Digital Applications in Archaeology and Cultural Heritage 41 (2026) e00534

Contents lists available at ScienceDirect

Digital Applications in Archaeology and Cultural Heritage

journal homepage: www.elsevier.com/locate/daach

When meshes lie formalising tacit knowledge for flaw recognition in cultural heritage photogrammetry

Mattia Sullini [1]

Department of Architecture, University of Bologna, Viale del Risorgimento, 2, 40136, Bologna, BO, Italy

ARTICLE INFO

Keywords:
Photogrammetry
3D mesh processing
Flaw recognition
Tacit knowledge
Cultural heritage digitisation
Semantic annotation
PCA analysis
Paradata integration

ABSTRACT

This article presents a methodological framework for the semantic characterisation of surface flaws in photogrammetric 3D models, developed within Cultural Heritage digitisation workflows. Rather than relying on predefined classifications or automatic detection, the method captures expert corrections during mesh cleaning and formalises them as traceable, interpretable annotations. Through a retro-projection mechanism, these interventions are mapped back onto the pre-corrected model, producing a semantic segmentation of flaws grounded in situated expertise. Each flawed region is analysed using mesh-derived indicators, such as PCA-based descriptors, edge length, and photogrammetric confidence, to identify recurring patterns of deviation. This process supports the construction of a frequency-based taxonomy of common artefacts, structured across low-, mid-, and high-frequency classes and illustrated with real-world examples. The taxonomy does not aim to provide exhaustive coverage or prescriptive correction strategies but serves as a proof of concept for encoding tacit judgement into a reproducible descriptive system. The approach is applied to a single-operator dataset and is presented as a foundational step toward reproducible flaw annotation and structured paradata integration. While further validation across operators and use cases is required, the method already offers a means to bridge procedural opacity and semantic clarity, particularly in workflows where dissemination assets derive from documentation-oriented models. By aligning mesh geometry, operator intervention, and scalar descriptors, the framework establishes a basis for more transparent assessment, future comparative validation, and potential integration into semi-automated classification systems. It also contributes to the ongoing effort to embed semantic meaning and interpretive traceability directly within 3D assets.

1. Introduction

1.1. Digitisation of CH

The digitisation of tangible Cultural Heritage serves two concurrent purposes: documentation and dissemination. While the former aims to construct a dense and traceable representation of the object's physical configuration, the latter prioritises accessibility, visual legibility, and platform compatibility (see Figs. 7–9).

These two purposes are not mutually exclusive. Nonetheless, in many institutional workflows, models generated for documentation are later adapted for dissemination through processes rarely subject to the same level of scrutiny. Derived assets are often treated as secondary outputs, requiring no further validation. Paradoxically, this undermines their communicative power, precisely where effectiveness and clarity are most critical.

As a result, rigour tends to be neglected even when documentation is the primary intent, due to the lack of structured methods for controlling and communicating surface quality throughout the transformation chain. The present work stems from this observation and proposes a descriptive framework to make such flaws recognisable, classifiable, and ultimately accountable. In particular, it redefines the correction phase (Fig. 1) not as a mere cleaning step, but as the epistemic moment in which flaws are detected, delimited, and semantically qualified.

The London Charter (2009) explicitly frames digital representations as scholarly outputs, requiring methodological clarity and full documentation of the processes involved. Complementing this perspective, the Sevilla Principles (2017) stress the epistemological responsibility of virtual archaeology, highlighting the need for plausible and well-argued interpretation of the available evidence as paradata (Bentkowska-Kafel and Denard, 2016). Yet, as noted by Koller et al. (2010), current infrastructures and tools often fail to operationalise these principles,

E-mail address: mattia.sullini2@unibo.it.
[1] **Address:** Via Franco Bolognese 41, 40,129, Bologna.

https://doi.org/10.1016/j.daach.2026.e00534
Received 24 July 2025; Received in revised form 17 March 2026; Accepted 13 April 2026
Available online 15 April 2026

especially in the early stages of 3D data processing and mesh inspection.

These limitations are particularly evident in low-cost or semi-automated workflows, where the lack of control over intermediate steps can obscure the provenance and quality of the resulting models (Remondino et al., 2012). Yet, this control can be partially restored by explicitly documenting the interpretive operations performed during defect remediation, transforming them into traceable and analysable actions within a structured framework, an approach that the present work assumes as its central methodological commitment.

As a form of measurement, every digital model, hereafter asset, represents an approximation of the physical object, hereafter item, and is inevitably a simplification, since perfect equivalence is unattainable. Environmental and operational factors further affect the process beyond instrumental limitations, introducing discrete, localised, and often non-systematic errors, which become embedded in the asset, making every digital transposition inherently flawed.

