simple mobile devices (e.g. smartphones) without the need for expensive professional equipment. Gaussian splatting enables the creation of lightweight digital twins that can be viewed in standard web browsers. This solution can be fairly accessible, despite some limitations regarding the visualisation parameters of the Gaussians and the size of data sets. However, there are still good prospects for improvement regarding these aspects. Compared to traditional methods, this technique democratises access to 3D modelling, making it available to museums and cultural institutions of all sizes and opening up new opportunities for digital engagement.

The use of these two digitisation methodologies varies depending on specific needs. For example, a general use case could involve enhancing an extended cultural landscape—like a literary park or an open-air museum—through simple immersive navigation of the entire space. Users would be able to explore the environment and focus on key details (e.g. a statue, a fountain, a work of art) by viewing these objects in high-quality 3D and closely examining their features.

In practice, a park or museum could easily create an immersive tour by connecting 360° photos, and the user could interact with the immersive environment through interactive points linked to 3D models (Gaussian splatting). This layered and combined approach improves the experience by balancing ease of navigation with the visual precision of 3D and allows for a simple, fast and relatively cheap way to enhance cultural heritage.

2.4 Metadata for traditional and digital CH objects

In the digitisation of cultural heritage, metadata plays a pivotal role and demands specific expertise. Firstly, transposing museographic information into a digital environment requires a rigorous transfer of existing metadata related to the digitised object, allowing for an enhancement of their informational value, for example, through semantic web technologies. Moreover, the digitisation process itself produces new metadata concerning the creation of the digital objects. This metadata is directly derived from the original information and constitutes an independent entity in its own right. For this reason, every step of the 3D model creation process should be documented to ensure reproducibility and clarity, considering that reality-based models can be influenced by human or software interpretation (Moore et al. 2022; Barzaghi et al. 2024a).

Having a well-defined and rigorous pipeline for metadata management factually increases the informative potential of digitised objects. It enables data visualisation, improved exploration and the tracking of provenance and changes over time, while also allowing the use of research products whose effectiveness has already been tested in previous case studies. Even though only a few workflows for managing metadata concerning 3D cultural heritage objects are openly available, we list below some relevant studies published on the topic and their proposed approaches.

Some early works explicitly depicting pipelines for these purposes date back more than a decade. Starting in 2010, the CARARE project allowed for the definition of a metadata schema for GLAM 3D digitised objects, which is required for their inclusion in the Europeana platform (D'Andrea and Fernie 2013). Following that, the 3D-ICONS contributed a substantial number of 3D models of archaeological and artistic items to Europeana, highlighting the need for a structured workflow to uniquely identify the items, make them publishable and harvestable on the platform through descriptive metadata and track the progression of the digitisation process via technical metadata (Guidi et al. 2013).

Around the same period, proposals emerged for the semantic exploitation of metadata related to 3D content. One such example is the development of an RDF-compliant methodology for harvesting Extensible 3D (X3D) web content through use of the Gleaning Resource Descriptions from Dialects of Languages (GRDDL), which enables semantic descriptions of real-world objects and their digital 3D reproductions referenced by URIs (uniform resource identifiers) (Flotyński 2013). The use of metadata to enhance the searchability of 3D digital items is also addressed in a study by Kolosova and Hermon (2013) focusing on real-time item visualisation and user experience in the X3DOM framework. In this case, metadata files associated with each object allow for an enhanced navigation experience through the content. Research concerning the AR/VR sector has further contributed to advancing the formalisation of metadata management procedures for 3D digitised CH objects. An example is the 5W1H-based metadata schema for context-aware AR applications, which links physical points of interest to AR contents through AR anchors and semantically defines the context through the five w-questions and specification of the interaction modality (Kim et al. 2015).

Today, new promising formalisations of metadata management pipelines are being proposed to exploit information about 3D digital items as linked data. A recent study concerning ontologies and metadata schemas for the capturing process of 3D heritage objects stressed the nature of digitisation as a metadata-generating activity at every stage, highlighting the need for tools and guidelines to properly record this information. The researchers produced a metadata schema based on the specific 3D scanners used in their case studies (Roman burial monuments in Trier, Assyrian cuneiform tablets and the preservation of antique wood samples), but they also openly published the obtained datasets on Zenodo (https://zenodo.org/), providing examples for accessible reproduction of the proposed workflow. They also extended the free and open-source GigaMesh Software Framework (https://gigamesh.eu/) to provide metadata about the 3D processing steps (Homburg et al. 2021). Most recent studies increasingly emphasise the importance of reusing existing standards rather than creating new purpose-specific ones for individual case studies. Amato et al. (2023), for instance, describe the procedure adopted for metadata management in two pilot studies—one on Samnite settlements and the other on the Roman Appia road—emphasizing the value of adopting application profiles that combine components from existing namespace schemas and optimise them for local applications.

In terms of reproducibility, the metadata related to objects, processes and provenance are pivotal in defining workflows for the cultural heritage domain (Barzaghi et al. 2024b). In this regard, a notable example is the tailored set of FAIR principles—introduced in 2018 specifically for archives, heritage libraries and museum collections—which emphasises the critical role of metadata (Koster and Woutersen-Windhouwer 2018). This approach defines three distinct levels of metadata: the object level (e.g. artefacts, datasets), the object metadata level (e.g. title, creator) and the metadata records level (e.g. the comprehensive set of metadata elements describing the object). The process/workflow data plays a key role in generating the metadata that accompanies the digital cultural heritage objects created during the digitisation process.

The definition of metadata should account for its model and representation schema, as these are crucial for data exchange and interoperability. In this regard, efforts have been made to provide guidelines that enhance data exchange across various levels—technological, semantic, organisational, legal and syntactic. This includes the development of a crosswalk for data models, vocabularies and aggregator guidelines (European Commission 2021). Crosswalk practices aimed at improving interoperability often involve making metadata available as Linked Open Data using a resource description framework (RDF) serialisation. A powerful tool for converting structured data into semantic triples is the RDF mapping language (RML) (Dimou et al. 2014). Given the non-trivial nature of the mapping task, RML provides detailed documentation and allows users to potentially avoid or significantly minimise direct interaction with code by defining RDF mapping files, in which they can specify the rules for converting JSON or CSV files into an RDF serialisation, using a data model of their choice. A more human-friendly option is also available, based on the use of a simple YAML file, whose low level of abstraction makes the procedure more accessible to less experienced users (Heyvaert et al. 2018). The Java codebase can be extended if the built-in functions do not fulfil user requirements for converting input datasets. To improve the factual reuse of this mapping approach, a tool named Morph-KGC (Arenas-Guerrero et al. 2024a) was developed in Python, broadening the potential user base and enabling the extension of the code with new user-defined functions (Arenas-Guerrero et al. 2024b).

For the long-term preservation of the resulting collections of information and maintenance of their trustworthiness, supporting the tracking of both object provenance information (OPI) and metadata record provenance information (MRPI) is pivotal. It is essential to capture changes in data and metadata over time, as digital representations of artefacts often evolve, either due to new findings or advancements in digitisation technologies. One possible practical approach to storing provenance information and tracking data changes is through the OpenCitations Data Model (OCDM) (Daquino et al. 2020). OCDM provides a framework for the comprehensive documentation of the entire lifecycle of digital artefacts using semantic web technologies. It captures snapshots when an entity is created, modified, merged or deleted and stores them in a provenance-named graph. Each snapshot includes information such as validity dates, responsible agents, primary sources and a link to the prior snapshot, ensuring a transparent and traceable record of data changes.

3. Acquisition and digitisation workflow

3.1. Managing objects and processes metadata

In the following sections, we describe the main activities in all the phases (fig. 1) that characterise the acquisition and digitisation process used to capture the cultural heritage objects (CHOs) featured in the temporary exhibition *The Other Renaissance: Ulisse Aldrovandi and the Wonders of the World*. The 301 CHOs featured in the exhibition were primarily from the permanent collection of Palazzo Poggi, with a smaller selection on loan from various cultural institutions (Balzani et al. 2024).

A wide array of technologies and methods were employed to develop and gather 3D data. In our case study, we chose to retain three distinct digital versions of each CHO, in addition to the raw material (RAW) generated during the acquisition phase (fig. 2):

- Processed raw model (RAWp): This is the preliminary output from the photogrammetry or scanner software after initial data processing but without any interpolation or geometry corrections.
- Digital cultural heritage object (DCHO): This version includes interpolation, gap filling and resolution of geometric issues, resulting in a refined and improved model.
- Optimised digital cultural heritage object (DCHOo): This version is optimised for seamless, real-time online interaction.

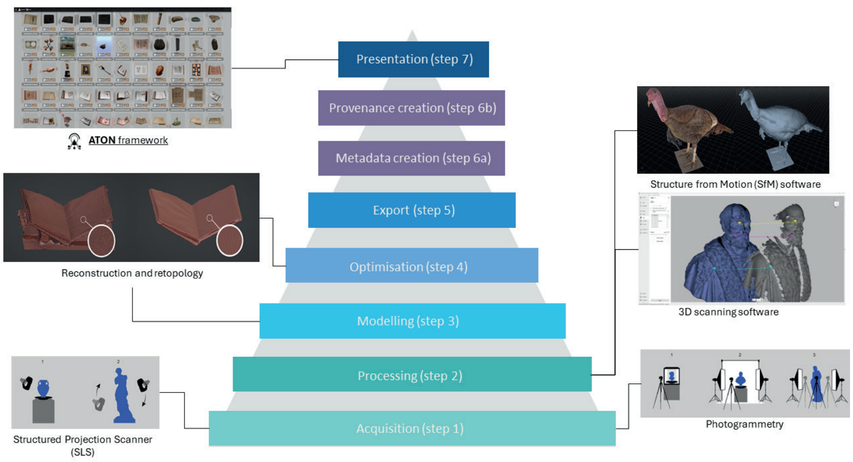

Figure 1: The acquisition and digitisation process.

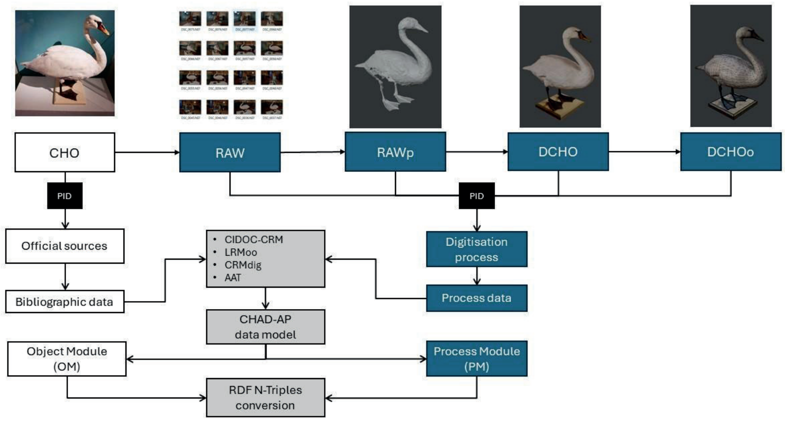

Figure 2: Parallel processes, digitisation and data management.

The metadata management process took place alongside the creation of 3D data, as illustrated in figure 2. In a digitisation workflow that aims to be reproducible, this process involves several prerequisites. First, objects (both physical and digital) must hold a unique global persistent identifier (PID). Second, the objects must be described with metadata, including the specification of their PIDs. Third, metadata records themselves must also have their unique PIDs. In addition, metadata are produced and managed at various levels, including metadata related to the exhibited objects and metadata resulting from the digitisation process itself.

Metadata production and management in our digitisation process started with data collection, which involved the creation of two tabular datasets. The first table was used to register cataloguing information (bibliographic data) about the CHOs. The second table focused on collecting information created during the acquisition and digitisation activities (process data), including the metadata accompanying the DCHOs created during digitisation (fig. 2). Table creation involved the definition of their structure, column names, expected cell data and controlled values for certain columns (with their values also aligned with existing authority files). While the bibliographic data table was structured according to official sources (such as the exhibition catalogue and curators' notes), the process data table was based on the steps that are part of the digitisation process and the attributes deemed relevant for tracking each step's provenance.

The acquisition phase, summarised in figure 1 and explained in section 3.2.1, aimed to capture CHOs and the RAW material needed to create the corresponding DCHOs. After acquisition, the generated raw data was reused throughout the remainder of the digitisation process in a sequence of activities, which could vary depending on the nature of the materials and the intended use of the digital files. However, the common procedure tended to involve the following phases:

- Processing (section 3.2.2), in which the RAW data produced during the acquisition phase is used to generate a processed raw model (RAWp) via automatic algorithms (with the human operator setting the input parameters);

- Modelling (section 3.2.3), in which the human operator resolves any topological issues in the RAWp, making, when needed, subjective interpretative contributions, resulting in the creation of the DCHO;
- Optimisation (section 3.2.4), which simplifies the DCHO for specific use purposes, creates an optimised DCHO (DCHOo);
- Export (section 3.2.5), in which the RAWp, DCHO and DCHOo data is converted into a specific format;
- Metadata creation (section 3.2.6), in which the metadata tables are exported in CSV format, and then converted into RDF N-Triples to create structured information for bibliographic and process data for CHO, RAW, RAWp, DCHO and DCHOo models;
- Provenance creation (section 3.2.7), in which provenance information is tracked and created for each metadata record;
- Presentation (section 3.2.8), in which DCHOos are moved from a local device or storage location to a web-based framework (e.g. ATON).

In this context, we drafted a specific and detailed document (Bordignon et al. 2024) describing the best practices tested within the current project. Our guidelines are aligned with those established by the Italian Ministry of Culture (MiC) (https://digitallibrary.cultura.gov.it/linee-guida/), with certain aspects having been simplified, as our guidelines are intended for purposes of valorisation and dissemination rather than documentation, as is the case for the MiC. Since the project in which this work was carried out is national in scope, our guidelines are currently available only in Italian, but an English version is in preparation. As detailed in sections 3.2 and 3.5, the workflow adopted in this project was specifically designed to align as closely as possible with the principles of open science, with particular emphasis on reproducibility, interoperability, the use of open platforms and technologies and the adoption of data formats that ensure long-term accessibility. While we acknowledge that companies involved in the digitisation of cultural heritage often follow their own well-established workflows, the proposed methodology is not intended as a one-size-fits-all solution. Instead, it offers a transparent and traceable framework that can add value in contexts where methodological clarity, workflow transparency, data sharing and cross-institutional collaboration are essential. The approach aims to facilitate the final dissemination and reuse of results at the time of publication, scheduled for thend of 2025 on Zenodo.

Each of the phases described above resulted in the production of distinct objects and their corresponding metadata, all of which needed to be properly recorded and described. Based on the bibliographic and process data collected in the tables, along with the metadata used to describe them, a semantic data model was developed to represent the CHOs and the various DCHO versions, as well as the processes involved in digitising the former into the latter, to support the reproducibility of the entire digitisation workflow. To achieve this, existing standards were reused, including the CIDOC Conceptual Reference Model (CIDOC-CRM) (http://www.cidoc-crm.org/cidoc-crm/) (Doerr 2003), the object-oriented Library Reference Model (LRMoo) (http://iflastandards.info/ns/lrm/lrmoo/) (Riva et al. 2022), the CRM digital extension (CRMdig) (http://www.cidoc-crm.org/extensions/crmdig/) (Doerr and Theodoridou 2011) and the Art & Architecture Thesaurus (AAT) (http://vocab.getty.edu/aat/) (Harpring 2010). In addition, the model's development and documentation followed the Simplified Agile Methodology for Ontology Development (Peroni 2017), which further contributes to the reproducibility and reusability of the overall process.

The result was the creation of the Cultural Heritage Acquisition and Digitisation Application Profile (CHAD-AP) (https://w3id.org/dharc/ontology/chad-ap) (Barzaghi et al. 2024c), an application profile encoded in OWL for describing as structured, machine-actionable data CHOs, DCHOs and the processes of acquisition and digitisation in the cultural heritage domain. CHAD-AP is structured around two main modules: the Object Module (OM) and the Process Module (PM).

The OM allows CHOs to be represented as constructs whose metadata is distributed along four conceptual layers, namely: 1. Work (the essence of the object, modelled as `lrmoo:F1_Work`), which is characterised by one or more titles (each an instance of `crm:E35_Title`), as well as relationships with other Works; 2. Expression (the intellectual realisation of the Work, modelled as `lrmoo:F2_Expression`), which is tied to the set of entities that contributed to the CHO's creation (usually instances of `crm:E7_Activity`, each with its own information, such as `crm:E39_Actor`, `crm:E55_Type`, etc.) as well as its content (including its subjects as instances of `crm:E73_Information_Object`); 3. Manifestation (the embodiment of an Expression in a physical carrier, modelled as `lrmoo:F3_Manifestation`), which describes the CHO's type (`crm:E55_Type`) and its related license (a semantic pattern centred around `crm:E73_Information_Object`); and 4. Item (the actual, physical CHO exemplar, modelled as `lrmoo:F5_Item`), which is accompanied by descriptive labels (identifiers and textual descriptions) and eventually linked to information related to its curation and conservation.

