TRACE-ing the Gaps: Mapping Interventions on Incomplete 3D Meshes

Luisa Ammirati[1], Alice Bordignon[1*], Federica Collina[2], Francesca Fabbri[1], Bruno Fanini[3], Daniele Ferdani[3], Maria Felicia Rega[3], Mattia Sullini[4]

[1]Department of Classical Philology and Italian Studies, University of Bologna, Via Zamboni 32, 40126, Bologna, Italy - luisa.ammirati2@unibo.it, alice.bordignon2@unibo.it, francesca.fabbri80@unibo.it
[2]Department of Cultural Heritage, University of Bologna, Via degli Ariani 1, 48121, Ravenna, Italy - federica.collina5@unibo.it
[3]Digital Heritage Innovation Lab, Institute of Heritage Science, National Research Council, Via Salaria km 29, 300, 00010, Montelibretti, Rome, Italy - bruno.fanini@cnr.it, daniele.ferdani@cnr.it, mariafeliciarega@cnr.it
[4]Department of Architecture, University of Bologna, Viale Del Risorgimento 2, 40136, Bologna, Italy - mattia.sullini2@unibo.it

Keywords: 3D Modelling, Digital Cultural Heritage, Hole-filling, 3D Digitisation, Paradata, FAIR.

Abstract

The 3D digital representations of physical objects are inevitably marked by gaps and imperfections arising from both technological limitations and contextual acquisition constraints. These shortcomings often necessitate manual or semi-automatic interventions, introducing elements of subjectivity into the modelling process. Despite the increasing use of 3D technologies across the Galleries, Libraries, Archives, and Museums (GLAM) sector, a standardised methodology for classifying and documenting mesh incompleteness is still lacking, undermining the transparency, reusability, and interoperability of such data. This paper addresses two core research questions: (RQ1) Which acquisition conditions can be identified as the main contributing factors to the occurrence of holes in 3D meshes? (RQ2) How can topological modifications resulting from the manual integration of holes in 3D cultural heritage models be automatically tracked and visualised? To address these questions, we propose a classification framework grounded in five case studies drawn from two digitisation campaigns conducted within the CHANGES project (Spoke 4). We further present a workflow for a post-hoc, method-agnostic identification approach for the automatic tracing and interactive visualisation of topological edits, implemented within a Web3D environment based on the ATON framework. Situated within a pipeline that adheres to FAIR principles, this methodology aims to enhance the interpretability and scholarly value of digital reconstructions by embedding provenance-aware metadata. Building on established approaches in critical architectural reconstruction, this work proposes a methodological extension tailored to the nuanced management of holes in cultural heritage 3D models.

1. Introduction

The 3D digital representation of real-world objects is inherently imperfect due to technological and data collection limitations, leading to the presence of gaps and inaccuracies widely recognised in the literature (Davis et al. 2002; Attene et al. 2013). These imperfections can range from local connectivity issues, where the mesh connectivity is compromised, to global topological errors, such as incorrect component counts, genus discrepancies, or orientation problems (Campen et al. 2012).

Geometric defects, in particular, pose more complex challenges that require case-specific analysis and nuanced solutions (Attene et al. 2013). These defects are referred to as holes and typically arise from insufficient surface sampling during the Surface Reconstruction phase (Sabbah et al. 2013). Holes may encompass areas of missing data of various dimensions, making their repair a challenging task (Attene et al. 2013).

Resolving these issues requires a combination of automated processes and manual interventions, inherently incorporating elements of subjectivity and authorial decision-making. Despite the growing adoption of 3D technologies in the Galleries, Libraries, Archives, and Museums (GLAM) sector, there is a lack of a shared methodology for classifying and analytically tracing mesh inaccuracies arising from acquisition conditions, leading to issues of data transparency and interoperability. Building on these premises, this study aims to address two complementary research questions:

RQ1: Which acquisition conditions can be identified as the main contributing factors to the occurrence of holes in 3D meshes?

RQ2: How can topological modifications resulting from the manual integration of holes in 3D cultural heritage models be automatically tracked and visualised?

To address these questions, the article is structured as follows. Section 2 introduces the most relevant related works on hole-filling techniques and the use of paradata to detect authorial interventions in the field of cultural heritage. Both research questions are explored in Section 3.

First, we propose a basic categorisation that links five representative case studies to the main problematic acquisition conditions encountered during two extended digitisation campaigns carried out within the CHANGES project, Spoke 4 (https://www.fondazionechanges.org/spoke-4/). This categorisation aims to identify and confirm potential patterns in the systematic presence of mesh gaps caused by acquisition techniques and environmental conditions.

Second, to address RQ2, the case studies serve as a basis for proposing a methodology for detecting manual integrations and visualising such data. In this context, to enable interactive visualisation, comparison, and highlighting of the obtained results, the Texture Rendering with Annotated Color Encoded (TRACE) is introduced, a small online Web3D tool developed on the ATON framework (https://aton.ispc.cnr.it/). Finally, in Section 4, we conclude the contribution by discussing the results obtained and the possible further improvement of the methodology to be implemented in the future.

Starting from an established FAIR-compliant pipeline developed within the project, this study seeks to improve 3D object reconstructions by enriching their informational depth, incorporating metadata that conveys data certainty and quality alongside geometric and texture details. Our proposal is, therefore, to adopt, through a process of methodological adaptation, the approach developed for documenting the integration of critical reconstructions of lost or non-extant architectural elements and to apply it to the remediation of gaps in digital models caused by intrinsic limitations in the data acquisition and processing workflow.

2. Related Work

Acquisition conditions strongly affect the quality of the outcome, influencing the geometric accuracy and, consequently, the presence of holes and noise levels (Ch'ng et al. 2019; Apollonio et al. 2021). Indeed, surveying artifacts in museums can be particularly challenging due to various constraints imposed by the environment—such as limited space and poor lighting conditions—which especially affect passive techniques, as they are more sensitive to ambient illumination than active sensors (Nicolae et al. 2014).