Research has developed techniques to identify and correct such errors, often by analysing local anomaly neighbourhoods. While effective for simple geometries, these approaches often fall short with the complex forms typical of Cultural Heritage. In such cases, expert intervention becomes essential: cleaning is performed manually or semi-automatically (Botsch et al., 2010), based on interpretive judgement and comparison with the item. Final accuracy thus depends not only on tools, but on domain-specific expertise.

However, in the absence of shared terminology and formalised classification of mesh defects, expert corrective actions remain undocumented, their rationale implicit. This work addresses that gap by reframing the correction phase not as a neutral refinement step, but as a critical interpretive stage. Here, flaws are not only removed but actively recognised, semantically delimited, and annotated through embedded segmentation.

To support this intent, mesh-derived indicators, such as curvature proxies, edge length, and photogrammetric confidence, serve as heuristic supports, enabling experts to ground their interventions in verifiable geometric evidence. This process yields a traceable map of the reasoning applied, allowing the extraction of recurring flaw patterns and supporting the construction of

a preliminary taxonomy. While not prescriptive, this taxonomy demonstrates how semantic segmentation can emerge from situated expertise and be anchored to measurable surface behaviour, laying the groundwork for reproducible annotation, interpretive transparency, and future validation strategies.

1.2. *From raw to readable: balancing documentation and dissemination*

The distinction between documentation and dissemination as a differentiation of scope is historically grounded in objective technological constraints, as to publish 3D content online or in interactive environments, models had to be aggressively simplified to ensure compatibility and responsiveness. While this rationale still underpins many workflows, advances in rendering pipelines, data transmission protocols, and Web3D platforms have significantly relaxed such limitations.

Technologies such as public 3D repositories based on WebGL/Web3D, metadata-rich formats like glTF (Sullini, 2026a)(, browser-accessible content, and remote streaming solutions are opening new scenarios not only for museum communication but also for collaborative research. Within this framework, maintaining a strict separation between scientific and disseminative assets becomes a limitation, hampering both the quality of communication and the opportunity to accelerate and democratise research.

Furthermore, even when dissemination is the primary objective, digitisation remains, de facto, a form of documentation. Digital models inevitably derive from measurement processes and therefore carry implicit claims about the physical objects they represent, but the epistemic qualification of the transformations that produced them becomes essential whenever such assets are expected to function as representations, that is, when they are cited, compared, reused, or evaluated against the objects they purport to depict.

This critique does not imply that dissemination-oriented campaigns are inherently flawed, nor that all digital models must adhere to the same epistemic standards as documentation-oriented assets. Indeed, digitisation processes vary widely in purpose, scale, and methodological ambition, and consequently models intended for storytelling, education, immersive interaction, or public engagement may legitimately prioritise perceptual and experiential qualities. In such contexts, geometric simplification is not only acceptable but often necessary, particularly when dissemination must occur through browser-based platforms, remote streaming infrastructures, or real-time Web3D environments, which impose strict limits in terms of polygon count, texture resolution, and file size. These constraints, however, do not remove the need for such transformations to remain methodologically controlled and explicitly documented.

Nonetheless, today dissemination-oriented assets are still routinely derived from models obtained directly from acquisition, which are generally assumed to be suitable for documentation because they preserve the original geometric information. However, the derivation process typically involves not only geometric optimisation and simplification with information loss, but also corrective operations and gap-filling interventions that disregard semantic and morphological integrity. Consequently, it becomes unclear which features of the original model have been preserved, discarded, distorted or simplified, and why. This opacity does not stem from geometric simplification itself, but primarily from the absence of traceable segmentation of correction interventions during mesh processing. Interpretive decisions, such as which flaws to remove and how, are rarely recorded, leaving

Fig. 1. The asset “Lophius Piscatorius” before (left) and, after (right) manual corrections. The corrected model is here shown untextured to show the complexity of the model's geometric features. Screencapture from Zbrush.

downstream users with visually convincing but epistemologically ungrounded assets.

On closer inspection, the issue of making the digitisation process epistemically transparent cannot be resolved simply by distinguishing between a full-resolution model intended for documentation and a dissemination model derived from it through simplification. Objectivity, in this sense, does not necessarily imply epistemic qualification.

Rather, it is methodological traceability that might provide a framework through which the transformations that digital assets undergo during refinement can be documented, interpreted, and evaluated, thereby strengthening not only scholarly transparency but also communicative reliability and supporting informed reuse and critical assessment across diverse contexts of application. In this sense, dissemination itself can achieve the same epistemic transparency required to sustain the documentary function of the full-resolution asset, notwithstanding the simplification or compression of its geometric information.