The PM focuses on defining a 3D digitisation workflow as a sequence of activities: the first, acquisition (`crmdig:D2_Digitization_Process`), involves the CHO at the Item level (the physical object being acquired) and produces the RAW file (`crmdig:D9_Data_Object`); the other activities, collectively referred to as "software activities" (`crmdig:D10_Software_Execution`), represent processing, modelling, optimisation, and so on. Both the acquisition and software activities involve other entities, including 3D data (`crmdig:D9_Data_Object`)—which passes from one activity to the other as either input or output (such as RAWp and DCHO)—people (`crm:E21_Person`) and organisations (`crm:E74_Group`) responsible for carrying out the activities, techniques (`crm:E55_Type`) and tools (`crmdig:D8_Digital_Device` and `crmdig:D14_Software`) involved and the activities' temporal information (`crm:E52_Time-Span`).

As detailed in section 3.2.6, once data collection was complete, the metadata tables were exported in CSV format and then converted into RDF N-Triples to create structured information for bibliographic and process data for each CHO, RAW, RAWp, DCHO and DCHOo, using the CHAD-AP data model as the vocabulary for semantic representation and the custom extension of the Morph-KGC software for converting the data into a knowledge graph based on said model.

3.2. Asset acquisition and digitisation workflow

The 301 CHOs selected for digitisation presented a wide range of characteristics in both geometric shapes and surface properties, which added further challenges to the existing time and space constraints.

It is widely acknowledged in the scientific literature and practical application that, regardless of the technique employed, the digitisation process typically struggles to capture (with sufficient accuracy) critical portions of CHOs—e.g. undercuts, occluded areas, thin elements and other geometric features—at frequencies approaching the sensitivity threshold of the instruments used. Moreover, the frequent presence of non-Lambertian materials, which reflect and scatter light differently depending on the viewing angle, can further hinder complete and adequate acquisition due to phenomena such as specular reflections, transparency, translucency, iridescence and anisotropy. These optical behaviours negatively affect scanner-based acquisition by interfering with sensor readability as well as photogrammetric acquisition since SfM and MVS algorithms rely on detecting consistent features across multiple images, features that appear differently in non-Lambertian materials depending on the perspective.

These adverse factors can lead to misreadings and undersampling or even significant gaps in the data, which is why the RAWp model cannot be considered objectively reliable in absolute terms, even when it appears complete. The refinement of the RAWp into a DCHO explicitly addresses this issue by correcting gaps and inconsistencies with the physical object through automatic, semi-automatic or, in the most critical cases, manual interventions.

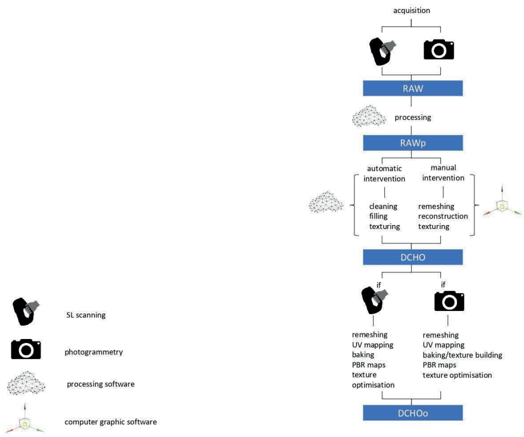

Figure 3: 3D data taxonomy established as research output.

Thus, even state-of-the-art acquisition does not guarantee the highest accuracy per se, and digitisation remains an inherently critical process. Also, the RAWp digital model is a representation of the physical object, just as the DCHO is, though this may be more intuitively evident in the latter. However, through a refinement process involving the removal of low-confidence areas and the integration of missing parts, the DCHO can achieve greater adherence to the original object, albeit at the cost of more intensive processing of the raw data. A necessary condition for this outcome, however, is that the refinement process be guided by solid and documentable evidence.

The CHOs we digitised included a variety of materials such as woodcuts, technical/scientific instruments, statues, specimens, archaeological artefacts and paper-based items like manuscripts, printed books and historical maps. Figure 3 presents the 3D data taxonomy obtained by testing this specific digitisation process. The schema serves as a reliable methodological reference model for future projects.

3.2.1. Step 1: Acquisition phase

The work began with a thorough analysis of the artefacts and the acquisition spaces. Managing such an extensive and diverse collection required the use of established techniques. Table 2 was created to represent the main types of objects and cases encountered during digitisation. The objects presented summarise the main problems and characteristics that influenced the choice of the acquisition methodology and the interventions in the following phases.

Consequently, the acquisition process employed reliable remote-sensing technologies frequently used in cultural heritage preservation such as 3D structured-light scanning (SLS) and digital photogrammetry (Bitelli et al. 2007;

Apollonio et al. 2021). Each CHO to be surveyed was initially examined and classified based on dimensions, geometry, surface irregularities, material properties, manoeuvrability and accessibility.

Based on these characteristics, the most suitable methodologies, problem mitigation strategies, tools and thresholds for geometry and image resolution were identified and established. For example, because the use of structured-light projection scanners requires strict adherence to a specific acquisition distance range, operators needed to be able to move around the CHO and maintain the correct distance. This was not always feasible when obstacles were present or space was limited. In cases where space was restricted or CHOs could not be moved, we opted for photogrammetry over SLS, as it is less constrained by acquisition (Ruiz et al. 2022).

Additionally, given the temporary nature of the exhibition, the limited time available for object acquisition made SLS a valuable tool for capturing large volumes of data quickly. We used all available equipment to achieve our objectives and gather satisfactory data in the shortest time possible. Generally, we prioritised the use of structured-light scanning for movable, medium-to-small CHOs (coins, woodcut tablets, etc.) with intricate details. Photogrammetry was used for both movable and immovable medium-to-large CHOs (specimens, manuscripts, scientific instruments, etc.). However, depending on the type of CHO being digitised, the teams could select from various scanner models suited to the CHO size.

This approach made it possible to adopt optimal geometric resolutions and image quality standards during the acquisition phase, tailored to the specific features of each artefact, while enabling research groups to proceed with the digitisation process simultaneously and efficiently.

SLS is an acquisition technique based on projecting calibrated blue or white light patterns onto the object's surface and obtaining its geometry in real-time (Geng 2011). For the CHOs acquired with this technology, proprietary software Artec Studio 15 Professional (https://www.artec3d.com/it/3d-software/artec-studio) and two Artec structured-light scanners with very high 3D point accuracy were used (Artec 3D 2023a; Artec 3D 2023b): Artec Space Scanner, to acquire small-to-medium objects (from 1-5 cm to 100 cm), and Artec Eva, to acquire medium-to-large objects (from 5-20 cm to more than 200 cm).

The first step was to set the "Geometry + Texture" scanner acquisition mode to enable simultaneous capture by the many dedicated cameras. Next, the most appropriate acquisition schemes (fig. 4) were identified as a function of a CHO's size, manageability and accessibility (for SLS, immovable CHOs stored in non-opening display cases were not considered):

1. For small-to-medium-sized movable CHOs, manual or turntable movement of the CHO itself was chosen;
2. For small-to-medium and medium-to-large immovable CHOs, the chosen approach was to move the operator around the CHO and use a ladder to reach even the highest and most difficult-to-acquire parts.

Table 2: Master data from five CHOs, representing the variety of objects and methods used in the proposed workflow.

	Object scale	Operating conditions	Acquisition Technology, Equipment and Devices	RAW data	RAWp data	DCHO	DCHOo
Rattle with Thevetia peruviana seeds 	• Small scale object, complex geometry • Bounding box diagonal: 11.34 cm	• **Surface complexity:** fragile object, without the possibility of being rotated, complex and articulated geometry • **Acquisition scheme:** capture from one side only, Lightbox with turntable, fixed camera on a tripod	**Photogrammetry** Panasonic DMC-LX100 CMOS 17.3 × 13mm 4112 × 3088px 4.19 µm 24–75 mm GSD: 0.003	• n. photos: 54 • n. target: 4	**Processing:** 3df Zephyr • Vertex: 212,315 • Polygons: 164,152 • Texel: 2,27 mm/pixel • Texture 8k	• Solving non-manifold issues • Authorial intervention: of the back of the seads and the joining threads • Texture 8k	• Automatic retopology in Instant Meshes • Vertex: 64.538 • Polygons: 65.025 • Texture 4k • PBR maps: 2k-4k
Adnotationes animalibus volume 	• Medium scale object, simple geometry • Bounding box diagonal: 46.43 cm	• **Surface complexity:** fragile object, without the possibility of being rotated, constant presence of museum staff needed • **Acquisition scheme:** non-movable objects, 2–4 continuous led lights on stands and cameras on tripods (moved around the object)	**Photogrammetry** Panasonic DMC-LX100 CMOS 17.3 × 13mm 4112 × 3088px 4.19 µm 24–75 mm GSD: 0.003	• n. photos: 38 • n. target: 8	**Processing:** 3df Zephyr • Vertex: 60.179 • Polygons: 119.983 • Texel :1,439 px/mm • Texture 8k	• Solving non-manifold issues • Authorial intervention: reconstruction of the back • Texture 8k	• Automatic retopology in Instant Meshes • Vertex: 31.927 • Polygons: 63.850 • Texture 4k • PBR maps: 2k-4k
Wooden Sawfish matrix 	• Medium scale object, simple geometry • Bounding box diagonal: 34.77 cm	• **Surface complexity:** movable object, good state of preservation	**Light Structured Scanner** Artec Space Spider2a 3D point precision: 0.05 mm 3D resolution: 0.1 mm Texture resolution: 1.3 MP Acquisition surface: 90x70 mm, 180x140mm Acquisition distance: 20-30 cm	• Fps: 6-7	**Processing:** Artec Studio 16 • Vertex: 1,614,083 • Polygons:3,228,162	• Solving non-manifold issues • Texture 16k-8k	• Automatic retopology in Instant Meshes • Vertex: 6684 • Polygons: 12456 • Texture 4k • PBR maps: 2k-4k
Quetzal bird embalmed model 	• Medium scale, non-Lambertian and unconventional materials • Bounding box diagonal: 135.3 cm	• **Surface complexity:** object in non-removable glass case, complex and articulated geometry, inconsistent surfaces (feathers) • **Acquisition scheme:** 2–4 continuous led lights on stands and cameras on tripods moved around the object	**Photogrammetry** Canon EOS 6D CMOS 36x24mm 5472x3648px 6.54 µm 50mm GSD: 0.008	• n. photos: 112 • n. target: 13	**Processing:** Metashape • Vertex:244, 979 • Volygons:489,135 • Texel: 0.001px/mm • Texture 8k	• Solving non-manifold issues • Authorial intervention: modelling of the bird • Alpha planes used for the feathers • Texture 8k	• Automatic retopology in Instant Meshes • Vertex:68,339 • Polygons: 84,271 • Texture 4k • PBR maps: 2k-4k
Narwhal horn 	• Large scale, simple geometry and texture, cylindrical shape • Bounding box diagonal: 281,51 cm	• **Surface complexity:** large, non-movable object, possibility to circle around the object	**Light Structured Scanner** Artec Eva 3D point precision: 0.1mm 3D resolution: 0.2mm Texture resolution: 1.3MP Acquisition surface: 214x148mm 536x371mm Acquisition distance: 40-100cm	• Fps:10	**Processing:** Artec Studio 16 • Vertex:1048049 • Polygons:2096094	• Solving non-manifold issues • Texture 16k-8k	• Automatic retopology in Instant Meshes • Vertex: 69224 • Polygons:131004 • Texture 4k • PBR maps: 2k-4k

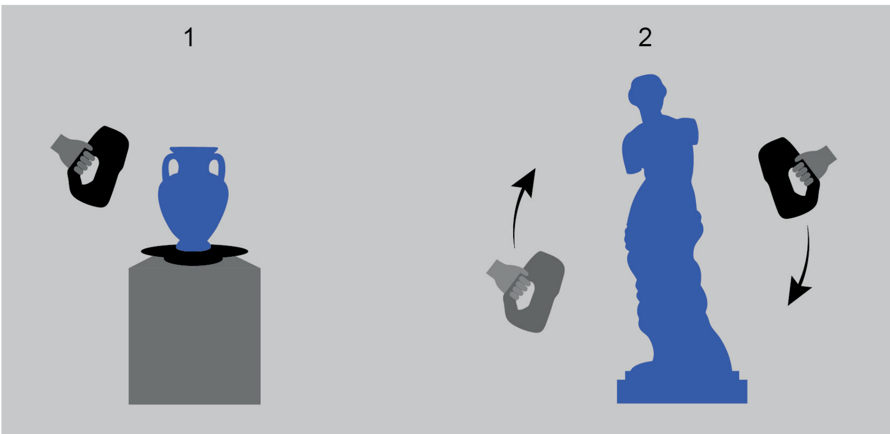

Figure 4: SLS acquisition schemes.

The CHOs were then captured following five key parameters: positioning the scanner at the appropriate distance and angle (both before and during the capture); setting the FPS (frames per second) according to each CHO's characteristics; evaluating "difficult to acquire" CHOs (shiny or reflective surfaces, or those with narrow, dark geometries) and adjusting scan sensitivity and brightness accordingly—while keeping in mind that higher sensitivity can increase noise; maintaining a valid area of overlap between consecutive scans; and acquiring a sufficient number of scans to cover 100% of the visible and acquirable surface of the CHO (McMillion 2022; Artec 3D 2020). In this case, the RAW consisted of a 3D point cloud stored in SCAN format (Artec 3D scan format).

For CHOs acquired through digital photogrammetry, various configurations and equipment were used based on the CHO's size, material and condition. The setup began with a controlled lighting environment, where diffuse light sources were arranged to minimise shadows and reflections, especially for non-Lambertian CHOs. Three configurations (fig. 5) were employed to adapt to the size and mobility of the CHOs:

1. For small movable CHOs, a lightbox, a turntable and a fixed camera on a tripod were used to ensure controlled lighting and consistent angles.
2. For medium-sized movable CHOs, the setup included two to four continuous LED lights on stands, a photographic studio backdrop, a turntable and a fixed camera on a tripod.
3. For large or immovable CHOs, two to four continuous LED lights on stands were positioned around the CHO. At the same time, the camera, mounted on a tripod, was moved around it to capture all necessary angles.

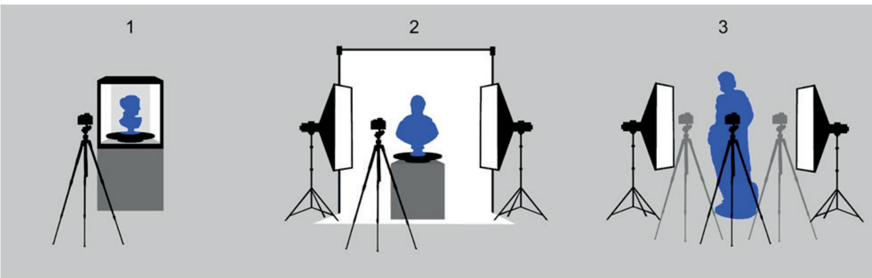

Figure 5: Photogrammetry acquisition schemes.

The ground sample distance (GSD) was calculated by determining the optimum sample resolution required to capture the sharpest detail on each sample. This calculation began with an assessment of the smallest feature that should be represented in the model. To achieve this, a ratio of 1:2 was applied, using Nyquist's theorem as a guide (Kraus 2011, 100) and consequently ensuring that the GSD was at least half the size of the smallest detail (i.e. if the minimum required feature size was 1 mm, the GSD was set to 0.5 mm) (Triglav Čekada et al. 2010, 7). This ratio allowed each pixel to represent a small enough portion of the surface to capture the desired level of detail and ensure that the texture generated from the images accurately represented the surface of the CHO. The GSD was determined by the distance of the camera to the CHO, the lens characteristics and the resolution of the camera sensor. By positioning the camera at an appropriate distance and choosing the correct lens and sensor resolution, the calculated GSD could be consistently achieved, allowing a balance between the level of detail required and manageable data size. This systematic approach ensured that the average GSD across all images met the project's accuracy and resolution requirements. However, as the CHOs are not flat and vary in shape, the GSD was approximate and often had to be adjusted. In cases where the CHO geometry made a precise 1:2 acquisition impossible, the GSD was occasionally set to a higher ratio, depending on the specific shape and surface complexity of the CHO in question.

For complete coverage of each CHO, we followed a 360° capture approach, moving around the CHO or using a turntable when possible at even intervals (around 15 degrees). This setup achieved an image overlap of 70-80%, providing sufficient data for accurate alignment and depth mapping during the processing phase. Camera settings were adjusted to capture sharp, high-quality images. Exposure, ISO and aperture were modified to balance brightness and detail: low ISO was used to reduce image noise, and a small aperture (f/11 to f/16) provided a greater depth of field, ensuring that all areas of the CHO remained in focus. Colour checkers were used to calibrate the white balance.