Furthermore, holes can result from factors like low reflectance, scanner placement limitations (Davis et al. 2002), or restricted accessibility to certain areas (Aldeeb and Hellwich, 2017). However, a systematic taxonomy of geometric issues tied specifically to acquisition methods is still lacking in the literature.

Regarding hole filling, existing literature predominantly emphasises automatic methods for gap closure (Santos et al. 2012; Sabbah et al. 2013; Aldeeb and Hellwich, 2017), while a systematic framework incorporating the role of human intervention in mesh reconstruction remains underexplored. Centin and Signoroni (2018) provide an interesting and detailed overview of the different types of automatic hole-filling techniques, along with explanations of the underlying algorithms.

Hole-filling and mesh completion are typically addressed through surface-based or volumetric reconstruction techniques. Surface-based methods, such as Poisson Reconstruction (Kazhdan et al. 2006), are robust but necessitate boundary matching, which becomes challenging for complex topologies. While they are effective for small holes, they may result in detail loss and artifacts, particularly in high-precision applications. Volumetric methods, like Octree-based reconstruction (Podolak and Rusinkiewicz, 2005), handle complex meshes more effectively but require re-meshing, which may lead to loss of detail due to voxelisation and contouring. Hybrid methods, such as the approach proposed by Centin and Signoroni (2018), combine volumetric guidance with surface-based filling techniques.

In addition to the classical algorithmic methods just mentioned that work exclusively based on 3D geometry, different approaches have been proposed that work on 2D projections of hole contours and solve them directly on the images by interpolation (Pèrez et al. 2021), or use the images to deform a provisional fill mesh as a solution to an energy-minimisation problem (Brunton et al. 2010).

In recent years, deep learning approaches have advanced the use of generative models to address surface completeness while disregarding topological information (Hermoza and Sipiran, 2018; Van Nguyen et al. 2024). Convolutional neural networks (CNNs) have also proven effective for hole repair, either by exploiting contextual cues from the surrounding geometry in the 2D domain (Wang et al. 2020) or by incorporating curvature-aware features of the mesh (Hernandez-Bautista and Melero, 2025).

Lastly, manual intervention by 3D modelers is considered essential, mainly when dealing with large and/or complex gaps (Campen et al. 2012; Sabbah et al. 2013; Apollonio et al. 2021), and this statement is still valid even with the aforementioned AI-based methods. In these cases, contextual interpretation and reconstruction significantly improve the precision of the results compared to purely mathematical or probabilistic solutions (Campen et al. 2012). This aspect, often underestimated by non-experts (Campen et al. 2012), is crucial in high-precision fields such as cultural heritage preservation.

Regardless of whether the hole- and gap-filling process is performed algorithmically, through AI inference, or manually, detailed tracking and documentation are essential. But this requirement becomes even more stringent in manual approaches, where the integration of new data involves critical reasoning in addition to technical mesh analysis, especially when dealing with geometries of scientific value (Apollonio et Giovannini, 2015; Apollonio et al. 2021a; Luengo et al. 2025).

Metadata and, more specifically, paradata play a crucial role in ensuring the accuracy and credibility of the model's visual output (ICOMOS General Assembly, 2017; Börjesson et al. 2020; Huvila, 2022). To address this need, Luengo et al. (2025) propose a standardised framework incorporating both paradata and metadata, which includes an entry for "Manual intervention/correction." This feature allows for tracking the extent of manual involvement, from full manual reconstruction to entirely automated processes, highlighting the importance of documenting the decision-making process in manual interventions.

Nevertheless, to effectively demonstrate the contribution of the author's intervention to the model, the inclusion of paradata in text form, while machine-readable, can be insufficient (Apollonio and Giovannini, 2015). A more comprehensive approach is needed to visually capture and communicate the extent of manual corrections made during the reconstruction process. This highlights the need for new methods that move beyond textual paradata, favouring more intuitive, media-based tools that align with visualisation and enhance accessibility for multidisciplinary teams, thus enabling more reliable data sharing (Apollonio and Giovannini, 2015).

Although there is a body of literature focused on documenting the reconstruction of 3D objects without a real-world counterpart, particularly historical reconstructions based on sources such as sketches, drawings, and photographs, mainly in the architectural field (Apollonio and Giovannini, 2015; Hauck and Kuroczyński, 2016; Apollonio et al. 2021a), the literature concerning reality-based models that exhibit holes—i.e., a mix of surfaces with high accuracy and parts with variable accuracy—remains scarce.

3. Materials and Methods

3.1 CHANGES project case studies

The case studies presented are part of the survey campaigns carried out within Spoke 4 of the CHANGES project, specifically: the pilot study dedicated to the creation of the digital twin of the temporary exhibition *The Other Renaissance: Ulisse Aldrovandi and the Wonders of the World* (https://site.unibo.it/aldrovandi500/en/mostra-l-altro-rinascimento) held in Palazzo Poggi Museum, focused on the

https://doi.org/10.5194/isprs-annals-X-M-2-2025-3-2025 |

naturalist Ulisse Aldrovandi; and the core project conducted at the Geological Collection Giovanni Capellini Museum, part of the University of Bologna museum system. The two acquisition campaigns included, respectively, a total of 301 objects (Balzani et al. 2024) and 87 objects digitised at the Capellini Museum. Both projects adopted an approach aligned with the FAIR principles of open science in terms of metadata management (Barzaghi et al. 2024) to develop a methodology that could be reproducible in other contexts (Barzaghi et al. 2024a; Barzaghi et al. 2024b).

Our research output prioritises the preservation of three 3D model versions to digitally represent each physical asset, as shown in Figure 1:

- Processed Raw Model (RAWp): the preliminary output obtained from the photogrammetry or 3D scanner software, based only on automatic and semi-automatic intervention to solve minor imperfections.
- Digital Cultural Heritage Object (DCHO): the complete and refined model obtained after solving geometry and texture issues.
- Optimised Digital Cultural Heritage Object (DCHOo): the optimised version for seamless, real-time online interaction on web-based platforms.