1.3. Epistemic qualification of digitisation workflows

In this sense, a shared epistemic qualification of both documentation and dissemination assets fully aligns with the stated purposes of the London Charter (2009), which "is concerned with the research AND dissemination of cultural heritage across academic, educational, curatorial and commercial domains" [capitalisation added by the author]. The relevance of this issue also extends to European data strategies, particularly the initiative to establish the European Collaborative Cloud for Cultural Heritage (ECCCH) (European Commission, 2022a), which aims to enable GLAM professionals and Cultural Heritage institutions to leverage the opportunities arising from the dematerialisation and cross-institutional circulation of digitised heritage. These developments are especially significant for tangible 3D heritage, whose digital surrogates are increasingly shared across institutional and national boundaries for research collaboration, data exchange, and dissemination.

Yet a project of this scale is not merely a matter of infrastructure, as dharing 3D digitisation results across member states presupposes interoperability, and interoperability in turn requires standardised formats, metrics, and methodologies for planning, executing, monitoring, and reporting digitisation campaigns. While metadata frameworks and repositories already exist, no shared standards currently govern the description and monitoring of the transformations that 3D digital assets undergo during processing, as documented by the VIGIE 2020/654 survey (European Commission, 2022b). Standardisation, however, cannot precede description; before formats and benchmarks can be meaningfully defined, the transformations that raw data undergo during processing, and the interpretive decisions embedded in them, must be made explicit and communicable.

This requirement is not abstract. In academic contexts, an unqualified model cannot be cited as evidence; in educational contexts, it cannot be trusted as a faithful representation; in curatorial contexts, it cannot support condition assessment or conservation decisions; in commercial contexts, it cannot underpin reproducible quality guarantees. Across all these domains, the absence of a qualified refinement process undermines both the credibility of the asset and the accountability of whoever produced it.

However, in the absence of shared terminology and a formalised classification of mesh defects, expert corrective actions remain undocumented and their rationale implicit. In such conditions, the transformations that digital assets undergo during refinement cannot be systematically analysed, compared, or evaluated. The transparency of these transformations therefore becomes a prerequisite for both scientific accountability and cross-institutional interoperability.

In this perspective, the correction phase cannot be understood as a neutral cleaning step. Rather, it constitutes the point at which interpretive decisions become accountable, communicable, and potentially interoperable across institutional contexts. Making these interpretive operations explicit contributes to the broader objective of establishing a digital continuum in which the actions and critical decisions involved in the digitisation process are documented within the models themselves (European Commission, 2022a).

Within this broader institutional context, the framework proposed in this work is intended primarily for operators involved in the production and refinement of 3D assets, while also providing researchers, curators, and heritage institutions with a means to assess and communicate the reliability of the resulting models, enabling more transparent reuse of digitised cultural heritage across the academic, educational, curatorial, and commercial domains identified by the London Charter.

To address this need, it introduces a mechanism through which interpretive operations performed during refinement can be explicitly recorded and later related back to the pre-correction state of the model, by embedding semantic segmentation, understood here as the manual tagging of surface regions according to their meaning or function, directly within the correction phase, marking the areas affected by expert intervention through polygroups before any optimisation occurs. These annotations can then be projected back onto the pre-correction mesh, making it possible to localise and analyse the corresponding flaws on the original geometry. In this way, both the pre-correction and post-correction meshes become semantically segmented, preserving a traceable correspondence between the regions affected by expert intervention and the underlying acquisition flaws. Resolving this dichotomy, between the procedural opacity of current workflows and the epistemic transparency that structured annotation can provide, is the core aim of the framework proposed here.

1.4. Structure of the paper

The paper is organised to progressively establish the conceptual and operational context of the proposed framework. The preceding sections have outlined the epistemic, methodological, and institutional background of CH digitisation workflows, while the following sections present the methodological structure and its technical implementation.

- Section 2**:** Situates the proposal within the current literature, first highlighting its synergy with the RAW–DCHO methodological framework and then reviewing the technical literature on acquisition defects and the analytical tools used to characterise mesh geometry, thereby defining the relevant technical domain.
- Section 3: Describes the proposed method. It first presents the procedure used to segment edited meshes and project the resulting segments back onto the pre-intervention meshes, then examines the analytical characterisation of mesh features associated with acquisition defects, and finally defines the parameterisation of the segmentation so that the semantic layer can be preserved and accessed outside the 3D authoring environment.
- Section 4: Presents the taxonomy of flaw types derived from expert feedback, as a demonstrator of the framework and a proof of concept for a formalised descriptive language.
- Section 5: Discusses the method's implications and limitations, along with parallel lines of research.
- Section 6: Concludes with potential devolpment of the method in its current form aimed to applications.