Finally, markers were strategically placed around the CHOs to serve as reference points during 3D reconstruction. These markers allowed for precise scaling and improved the alignment and positioning of images within the photogrammetric software, ensuring the resulting 3D model accurately reflected the artefact's geometry. To achieve this, different homemade marker configurations were used according to the acquisition configuration:

- Scale bars: Rigid bars with coded markers placed at known distances, which provided reference points for scaling.
- Grids: Coded markers printed on a grid and evenly distributed, with known coordinates for each marker. The grid was placed beneath the CHO or attached to the turntable.

Although a metrology kit would have achieved sub-millimetre accuracy, the homemade marker setup provided sufficient accuracy (around 1 mm) to meet the project's requirements.

During the acquisition phase, various cameras and lenses were selected based on average GSD, CHO dimensions and accessibility, with a preference for full-frame cameras equipped with standard lenses (35-50 mm). However, in some cases where accessibility posed challenges, such as with immovable CHOs or CHOs in fixed display cases, alternative solutions were required. These included using a panoramic camera (Insta360) or wide-angle lens to capture areas that were otherwise difficult to reach.

3.2.2. Step 2: Processing phase

Once the RAW scans had been captured from all necessary angles and in sufficient number, they were refined and processed using software tools (fig. 6) in preparation for the creation of the 3D model. As a first step, the original scans (RAW) were exported to create a backup in case of loss or the need to reprocess a 3D model. Before starting the generation of the RAWp models, the scans were reviewed to detect any acquisition errors. Next, individual scans were cleaned of extraneous object elements captured during the acquisition (such as the operator's hands, walls, the surface on which the object rests and other irrelevant objects) using the `erase` tool offered by the software (Artec 3D 2020, 112-116). After refining the scans, those useful for alignment were selected, and `rigid alignment` was performed in one of two ways: manually, by dragging and positioning them as closely as possible to each other (drag-and-drop method), or by using homologous points, where at least three corresponding points were identified in both scans and the `align` command was executed (Artec 3D 2020, 117-124). After finishing the alignment of all the scans, the next step involved the application of two algorithms in succession: global registration, which transformed the surfaces of each frame to correctly align them with each other (Artec 3D 2020, 130), and outlier removal, which worked as a filter, automatically removing extraneous fragments from the main surface (Artec 3D 2020, 112-113). As a final step, the scans were merged using the `sharp fusion` command without applying automatic gap closure to obtain the most detailed RAWp (Artec 3D 2020, 137-140). The RAWp obtained in this phase was not textured; for models acquired with structured-light scanners, the texturing process took place during the subsequent modelling phase (section 3.2.3).

The processing of photogrammetric data, on the other hand, involved several stages (fig. 3). Since the quality of the surface reconstruction in the photogrammetric workflow is significantly influenced by the quality of the dataset used, it was essential to define a standard methodology to address potential issues during data elaboration (De Paolis et al. 2020).

Before processing the dataset, an audit was carried out to verify whether any adjustments to the RAW files were necessary in dedicated photo editing software (fig. 7). The main software used included Adobe Lightroom (https://www.adobe.com/it/products/photoshop-lightroom.html), Camera Raw (https://adobe-camera-raw.it.softonic.com/) and open-source Raw Therapee (https://www.rawtherapee.com/), for correcting exposure, sharpness, highlights, shadows and white balance. The edited images were then exported in TIFF using LZW compression or JPEG format.

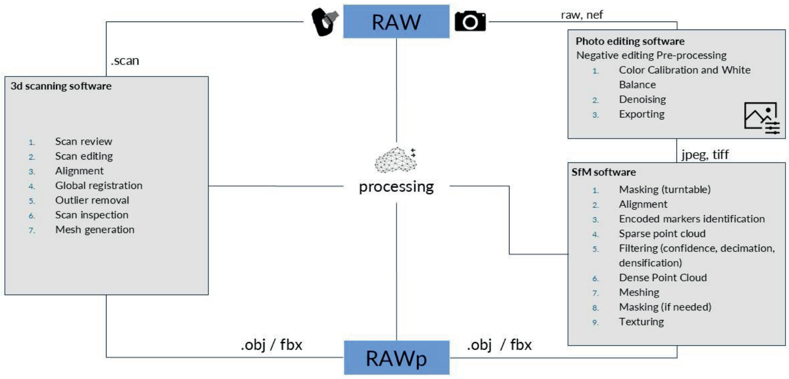

Figure 6: Processing phase from RAW to RAWp, depending on the acquisition technique.

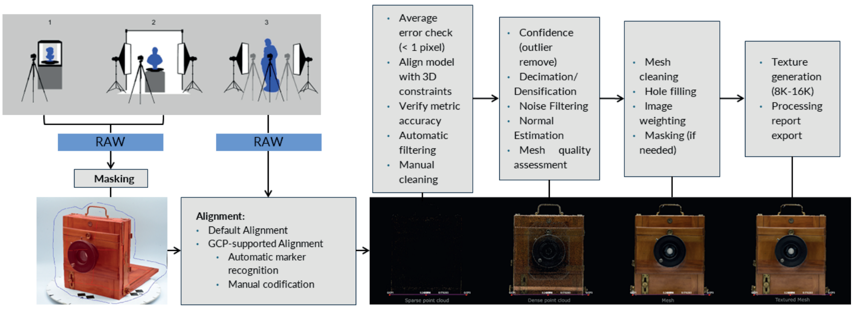

Figure 7: From RAW data to textured mesh in 3DF Zephyr data processing.

After performing a quality check on the images, the dataset is loaded into the structure-from-motion software. For this phase, the main software tools used were Agisoft Metashape (https://www.agisoft.com/), 3DF Zephyr (https://3dflow.net/it/software-di-fotogrammetria-3df-zephyr/) (table 2) and Meshroom, which is based on the open-source AliceVision (https://alicevision.org/#meshroom) framework, currently under testing. For images acquired using a turntable (acquisition schemas 1-2 in fig. 5), it was necessary to apply automatic or manual masks to exclude the background and isolate the object (fig. 7). Masking refers to the process of manually or automatically defining the areas of an image that should be ignored during feature detection and matching. This technique is commonly used to exclude unwanted elements such as moving objects, background noise or reflections that could interfere with the accuracy of the 3D reconstruction. Masking can be particularly useful when working with turntable-based image sets, where the background remains static, or in scenarios where specific parts of an image need to be isolated to improve alignment and point cloud generation (Brandolini and Patrucco 2019; Farella et al. 2022b).

Next, the image alignment process was carried out. In this phase, the software identifies key points—distinctive features within individual images—and establishes tie points, which are matches of those features across multiple images to calculate the relative position of the cameras. This process generates a preliminary low-density point cloud (sparse point cloud) that represents the approximate geometry of the object. In this phase, it was important to ensure that all images were properly aligned by examining the key points and alignment report. To proceed, the average error of < 1 pixel was considered the optimal value to ensure high accuracy in the following steps. To achieve greater accuracy in camera alignment, ground control points were used, through either automatic recognition or manual assignment of coded markers within the images (Remondino 2011). The known coordinates of the markers were used when available, otherwise manually input distances are applied to align the photogrammetric model using 3D constraints. After setting specific targets as control points and checkpoints, the alignment was performed to scale and the model was oriented within a local reference framework. Following alignment, the checkpoints were used to verify the metric accuracy of the model. The difference between the actual coordinates of the checkpoints and those calculated by the software indicated the model's error: in this case, an RMSE error of less than 1-2 mm was considered the optimal value for medium- to large-sized object models. In the case of a lack of partial alignment within problematic datasets (such as with thin or flat objects like coins), it was necessary to process data by creating different workspaces, aligning them and subsequently merging the separated point clouds before proceeding to mesh creation. The merging was done by manually placing common distinctive points on the images to help the photogrammetry software recognise and subsequently align them.

Once the workspace merging was complete, a comprehensive dense point cloud was computed. This process involved creating depth maps, which represent the distance of each pixel from the camera (Lu et al. 2018). Next, any outliers generated by reflective surfaces or incorrectly acquired areas were removed. In complex situations, external software such as CloudCompare (https://www.danielgm.net/cc/) would be recommended for more thorough cleaning (Demetrescu et al. 2020).

Based on the points in the dense point cloud, the mesh was generated. In this context, the range of polygons was chosen depending on the geometric complexity and size of the mesh involved (see table 2). The level of detail is a decisive factor in creating a satisfying normal map, starting from the RAWp geometry for the DCHO at a later stage.

Finally, a single texture or multiple textures were generated. The texture was an image mapped onto the surface of the 3D model to faithfully replicate the colours, details and surface features of the real object. In this step, the photogrammetry software used the images acquired during the survey to project the surface details of the object onto the 3D model. This was achieved by associating the 3D coordinates of the model points with their positions in the 2D images using the UV mapping technique, which associated three-dimensional mesh coordinates with two-dimensional image coordinates. Each polygon in the mesh was developed (unwrapped) onto a flat surface and then textured. For texture building, software employing different blending techniques was used to eliminate any seam lines or issues between the images (Dostal and Yamafune 2018). Problematic datasets underwent an additional masking process (fig. 7) specifically designed for the texturing phase, excluding any reflections on the object surface within the images. Furthermore, where needed, an image weighting process was implemented to ensure a sharp texture while minimising data loss. Among the options for texture generation, the desired resolution for the image applied to the model could be specified: for the RAWp, a high-resolution texture (16K or 8K) was generated to ensure that a high level of both geometric and texture detail was used as a starting point for subsequent interventions. Finally, if users want to download a processing report to keep track of key operations performed within the software and ensure more transparent data manipulation tracking (Moore et al. 2022), they can do so in 3DF Zephyr or Metashape.

3.2.3. Step 3: Modelling phase

The RAWp is the initial output from the scanner or photogrammetry software, representing the data processing result without any interpolation or geometry correction, except for rough data cleaning. To produce a complete DCHO without topological issues, two primary methodologies were applied: automatic tools were used to fill small gaps, while manual interventions within 3D modelling software addressed larger gaps and significant non-manifold issues.

For data obtained from SLS, we considered “modelling” to be the stage in which the RAWp was adjusted to obtain a closed, textured model (DCHO).

Sometimes the complexity of an object’s shape or the scanning conditions prevented the capture of all necessary data, resulting in local errors, such as gaps, in the model. These errors were addressed directly within the Artec software prior to subsequent export. Using the `fix holes` tool, the software automatically closed small gaps and non-complex shapes. For holes that were too large or shapes that were too complex for the algorithm to close correctly, manual interventions were used to reduce and modify these gaps (e.g. the `bridge` tool) (Artec 3D 2020, 141–151).

For holes of high complexity, the subsequent closure was performed within 3D modelling software such as Blender. Once all the holes in the 3D model were closed, the reconstructed surfaces were assessed for consistency with the surrounding geometry. If any edges appeared too visible or poorly integrated, they were smoothed to create a more seamless and coherent surface across the 3D object (Artec 3D 2020, 145).

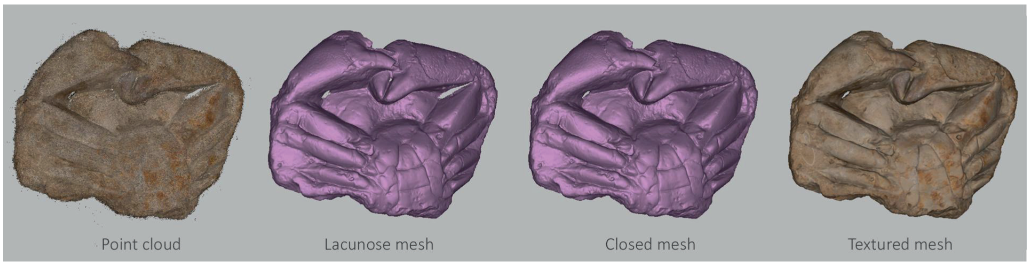

Figure 8: Different SLS phases from point cloud to textured mesh in Artec Studio 15 Professional.

Figure 9: Flying fish CHO, Museum of Palazzo Poggi, Bologna. Photograph by Alice Bordignon.

The final step in obtaining the DCHO was texturing. After selecting all the scans needed to reconstruct the object's texture and enabling the `inpaint missing texture` option, which allows the algorithm to fill regions lacking texture by diffusing information from neighbouring areas, the maximum output resolution achievable by the software (16K) was set. The texture application algorithm was then executed (Artec 3D 2020, 155-162). In cases where the automatic application of texture to missing areas was unsatisfactory, manual texture correction was performed outside the scanning software, using the `texture paint` tool in Blender, Photoshop (https://www.adobe.com/uk/products/photoshop.html) or open-source GIMP (https://www.gimp.org/) (Whitt 2023). The phases described are summarised in figure 8.

For the RAWp obtained from the SfM photogrammetry software, the process was intended to create a closed and complete DCHO free from geometry issues.

The most common scenario involved the RAWp containing gaps of varying sizes. For small gaps, mesh optimisation and cleaning tools were used to optimise the mesh and correct errors. The mesh was sometimes exported to perform minor optimisation tasks in external software, such as MeshLab (https://www.meshlab.net/) (Cignoni et al. 2011) or CloudCompare. The cleaning tools removed unnecessary or incorrectly generated triangles, while non-manifold geometry was identified and corrected to ensure proper mesh structure. Surface smoothing helped eliminate minor distortions, while gap-filling algorithms addressed holes or inconsistencies, creating a continuous surface. As was the case for scanned RAWp, when it was impossible to solve gaps or non-manifold issues with automatic tools, a manual intervention was applied using Blender. To facilitate manual editing of complex models, an initial geometric simplification was performed to achieve a polygon count suitable for the processing power of the hardware in use.

Figure 10: The DCHO of the flying fish as displayed in the exhibition (left); the DCHO as the living, pre-taxidermied reinterpretation of the fish prior to becoming a CHO (right).

Figure 11: Difference between RAWp and DCHO of the microscope CHO.

Topological issues that needed to be addressed (mainly holes, non-manifold edges and overlapping faces) depended on several factors, including the conditions of the survey environment (e.g. shadows), the tools and methodologies used during acquisition (e.g. metrological kits) and CHO features (e.g. occluding geometries). For example, certain items, like manuscripts, could not be handled due to their fragile condition, restricting data acquisition to their front and sides. Manual editing involved authorial interventions to the mesh, aimed at restoring missing or altered details with the utmost attention to the reliability and fidelity of the original asset (table 2). These adjustments were executed using a combination of cutting tools and free-form modelling brushes, enabling the creation of geometries that closely matched the original form. In some cases, we chose to add an additional layer of interpretation to the DCHO, illustrating the transformation and potential changes the asset underwent over time before becoming a CHO. In this context, an illustrative case is that of the flying fish CHO (fig. 9). The fins, preserved through embalming, were in a very degraded and delicate state, giving a misleading impression of the animal's natural appearance. After creating the DCHO from the displayed CHO, an additional DCHO was generated, replacing the embalmed fins with a representation closely resembling the fins of a live fish of the same species, thus depicting its appearance prior to preservation as a CHO (fig. 10). In both cases, orthographic photos of the fins were used and projected onto a plane with an active alpha channel.

Among the DCHOs, some exceptional cases required a complete remodelling of the asset due to its highly intricate geometries. When acquired through photogrammetry, these complex structures were subject to interpolation and generation of low-confidence geometry, resulting in a distorted appearance. This intervention was necessary for the microscope (fig. 11) due to its mix of non-Lambertian materials (glass, metal) and complex small-scaled geometries that were challenging to capture with the available technologies. In the RAWp model, these features were automatically interpolated, leading to an inaccurate and distorted representation of the small mechanisms (gears, levers, knobs, screws) of the asset. The object was then remodelled in Blender, based on the RAWp, and the original textures were preserved through texture baking. Metalness and roughness maps were created to restore the brass surface characteristics that had been lost during acquisition.

3.2.4 Step 4: Optimisation phase

The DCHO is the model resulting from interpolation and the resolution of geometric issues generated from the processing of scanned or photogrammetric data. The final step in optimising the DCHO was to obtain a performant DCHOo suitable for web applications and real-time interaction while retaining an adequate level of detail and visual quality.

The DCHO went through a re-meshing process—which could be done automatically or manually—to achieve a quad-dominant topology.

An initial automatic remeshing often served as a visual guide in manual retopology. Tools for snapping and constraint-based editing helped align vertices and edges to the underlying geometry, particularly in areas with complex curvature. After guiding polygon strips were established, they were extended and connected using bridging techniques along principal curvature directions. This iterative process continuously evaluated connectivity and curvature and often integrated sculpting tools to refine intricate details.