The preservation of these three different models ensures that the transformations applied to the model during all stages are fully traceable.

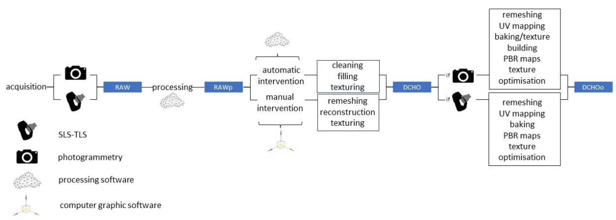

Figure 1. 3D data taxonomy established as research output.

The process began with an evaluation of acquisition methodologies. Since both campaigns took place in museum settings, the acquisition had to be carried out within strict time constraints, making use of the days when the spaces were closed to visitors. The acquisition processes were also impacted by operational and spatial limitations. Specifically, the inability to move certain items from their display cases for conservation and the lighting conditions, not always adjustable. This contributed to datasets that were sometimes incomplete, requiring the use of automatic, semi-automatic, and manual hole-filling interventions by the operator.

The initial collection of 3D raw data (RAW) was conditioned by the objects' heterogeneity in terms of shape, size, and material. In Aldrovandi's digitisation campaign, the heterogeneity of the objects involved a wide range of typologies, including natural specimens, printed and manuscript volumes, scientific instruments, woodcuts, dioramas, nautical charts, archaeological findings, models and artworks such as paintings and sculptures. In the context of the acquisitions carried out for the Capellini collections, the objects to digitise included primarily fossils, bones, and plaster casts, all belonging to the paleontological domain.

The digitisation process involves several phases to provide the digital representation of each physical asset: 1) the **acquisition phase**, where the most suitable technique among photogrammetry and scanning technologies is selected to capture the object’s geometry and texture; 2) the **processing phase**, where small holes or gaps in the raw data are filled through interpolation using Software from Motion (SfM) or 3D scanning software. In this phase, human operators define the input parameters to solve minor imperfections, resulting in the RAWp. Issues that can’t be fully resolved during processing are addressed in 3) the **modelling phase**, where the operator manually corrects topological errors using 3D modelling software to obtain the DCHO. If the missing areas are extensive, orthographic photographs are used for reference; otherwise, the operator relies on their knowledge, visual perception, and tactile experience with the object to complete the model; 4) the **optimisation phase**, where the DCHO undergoes simplification and remeshing to achieve the DCHOo; 5) the **export phase**, where the various versions of the model (RAWp, DCHO, DCHOo) are converted into specific formats suitable for archiving and sharing; 6) a **metadata and provenance creation** phase, during which structured information is generated for each physical and digital asset, and provenance details are tracked and recorded for every metadata entry; 7) finally, in the **presentation phase**, the DCHOo is uploaded to the web platform ATON (Fanini et al. 2021), enabling online access and user interaction with the 3D content.

For a detailed description of each stage of this workflow and the management of the resulting data, see previous work (Balzani et al. 2024; Barzaghi et al. 2024, Barzaghi et al. 2024a, Barzaghi et al. 2024b).

3.2 Tracing integrations

During the acquisition campaigns of the objects, we encountered significant challenges due to various factors, which can be summarised into four main categories:

a) **Problematic lighting conditions**: the lighting in the exhibition space was not optimised for digital capture, resulting in cast shadows, reflections and uneven illumination that interfered with accurate data capture.

b) **Geometric complexity**: objects with intricate shapes, undercuts, cavities, or fine surface details are harder to capture accurately, especially if occluded parts are not accessible from all angles and if the mesh sampling reflects the frequency of the geometric features (Lévy and Zhang, 2010).

c) **Immovability of the object**: many of the objects could not be removed from their display cases, due to conservation restrictions, fragility, or size, limiting the angles and proximity from which they could be scanned or photographed. Furthermore, the display cases themselves often posed a physical barrier to data acquisition

d) **Non-Lambertian characteristics**: materials, such as dielectrics (e.g., glass, porcelain) and conductors (e.g., metals), exhibit specular reflection properties that vary depending on the viewer’s position. These materials reflect light in non-uniform ways, which can introduce inconsistencies in photogrammetric workflows and hinder accurate interpretation in scans based on structured light (SLS). In contrast, Lambertian surfaces appear evenly lit from all viewing angles and, therefore, do not present such issues.

In this context, we will discuss 5 representative case studies of these four conditions (Tab. 1), applying a common methodology

https://doi.org/10.5194/isprs-annals-X-M-2-2025-3-2025

conceived to highlight the mesh modifications due to manual and semi-automatic reconstruction made on the RAWp to obtain the DCHO and the DCHOo.

Problematic condition	Representative case studies	Acquisition techniques
a	CS1	Photogrammetry
b	CS2	Structured Light Scanner
c	CS3	Photogrammetry
d	CS4	Photogrammetry
	CS5	Structured Light Scanner

Table 1. Problematic acquisition conditions and representative case studies.

Once the RAW data has been generated, it is processed within the SfM or 3D scanning software to obtain the RAWp. During this process, one possible step is to use an initial automatic hole-filling procedure directly within the software, following an analysis of the gaps. However, this option is generally effective only for smaller holes, while it proves unsatisfactory for larger gaps, which require more extensive intervention (Apollonio et al. 2021).

Following processing, the RAWp model is imported into a 3D modelling software, where further refinement is achieved through a combination of semi-automatic and manual modifications.

A key step in the post-processing phase entails addressing the excessive geometry present in areas where alignment errors are most pronounced and the elimination of redundant and misaligned geometry. Semi-automatic techniques are then employed to fill the connectivity errors and solve non-manifold issues. Subsequently, sculpting operations are performed to slightly soften the excessively sharp edges of the newly added geometry. In some cases, when a dynamic topology tool (i.e. *Dyntopo* in Blender) is used during the process to perform real-time remeshing of manual integrations, an initial form of tracking can be obtained on the texture map as a side effect of the UV destruction that occurs alongside dynamic topology. This feature allows for dynamic tessellation of the mesh by locally subdividing or collapsing geometry based on the sculpting strokes, enabling more detailed modelling without predefined topology.