2. Background and related works

2.1. Towards a unified master asset for documentation and dissemination

To overcome the aforementioned dichotomy between documentation and dissemination, raw assets should be cleaned and optimised through a controlled process that preserves both geometric consistency and interpretive transparency. In such a framework, the master asset should no longer coincide with the raw acquisition model but rather take the form of a cleaned and optimised version that combines technical

robustness with visual usability, while also preserving an embedded record of the interpretive decisions made during the correction process. In Barzaghi et al. (2025), this principle was already addressed through a structured revision pipeline based on four stages (Fig. 2).

- *RAW*: refers to all unprocessed data directly obtained from the instruments.
- *RAWp:* (RAW processed) refers to the data transformed into a textured mesh and then globally cleaned through automated procedures to remove well-known and clearly characterised defects.
- *DCHO*: (Digital Cultural Heritage Object) is the model corrected for errors and defects and supplemented in its gaps through the application of inferential and critical corrective interventions conducted by operators. In the original framework the simplification might be contextual to correction.
- *DCHOo*: (DCHO optimised) is the final optimised, simplified and metadata-enriched model, prepared for publication in glTF format.

In the present proposal, this structure is retained with one critical variation, where DCHO is required to preserve the full data content of RAWp, with simplification strictly deferred to the DCHOo phase and performed only as required by the target application. This enforces a rigorous separation between cleaning and optimisation, making DCHO the definitive Master asset it was originally intended to be, fully suitable for documentation, while dissemination-oriented assets can be derived from it through application-dependent optimisation and simplification (DCHOo). This condition also enables a clearer separation between the correction phase and the subsequent optimisation and simplification that transforms DCHO into DCHOo, thereby introducing an additional epistemic distinction within the workflow alongside the one that differentiates RAWp and DCHO along the axis of agency, automated in the former and interpretive in the latter.

Within this structure, however, a key bottleneck persists in the transition from RAWp to DCHO, namely, the need to systematically identify, classify, and eliminate the aforementioned acquisition-induced defects, which are often addressed heuristically and without formal traceability. By contrast, the subsequent simplification stage relates to the distinct problem of managing the discretisation error introduced by geometric reduction and therefore falls outside the scope of the present discussion.

Previous developments have partially addressed this gap in the identification and correction of defects emerging from ambiguous photogrammetric data during surface reconstruction. In Bordignon et al. (2026), the correction of mesh voids was explicitly annotated to preserve interpretive traceability, while Ammirati et al. (2025a) introduced a projection-based method to detect such gaps by comparing RAWp and DCHO geometries. Furthermore, the classification of corrective interventions and the annotation of their epistemic impact on mesh geometry have been explored in previous work (Sullini, 2026b)(), which introduced a reliability-based segmentation of mesh elements grounded in the type of corrections applied. The present article extends this approach by generalising the framework to encompass corrective processes across the RAWp–DCHO transition and by formalising their integration within the established RAW → DCHOo pipeline.

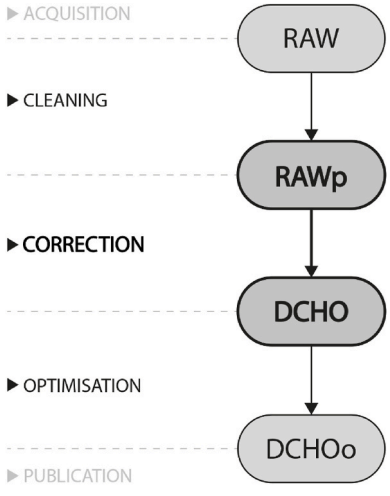

Fig. 2. The adopted workflow for mesh generation and optimisation in CH digitisation pipelines. The main steps follow Barzaghi et al. (2025), while the intermediate phases have been re-elaborated.

Building on these foundations, the present article generalises the approach to encompass all corrective operations. It introduces a unified system that captures tacit operator knowledge during the editing phase, annotates each intervention at the DCHO level, and projects these annotations back onto the RAWp mesh, where the original flaws are segmented, localised, and analysed using mesh-based indicators.

The resulting framework enables a reproducible classification of flawed regions, not by pre-correction inspection, but through interpretive feedback grounded in actual expert practice. These correspondences support a double semantic stratification: RAWp is structured by defect type, and DCHO by correction type, allowing both assets to be semantically and metrically aligned. This alignment forms the basis for the proposed taxonomy, which organises flaw categories according to morphological traits and mesh-derived metrics, providing a formal language for describing imperfections across diverse CH digitisation scenarios.