In the case of automatic remeshing, the software Instant Meshes (Jakob et al. 2015) was used. This process ensured a uniform, quad-based topology and reduced polygon density to an optimised level based on the features of the objects (table 2). The mesh was checked for defects, adjustments were made where necessary and complex parts were verified to ensure they were not overly simplified compared to their original structure. Where needed, surface defects were addressed with sculpting tools.

Despite manual retopology being the best method for achieving mesh control on details and highly optimised geometric models, it was undoubtedly the most time-consuming approach. Given the project's limited timeline and scope, automated re-meshing was preferred as standard procedure for most assets and manual interventions were made in cases where the automatic tool could not achieve satisfactory results in the thinner or more complex parts of the mesh. This choice still provided an adequate level of performance for web-based publication and real-time interaction, balancing efficiency and sufficient quality given the project's goals. The diagram in figure 12 outlines the optimisation workflow, highlighting the differences between manual and automatic remeshing. It also compares the UV maps of RAWp and DCHOo.

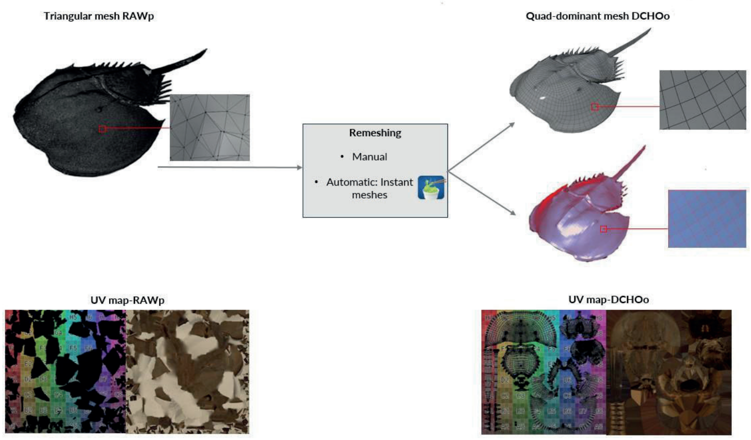

Figure 12: Comparison between automatic retopology using Instant Meshes and manual retopology on the horseshoe crab CHO, alongside the differences in UV maps and textures before and after the intervention.

Following remeshing, a new UV map was needed to preserve texture integrity. The new UV map was created manually, with seams carefully placed in key areas to ensure precise and controlled texture mapping (fig. 9). Following seam selection, the model underwent unwrapping, where it was “opened” along the seams to produce UV islands or separate polygon groups representing sections of the mesh. After the unwrapping, UV layout adjustments were applied to ensure no UV islands overlapped and the texel density was adjusted to increase the resolution in the areas of the model that required higher visual detail. Texel density is a measure that represents the resolution of a texture in relation to the size of a 3D object and is directly influenced by the GSD of the photos used to generate the texture. More specifically, it indicates how many “texels” (single units of texture space) are applied per unit area of the object’s surface.

This approach is beneficial for objects with prioritised details, such as books, allowing for the enlargement of specific UV faces to capture fine details and achieve a well-organised texture layout (Balzani et al. 2024).

After creating a balanced UV set and prioritising the most relevant parts of the mesh, texture building was performed. For 3D models obtained through photogrammetric techniques, SfM software could be used to re-project the textures obtained during the generation of the RAWp onto the DCHOo using the new UV map. Alternatively, a new texture could be generated using photos from the original dataset.

For 3D models obtained from structured-light scanners, textures were baked from the DCHO onto the DCHOo in Blender. The baking process involves simulating the characteristics of the high-poly mesh model (such as geometry, depth, colour and texture) and transferring them onto the low-poly mesh model (DCHOo) (Apollonio et al. 2021).

In some cases, where the texture presented gaps or visible UV seams, issues were resolved by directly editing the image using software like Photoshop or GIMP. Additional modifications were applied directly in Blender using the `texture paint` tool. Once a complete sharp texture was obtained, physically based rendering (PBR) maps (e.g. roughness, metalness, normal, and ambient occlusion) (Farella et al. 2022b) were created, depending on the specific cases. The main aim was to achieve a realistic simulation of light–surface interaction. Roughness maps define surface smoothness or roughness, regulating reflection intensity: white indicated a rough surface, while black represents a smooth one. Metalness maps distinguish between metallic (white) and non-metallic (black) materials, enabling standardised reflection management (Farella et al. 2022b). Normal maps simulate detailed textures on low-poly models by transferring high-poly geometry through a “baking” process using software like Blender or

Substance Painter. Alternatively, normal maps can be exported directly from the photogrammetry software for models obtained through photogrammetric acquisition. Finally, ambient occlusion maps improve realism by simulating soft shadows in concave areas, improving depth perception. Each map can be created or refined in software tools like Photoshop, GIMP or Materialize (https://materializecss.com/) or in advanced texturing tools such as the TexTools plugin (https://github.com/franMarz/TexTools-Blender), enabling precise control over the final appearance of the 3D model.

Table 3: Multi-Parametric Optimisation Matrix for LOD Comparison.

LOD Level	Texture Resolution	Approx. Poly Count	Approx. File Size	Format	Performance Impact	Use Case
Full Detail	8192×8192	>100K polygons	High (e.g. 50+ MB)	.obj or .fbx	Maximum visual fidelity; highest GPU/CPU load	RAWp
High Detail	8192×8192	50K–100K polygons	Medium (e.g. 20–50 MB)	.obj or .fbx	Reduced load while retaining all the details; balance between performance and visual fidelity	DCHO
Medium Detail	4096×4096	10K–50K polygons	Low-Medium (e.g. 10–20 MB)	.gltf	Noticeable performance gain; acceptable detail loss compensated by PBR maps	DCHOo
Low Detail	4096×4096 / 2048×2048	<10K polygons	Low (e.g. <10 MB)	.gltf	Minimal resource usage; lowest visual fidelity	Placeholders

The DCHOo model was thus designed to balance visual fidelity with performance, making it ideal for applications requiring rapid rendering and reduced resource demands, such as the ATON framework. At this stage, to enable even more efficient performance by the DCHOo, the textures were optimised to 4K, with PBR maps resized to 4K or 2K, depending on specific cases; multiple textures associated with a single asset were baked, resulting in a single 4K map for each DCHOo (table 2).

Table 3 presents a multi-parametric optimisation matrix for a level of detail (LOD) comparison, designed to facilitate LOD comparison in 3D asset workflows. By evaluating key technical parameters—such as texture resolution, polygon count, file size and performance impact—across varying levels of detail, the scheme provides a structured framework for selecting the most suitable LOD configuration based on specific project requirements.

3.2.5. Step 5: Export phase

Each 3D model obtained by the workflow was exported in a specific format. Both the RAWp and the DCHO models were exported from Blender in OBJ or FBX formats, allowing for the preservation of the maximum texture resolution and ensuring that visual fidelity was maintained in the final output. OBJ, an open-source format, was the first choice, but for complex 3D models that required features not supported by OBJ, such as multiple UVs and vertex weights, FBX was used.

The DCHOo model was exported exclusively in the glTF format, which is particularly advantageous for sharing and the interactive navigation of 3D objects within Web3D applications. glTF is a standardised format that facilitates efficient transmission and rendering of 3D models, making it ideal for real-time visualisation in Web3D applications. It supports various features, including PBR materials (Robinete et al. 2018).

Throughout the export process, we adhered to a set of specific guidelines (https://osiris.itabc.cnr.it/aton/index.php/tutorials/creating-3d-content-for-aton/exporting-3d-models-from-blender/) tailored to the Web3D platform employed in this project. These instructions were developed to ensure compatibility and optimal performance within the target application, addressing considerations such as geometry optimisation, material settings and texture mapping. By following these guidelines, we aimed to enhance the usability and accessibility of the 3D models in a web context, thus supporting interactive engagement and exploration by users. This approach underscores the importance of format selection and export procedures in the development of high-quality 3D assets for digital environments.

3.2.6. Step 6a: Metadata creation phase

The creation of metadata is a process composed of several substeps, which are partially independent of and partially overlapping with other phases of the workflow. The process is structured such that the data creation and gathering can be simple, intuitive and collaborative. Concerning the choice of tool for the task, Google Sheets is a valuable option for real-time joint contributions; the importance of this feature is also related to guaranteeing the coherence of the values in the fields. It is possible to divide the metadata into two main groups: bibliographic or catalographic data about the physical cultural heritage objects included in the exhibition, and the process data about the acquisition and digitisation process leading to the creation of the related RAW, RAWp, DCHO and DCHOo models. These two datasets are collected in two separate spreadsheets in tabular format, where the fields are clearly defined to allow a metadata crosswalk and conversion to RDF according to a chosen application profile. For the data collection phase, the use of template tables, downloadable from Zenodo

(Moretti, 2024), is recommended. Each template is accompanied by a short CSV sample dataset as a guide for the insertion of field values. To proceed with the subsequent steps in the metadata management process, users should export the tables in CSV format. For compliance with open science principles, the datasets should be published on platforms that allow open access to them, such as Zenodo.

In our case study, to perform the conversion, the open-source tool Morph-KGC (https://github.com/dharc-org/morph-kgc-changes-metadata) was extended to meet the requirements imposed by the defined workflow. Its main components are:

1. Mapping files. A YARRRML (https://rml.io/yarrrml/) file is compiled for each of the datasets, defining the mapping rules for converting the input information in RDF triples, according to the application profile. Once the rules are defined, the mapping file can be reused to convert data in any CSV following the same structure as the initial one.
2. Configuration file. This is an .ini file in which the parameter values for the conversion of the input datasets are set. The README.md file of the software extension provides further guidance about the field values to define. The file is divided in sections introduced by their key, stated in squared brackets. The general section contains both mandatory and optional parameters, including the name of the output file, its format and the directory in which the file is going to be stored. Each dataset to be converted—in this case two—has its own subsection, in which the path to the correct mapping file is declared, along with the input path.
3. User-defined functions. Morph-KGC provides built-in functions (https://rml.io/docs/rmlmapper/default-functions/#combinin) in Phyton. This approach follows the model originally adopted in RML, where Java functions are integrated into the official RMLMapper (https://github.com/RMLio/rmlmapper-java) to address standard conversion tasks. In the same file, it is possible to add user-defined functions for handling specific needs determined by more complex input dataset structures, although the coverage of the built-in functions is quite extensive and addresses the majority of the issues encountered by the average user. These are declarative transformation functions implemented through the RML function mapping language (RML-FNML) (Arenas-Guerrero et al. 2024b).
4. Launch script. A Python file executes the process. In addition to orchestrating the conversion process, this file can also be exploited to perform some data preprocessing, which can make the mapping activities smoother by allowing the reuse of predefined built-in functions instead of defining new user-defined ones to address specific needs. Similarly, although it should not be necessary, some post-processing tasks can be performed on the produced output if minor formal inconsistencies are spotted. This latter procedure is discouraged, however, since Morph-KGC is extensively documented such that the user will benefit from the multiple conversion solutions offered.

3.2.7. Step 6b: Provenance creation phase

The provenance creation phase systematically tracks metadata record changes using the OpenCitations Data Model (OCDM) (Daquino et al. 2020). This approach creates detailed "snapshots" whenever an entity undergoes creation or modification, preserving each snapshot in a dedicated provenance graph to document the entity's state at that specific moment.

The data model establishes clear relationships between entities and their snapshots using semantic properties: snapshots (categorised as `prov:Entity`) connect to their source entities via `prov:specializationOf`. Critical temporal data is captured through `prov:generatedAtTime` and `prov:invalidatedAtTime` timestamps, marking creation and invalidation points, respectively. Attribution and accountability are maintained by recording responsible parties through `prov:wasAttributedTo`, while `prov:hasPrimarySource` establishes information lineage back to original sources. The evolution chain between snapshots is preserved using `prov:wasDerivedFrom` relationships.

OCDM enhances the base provenance ontology (PROV-O) (Lebo et al. 2013) capabilities by introducing `oco:hasUpdateQuery` (Peroni et al. 2016). This extension enables detailed tracking of RDF graph modifications through SPARQL `INSERT` and `DELETE` operations, making it possible to reconstruct previous entity states by reversing subsequent changes.

Incorporating this provenance tracking system into the broader metadata creation workflow significantly enhanced FAIR principle compliance for the DCHOs and their associated data. The comprehensive documentation of changes created a reliable foundation for research applications by ensuring authenticity and enabling confident reuse across different contexts.

3.2.8. Step 7: Presentation phase

In order to make DCHOs accessible on the web, the final step is to employ Web3D technologies to offer an interactive presentation layer to end users. In this phase, optimised DCHOs are transferred from local storage to dedicated server-side storage and converted into Web3D interoperable standards (see section 2.2).

In our case, the open-source framework ATON—designed and developed by CNR-ISPC in 2016 (Fanini et al. 2021)—was adopted. ATON is built on large open-source frameworks such as Three.js (https://threejs.org/) and Node.js (https://nodejs.org/) and on solid web standards.

Single 3D items can be uploaded by different professionals into a dedicated presentation node, using a multitude of approaches to populate digital collections.

In our case, the open-source cloud service NextCloud (https://nextcloud.com/) was adopted for the data layer, allowing easy management by the 3D modelling and content creation team.

ATON uses scene descriptors, similar to the Smithsonian's Voyager SVX (https://smithsonian.github.io/dpo-voyager/document/), to reference 3D items in collections and their hierarchies, with lighting setups, viewpoints, semantic annotations and much more. Each scene in ATON has a unique (alphanumeric) ID, thus providing persistent identifiers: this is ideal for 3D galleries in online virtual museums (VMs) or digital libraries. Published 3D scenes can also be accessed or referenced by any ATON Web3D/WebXR application deployed on the same instance.

3.3 Toward the digital twin

3.3.1 Environment 360° documentation

Given the exhibition's temporary nature, a crucial aspect of effective post-processing involved meticulously documenting the curated spaces and spatial placement of objects to preserve the intended narrative from a curatorial standpoint. Before returning certain pieces to their institutions of origin, we opted to create a basic mock-up using 360° panoramic views of the exhibition. This approach provided a common reference point for the creation of the digital twin. Using 360° panoramas to create virtual tours proved to be an efficient solution; moreover, this technique has gained traction among museum institutions in recent years for a variety of applications, including the enhancement and documentation of cultural heritage (Wu 2022).

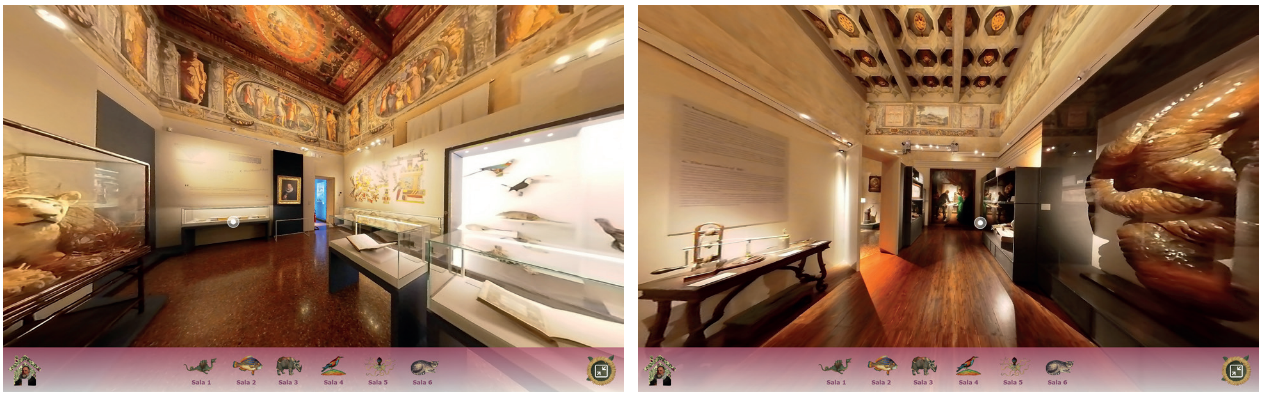

Figure 13: Example of the UI of the virtual tour navigation mode designed by the Skin Editor tool.

For the Aldrovandi exhibition, we used 360° spherical or equirectangular images captured in each room, as this proved to be the most effective solution for providing comprehensive information. As high-resolution images were unnecessary for our purposes, we opted for a quicker method using the Insta360 ONE X2, a dual-lens 360 action camera designed to capture panoramic images and videos up to 5.7K. Our plan involved capturing multiple panoramas in each room, enabling a virtual tour that follows the standard visitor path and aligns with the audio guide available at the exhibition. During image capture, we were able to instantly review the results on a tablet via the Insta360 Studio mobile app, connected to the camera over Wi-Fi. This approach took less time and helped prevent potential errors such as visitor presence in photos. The images were then edited and exported at the highest available resolution, resulting in 360° images of 6080 × 3040 pixels. Each individual panorama was uploaded to a shared cloud platform, and the various shots were interconnected to simulate guided, one-way navigation through the spaces. This process, which created the interactive virtual tour, was completed using Pano2VR (https://ggnome.com/pano2vr/). The primary function of this tool (Cao 2022) is to transform panoramic images into web-compatible formats (HTML5/CSS3) suitable for various platforms. By importing panoramic images into the software, users can link them together to recreate the exhibition's path. The tour can be enhanced with interactive features such as photographic hotspots, directional audio, informational pop-ups and video elements. Additionally, the user interface is customisable using simple CSS through the Skin Editor (fig. 13), allowing for tailored formatting and visual styling.