In a previous work (Bordignon et al. 2025; Bordignon et al. 2025a), we described a method for tracking the integrations applied to incomplete meshes by continuously and contextually updating a semantic vertex color map. In the present work, we propose a generalisation of the concept of communicating the degree of confidence associated with modifications and integrations without relying on continuous tracking. Instead, we introduce a post-hoc, method-agnostic identification approach.

In the case studies presented here, although the processes were carried out using the tracking-based method, we apply the generalised approach to serve a dual purpose: to test its applicability and to validate it by comparing its results with the data we tracked during processing. Thus, while the case study descriptions implicitly assume live tracking of all modifications, the final comparison remains agnostic to the parallel tracking sub-process, allowing for a post-hoc assessment independent of continuous modification tracking, thereby increasing resilience to errors, inconsistencies across file formats, or data loss.

To systematically detect the geometry modifications applied to the model, a false-colour texturing methodology is adopted. The RAWp is assigned a uniform colour texture to standardise the base appearance. Meanwhile, the mid-poly DCHO obtained from 3D modelling interventions is assigned an empty 4K texture. Following verification of the precise overlap between the two models, the Diffuse component is baked. In this context, depending on the specific modifications to be detected, the Extrusion and Max Ray Distance parameters can be adjusted. While Extrusion refers to the outward expansion of the low-poly mesh along its normals to ensure proper ray intersection during baking, the Max Ray Distance defines the maximum length a ray can travel to locate corresponding details on the high-poly mesh.

To detect all types of modifications made to the RAWp mesh—from minor manual interventions to resolve non-manifold issues, to remeshing, and systematic work on major holes—the value of 0.01 m proved to be optimal for both parameters. Starting from this value, it is possible—depending on the accuracy of the resulting mesh—to gradually adjust the parameters to enable selective detection: for instance, in cases where only major gaps are to be identified, rather than the overall mesh topology (Fig. 2). The false-color baking was planned to obtain a color-coded map (referred to as the *Refinement Trace Map* from this point forward), where render modifications and integrations on the model are identified in white (1) and the areas of the mesh that align perfectly with the original in black (0).

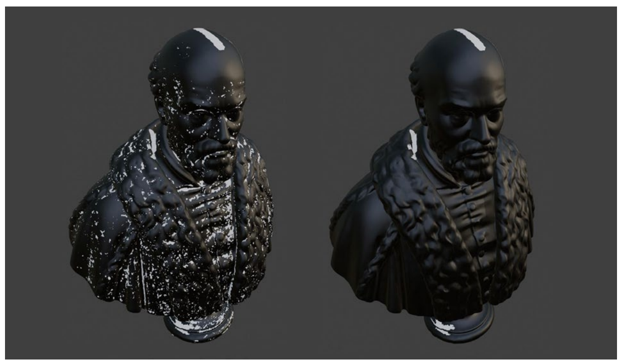

Figure 2. Example of selective detection on the 3D model of the bust of Ulisse Aldrovandi.

Similarly, to trace areas that, due to the presence of holes, lack sufficient data to generate a texture during reprojection, a secondary baking process is performed on the original texture map to highlight the areas where the texture will be completed using editing software. Indeed, while on one hand the areas where geometry has been reconstructed remain "empty", the same applies to certain regions where the baking process fails to transfer the original texture accurately.

This methodology enables clear differentiation between the original and the reconstructed areas, facilitating the detection of modifications throughout the post-processing workflow. The case studies presented in the following sections aim to provide concrete examples of the systematic implementation of this methodology for detecting manual interventions and visualising them, thereby further reinforcing the transparency of the initial pipeline presented.

3.2.1 Problematic lighting conditions: The mammoth case study (CS1) examines the digitisation of a specimen from the Giovanni Capellini Museum's collection, highlighting challenges encountered during its data acquisition under suboptimal lighting conditions using photogrammetry techniques (Nicolae et al. 2014).

The subject is a medium-to-large-sized mammoth cast, crafted from painted plaster and supported with an internal wrought iron framework. The item's fixed position at the intersection of a room and an adjoining corridor restricted operator mobility during data collection. The museum's lighting system lacked adjustable intensity controls. It primarily consisted of direct spotlights in both the main room and the corridor, leading to significant issues with contrast and reflectance on the object's surface. Additionally, the system did not allow the selective activation of light sources in the environment involved.

To mitigate reflectance problems, the environment's lighting was turned off. Natural light, diffused through a curtained window, adequately illuminated the model's front but left the rear insufficiently lit, resulting in uneven illumination. Available artificial lights lacked polarising filters and were inadequate for properly illuminating the model's rear, causing surface distortions.

Data acquisition was performed using a NIKON D750 camera, supported by a ladder to reach necessary heights. Camera settings were maintained at ISO 400, f/16 aperture, 48mm focal length, and a 6-second exposure time. The resulting RAW dataset was initially processed in RawTherapee (https://www.rawtherapee.com/) to address underexposure and white balance issues before being processed in 3DF Zephyr (https://www.3dflow.net/it/). Despite placing coded markers around the model, the software failed to align the rear and portions of the upper sections, leading to significant gaps and areas of low confidence in the RAWp.

Furthermore, the object's non-Lambertian surface characteristics, combined with challenging lighting conditions, contributed to generating a noisy mesh. The manual modelling of the missing parts was done in Blender (https://www.blender.org/), using the photographs of the rear as references to reconstruct and sculpt the model as faithfully as possible to the physical object (Fig. 3).