This multi-stage pipeline progressively introduces semantic control: from automatic filtering in RAWp to expert-guided correction and annotation in DCHO, where flaws are manually addressed and, ideally, segmented in a traceable way.

2.2. *Photogrammetry-based surface reconstruction*

The assets entering this pipeline are most commonly produced through photogrammetry, widely used for the 3D digitisation of Cultural Heritage (CH) tangible assets, as it offeris a flexible and cost-effective alternative to active scanning technologies (Moyano et al., 2023). When acquisition is performed under controlled conditions, the method yields metrically accurate and visually rich 3D models (Menna et al., 2016; Nocerino et al., 2014);). However, such conditions are rarely achievable in CH practice, particularly during large-scale digitisation campaigns of temporary museum collections, as noted by Balzani et al. (2024) and reflected in the forthcoming Guidelines for the 3D digitisation of Historical, Artistic and Museum Heritage (ParteciPa, 2025).

Even when acquisition constraints are acknowledged, factors such as access limitations, suboptimal lighting, surface reflectivity, or time restrictions often result in irregular image sampling, occlusions, and misalignments. This is particularly common in digitisation campaigns conducted by small and medium-sized museums and other GLAM institutions, where the widespread availability of low-cost equipment (Kirchhöfer et al., 2011; Leksell, 2018) and user-friendly pipelines (Remondino et al., 2012; Rahaman and Champion, 2019) has significantly lowered the barriers to entry. This tendency is especially evident in photogrammetric pipelines (Apollonio et al., 2021; Farella et al., 2022), which can be implemented with relatively inexpensive acquisition setups and open-source reconstruction software, and for which several adaptations of professional workflows have been proposed to support low-cost digitisation environments (Sebar et al., 2021). While this democratisation has expanded the adoption of 3D digitisation practices, it also reduces the degree of control that can be exercised over the aforementioned adverse factors, and photogrammetry, as an indirect reconstruction method, is particularly vulnerable to their consequences.

Errors originating already in the point cloud, where incorrect correspondences or insufficient sampling may survive standard filtering, are often propagated and sometimes amplified during mesh

reconstruction, introducing discrete and non-systematic errors that manifest as local morphological anomalies (Berger et al., 2016; Frost et al., 2023). Such artefacts may be perceptible to expert operators but remain difficult to detect algorithmically due to their variability and lack of formal characterisation.

Importantly, not all local anomalies are flaws: some correspond to legitimate object features, especially at medium or high spatial frequencies. By contrast, flaws are artefacts that both violate expected morphological continuity and cannot be attributed to actual geometry. This ambiguity makes the distinction between flaw and feature non-trivial, reinforcing the need for a semantic framework grounded in expert correction. Such a framework must support the annotation and classification of defects as they are detected and resolved, thus enabling informed, communicable, and reproducible assessments.

It should be noted that active sensing technologies, which directly measure point positions in space and often implement onboard rejection of anomalous readings (Beraldin et al., 2015; Pacheco et al., 2014), are inherently less prone to the hallucination phenomena described above, as they do not rely on indirect surface inference. However, their cost, operational constraints, and reduced flexibility in field conditions make photogrammetric pipelines the preferred solution in most CH contexts (Leksell, 2018; Wachowiak and Karas, 2009), and the present discussion is therefore confined to the specific challenges of indirect reconstruction.

2.3. Validation in 3D digitisation

Confining the discussion to indirect reconstruction has a direct methodological implication: 3D models produced through photogrammetry are, from a strictly metrological standpoint, locally uncertain. They may contain geometric features that are in fact morphological configurations with no physical correspondence on the actual object, arising not from noise in real measurements but from structural misinterpretation during surface estimation (Fig. 3) that, in the present work, are referred to as morphological hallucinations.

Given these premises, metrological precision, while desirable, is not sufficient. This intrinsic uncertainty indeed limits the effectiveness of purely numerical validation and highlights the need for interpretive frameworks that can operate within the photogrammetric pipeline itself and that are capable of identifying, classifying, and communicating the presence of surface anomalies in a structured and reproducible way.

In this light, the reliability of a photogrammetric model depends not only on the quality of its input data, but above all on the transparency of how its surface is examined, interpreted, and refined. Structured annotation of flaws conducted during expert-led correction thus becomes a critical validation layer: more robust than global metrics, more communicable than aggregate errors, and particularly well-suited to the specific challenges of CH digitisation.