3.3.2 Environment 3D modelling

One of the final stages in creating a digital twin for the temporary exhibition involves the meticulous reconstruction of the display spaces for each exhibit item. This process includes recreating not only the walls, floors and ceilings of the exhibition hall but also the precise replication of display cases, panels and captions accompanying each item in the physical exhibition. The re-creation of the rooms is a multi-faceted process that requires various techniques and specialised software to achieve a realistic digital replica.

The first step was the photogrammetric acquisition of the ceiling and parts of the walls where frescoes are preserved. The RAW data was acquired using photogrammetric techniques. Images obtained from this process were processed in 3DF Zephyr software. The RAWp obtained (fig. 11) was then imported into Blender, where it was rotated and scaled to align correctly with the intended exhibition layout. This initial adjustment allowed us to proceed to the next steps of detailed modelling and optimisation.

Thus, manual modelling was performed in Blender to reconstruct the ceiling and walls to achieve a mesh as low-poly as possible. The manual modelling process used the RAWp as a foundational reference. Since ceilings are composed by assembling repetitive elements such as trusses and decorations, the ceiling was reconstructed by modelling a smaller section that was later duplicated using the `mirror` and `clone` modifiers in Blender to achieve symmetry, though the mirror modification was only finalised at the end of the baking process.

For this step, the geometry was modelled as perfectly regular and without modifications to simplify the UV mapping and to get identical UV islands so that at a later stage, eventual missing parts could be sampled by other instances of the same module, at which point simple deformations could be applied both to match the real geometry of the design and to account for deformations such as truss bending and similar.

Next, the texture and normal maps were baked, following a process similar to the one used during the optimisation phase of the assets (see section 3.2.4). The diffuse and normal maps were both baked directly in Blender. The normal map, in particular, was used to recover fine details and subtle reliefs in the ceiling decoration that had been lost during the modelling phase. Any gaps or imperfections in the resulting model were addressed using Blender's `texture paint` tools or in GIMP. This was possible because the maps had been built from a regularised and straightened version of the ceilings such that the resulting UV islands were regular and allowed for direct editing within software that processes raster maps.

After completing the ceiling (fig. 14), the next step was to recreate the room's walls by starting with a basic geometric shape and adjusting it to match the dimensions of the DCHO. Once aligned, a precise UV map was generated to enable the baking of a new texture. This texture was then refined and enhanced in GIMP to better fit the walls, improve colour quality and fill in any missing sections. The missing parts were added through a cut-and-paste process, utilising orthogonal photos taken during the photogrammetric acquisition phase. This final process is particularly useful when it is not possible to achieve new texturing through reprojection in photogrammetric software or by baking, especially when there are missing sections in the geometry. In this case, it could also be necessary to add synthetic material to the missing parts and finalise the texture with the final baking, as we will see later.

The second phase in recreating the exhibition environment involved manually constructing each display case and exhibition panel. To achieve the highest level of accuracy, we relied on the floor plans and elevations provided in the original museum exhibition design. These documents included detailed specifications of room dimensions, as well as the height, depth, shape and content of the display cases. For this process, photographs and videos taken during the acquisition phases were essential in understanding the position of each element in the scene, allowing for a faithful reproduction of the original setup.

At the end of the creation process, each element within the scene had been associated with its own synthetic material, accurately reflecting those present in the physical exhibition. For example, the metal material for the display cases was created either using an online resource like Poly Haven (https://polyhaven.com/), which provides realistic textures for download, or appropriately applying colours and textures acquired from the actual environment and setup, such as the Venetian floor material or the displays with their graphics. Since synthetic materials define the behaviour of the material itself when interacting with light, the rendered colours, reflections, refractions and other complex behaviours are locally dependent on the synthetic illumination.

Figure 14: RAWp of the ceiling imported into Blender.

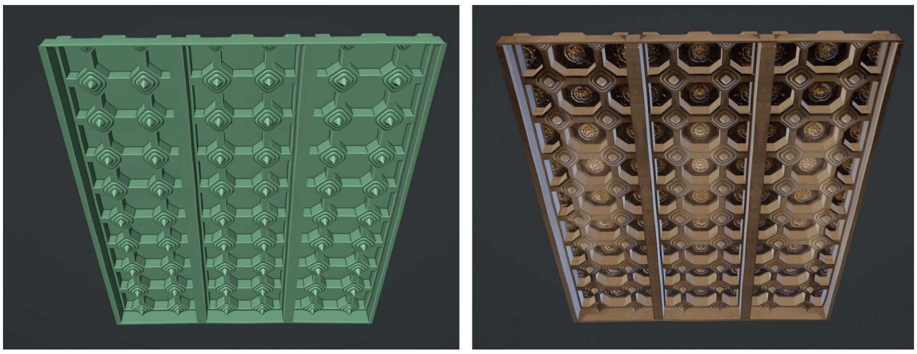

Figure 15: Modelling and texturing the ceiling.

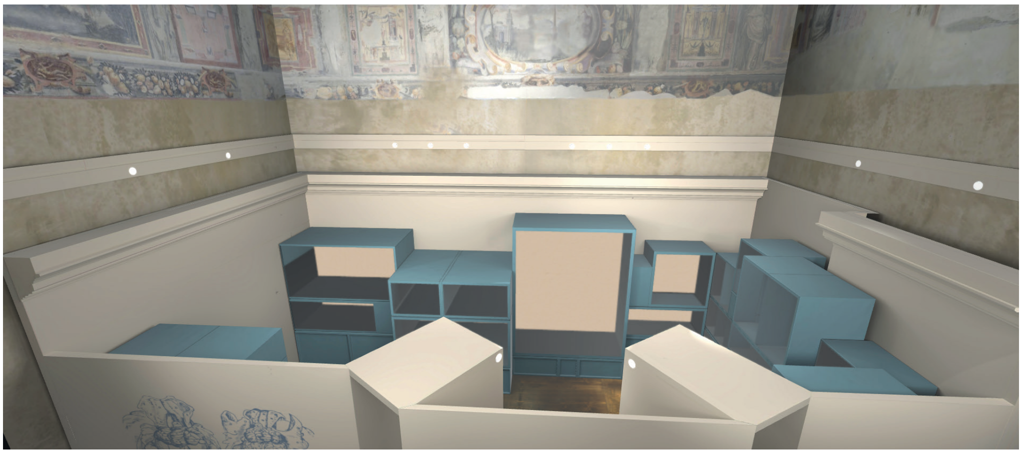

Figure 16: View of Room 6 exported in glTF format after baking the lighting.

In the final baking process (Apollonio and Giovannini 2015), these elements were unified, resulting in a single, cohesive texture that accounted for the resulting interaction between the original materials under the given illumination in the synthetic scene, which is designed to mimic the real one. This strategy overcomes the limitations of current Web3D applications, which cannot make such complex calculations in real-time. The baking process allows these interactions to be pre-calculated and to re-applied to the objects as a simple colour texture, significantly enhancing photorealism while maintaining real-time performance.

Glass surfaces, such as those of the display cases, required specific adjustments to replicate the transparency, reflectivity and refraction that would allow viewers to see the objects inside. Achieving this involved using different parameters: the displacement settings were modified by incorporating a noise texture to add subtle surface variations, and both the alpha and roughness channels were adjusted to reach the right balance between clarity and natural light diffusion. In this way, the final rendering achieved the desired glass transparency. This process contributes to a more immersive scene, where the glass surfaces not only look visually realistic but also interact naturally with the light and surrounding environment.

Accurate lighting placement was essential for finalising the recreation of the exhibition's digital environment. Each light source was positioned to simulate the actual setup in the original exhibition space, including direction, intensity and colour temperature, which were adjusted to match those of the physical room. Once configured, the lighting was baked into the scene, capturing shadows and reflections to maintain a high level of realism without overloading the scene. To bake the room's lights and shadows directly onto the texture, the room was divided into several sections (e.g. ceiling, walls, floor, panels and display cases). For elements that did not require roughness, reflections, transparency or other specific characteristics, a diffuse bake was performed, allowing the texture to incorporate the room's lighting effects in a uniform way. For elements with transparent or reflective parts (such as display cases), a combined bake was used to retain those physical characteristics. After the baking process, the room was exported as a single glTF file with a reduced number of 4K textures (fig. 13).

3.3.3 Final optimisation

After recreating the spaces of the temporary exhibition, it was necessary to produce a further optimised version of the environment and each asset within it, combining all individually generated elements into a single scene to enable fluent user navigation. The optimisation was carried out simultaneously for each of the six rooms that hosted the exhibition.

As for the environment, no geometric optimisation was needed: the output of the manual modelling for each room was a low-poly mesh created from simple shapes. The situation was different, however, for the textures. After light-baking, which was done to recreate shadows and ambient lighting directly on the textures, an additional baking process was required to merge the room's multiple textures into as few maps as possible, at 2K or 4K resolution each. The goal of baking is to optimise the performance of the 3D model as much as possible while maintaining acceptable resolution quality depending on the contexts of reuse.

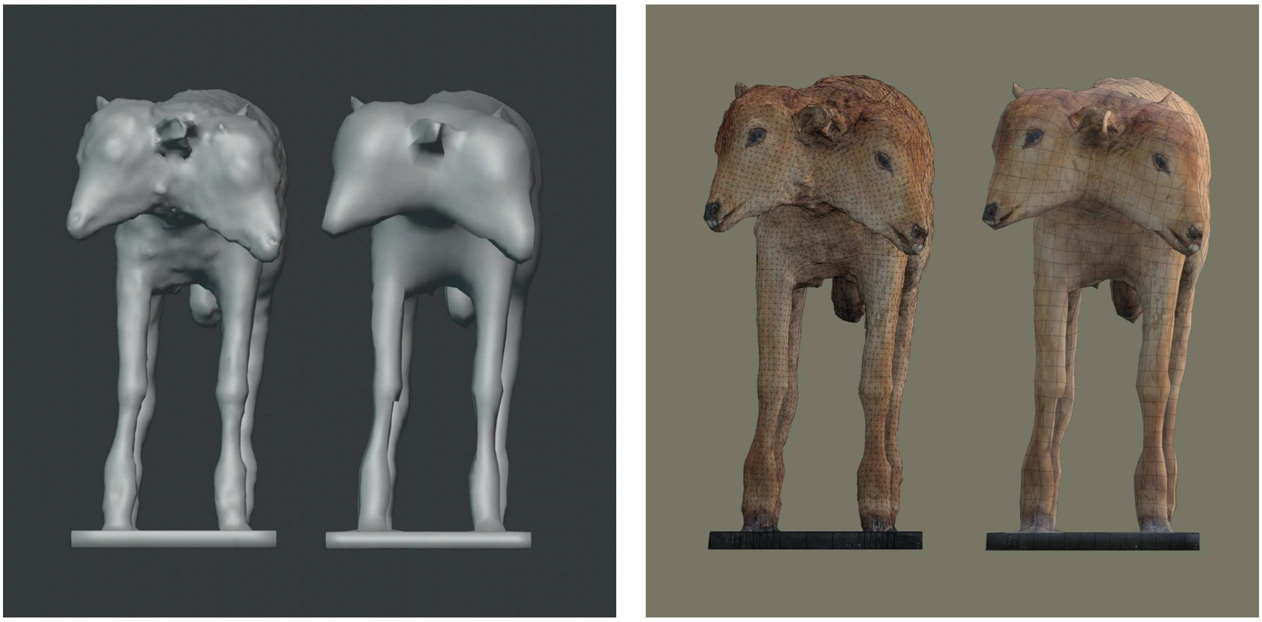

Figure 17: Example of PL starting from the re-meshing of the DCHOo.

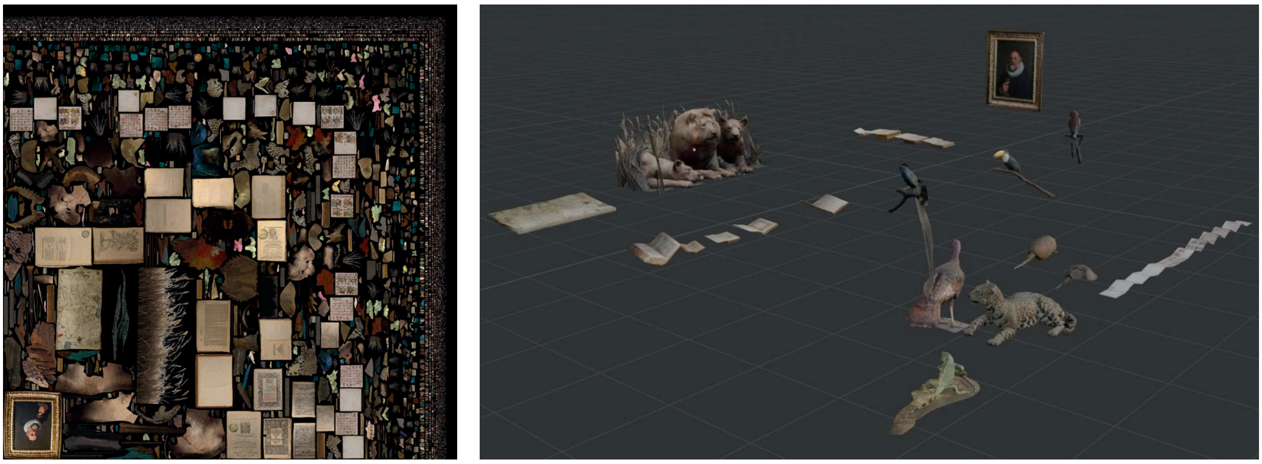

Figure 18: Texture map (left) obtained from the texture baking of all PL in Room 1 (right).

On the asset side, however, each DCHOo was optimised from both a geometric and texture perspective to obtain a placeholder model (PL). This additional version was needed to create a performant immersive environment for the user's exploration, where all the elements coexist in the same scene. The geometric optimisation to obtain the PL was done through two main methodologies: polygon decimation using Blender or remeshing using Instant Meshes (Jakob et al. 2015). In both cases, the goal was to reduce the number of polygons as much as possible (fig. 17). While internal decimation did not affect the UV map, remeshing required a new UV mapping, followed by transferring the DCHOo texture to PL through baking.

For texture optimisation, the DCHOo textures for each room were optimised into a single 4K map through baking, then assigned as the only material of a single mesh, including all DCHOo. For example, if the first room contained 21 assets, each with a 4K texture, the final optimisation output involved merging the 21 meshes into a single mesh. This mesh was assigned a single material with one 4K map, which was created by baking the 21 pre-existing maps into one (fig. 18). The overall final optimisation process was conceived to create optimal conditions to facilitate the web-based and immersive implementation for real-time performances.

3.3.4 Web-based and immersive implementation

A digital twin application was designed to create a fully explorable virtual replica of the temporary exhibition. To achieve a minimum viable product, core features were defined, including user navigation, high-quality exhibit exploration and a clear presentation of each exhibit's metadata. A functional prototype was developed following a UI/UX study.

The exhibition consists of six connected rooms, navigable one by one through interactive doors. The user journey begins at a specific starting point in the first room, with a welcome screen launching an audio guide to introduce the overall exhibition. Within the virtual environment, the user can navigate the exhibition space freely using the teleport tool (Bozgeyikli et al. 2016), viewing all the objects placed in cases and the graphics displayed on the room's walls. When interacting with a single object, the user can explore it in detail, viewing the high-quality version and metadata.

A prototype was designed as a web application accessible on handheld devices and VR headsets without installation, offering an embodied exploration in virtual reality and a flexible interface via desktop/tablet/touchscreen devices. Built using ATON, the web app takes advantage of the framework's web-based, VR-ready 3D features. Although the process is streamlined by using ATON's API, ensuring a convenient trade-off between performance and quality presented technical challenges involving the optimisation of 3D assets and a tailored 3D scene management strategy.

With a scalable approach in mind, the browser's memory usage was optimised by loading in the virtual environment one room at a time to avoid overloading the client's device and managing an un-predetermined set of exhibition rooms. To allow users to view the entire collection within the room and appreciate each object in detail, two complementary exploration modes were implemented:

1. Room exploration: In this mode, the user can explore a scene with two main elements—the room (with its materials and textures) and the collection of exhibits. Low-poly versions of all objects are used, with meshes merged and textures baked into one or two files. This optimisation reduces draw calls and texture loading, speeding up load times and allowing more exhibits to be managed within each room.
2. Exhibit inspection: From room exploration mode, users can click on an object to view it in detail. Selection is triggered by simple clickable cubes, moving the user to a close-up viewpoint. After this transition, the low-poly room and exhibit collection are unloaded to free up memory, and a high-quality version of the selected object is displayed. The positions of the low- and high-poly versions are carefully aligned to maintain visual consistency, while metadata for the exhibit appears on the side of the screen.