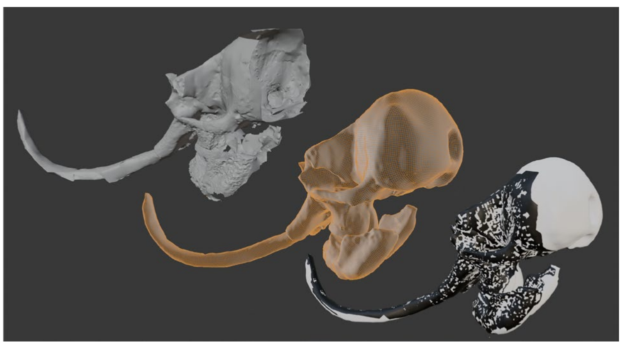

Figure 3. From left, the RAWp, the DCHO and the Refinement Trace Map of CS1.

3.2.2 Immovability of the object: The Mosasaurus skull (CS2) belongs to the Giovanni Capellini Museum's collection and was a small to medium-sized object with a dark, uneven and slightly reflective surface due to the preservative in which it had been immersed during restoration.

The find was displayed close to a wall and resting on a stand that maintained its correct anatomical position. Given the physical characteristics of the object, the impossibility of moving the element due to its weight, and the difficulty of acquisition due to its location, it was decided to acquire the Mosasaurus using an Artec Space Spider structured light projection scanner. This type of scanner allows acquisitions to be made at a close distance, plus adjusting the instrument sensitivity to overcome the problems caused by surface refraction.

The object and its base were acquired at 4 FPS, with texture data recorded every 3rd frame, and scanner sensitivity set to a low-medium level. The obtained scans were processed with Artec Studio 15 software (https://www.artec3d.com/).

The resulting mesh showed some small holes due to the surface irregularity and two huge ones in the lower portion of the skull at the exhibition stand and the support surface. Smaller and less complex holes were semi-automatically closed by the software, while the two larger ones required manually bridging their edges to reduce surface area and facilitate closure. Finally, the reconstructed surfaces have been chamfered to smooth out the difference between these and the surrounding surface (Fig. 4).

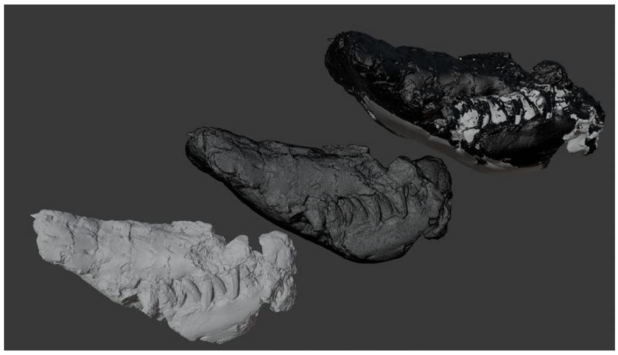

Figure 4. From left, the RAWp, the DCHO and the Refinement Trace Map of CS2.

3.2.3 Geometric complexity: The Hammerhead Shark (CS3) was featured in the exhibition dedicated to Ulisse Aldrovandi, alongside other zoological fish specimens distinguished by their complex geometry.

It posed several challenges during the digital acquisition process, primarily due to its intricate shape and fixed position on a support structure from which it could not be removed. Photogrammetry was chosen as the most suitable acquisition technique, as the inability to reposition the specimen would have restricted the range of scanning angles, resulting in occlusions and incomplete data capture with structured-light scanning.

The acquisition was conducted using a Sony Alpha 7i, placing the model inside an open white box and covering the support structure with a white cloth to minimise reflections. Camera settings were maintained at ISO 800, f/10 aperture, 28mm focal length, and a 1.6-second exposure time. The support was later reconstructed through manual 3D modelling in Blender.
Following the photogrammetric reconstruction in 3DF Zephyr, the processed model exhibited numerous geometric holes, particularly in the thinner regions of the body, such as the head, fins and tail. These holes were primarily attributed to small cavities, occlusions and difficulties in achieving sufficient image

overlap in these areas. Additional holes were observed in the lower portion of the body, caused both by the difficulty of capturing it with the camera lens and the features of the fish's skin, which disrupted texture reconstruction. The interplay of these factors led to challenges in achieving an accurate and complete representation of the specimen's surface characteristics. The holes were closed using a combination of semi-automatic and manual operations, all carried out in Blender.

Nevertheless, the baking process highlights not only the regions where new geometry has been introduced but also those where texture reconstruction proves particularly challenging. A notable example is the junction between the right fin and the body of the hammerhead shark, where a chromatic inconsistency is observed. This anomaly likely arises from difficulties in re-establishing a seamless texture transition in that specific area (Fig. 5).

Figure 5. From the bottom, the RAWp, the DCHO and the Refinement Trace Map of CS3.

3.2.4 Non-Lambertian characteristics: The *Linum* botanical model (CS4), part of the Brendel collection (https://sma.unibo.it/it/il-sistema-museale/orto-botanico-ed-erbario/collezioni/i-modelli-botanici-brendel), historically used for educational purposes in academic institutions, was selected for digital acquisition as part of the exhibition dedicated to Ulisse Aldrovandi.

The model, representing the flower of *Linum usitatissimum*—a plant belonging to the Linaceae family—is characterised by a high level of accuracy in relation to the actual flower and was originally constructed using a variety of materials, primarily wood, metal, and wax. The presence of a reflective, multi-material surface posed significant challenges during its acquisition.

The acquisition was conducted using a Nikon D3300 digital camera equipped with a 24 mm focal length lens. Image capture settings included an aperture of f/11, an exposure time of 1/11s, and ISO set to 400. The model was placed on a turntable to facilitate full rotational coverage during image capture. The image dataset was processed using 3DF Zephyr; due to the reflective surface, several alignment issues caused major holes. The resulting reconstruction exhibited significant deficiencies, particularly in the inner and outer regions of the petals. The initial RAWp model contained substantial holes and interpolated areas, especially on the lower portion of the bud.

Subsequent post-processing was carried out in Blender. Missing anatomical features, including the stem and internal pistils, were reconstructed. These elements, originally represented by elongated geometries, had been subject to distortion and interpolation during the processing phase.