Fig. 3. The asset “Lofoforo” presenting various Surface Reconstruction hallucinations, such as the apparent holes, which are indeed the result of the thinness of the feathers, i.e. the ambiguity of the position of the points on the front and back faces. Screencapture from Zbrush.

2.4. Connectivity

The identification and treatment of such anomalies, however, remains poorly systematised: although photogrammetry workflows are increasingly consolidated, literature still focuses on isolated issues such as hole closure (Pérez et al., 2016) or filtering of noise and topological errors (Botsch et al., 2010; Zhou et al., 2022). To the best of author's knowledge, a comprehensive framework to identify, classify, and manage the full range of flaws, especially in CH contexts with complex geometries and suboptimal acquisition, is lacking.

This gap reflects the limited availability of tools specifically designed for Cultural Heritage, as most software, hardware, and algorithms are adapted from adjacent fields, such as reverse engineering (Barazzetti, 2020) and entertainment industries (Cipriani and Fantini, 2017). While the latter contribute mainly to optimisation and visual enhancement, flaw correction often relies on methods developed for man-made hard-surface objects, where geometry is simpler and boundaries clearly defined, thus favouring automated correction pipelines, where connectivity issues are fixed algorithmically.

In facts, these meshes are typically repaired using tools that operate at the level of connectivity (Botsch et al., 2010; Sorgente et al., 2023), i.e., the structure linking vertices, edges, and polygons. Such methods are effective against high-frequency issues such as non-manifold geometry, outliers, open borders, degenerate polygons, but fail on discrete or large-scale flaws that are typical in CH assets with complex or critical surfaces, where defects may not break topological rules yet still require correction guided by semantic and morphological criteria.

2.5. Holes and algorithmic hallucinations

A relevant exception to the aforementioned considerations are holes, which are among the most studied flaws in 3D digitisation workflows and the only category to receive sustained attention in both technical and CH literature. Yet this focus remains fragmented, limited to cause classification (Nicolae et al., 2014) or procedural repairs (Pérez et al., 2016), rather than to a comprehensive analysis of their role, semantics, and treatment in CH contexts.

Holes typically arise from insufficient or uneven sampling caused by occlusion, reflectance, or surface inaccessibility, varying in shape and scale and compromising geometry, texture, or both. Their presence, however, is algorithm-dependent: while methods such as Poisson reconstruction do not generate holes in the topological sense, as they return closed or contiguous surfaces even over poorly sampled regions, this is accomplished by patching those areas with low-confidence and low-sampled interpolated geometry, effectively substituting the topological condition of a hole with a patch of geometrically unreliable surface. Nonetheless, these patches are often removed during post-processing, and the resulting absence of geometry produces the same topological condition as a hole, albeit with different semantic implications.

Algorithmic hallucinations, similarly, when not corrected through deformation while preserving mesh connectivity, are commonly removed once identified, thus recreating the same topological condition reported for patch removal. Yet their identification is not straightforward, as geometry remains continuous while exhibiting localised distortions, morphological, cumulative, or patterned, that alter form without generating missing data. Consequently, these artefacts often escape immediate detection, despite their impact on visual plausibility and metric reliability.

In practice, many such flaws are well known to experienced operators, who recognise them empirically in the course of established correction routines, reflecting what Nonaka (1994) described as tacit

knowledge. However, these insights rarely translate into formal classifications. A comparable attention to the causes of acquisition defects is also found in the *Guidelines for the 3D Digitisation of Historical, Artistic and Museum Heritage* (ParteciPa, 2025), although there the focus is primarily on their aetiology and on mitigation strategies during the acquisition phase. Issues such as insufficient sampling in high-curvature regions, surface bridging, photometric misinterpretation, or local shading artefacts are widespread, yet typically fall outside structured analysis frameworks or are absorbed into global quality metrics (Remondino et al., 2017).

This reveals a persistent disconnection between operational knowledge and theoretical formalisation. Without explicit characterisation, the legitimacy of correcting such flaws remains ambiguous, and their documentation is often neglected. By embedding both semantic segmentation and correlation with mesh analysis indicators within the correction process, the present framework aims to convert situated expertise into structured descriptions, enabling the emergence of a consistent, empirically grounded taxonomy.

2.6. Sampling and photogrammetric indicators

Indeed, hallucinations are not expressed solely at the morphological level, but also in the sampling structure of the reconstructed surface. Empirical inspection of problematic reconstructions shows that hallucinated morphologies are often preceded by unstable or inconsistent sampling patterns that already emerge in the point cloud. More specifically, clusters of spatially inconsistent samples may form because these samples are not recognised as outliers and therefore survive the filtering stage. During surface reconstruction these samples are nevertheless incorporated into the mesh, but must be geometrically connected to neighbouring regions supported by consistent observations. In practice, this process behaves similarly to the treatment of small gaps: the reconstruction algorithm interpolates across these inconsistencies, effectively resolving them as local anomalies in sampling density.