This approach balances graphic quality with performance, providing an efficient user experience. Core features were tested with real users, revealing that the switch gesture between room exploration and exhibit inspection mode works without disorientation, even if a short delay is needed to load the asset. However, this gesture can cause motion sickness in VR, where maintaining a consistent user position is essential. This issue warranted an alternative design of the VR logic, which is currently being developed.

Future improvements will include integrating a MELODY-based API for metadata and experimenting with level of detail (https://threejs.org/docs/#api/en/

objects/LOD), 3D tiles (https://cesium.com/why-cesium/3d-tiles/) and other techniques to improve performance and reduce optimisation time for 3D assets. Real-time lighting and advanced post-processing effects are also under consideration for enhanced visual quality.

3.4. Software for human-readable documentation and narratives

To quantitatively analyse and visualise aggregated data, we used MELODY, an open-source, web-based platform designed to visualise and narrate Linked Open Data (Renda et al. 2023a). This tool serves as a bridge between technical data exploration and user-friendly storytelling, where narratives can capture historical contexts, provenance and thematic elements. MELODY was initially designed using a hybrid methodology that combines eXtreme design with design thinking, guiding the tool's development from ontology-driven data requirements to a UI/UX structure that supports a diverse range of storytelling elements (Renda et al. 2023b). This method not only allows MELODY to address complex cultural heritage use cases but also ensures flexibility for other domains where structured storytelling is needed.

MELODY allows users to explore one or more Linked Open Data sources via SPARQL queries. The results of these queries are linked to various user interface (UI) components, such as charts, maps, text searches and tables. These components can be selected, added, rearranged and displayed on a canvas, allowing the data story creator to alternate between visualisations and curated text, thus contextualising visual aspects and offering interpretative insights. The final data story, as well as each user interface (UI) component, can be exported and embedded in other web pages. Data storytelling options are available to data curators, who can select appropriate (dynamic) visualisations to be included in their narrative, and to end users, who can explore the chosen dataset via the MELODY online platform (https://projects.dharc.unibo.it/melody/) and publish their own data stories in a dedicated catalogue (https://melody-data.github.io/stories/).

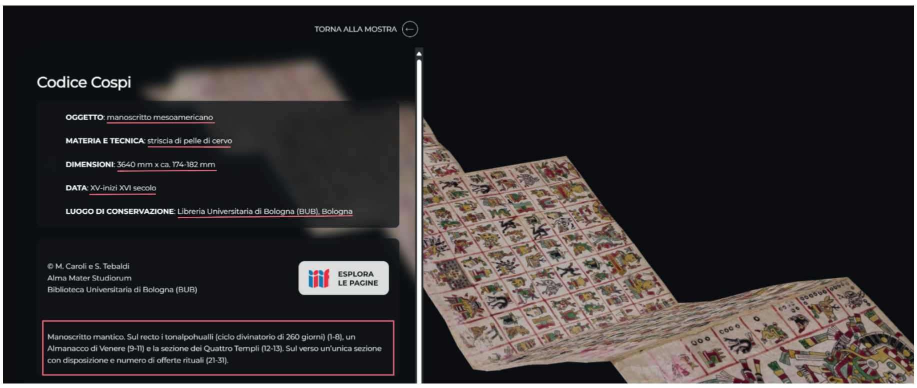

Figure 19: A mock-up example of a MELODY API's response integrated in the sidebar of a 3D object.

Our choice of MELODY was guided by its open-access nature, with both its code and documentation available online (https://github.com/polifonia-project/dashboard), and by its use of standard technologies and frameworks, making it easily extendable and reusable across different contexts. MELODY is developed as a Python application using the Flask framework (https://flask.palletsprojects.com), configured through a JSON file. This file contains the background information and SPARQL queries necessary to retrieve data for the data stories, which are presented in a web interface served as HTML pages. Each data story comprises React components (https://react.dev/) that can be combined as desired during the data story creation process.

To address the needs of our pilot study, we are extending MELODY to include a new API to enhance interoperability and reuse. Specifically, the new API will allow direct communication with ATON. As discussed in section 3.3, within the ATON framework, each cultural object is visualised in a 3D environment. Upon selecting an object, ATON sends an HTTP request containing the object's unique global persistent identifier, along with configuration parameters. These parameters, provided as a JSON object, specify which information to display, including SPARQL queries for data retrieval, textual descriptions with object-specific data, visual elements and additional files for styling and interactivity. The API processes these parameters and delivers a customised HTML response, which is displayed as complementary information to the object (fig. 19).

In the mock-up in figure 19, the highlighted sections represent dynamic content that changes based on the object identifier. This includes details such as the type of object, materials and techniques used, historical period, place of conservation and a brief description. The goal is to provide visitors with comprehensive metadata about the object and insights into the digitisation process. This approach ensures that users interacting with the digital twin not only learn about the artefact but also gain an understanding of the technical and methodological steps involved in its digital capture.

3.5 Interoperability and sustainable reuse with open technologies and standard formats

The final goal of the project was to use open technologies and software at each stage of the process to maximise the workflow's adaptability for reuse in various scenarios, both within and beyond Spoke 4.

While metadata management did not present challenges in this regard, handling 3D models did. We carefully selected open-source software for every feasible step; however, proprietary software was necessary for certain specific tasks where open-source options did not yield satisfactory results.

To avoid dependency on proprietary software, we prioritised the use of standard formats for all research data types generated. For each type, these formats included: glTF, GLB, OBJ, MTL, PNG, JPG, TIFF and E57 (among others) for 3D models; PNG, TIFF, RAW and JPG for images, MOV and MP4 for video and MP3 for audio. These format choices were documented and, in some cases, guided by the project's data management plan (Gualandi and Peroni 2024).

We advocated for the adoption of glTF as the primary 3D format, an open standard designed for interactive Web3D applications, to ensure strong interoperability with current 3D platforms and services. This choice also supports the reusability and integration of licensing information directly within the format (Robinete et al. 2018).

Attention was also paid to the compatibility of the chosen formats with the existing technological ecosystem, particularly with widely used online repositories and Web3D viewers. This ensured that the processed assets could be optimally employed in the chosen framework (ATON) and the target application, while also allowing for the potential publication of assets independently on existing platforms. The selection of formats was not based solely on an evaluation of technical criteria—such as compactness or supported features—but also considered the current level of dissemination and documentation.

The final stages of the project are still underway. At the time of writing, the 3D models, along with their associated data and metadata, have yet to be deposited in a repository for long-term preservation. The majority of them will be published with a CC0 license to maximise reuse, while others will be published with more restrictive licenses (e.g. CC-BY-NC), depending on the permissions from the museums involved. For the purpose of this article, we have published on Zenodo the RAWp 3D model (Collina 2025) and the DCHO and DCHOo 3D models (Bordignon and Collina 2025) of one of the objects included in the digital twin, i.e. the bust of Ulisse Aldrovandi. The team has opted to use Zenodo, a general-purpose repository, at least temporarily, until a more specialised platform becomes available. Although Zenodo is not tailored specifically for 3D or cultural heritage data and metadata, it provides a DOI for each deposited item, supports high-level metadata schemas (DataCite Metadata Schema, Dublin Core) and is well-regarded within the research community. Zenodo was chosen for its alignment with open science principles, its familiarity among all project partners, its independence from the institutions involved and its capability to host a dedicated community for the CHANGES–Spoke 4 project. This choice allows for the consolidation of all project outputs under one umbrella rather than dispersing them across multiple repositories.

4. Discussion

4.1 Towards an open, reproducible and transparent workflow

Moving forward, a key focus of this study will be enhancing the informational complexity associated with the geometry of 3D objects. This involves not only capturing and representing geometric and texture details but also incorporating metadata that provides a nuanced understanding of the accuracy and quality of the reconstructions. This focus is particularly important because the process of digitisation is inherently imperfect: due to technological limitations or data collection constraints, reconstructed objects often exhibit gaps, inaccuracies or poorly inferred geometric features.

To address these imperfections, it is imperative to complement automatic reconstruction with the manual interventions that have been briefly outlined in the workflow in section 3. However, manual integration carries a degree of subjectivity and arbitrariness, even when it is guided by supplementary references or supporting evidence. This blending of automated and manual processes is well illustrated by the distinction between RAWp and DCHO, which underscores a systematic approach grounded first in objective data capture and subsequently in informed interpretative efforts. Such a framework aligns closely with practices in archaeological virtual reconstructions, where integrating missing elements demands a delicate balance of data-driven accuracy and informed conjecture.

The proposed future direction builds on the consolidated findings of the scientific concerning virtual archaeological reconstructions. Grounded by foundational guidelines such as the London Charter (2009), which outlines principles for the use of 3D visualisations in cultural heritage research and preservation, and the Sevilla Principles (Bendicho 2013), which extend these guidelines specifically to digital and computer-based representations, this work will draw on a well-defined methodological and ethical framework.

A critical aspect of future developments involves the structured documentation of the processes, interpretive decisions and the contextual reasoning that underpin the integration of missing parts into 3D reconstructions. This aligns with the concept of paradata, first introduced in the context of the London Charter (2009) and subsequently expanded in *Paradata and Transparency in Virtual Heritage* (Bentkowska-Kafel and Denard 2016). Rather than offering a single theoretical definition, this body of work presents paradata as a pragmatic framework to account for and communicate the interpretive dimension of digital reconstruction processes. It has been effectively employed in applied case studies, such as Apollonio and Giovannini (2015), to express the degrees of uncertainty and rationale behind specific modelling choices.

In our context, the directly acquired portions of the object serve as a form of empirical reference, while the reconstructed areas—modelled with varying levels of plausibility—require explicit documentation of both the procedures and the interpretive assumptions involved. While there is no fixed ontology of paradata, its flexible structure makes it particularly suitable for addressing these situations, provided that it is used not as a label but as a set of transparent practices embedded in the workflow.

In line with the aim of ensuring maximum workflow alignment with the FAIR principles, the use of open-source software and formats should be improved by updating the workflow as plausible alternatives emerge. Concerning the formats, while glTF is both an open format and currently the best choice for delivering 3D assets to Web3D applications, there is still no viable free alternative to FBX for authoring 3D assets requiring advanced features that are not supported by OBJ. DAE failed to fill this gap since its adoption remained limited and continues to decline, while USD shows promise but has not yet fully overcome key barriers to wider adoption: its inherent complexity, partial standardisation across software tools, potential performance issues due to dependencies, steep learning curve and integration difficulties with legacy systems. Additionally, the need for ongoing support and customisation further complicate its seamless integration into existing workflows.

Finally, regarding software, the GLAM sector's need to adopt increasingly accessible digitisation workflows (both in terms of cost and technical skills) makes the use of open-source software at every stage of the process a significant consideration. For most stages of the workflow, open-source alternatives appear suitable; however, data processing remains an exception. As noted, in this stage, open-source options in both scanning and photogrammetry currently do not provide functions and results comparable to those of proprietary software (Rahaman and Champion 2019). In our specific case, we tested Meshroom as an open-source alternative for processing photogrammetric data, confirming the observations in the literature. As for structured-light scanner digitisation, there is potential to explore open-source alternatives in the future (e.g. Open Scan, 3DUNDERWORLD-SLS), with the aim of enhancing interoperability for equipment that is already difficult for the GLAM sector to access.

In this context, a commercially available software program, Reality Capture, now offers a free version that is highly relevant for both photogrammetry and LiDAR data. This platform allows users to combine photographs with LiDAR data or meshes from 3D scanners, leveraging both input types to create highly accurate models. At the moment, Reality Capture does not natively support data from structured-light scanners: the data may need to be processed or converted into a compatible format before it can be imported. The free version of Reality Capture is available for students, educators, hobbyists, and companies with an annual gross revenue under $1 million USD. Given these advantages, we are considering a shift towards this type of alternative (at least for photogrammetry data processing), which combines processing speed, precision and functionality, offering results comparable to proprietary software used in our pipeline.

4.2. Embedding metadata into DCHO and DCHOo

As previously described, specific 3D formats involved in our pipeline—especially open options (OBJ and glTF)—allow extensibility and enrichment. This is particularly important for custom metadata and the general injection of attributes directly into the 3D format. The need to embed semantic information describing virtual products has become increasingly crucial in recent years, especially in the 3D commerce sector. 3D assets that contain descriptive and administrative metadata such as product descriptions, details on intellectual property rights, creation and modification dates and other detailed authoring histories enable the management of 3D virtual product catalogues as well as the sharing of assets between vendors, retailers and end-user platforms at industrial scale.

Within delivery formats such as glTF, this is being addressed by Khronos extensions such as *KHR_xmp_json_ld* (https://github.com/KhronosGroup/glTF/blob/main/extensions/2.0/Khronos/KHR_xmp_json_ld/README.md), which adds support for XMP (Extensible Metadata Platform) (ISO 16684-1) metadata for glTF. This replaces the original *KHR_xmp* extension proposed by Adobe in 2019. Metadata is used to transmit information such as, among other things, the attribution, licensing and creation date associated with the glTF asset and has no normative effect on the asset's appearance and rendering. Other platforms, such as SketchFab, may also inject custom data into the "extras" entry of the glTF model, such as author(s) and licensing as well as the original persistent identifier on the platform.

A few open-source frameworks, such as ATON, are capable of extracting such data at runtime and exposing it at the presentation level. Specifically, ATON introduced support in 2022 for both custom schema in the glTF "extras" attribute (thus showing embedded information in glTF models downloaded from SketchFab, for instance) and *KHR_xmp_json-ld*. To further assist content creators and publishers, ATON also offers asset injectors operating on JSON formats, such as glTF or Cesium 3D Tiles.

4.3. Beyond the valorisation and preservation of CH

Although the current workflow already involves extensive collaboration among experts from various disciplines and institutions, it is essential to focus on the sustainability and reusability of a project of this size. While the primary goal here is to enhance cultural heritage, there is nothing preventing a digital twin from being adapted for other purposes as a tool for comparative studies.

The sustainability of reusing the knowledge base created with the digital twin lies in its adaptability: the digital twin becomes a structured container where the various components are replaceable, interchangeable and modifiable according to different objectives. This flexibility is achievable only if the workflow used is replicable and well-defined.

In the context of Spoke 4, a recent example of this adaptability involved using the first two rooms of the digital twin for both usability testing and psychological research. A team of psychologists, designers and digital heritage experts collaborated to: gather requirements for defining features that capture user behaviour in 3D and VR environments; develop a prototype incorporating these features for a psychological study on navigation styles and spatial orientation in museum or exhibition settings; test the prototype with psychologists acting as operators during lab sessions with real users (fig. 20); and analyse the results to enhance the prototype's usability and functionality for broader applications beyond the pilot phase (Massidda et al. 2024).

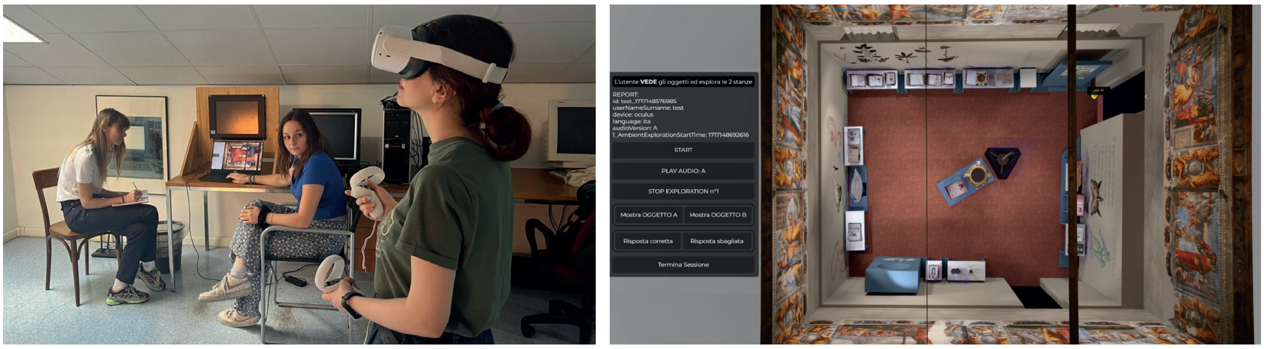

Figure 20: Observation of a subject conducting a VR session (left) and ATON virtual environment, with the admin panel (right).