Manual reconstruction was performed to design these incomplete features using regular geometries consistent with the original physical model. Afterwards, the model was reimported into 3DF Zephyr –maintaining its original local coordinates–and the texture was reprojected on the reconstructed model. Even after reprojection, the texture still showed holes, which were resolved through editing operations performed in Blender using the "clone" feature of the texture painting tool, as well as through manual refinement in Photoshop (https://www.adobe.com/it/products/photoshop.html).

The model of the *Linum* required substantial manual intervention to produce a version suitable for dissemination. The reflective metallic texture, along with the complex geometry composed of thin and easily distorted elements, were the main factors contributing to the gaps and inconsistencies in the model (Fig. 6).

Figure 6. From left, the RAWp, the DCHO and the Refinement Trace Map of CS4.

The Aldrovandi bust (CS5), part of the permanent collection of Palazzo Poggi Museum, is a medium-sized object in white marble whose surface is light and pretty reflective due to material properties. The bust is displayed on a pedestal and brightened by several spotlights.

Given the size and the difficulty of acquisition due to its height, it was decided to acquire the bust using an Artec Eva structured light projection scanner. The object was acquired at 8 FPS, with texture recorded every 3rd frame; then, the obtained scans were processed with Artec Studio 15 software.

Although the marble surface was visually opaque, the material, affected by the spotlights, prevented the correct acquisition of some areas, such as part of the base of the torso and the top of the head, where gaps were created; the remaining holes, however, were located at the undercuts of the robe. Therefore, the resulting RAWp model was mostly complete due to the Artec Eva scanner's automatic sensitivity adjustment, which allows even objects whose surface is reflective or translucent to be easily captured. Given the regularity and simplicity of the gaps, these were closed semi-automatically within the Artec software (Fig. 7).

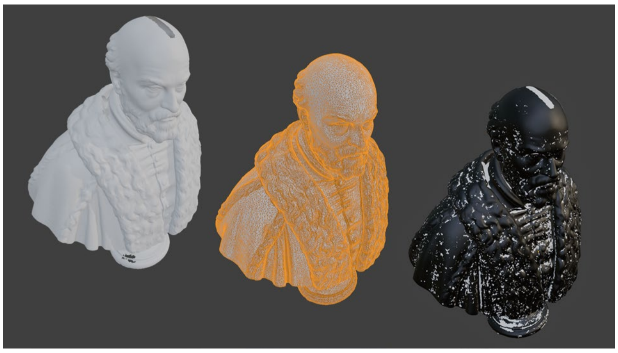

Figure 7. From left, the RAWp, the DCHO and the Refinement Trace Map of CS5.

3.3 TRACE Web3D tool

To provide a way to interactively visualise, compare and highlight specific areas involved in the described processes, a small online Web3D tool, called TRACE (https://aton.ispc.cnr.it/a/trace/), was developed.

The tool is based on the ATON framework developed by ISPC-CNR (Fanini et al. 2021): an open-source project that offers several features, including cross-device presentation (from mobile devices, workstation, up to XR devices); physically-based rendering (PBR); various XR presentation modes; semantic annotations and much more.

The open-source Web3D TRACE tool (https://github.com/phoenixbf/trace) was developed using the "plug&play" architecture of the ATON framework: this model allows to quickly craft javascript web-apps with custom logic and user interface, while accessing core ATON features - for more details on the app architecture see (Fanini et al. 2021, section 3.3).

Similarly to previous projects, this approach allowed to accelerate the development and deployment of the interactive Web3D application. Specifically, it was designed to interactively present a given 3D item (a glTF 3D model with standard maps, like base color, normal map, roughness, metalness, etc.) and associated Refinement Trace Maps. TRACE can access such data managed on a centralised collection on cloud storage. This is possible thanks to specific ATON features related to 3D collections, already providing routines to access items (3D models) from online locations and use them for presentation or scene composition - for more details, see (Fanini et al. 2021, section 3.2).

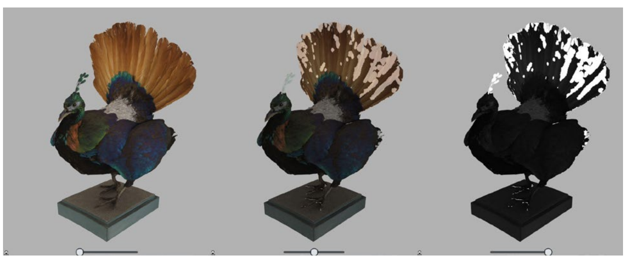

Figure 8. Testing the TRACE tool on the 3D model of the taxidermied Himalayan Monal.

Custom functionalities operating on item materials allow the control of the visualisation and inspection of color-coded information, using a slider (Fig. 8). Specifically, a custom fragment shader was used to control the base color of the 3D model and associated Refinement Trace Maps (see github repository of the TRACE tool). The designed routines also allow handling 3D items with multiple atlases associated and multiple Refinement Trace Maps.

4. Discussion and conclusion

Building upon an established pipeline, this article aims to address two specific research questions (RQ1 and RQ2) based on the results obtained and the taxonomy (Fig. 1) developed within the project framework.

In response to **RQ1**, the proposed clustering (Table 1) serves to identify the main problematic conditions encountered during the acquisition campaigns, offering concrete implementation examples that illustrate these challenges. At the same time, the presentation of the case studies is closely tied to the hole-filling process, which often requires extensive manual intervention, as demonstrated in the case studies presented. Regardless of whether the hole- and gap-filling process is carried out algorithmically, through AI inference, or manually, thorough tracking and documentation remain essential in the research ecosystem. Within this context, the present contribution does not advocate for manual over automated interventions. Rather, it introduces one possible approach capable of retrospectively tracking modifications in a manner that remains as neutral as possible with respect to the software used.