This observation suggests that hallucinated geometries may have a detectable signature not only in morphology but also in the sampling structure of the mesh. Consequently, it becomes relevant to examine which existing indicators may function as diagnostic cues for such algorithmic interventions.

Among the available cues concerning density, average *Edge Length* is an effective proxy for acquisition sampling, already used for the characterisation of digitised assets in the GLAM sector (Apollonio et al., 2021). While smooth regions require fewer samples to represent the underlying geometry and are therefore often reconstructed with lower mesh densities, areas with high curvature or complex features demand denser and more structured sampling. When the local density is insufficient to resolve features at a given spatial frequency, reconstruction fails to capture them accurately. In fact, a feature of scale d must be sampled with spacing smaller than d/2 in order to be reconstructed as geometry rather than noise, as stated by the Nyquist–Shannon theorem. This condition is often unmet in real-world acquisitions, especially in occluded or critical regions.

In light of the considerations regarding sampling density, it emerges that algorithmic interventions leave detectable traces on the reconstructed surface when they resolve the latent form implied by the acquired samples. In particular, hallucinations tend to arise in underdetermined or unstable sampling domains, where the available observations provide insufficient support to uniquely constrain the reconstructed geometry. In photogrammetric pipelines, this condition is expressed by an indicator that reflects the robustness of the relationship between the available local observations and the reconstruction solution derived from them: Confidence.

Technically, *Confidence* refers to a scalar field assigned to each vertex or face, which reflects the stability of the reconstruction process rather than a geometric property, and is typically estimated from factors such as image coverage, angular redundancy, and depth consistency. Although confidence is derived from internal reconstruction parameters, it does not directly measure the geometric accuracy of the resulting model. Hence, while low-confidence areas often correspond to occlusions, poor sampling, or suboptimal viewing angles, conditions that also lead to frequent reconstruction flaws, high-confidence regions, despite being well supported by the dataset, do not alone guarantee accuracy, as systematic errors such as poor photometric quality, repeated textures, or alignment faults can still produce flaws in geometrically stable areas.

2.7. Morphometric indicators

Beyond their densitometric signature, hallucinated morphologies are often perceived by operators as autonomous entities; in practice, however, they can be interpreted as polymorphic expressions of a locally simpler underlying structure, resulting from algorithmic interventions that reconnect clusters of inconsistent samples displaced with respect to the consistent ones. In addition to the densitometric characteristics discussed above, it is relevant to observe that these phenomena are also accompanied by morphometric anomalies. Since reconstruction algorithms tend to treat such stitches as local gaps, the variation of curvature within the intervention generally assumes a monotonic behaviour.

Yet, while the locally observable morphological pattern may appear relatively simple, its articulation in space is not, as the overall morphology of hallucinations is ultimately determined by the geometry of the underlying inconsistent sample cluster.

As a result, their expression may span multiple scales, from local irregularities and microgeometric distortions to larger disruptions of surface continuity. In mesh processing, this variation is commonly interpreted through the notion of frequency, a concept derived from signal processing and formalised via spectral analysis tools (Zhang et al., 2007) that decompose surface variation by scale, supporting operations such as smoothing, simplification, or feature-aware filtering, and enabling a multiscale reading of geometric behaviour.

In the present framework, frequency is used in a qualitative sense to articulate both the behaviour of the geometric indicators introduced in the following section and the morphology of defects. It also underpins the proposed flaw taxonomy, where defects are grouped by their dominant spatial frequency.

Although frequency-based analysis operates on the intrinsic 2D manifold of the mesh, the surface itself is realised in 3D space through a parametric immersion (Do Carmo, 2016), which induces on the parametric domain (u,v) a two-dimensional metric inherited from the Euclidean metric of the ambient 3D space, thereby allowing geometric measurements performed in the 2D parametrisation to remain valid for the surface in 3D. The local shape of the surface can thus be described by geometric quantities such as curvature, whose definitions originate in the smooth setting and must therefore be adapted to the discrete domain of polygonal meshes. In fact, curvature is formally defined within the framework of differential geometry, which assumes smooth and continuous surfaces and where it is a local property derived from surface normals and second-order derivatives. The curvature of a surface varies depending on the direction along which it is evaluated; hence, the most relevant quantities are the principal curvatures, that is, the maximum and minimum values at each point. Notably, their corresponding directions are orthogonal and provide an effective description of local surface behaviour.