5. Conclusions

We have presented a reproducible workflow for structuring acquisition and digitisation campaigns as well as digital twin creation in cultural heritage institutions (e.g. museums) and events (e.g. exhibitions). The workflow was derived from the work conducted on a pilot study, i.e. the temporary exhibition entitled *The Other Renaissance: Ulisse Aldrovandi and the Wonders of the World*, which ended in May 2023. We have detailed the materials, methods and tools used for acquiring and digitising 3D cultural heritage artefacts, answering RQ1. Also, in light of the needs identified within the GLAM context regarding the acquisition and digitisation processes, we have discussed possible practices and future adoptions aimed at improving the workflow's openness, accessibility, transparency, reproducibility, reusability and sustainability, thus answering RQ2.

One of the main goals of this activity was to establish a shared and reproducible process for all acquisition and digitisation campaigns conducted in the context of the project CHANGES–Spoke 4, involving nine case studies developed within the project and implemented with the support of several companies. We are collaborating with institutions across Italy: Museo Egizio in Turin, several museum networks in Italian universities (University of Bologna, University of Ferrara, University of Turin), Centro Studi e Archivio della Comunicazione (CSAC) in Parma, Ente morale Istituto Suor Orsola Benincasa in Naples, Carlo Levi and Grazia Deledda Literary Parks in Aliano and Galtellì, and Reggia di Caserta in Caserta.

A complementary goal of the work presented in this paper was to foster a culture of accountability and reproducible research by documenting a time-consuming and specialised research workflow in the area of cultural heritage. Adopting principles like FAIR and guidelines for documenting research methodologies in the humanities is a positive step forward and may increasingly become the norm in the academic setting.

References

Amato, Alba, Dario Branco, Salvatore Venticinque, Giuseppina Renda, and Sabrina Mataluna. 2023. "Metadata and Semantic Annotation of Digital Heritage Assets: A Case Study." In *2023 IEEE International Conference on Cyber Security and Resilience (CSR)*, 516–22. https://doi.org/10.1109/CSR57506.2023.10224935.

Apollonio, Fabrizio, and Elisabetta Giovannini. 2015. "A Paradata Documentation Methodology for the Uncertainty Visualization in Digital Reconstruction of CH Artifacts". *SCIRES-IT* 5 (July):1–24. https://doi.org/10.2423/i22394303v5n1p1.

Apollonio, Fabrizio Ivan, Filippo Fantini, Simone Garagnani and Marco Gaiani. 2021. "A Photogrammetry-Based Workflow for the Accurate 3D Construction and Visualization of Museums Assets." *Remote Sensing*, 13 (3). https://doi.org/10.3390/rs13030486.

Arenas-Guerrero, Julián, David Chaves-Fraga, Jhon Toledo, María S. Pérez, and Oscar Corcho. 2024a. "Morph-KGC: Scalable Knowledge Graph Materialization with Mapping Partitions." *Semantic Web* 15 (1): 1–20. https://doi.org/10.3233/sw-223135.

Arenas-Guerrero, Julián, Paola Espinoza-Arias, José Antonio Bernabé-Diaz, Prashant Deshmukh, José Luis Sánchez-Fernández, and Oscar Corcho. 2024b. "An RML-FNML Module for Python User-Defined Functions in Morph-KGC." *SoftwareX* 26 (May). https://doi.org/10.1016/j.softx.2024.101709.

Artec 3D. 2020. "Artec Studio 15 User Guide." Artec 3D. https://docs.artec-group.com/as/15/en/_downloads/25ad86ff3ec13dc8e-ae3208423c3eb81/Manual-15-EN.pdf.

Artec 3D. 2023a. "Artec 3D Eva." Artec 3D. https://cdn.artec3d.com/content-hub-files/artec-eva-brochure-eng-2023.pdf.

Artec 3D. 2023b. "Artec 3D Space Spider." Artec 3D. https://cdn.artec3d.com/content-hub-files/artec.s.spider-b-a.8-web-en-nop.pdf.

Baker, Monya. 2016. "1,500 Scientists Lift the Lid on Reproducibility." *Nature* 533 (7604): 452–54. https://doi.org/10.1038/533452a.

Balzani, Roberto, Sebastian Barzaghi, Gabriele Bitelli, Federica Bonifazi, Alice Bordignon, Luca Cipriani, Simona Colitti, et al. 2024. "Saving Temporary Exhibitions in Virtual Environments: The Digital Renaissance of Ulisse Aldrovandi—Acquisition and Digitisation of Cultural Heritage Objects." *Digital Applications in Archaeology and Cultural Heritage* 32 (March):e00309. https://doi.org/10.1016/j.daach.2023.e00309.

Barzaghi, Sebastian, Alice Bordignon, Bianca Gualandi, and Silvio Peroni. 2024a. "Thinking Outside the Black Box: Insights from a Digital Exhibition in the Humanities." Preprint, submitted April 10, 2024, https://doi.org/10.48550/ARXIV.2402.12000.

Barzaghi, Sebastian, Alice Bordignon, Bianca Gualandi, Ivan Heibi, Arcangelo Massari, Arianna Moretti, Silvio Peroni, and Giulia Renda. 2024b. "A Proposal for a FAIR Management of 3D Data in Cultural Heritage: The Aldrovandi Digital Twin Case." *Data Intelligence*, 6 (4): 1190-221. https://doi.org/10.3724/2096-7004.di.2024.0061.

Barzaghi, Sebastian, Ivan Heibi, Arianna Moretti, and Silvio Peroni. 2024c. "Developing Application Profiles for Enhancing Data and Workflows in Cultural Heritage Digitisation Processes." Preprint, submitted August 2, https://doi.org/10.48550/arXiv.2404.12069.

Bendicho, Víctor Manuel López-Menchero. 2013. "International Guidelines for Virtual Archaeology: The Seville Principles." In *Good Practice in Archaeological Diagnostics*, edited by Cristina Corsi, Božidar Slapšak, and Frank Vermeulen, 269–83. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-01784-6_16.

Benítez, Anto J., Xose Prieto Souto, Manuel Armenteros, Esteban M. Stepanian, Rubén Cantos, Miguel García-Villaraco, Jaime Solano, and Ángela Gómez Manzanares. 2022. "Multi-Camera Workflow Applied to a Cultural Heritage Building: Alhambra's Torre de La Cautiva from the Inside." *Heritage* 5 (1): 21–41. https://doi.org/10.3390/heritage5010002.

Bentkowska-Kafel, Anna, and Hugh Denard, eds. 2016. *Paradata and Transparency in Virtual Heritage*. London: Routledge. https://doi.org/10.4324/9781315599366.

Bitelli, Gabriele, Valentina A. Girelli, Fabio Remondino, and Luca Vittuari. 2007. "The Potential of 3D Techniques for Cultural Heritage Object Documentation." In *Proceedings SPIE,* Vol 6491, Videometrics IX, edited by J.-Angelo Beraldin, Fabio Remondino, and Mark R. Shortis, 64910S. San Jose, CA, USA. https://doi.org/10.1117/12.705012.

Bolognesi, C. M., and V. Manfredi. 2024. "Optimization of VR Application in Texturing Cultural Heritage." *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XLVIII-2/W4-2024 (February): 73–78. https://doi.org/10.5194/isprs-archives-XLVIII-2-W4-2024-73-2024.

Bolton, Alexandra, Lorraine Butler, Ian Dabson, Mark Enzer, Matthew Evans, Tim Fenemore, Fergus Harradence, et al. 2018. "Gemini Principles." CDBB. https://doi.org/10.17863/CAM.32260.

Bordignon, Alice, Sebastian Barzaghi, Federica Collina, Francesca Fabbri, Bruno Fanini, Daniele Ferdani, Nicola Marinello, Arianna Moretti, Maria Felicia Rega, and Mattia Sullini. 2024. "Guidelines for the Digitisation of Museum and Art Collections: Guidelines for Digitisation of Collections' Items: Guidelines for 360° Virtual Tours: Guidelines for Digitisation of Architecture-Scale Objects". Zenodo. Version v1. https://doi.org/10.5281/ZENODO.14249935.

Bordignon, A., and Collina, F. 2025. Busto Aldrovandi—DCHO and DCHOo models [Graphic]. Zenodo. Version v1. https://doi.org/10.5281/zenodo.15389506.

Bozgeyikli, Evren, Andrew Raij, Srinivas Katkoori, and Rajiv Dubey. 2016. "Point & Teleport Locomotion Technique for Virtual Reality." In *Proceedings of the 2016 Annual Symposium on Computer-Human Interaction* *in Play*, 205–16. Austin Texas USA: ACM. https://doi.org/10.1145/2967934.2968105.

Brandolini, Filippo, and Giacomo Patrucco. 2019. "Structure-from-Motion (SFM) Photogrammetry as a Non-Invasive Methodology to Digitalize Historical Documents: A Highly Flexible and Low-Cost Approach?" *Heritage* 2 (3): 2124–36. https://doi.org/10.3390/heritage2030128.

Cao, Kai. 2022. "Development and Design Case Function Comparison of Panoramic Roaming System of Virtual Museum Based on Pano2VR." *Mobile Information Systems* 2022 (July):1–11. https://doi.org/10.1155/2022/7363221.

Chai, Kedong, and Yue Li. 2022. "Cultural Heritage Assets Optimization Workflow for Interactive System Development." In *2022 IEEE 46th Annual Computers, Software, and Applications Conference (COMPSAC)*, 1822–27. Los Alamitos, CA, USA: IEEE. https://doi.org/10.1109/COMPSAC54236.2022.00290.

Champion, Erik, and Hafizur Rahaman. 2020. "Survey of 3D Digital Heritage Repositories and Platforms." *Virtual Archaeology Review*, 11(23): 1–15. https://doi.org/10.4995/var.2020.13226.

Chapinal-Heras, Diego, Carlos Díaz-Sánchez, Sergio España-Chamorro, Natalia Gómez-García, Lucía Pagola-Sánchez, Manuel Parada López De Corselas, and Manuel Elías Rey-Álvarez Zafiria. 2023. "Photogrammetry and 3D Modelling in University Teaching. A Case Study Applied to the History Degree." *Journal of Classics Teaching* 24 (48): 133–42. https://doi.org/10.1017/S2058631023000429.

Cignoni, P., G. Ranzuglia, M. Callieri, M. Corsini, F Ganovelli, N Pietroni, and M. Tarini. 2011. MeshLab [Software]. (Cross platform). https://hdl.handle.net/2434/625490.

Collina, F. (2025). Busto Aldrovandi—RAWp model [Graphic]. Zenodo, Version v1. https://doi.org/10.5281/zenodo.15389391.

D'Andrea, Andrea, and Kate Fernie. 2013. "CARARE 2.0: A Metadata Schema for 3D Cultural Objects." In *2013 Digital Heritage International Congress (DigitalHeritage)*, 2: 137–43. https://doi.org/10.1109/DigitalHeritage.2013.6744745.

Daquino, Marilena, Silvio Peroni, and David Shotton. 2020. "The OpenCitations Data Model." figshare. https://doi.org/10.6084/M9.FIGSHARE.3443876.V7.

De Paolis, Lucio Tommaso, Valerio De Luca, Carola Gatto, Giovanni D'Errico, and Giovanna Ilenia Paladini. 2020. "Photogrammetric 3D Reconstruction of Small Objects for a Real-Time Fruition." In *Augmented Reality, Virtual Reality, and Computer Graphics*, eds. Lucio Tommaso De Paolis and Patrick Bourdot, 12242: 375–94. Lecture Notes in Computer Science. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-030-58465-8_28.

Demetrescu, Emanuel, Enzo d'Annibale, Daniele Ferdani, and Bruno Fanini. 2020. "Digital Replica of Cultural Landscapes: An Experimental Reality-Based Workflow to Create Realistic, Interactive Open World Experiences." *Journal of Cultural Heritage* 41 (January): 125–41. https://doi.org/10.1016/j.culher.2019.07.018.

Dimou, Anastasia, Miel Vander Sande, Pieter Colpaert, Ruben Verborgh, Erik Mannens, and Rik Van de Walle. 2014. "RML: A Generic Language for Integrated RDF Mappings of Heterogeneous Data." In *Proceedings of the 7th Workshop on Linked Data on the Web*, ed. Christian Bizer et al., 1184.

Doerr, Martin. 2003. "The CIDOC Conceptual Reference Module: An Ontological Approach to Semantic Interoperability of Metadata." *AI Magazine* 24 (3): 75–75. https://doi.org/10.1609/aimag.v24i3.1720.

Doerr, Martin, and Maria Theodoridou. 2011. "CRMdig: A Generic Digital Provenance Model for Scientific Observation." Paper presented at 3rd USENIX Workshop on the Theory and Practice of Provenance, Heraklion, Crete, Greece, June 20-21. https://cidoc-crm.org/sites/default/files/CRMdig-TAPP11.pdf.

Dostal, Christopher, and Kotaro Yamafune. 2018. "Photogrammetric Texture Mapping: A Method for Increasing the Fidelity of 3D Models of Cultural Heritage Materials." *Journal of Archaeological Science: Reports* 18 (April): 430–36. https://doi.org/10.1016/j.jasrep.2018.01.024.

Ekengren, Fredrik, Nicolò Dell'Unto, Matteo Callieri, Paolo Derudas, Daniele M. Campanaro, Giulia Naponiello, Åsa Ottosson Berggren, Erik Arvidsson, and Robert Ringdahl. 2024. *Dynamic Collection Plus: A 3D Web Infrastructure Designed to Support Higher Education and Research in Archaeology*. Web publication. Department of Archaeology and Ancient History, Lund University. https://dyncolldev.ht.lu.se/plus/.

European Commission. Directorate General for Communications Networks, Content and Technology. 2022. *Study on Quality in 3D Digitisation of Tangible Cultural Heritage: Mapping Parameters, Formats, Standards, Benchmarks, Methodologies, and Guidelines—Executive Summary*. Publications Office. https://data.europa.eu/doi/10.2759/581678.

European Commission. Directorate General for Research and Innovation, EOSC Executive Board. 2021. *EOSC Interoperability Framework: Report from the EOSC Executive Board Working Groups FAIR and Architecture*. Publications Office. https://data.europa.eu/doi/10.2777/620649.

Fanini, Bruno, Daniele Ferdani, Emanuel Demetrescu, Simone Berto, and Enzo d'Annibale. 2021. "ATON: An Open-Source Framework for Creating Immersive, Collaborative and Liquid Web-Apps for Cultural Heritage." *Applied Sciences* 11 (22): 11062. https://doi.org/10.3390/app112211062.

Farella, E. M., L. Morelli, E. Grilli, S. Rigon, and F. Remondino. 2022a. "Handling Critical Aspects in Massive Photogrammetric Digitization of Museum Assets." *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XLVI-2/W1-2022 (February): 215–22. https://doi.org/10.5194/isprs-archives-XLVI-2-W1-2022-215-2022.

Farella, Elisa Mariarosaria, Luca Morelli, Simone Rigon, Eleonora Grilli, and Fabio Remondino. 2022b. "Analysing Key Steps of the Photogrammetric Pipeline for Museum Artefacts 3D Digitisation." *Sustainability* 14 (9): 5740. https://doi.org/10.3390/su14095740.

Flotyński, Jakub. 2013. "Harvesting of Semantic Metadata from Distributed 3D Web Content." In *2013 6th International Conference on Human System Interactions (HSI)*, 193–200. https://doi.org/10.1109/HSI.2013.6577822.

Gabellone, Francesco. 2022. "Digital Twin: A New Perspective for Cultural Heritage Management and Fruition." *Acta IMEKO* 11 (1): 7 pp.-7 pp. https://doi.org/10.21014/acta_imeko.v11i1.1085.

Geng, Jason. 2011. "Structured-Light 3D Surface Imaging: A Tutorial." *Advances in Optics and Photonics* 3 (2): 128–60. https://doi.org/10.1364/AOP.3.000128.

Grieves, Michael, and John Vickers. 2017. "Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems." In *Transdisciplinary Perspectives on Complex Systems: New Findings and Approaches*, edited by Franz-Josef Kahlen, Shannon Flumerfelt, and Anabela Alves, 85–113. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-38756-7_4.

Gu, Qing, Kyriakos Herakleous, and Charalambos Poullis. 2016. "3DUNDERWORLD-SLS: An Open-Source Structured-Light Scanning System for Rapid Geometry Acquisition." Preprint, submitted August 20. http://arxiv.org/abs/1406.6595.

Gualandi, Bianca, and Silvio Peroni. 2024. "Data Management Plan: Second Version." Zenodo, Version 1.0. https://doi.org/10.5281/ZENODO.10727879.

Guidi, G., L. L. Micoli, S. Gonizzi, P. Rodriguez Navarro, and M. Russo. 2013. "3D Digitizing a Whole Museum: A Metadata Centered Workflow." In *2013 Digital Heritage International Congress (DigitalHeritage)*, 307–10. Marseille, France: IEEE. https://doi.org/10.1109/DigitalHeritage.2013.6744768.

Harpring, Patricia. 2010. "Development of the Getty Vocabularies: AAT, TGN, ULAN, and CONA." *Art Documentation: Journal of the Art Libraries Society of North America* 29 (1): 67–72. https://doi.org/10.1086/adx.29.1.27949541.