Addressing **RQ2**, the proposed methodological extension involves the generation of Refinement Trace Maps, which are color-coded maps derived from a semi-automated process, intended for real-time visualisation and comparison using a 3D web-based tool. While it is relatively straightforward to define a standard parametric threshold for detecting changes between different mesh versions, it remains challenging to establish a consistent value for the selective identification of hole types, as this may vary depending on the initial mesh's level of accuracy. At present, the generated Refinement Trace Maps are limited to offering a binary value, distinguishing between the presence (0) and absence (1) of data. Looking forward, the goal is to develop more nuanced color-coded maps, where each color corresponds to a semantic value indicating the nature and cause of the integration—such as areas of low confidence, fully reconstructed sections, or minor manual adjustments. These color-coded outputs are envisioned as structured and machine-readable information layers.

This approach opens up new possibilities for auditing and reviewing digital restoration workflows, helping institutions and professionals maintain a higher degree of accountability and control over the digital lifecycle of cultural heritage assets. It also supports better-informed decisions regarding the future use and sharing of data, including scholarly analysis, virtual exhibitions, and long-term digital preservation.

In conclusion, the objective of this contribution, along with its prospective developments, is to further enhance the transparency of a well-established methodology and to improve 3D object reconstructions by enriching them with greater informational depth, integrating metadata that communicates data certainty and quality alongside geometric and textural attributes.

Acknowledgements

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/):
Luisa Ammirati: Resources, Writing – original draft, Writing – review & editing
Alice Bordignon: Project Administration, Conceptualization, Methodology, Writing – original draft, Writing – review & editing
Federica Collina: Methodology, Writing – original draft, Writing – review & editing
Francesca Fabbri: Methodology, Investigation, Visualisation, Writing – original draft, Writing – review & editing
Bruno Fanini: Data curation, Software, Writing – original draft
Daniele Ferdani: Methodology, Writing – review & editing
Maria Felicia Rega: Methodology, Investigation, Writing – original draft, Writing – review & editing
Mattia Sullini: Supervisor, Conceptualization, Methodology, Writing – review & editing

References

Aldeeb, N.H., Hellwich, O., 2017. Detection and Classification of Holes in Point Clouds, in: Proceedings of the 12th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications. Presented at the International Conference on Computer Vision Theory and Applications, SCITEPRESS - Science and Technology Publications, Porto, Portugal, pp. 321–330. https://doi.org/10.5220/0006296503210330

Apollonio, Fabrizio Ivan, Fantini, F., Garagnani, S., Gaiani, M., 2021. A Photogrammetry-Based Workflow for the Accurate 3D Construction and Visualization of Museums Assets. Remote Sensing 13, 486. https://doi.org/10.3390/rs13030486

Apollonio, Fabrizio I., Fallavollita, F., Foschi, R., 2021a. The Critical Digital Model for the Study of Unbuilt Architecture, in: Niebling, F., Münster, S., Messemer, H. (Eds.), Communications in Computer and Information Science. Springer International Publishing, Cham, pp. 3–24. https://doi.org/10.1007/978-3-030-93186-5_1

Apollonio, F.I., Giovannini, E.C., 2015. A paradata documentation methodology for the Uncertainty Visualization in digital reconstruction of CH artifacts. SCIRES-IT - SCIentific RESearch and Information Technology 5, 1–24. https://doi.org/10.2423/i22394303v5n1p1

Attene, M., Campen, M., Kobbelt, L., 2013. Polygon mesh repairing: An application perspective. ACM Comput. Surv. 45, 1–33. https://doi.org/10.1145/2431211.2431214

Balzani, R., Barzaghi, S., Bitelli, G., Bonifazi, F., Bordignon, A., Cipriani, L., Colitti, S., Collina, F., Daquino, M., Fabbri, F., Fanini, B., Fantini, F., Ferdani, D., Fiorini, G., Formia, E., Forte, A., Giacomini, F., Girelli, V.A., Gualandi, B., Heibi, I., Iannucci, A., Manganelli Del Fà, R., Massari, A., Moretti, A., Peroni, S., Pescarin, S., Renda, G., Ronchi, D., Sullini, M., Tini, M.A., Tomasi, F., Travaglini, L., Vittuari, L., 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, e00309. https://doi.org/10.1016/j.daach.2023.e00309

Barzaghi, S., Bordignon, A., Gualandi, B., Heibi, I., Massari, A., Moretti, A., Peroni, S., Renda, G., 2024. A Proposal for a FAIR Management of 3D Data in Cultural Heritage: The Aldrovandi Digital Twin Case. DI 6, 1190–1221. https://doi.org/10.3724/2096-7004.di.2024.0061

Barzaghi, S., Bordignon, A., Collina, F., Fabbri, F., Fanini, B., Gualandi, B., Heibi, I., Mariniello, N., Massari, A., Moretti, A., Peroni, S., Pescarin, S., Rega, M.F., Renda, G., Sullini, M., 2024a. A reproducible workflow for the creation of digital twins in the cultural heritage domain. https://doi.org/10.5281/zenodo.14170838

Barzaghi, S., Bordignon, A., Gualandi, B., Peroni, S., 2024b. Thinking Outside the Black Box: Insights from a Digital Exhibition in the Humanities. https://doi.org/10.48550/ARXIV.2402.12000

Bordignon, A., Collina, F., Fabbri, F., Ferdani, D., Rega, M.F., Sullini, M., 2025. Balancing Photorealism and Transparency: A FAIR Approach to 3D Digitisation for Cultural Heritage documentation and dissemination. https://doi.org/10.5281/zenodo.14638115

Bordignon, A., Collina, F., Fabbri, F., Ferdani, D., Rega, M.F., Sullini, M., 2025a. Visual documentation of interpretative processes with vertex colour maps in 3D Modelling for Cultural Heritage. https://doi.org/10.5281/ZENODO.15210329

Börjesson, L., Sköld, O., Huvila, I., 2020. Paradata in Documentation Standards and Recommendations for Digital Archaeological Visualisations. Digital Culture & Society 6, 191–220. https://doi.org/10.14361/dcs-2020-0210

Brunton, A., Wuhrer, S., Chang, S., Bose, P., Demaine, E., 2010. FILLING HOLES IN TRIANGULAR MESHES USING DIGITAL IMAGES BY CURVE UNFOLDING. Int. J. Shape Model. 16, 151–171. https://doi.org/10.1142/S0218654310001328

Campen, M., Attene, M., Kobbelt, L., 2012. A Practical Guide to Polygon Mesh Repairing. The Eurographics Association. https://doi.org/10.2312/CONF/EG2012/TUTORIALS/T4

Centin, M., Signoroni, A., 2018. Advancing mesh completion for digital modeling and manufacturing. Computer Aided Geometric Design 62, 73–90. https://doi.org/10.1016/j.cagd.2018.03.016

Ch'ng, E., Cai, S., Zhang, T.E., Leow, F.-T., 2019. Crowdsourcing 3D cultural heritage: best practice for mass photogrammetry. Journal of Cultural Heritage Management and Sustainable Development 9, 24–42. https://doi.org/10.1108/JCHMSD-03-2018-0018

Davis, J., Marschner, S.R., Garr, M., Levoy, M., 2002. Filling holes in complex surfaces using volumetric diffusion, in: Proceedings. First International Symposium on 3D Data Processing Visualization and Transmission. Presented at the First International Symposium on 3D Data Processing Visualization and Transmission, IEEE Comput. Soc, Padova, Italy, pp. 428–861. https://doi.org/10.1109/TDPVT.2002.1024098

Fanini, B., Ferdani, D., Demetrescu, E., Berto, S., d'Annibale, E., 2021. ATON: An Open-Source Framework for Creating Immersive, Collaborative and Liquid Web-Apps for Cultural Heritage. Applied Sciences 11, 11062. https://doi.org/10.3390/app112211062

Hermoza, R., Sipiran, I., 2018. 3D Reconstruction of Incomplete Archaeological Objects Using a Generative Adversarial Network. https://doi.org/10.48550/arXiv.1711.06363

Hernández-Bautista, M., Melero, F.J., 2024. SR-CurvANN: Advancing 3D Surface Reconstruction through Curvature-Aware Neural Networks. https://doi.org/10.48550/arXiv.2407.17896

Huvila, I., 2022. Improving the usefulness of research data with better paradata. Open Information Science 6, 28–48. https://doi.org/10.1515/opis-2022-0129

ICOMOS, 2017. The Seville Principles—International Principles of Virtual Archaeology. Fórum Internacional de Arqueología Virtual, ICOMOS General Assembly. https://icomos.es/wp-content/uploads/2020/06/Seville-Principles-IN-ES-FR.pdf (accessed 16 April 2025).

Kazhdan, M., Bolitho, M., Hoppe, H., 2006. Poisson Surface Reconstruction. Eurographics Symposium on Geometry Processing. https://doi.org/10.2312/SGP/SGP06/061-070

Kuroczyński, P., Hauck, O., Dworak, D., 2016. 3D Models on Triple Paths - New Pathways for Documenting and Visualizing Virtual Reconstructions, in: Münster, S., Pfarr-Harfst, M., Kuroczyński, P., Ioannides, M. (Eds.), 3D Research Challenges in Cultural Heritage II, Lecture Notes in Computer Science. Springer International Publishing, Cham, pp. 149–172. https://doi.org/10.1007/978-3-319-47647-6_8

Lévy, B., Zhang, H. (Richard), 2010. Spectral mesh processing, in: ACM SIGGRAPH 2010 Courses. Presented at the SIGGRAPH '10: Special Interest Group on Computer Graphics and Interactive Techniques Conference, ACM, Los Angeles California, pp. 1–312. https://doi.org/10.1145/1837101.1837109

Luengo, P., Varela, G.B., Pla, S.C., Jover, A.C., Sosa, A.S., Codina, D.V., 2025. Toward a Trustable Digitisation of Built Heritage: The Role of Paradata, in: Ioannides, M., Baker, D., Agapiou, A., Siegkas, P. (Eds.), 3D Research Challenges in Cultural Heritage V: Paradata, Metadata and Data in Digitisation. Springer Nature Switzerland, Cham, pp. 139–150. https://doi.org/10.1007/978-3-031-78590-0_12

Nicolae, C., Nocerino, E., Menna, F., Remondino, F., 2014. Photogrammetry applied to Problematic artefacts. Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci. XL–5, 451–456. https://doi.org/10.5194/isprsarchives-XL-5-451-2014

Pèrez, E., Salamanca, S., Merchan, P., Adan, A., 2021. Repairing 3D Models Obtained From Range Sensors. IEEE Access 9, 43474–43493. https://doi.org/10.1109/ACCESS.2021.3061525

Podolak, J., Rusinkiewicz, S., 2005. Atomic Volumes for Mesh Completion, in: Proceedings of the 3rd Eurographics Symposium on Geometry Processing (SGP'05). Vienna, Austria. https://dl.acm.org/doi/10.5555/1281920.1281926

Sabbah, K., Ebied, H., Shedeed, H.A., Salem, A.-B.M., 2013. Hole-Filling Techniques for 3D-Objects: A Comparative Study. Presented at the Conference: 4th International Conference on Information and Communication systems (ICICS 2013).

Santos Junior, J.D.O., Vrubel, A., Bellon, O.R.P., Silva, L., 2012. 3D reconstruction of cultural heritages: Challenges and advances on precise mesh integration. Computer Vision and Image Understanding 116, 1195–1207. https://doi.org/10.1016/j.cviu.2012.08.005

Van Nguyen, S., Le, Son Thanh, Tran, M.K., Le, Sach Thanh, 2025. Filling the Holes on 3D Heritage Object Surface Based on Automatic Segmentation Algorithm. Expert Systems 42, e13749. https://doi.org/10.1111/exsy.13749

Wang, X., Xu, D., Gu, F., 2020. 3D Model Inpainting Based on 3D Deep Convolutional Generative Adversarial Network. IEEE Access 8, 170355–170363. https://doi.org/10.1109/ACCESS.2020.3024288