Hernández-Muñoz, Óscar. 2023. "Analysis of Digitized 3D Models Published by Archaeological Museums." *Heritage*, 6(5), 3885–3902. https://doi.org/10.3390/heritage6050206.

Heyvaert, Pieter, Ben De Meester, Anastasia Dimou, and Ruben Verborgh. 2018. "Declarative Rules for Linked Data Generation at Your Fingertips!" In *The Semantic Web: ESWC 2018 Satellite Events*, edited by Aldo Gangemi, Anna Lisa Gentile, Andrea Giovanni Nuzzolese, Sebastian Rudolph, Maria Maleshkova, Heiko Paulheim, Jeff Z. Pan, and Mehwish Alam, 213–17. Lecture Notes in Computer Science. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-98192-5_40.

Homburg, Timo, Anja Cramer, Laura Raddatz, and Hubert Mara. 2021. "Metadata Schema and Ontology for Capturing and Processing of 3D Cultural Heritage Objects." *Heritage Science* 9 (1): 91. https://doi.org/10.1186/s40494-021-00561-w.

Hutson, Matthew. 2018. "Artificial Intelligence Faces Reproducibility Crisis." *Science* 359 (6377): 725–26. https://doi.org/10.1126/science.359.6377.725.

Jakob, Wenzel, Marco Tarini, Daniele Panozzo, and Olga Sorkine-Hornung. 2015. "Instant Field-Aligned Meshes." *ACM Transactions on Graphics* 34 (6): 1–15. https://doi.org/10.1145/2816795.2818078.

Jones, Christine A., and Elizabeth Church. 2020. "Photogrammetry Is for Everyone: Structure-from-Motion Software User Experiences in Archaeology." *Journal of Archaeological Science: Reports* 30 (April): 102261. https://doi.org/10.1016/j.jasrep.2020.102261.

Kerbl, Bernhard, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. "3D Gaussian Splatting for Real-Time Radiance Field Rendering." Preprint, submitted August 8, 2023. https://doi.org/10.48550/ARXIV.2308.04079.

Kim, Eunseok, Jungwha Kim, and Woontack Woo. 2015. "Metadata Schema for Context-Aware Augmented Reality Applications in Cultural Heritage Domain." In *2015 Digital Heritage*, 2: 283–90. https://doi.org/10.1109/DigitalHeritage.2015.7419509.

Kolosova, Alexia, and Sorin Hermon. 2013. "Metadata Enhanced 3D Content Search for Real-Time Visualization of 3D Digital Assets." In *2013 Digital Heritage International Congress (DigitalHeritage)*, 2: 173–76. https://doi.org/10.1109/DigitalHeritage.2013.6744750.

Koster, Lukas, and Saskia Woutersen-Windhouwer. 2018. "FAIR Principles for Library, Archive and Museum Collections: A Proposal for Standards for Reusable Collections." *Code4Lib Journal*, no. 40. https://journal.code4lib.org/articles/13427.

Kraus, Karl. 2007. *Photogrammetry: Geometry from Images and Laser Scans*, translated by Ian Harley and Stephen Kyle. 2nd ed. Berlin: De Gruyter Brill. https://doi.org/10.1515/9783110892871.

Kritzinger, Werner, Matthias Karner, Georg Traar, Jan Henjes, and Wilfried Sihn. 2018. "Digital Twin in Manufacturing: A Categorical Literature Review and Classification." *IFAC-PapersOnLine*, 16th IFAC Symposium on Information Control Problems in Manufacturing INCOM 2018, 51 (11): 1016–22. https://doi.org/10.1016/j.ifacol.2018.08.474.

Lebo, T., S. Sahoo, and D. McGuinness. 2013. PROV-O: The PROV Ontology'. World Wide Web Consortium Recommendation, April 30. http://www.w3.org/TR/2013/REC-prov-o-20130430/.

Lentz, Rick William, Donald Paul Brutzman, and Michalis Kamburelis. 2021. "X3D and glTF Model Differencing for Conversions, Comparison and Conformance Testing." In *The 26th International Conference on 3D Web Technology*, 1–7. Pisa Italy: ACM. https://doi.org/10.1145/3485444.3493230.

Lo Brutto, Mauro, and M.G. Spera. 2011. "Sperimentazione Di Procedure Automatiche in Fotogrammetria Close-Range per Il Rilievo Di Beni Culturali." In *Atti 15a Conferenza Nazionale*. Colorno (Parma).

Lu, Yuhang, Jun Zhou, Jing Wang, Jun Chen, Karen Smith, Colin Wilder, and Song Wang. 2018. "Curve-Structure Segmentation From Depth Maps: A CNN-Based Approach and Its Application to Exploring Cultural Heritage Objects." *Proceedings of the AAAI Conference on Artificial Intelligence* 32 (1). https://doi.org/10.1609/aaai.v32i1.12306.

Luhmann, Thomas, Stuart Robson, Stephen Kyle, and Jan Boehm. 2023. *Close Range*

Photogrammetry and 3D Imaging. Caithness, Scotland: Walter de Gruyter GmbH & Co KG. https://www.degruyterbrill.com/document/doi/10.1515/9783110607253/html.

Malik, Umair Shafqat, and Gabriele Guidi. 2018. "Massive 3D Digitization of Sculptures: Methodological Approaches for Improving Efficiency." *IOP Conference Series: Materials Science and Engineering* 364 (June): 012015. https://doi.org/10.1088/1757-899X/364/1/012015.

Massidda, Marcello, Alice Bordignon, Francesca Fabbri, Raffaella Nori, Laura Piccardi, Laura Travaglini, Manuele Veggi, and Sofia Pescarin. 2024. "Towards an Experiment Planner for Cognitive Studies in Virtual Heritage Environments. A Pilot Study." In *GCH2024–Eurographics Workshop on Graphics and Cultural Heritage*. https://doi.org/10.2312/GCH.20241261.

McMillion, Matthew. 2022. '"How Does Structured-Light 3D Scanning Work?" Artec 3D Learning Center (blog). 2022. https://www.artec3d.com/learning-center/structured-light-3d-scanning.

Medina, Joshua J., James M. Maley, Siddharth Sannapareddy, Noah N. Medina, Cyril M. Gilman, and John E. McCormack. 2020. "A Rapid and Cost-Effective Pipeline for Digitization of Museum Specimens with 3D Photogrammetry." *PLoS ONE* 15 (8): e0236417. https://doi.org/10.1371/journal.pone.0236417.

Menna, F., E. Nocerino, D. Morabito, E. M. Farella, M. Perini, and F. Remondino. 2017. "An Open Source Low-Cost Automatic System for Image-Based 3D Digitization." *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XLII-2/W8 (November): 155–62. https://doi.org/10.5194/isprs-archives-XLII-2-W8-155-2017.

Menna, F., E. Nocerino, F. Remondino, M. Dellepiane, M. Callieri, and R. Scopigno. 2016. '3D Digitization of an Heritage Masterpiece—A Critical Analysis on Quality Assessment." *ISPRS—International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XLI-B5 (June): 675–83. https://doi.org/10.5194/isprsarchives-XLI-B5-675-2016.

Milosz, Marek, Jacek Kęsik, and Jerzy Montusiewicz. 2020. "3D Scanning and Visualization of Large Monuments of Timurid Architecture in Central Asia—A Methodical Approach". *Journal on Computing and Cultural Heritage* 14 (1): 1–31. https://doi.org/10.1145/3425796.

Moore, Jennifer, Adam Rountrey, and Hannah Scates Kettler. 2022. *3D Data Creation to Curation: Community Standards for 3D Data Preservation*. ALA. https://www.alastore.ala.org/content/3d-data-creation-curation-community-standards-3d-data-preservation.

Moretti, Arianna. 2024. Modelli Spreadsheet CHANGES—Acquisizione e Oggetti (1.0) [Data set]. Zenodo. Version 1.0 https://doi.org/10.5281/zenodo.14277220

Niccolucci, Franco, Béatrice Markhoff, Maria Theodoridou, Achille Felicetti, and Sorin Hermon. 2023. "The Heritage Digital Twin: A Bicycle Made for Two. The Integration of Digital Methodologies into Cultural Heritage Research." *Open Research Europe* 3: 64. https://doi.org/10.12688/openreseurope.15496.1.

Niven, Daniel J., T. Jared McCormick, Sharon E. Straus, Brenda R. Hemmelgarn, Lianne Jeffs, Tavish R. M. Barnes, and Henry T. Stelfox. 2018. "Reproducibility of Clinical Research in Critical Care: A Scoping Review." *BMC Medicine* 16 (1): 26. https://doi.org/10.1186/s12916-018-1018-6.

Papadopoulos, Costas, Kelly Gillikin Schoueri, and Susan Schreibman. 2025. "And Now What? Three-Dimensional Scholarship and Infrastructures in the Post-Sketchfab Era." *Heritage* 8(3): 99. https://doi.org/10.3390/heritage8030099.

Peels, Rik. 2019. "Replicability and Replication in the Humanities." *Research Integrity and Peer Review* 4: 2. https://doi.org/10.1186/s41073-018-0060-4.

Peels, Rik, and Lex Bouter. 2018. "The Possibility and Desirability of Replication in the Humanities." *Palgrave Communications* 4: 95. https://doi.org/10.1057/s41599-018-0149-x.

Peinado-Santana, Sara, Patricia Hernández-Lamas, Jorge Bernabéu-Larena, Beatriz Cabau-Anchuelo, and José Antonio Martín-Caro. 2021. "Public Works Heritage 3D Model Digitisation, Optimisation and Dissemination with Free and Open-Source Software and Platforms and Low-Cost Tools." *Sustainability* 13 (23): 13020. https://doi.org/10.3390/su132313020.

Peng, Roger. 2015. "The Reproducibility Crisis in Science: A Statistical Counterattack". *Significance* 12 (3): 30–32. https://doi.org/10.1111/j.1740-9713.2015.00827.x.

Peroni, Silvio. 2017. "A Simplified Agile Methodology for Ontology Development." In *OWL: Experiences and Directions—Reasoner Evaluation*, edited by Mauro Dragoni, María Poveda-Villalón, and Ernesto Jimenez-Ruiz, 10161:55–69. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-54627-8_5.

Peroni, Silvio, David Shotton, and Fabio Vitali. 2016. "A Document-Inspired Way for Tracking Changes of RDF Data. The Case of the OpenCitations Corpus." In *Detection, Representation and Management of Concept Drift in Linked Open Data*, edited by Laura Hollink, Sándor Darányi, Albert Meroño Peñuela, and E. Kontopoulos, 26–33. Bologna: CEUR Workshop Proceedings. http://ceur-ws.org/Vol-1799/Drift-a-LOD2016_paper_4.pdf.

Pescarin, Sofia, d'Annibale, Enzo, Fanini, Bruno, and Ferdani, Dantiele. 2018. "Prototyping on Site Virtual Museums: The Case Study of the Co-design Approach to the Palatine Hill in Rome (Barberini Vineyard) Exhibition." In *2018 3rd Digital Heritage International Congress (DigitalHERITAGE) held jointly with 2018 24th International Conference on Virtual Systems & Multimedia (VSMM 2018)* IEEE. https://ieeexplore.ieee.org/abstract/document/8810135.

Rahaman, Hafizur, and Erik Champion. 2019. "To 3D or Not 3D: Choosing a Photogrammetry Workflow for Cultural Heritage Groups." *Heritage* 2(3): 1835–51. https://doi.org/10.3390/heritage2030112.

Rahaman, Hafizur, Erik Champion, and Mafkereseb Bekele. 2019. "From Photo to 3D to Mixed Reality: A Complete Workflow for Cultural Heritage Visualisation and Experience." *Digital Applications in Archaeology and Cultural Heritage* 13 (June):e00102. https://doi.org/10.1016/j.daach.2019.e00102.

Raimo, Nicola, Ivano De Turi, Alessandra Ricciardelli, and Filippo Vitolla. 2022. 'Digitalization in the Cultural Industry: Evidence from Italian Museums'. *International Journal of Entrepreneurial Behavior & Research* 28 (8): 1962–74. https://doi.org/10.1108/IJEBR-01-2021-0082.

Remondino, Fabio. 2011. 'Heritage Recording and 3D Modeling with Photogrammetry and 3D Scanning'. *Remote Sensing* 3 (6): 1104–38. https://doi.org/10.3390/rs3061104.

Renda, Giulia, Marilena Daquino, and Valentina Presutti. 2023a. 'Melody: A Platform for Linked Open Data Visualisation and Curated Storytelling'. In *Proceedings of the 34th ACM Conference on Hypertext and Social Media*, 1–8. HT '23. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/3603163.3609035.

Renda, Giulia, Marco Grasso, and Marilena Daquino. 2023b. "From Ontology Design to User-Centred Interfaces for Music Heritage." In *La Memoria Digitale: Forme Del Testo e Organizzazione Della Conoscenza*, edited by Emmanuela Carbé, Gabriele Lo Piccolo, Alessia Valenti, and Francesco Stella, 168–72. Siena. https://doi.org/10.48550/arXiv.2306.12973.

Riva, Pat, Maja Žumer, and Trond Aalberg. 2022. "LRMoo, a High-Level Model in an Object-Oriented Framework." IFLA WLIC Dublin, October. https://repository.ifla.org/server/api/core/bitstreams/8b-6f71d0-e53c-4f41-952f-8d3a04372eae/content.

Robinete, Fabrice, Rémi Arnaud, Tony Parisi, and Patrick Cozzi. 2018. "glTF: Designing an Open-Standard Runtime Asset Formati." In *GPU Pro 360 Guide to 3D Engine Design*, edited by Wolfgang Engel, 243–60. A K Peters/CRC Press. https://www.taylorfrancis.com/chapters/edit/10.1201/9781351172486-20/gltf-designing-open-standard-runtime-asset-format-fabrice-robinet-re-mi-arnaud-tony-parisi-patrick-cozzi-fabrice-robinet-re%C2%B4mi-arnaud-tony-parisi-patrick-cozzi.

Ruiz, Rafael Melendreras, Ma Teresa Marín Torres, and Paloma Sánchez Allegue. 2022. "Comparative Analysis Between the Main 3D Scanning Techniques: Photogrammetry, Terrestrial Laser Scanner, and Structured Light Scanner in Religious Imagery: The Case of the Holy Christ of the Blood." *Journal on Computing and Cultural Heritage* 15 (1): 1–23. https://doi.org/10.1145/3469126.

Sapirstein, Philip. 2018. "A High-Precision Photogrammetric Recording System for Small Artifacts." *Journal of Cultural Heritage* 31 (May): 33–45. https://doi.org/10.1016/j.culher.2017.10.011.

Sebar, Leila Es, Sabrina Grassini, Marco Parvis, and Luca Lombardo. 2021. "A Low-Cost Automatic Acquisition System for Photogrammetry." *2021 IEEE International Instrumentation and Measurement Technology Conference (I2MTC)*, 1–6. Glasgow, United Kingdom: IEEE. https://doi.org/10.1109/I2MTC50364.2021.9459991.

Statham, Nataska. 2019. "Scientific Rigour of Online Platforms for 3D Visualization of Heritage." *Virtual Archaeology Review* 10 (20): 1–16. https://doi.org/10.4995/var.2019.9715.

The London Charter. "The London Charter for the Computer-Based Visualisation of Cultural Heritage." February, 2009. https://londoncharter.org/.

Triglav Čekada, Mihaela, Fabio Crosilla, and Mojca Kosmatin Fras. 2010. "Teoretična Gostota Lidarskih Točk Za Topografsko Kartiranje v Največjih Merilih." *Geodetski Vestnik* 54 (03): 403–16. https://doi.org/10.15292/geodetski-vestnik.2010.03.403-416.

Vuković, M., J. Balen, H. Potrebica, M. Mađerić, and V. Španiček. 2022. "3D Digitization of Museum Artefacts within the Interreg Iron Age Danube Project." *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XLIII-B2-2022 (May):1159–65. https://doi.org/10.5194/isprs-archives-XLIII-B2-2022-1159-2022.

Webb, E. K., S. Robson, and R. Evans. 2020. "Quantifying Depth of Field and Sharpness for Image-based 3D Reconstruction of Heritage Objects." *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XLIII-B2-2020 (August): 911–18. https://doi.org/10.5194/isprs-archives-XLIII-B2-2020-911-2020.

Whitt, Phillip. 2023. *Beginning Photo Retouching and Restoration Using GIMP: Learn to Retouch and Restore Your Photos Like a Pro*. 2nd ed. Berkeley, CA: Apress.

Wu, Shan. 2022. "Research on the Application of Computer Virtual Reality Technology in Museum Cultural Relics Exhibition Hall." *2022 IEEE International Conference on Advances in Electrical Engineering and Computer Applications (AEECA)*, 1049–53. Dalian, China: IEEE. https://doi.org/10.1109/AEECA55500.2022.9919104.