But when a surface is represented as a polygonal mesh, differential tools are no longer directly applicable. A first remedy consists in local surface fitting, which enables the application of differential analysis on a fitted continuous surface such as quadrics (Cipriano et al., 2009) or osculating jets (Cazals and Pouget, 2005) within a fixed radius. An alternative uses advanced tools adapted to discrete settings, such as discrete shape operators and the cotangent Laplacian (Meyer et al., 2003), but while formally accurate, these are computationally demanding and not natively supported in CH digitisation software like MeshLab or CloudCompare.

For morphological purposes, however, proper curvature is often less relevant than curvature variation across a neighbourhood, which can be captured by compact descriptors. These descriptors, and especially their distribution across the surface, support the identification of salient geometric features at different spatial frequencies, so that by analysing how shape-related quantities vary at increasing scales, it becomes possible to distinguish low-frequency deformations from high-frequency details. Such multiscale morphometric descriptors are widely used in mesh analysis, feature extraction, remeshing, and feature-driven smoothing (Pauly et al., 2003).

A common approach to extract such descriptors is Principal Component Analysis (PCA) (Jolliffe and Cadima, 2016), which provides robust indicators of shape variation and is both well-established and computationally accessible. At each vertex, PCA computes the covariance matrix of a local neighbourhood, typically defined either by a fixed spatial radius or by a k-nearest neighbours (kNN) criterion. The eigenvectors of this matrix define the principal directions of local spatial variance, while the corresponding eigenvalues quantify the magnitude of variation along these directions. Together, they are correlated with the local geometric structure of the surface, but in regular regions and under dense sampling, the direction of minimal variance tends to align with the local normal, while the other two span the tangent plane and approximate the principal curvature directions. In the limit case of infinitesimal radius and infinitely dense sampling, this decomposition converges to the differential frame defined by the surface normal and the principal directions of curvature.

Moreover, PCA acts as a low-pass approximation of curvature-like behaviour near the kernel scale, so that surface features with higher spatial frequencies are effectively suppressed, making PCA sensitive only to variation within the kernel's resolution, as the resulting eigenstructure encodes the smoothed geometric behaviour of the surface at that scale. In a multiscale framework, this becomes an asset rather than a limitation, as it enables the extraction of curvature-like descriptors that reflect the latent morphology of the object across diverse levels of detail. Scalar indicators derived from PCA, such as planarity, sphericity, or anisotropy, provide robust, noise-tolerant and multiscale cues for morphological classification, and are therefore adopted for the multiscale analysis of the mesh, both to detect the local signatures of algorithmic stitching interventions and to describe their spatial configuration within the hallucinated morphology as a whole.

These analytical tools, however, acquire their diagnostic value only when situated within a structured interpretive practice, one in which the operator's corrective decisions provide the semantic grounding that indicators alone cannot supply.

3. Methodology

3.1. Epistemological aspects

In photogrammetric workflows, correction is often treated as a neutral post-processing step. However, in Cultural Heritage digitisation, where shape plausibility is context-dependent and reconstruction flaws may resemble legitimate features, correction becomes an interpretive act, as the decision to intervene implies a judgment about what should or should not be present. Each correction thus encodes an implicit classification: it isolates a flawed region, delimits its extent, and modifies or reconstructs it according to an internal plausibility model informed by domain knowledge, visual inspection, and contextual cues.

Rather than preceding correction, flaw recognition typically emerges through it. The operator may begin without a complete map of defects, progressively identifying anomalies as familiarity with the model increases. This process relies not only on geometry, but also on texture analysis, as chromatic deformations, background leakage, or projection inconsistencies often signal problematic regions even before geometric inspection. Crucially, this evaluation is intrinsically binary: a surface is either accepted or modified. At scale, these discrete decisions define a de facto segmentation of the model into reliable and flawed regions.

This non-intrusive configuration is essential. The framework does not alter operator behaviour but retrospectively captures decisions already enacted, ensuring that classification emerges from actual expert practice rather than from prescriptive modelling. It functions as an epistemic observer that, rather than prescribing correction, retroactively structures the interpretive decisions made by the operator, transforming them into a traceable semantic layer, since each intervention embeds a prior diagnosis of local plausibility, albeit non-formalised. Once segmented and projected back onto the unedited model, these corrections can be analytically interpreted through mesh-based indicators such as curvature, sampling density, and reconstruction confidence, which provide geometric support without constraining operator evaluation.

As segmentation becomes available, it becomes possible to investigate whether recurring patterns emerge in the combination of geometric and morphometric indicators associated with corrected regions. If such patterns prove consistent, they may serve as proxies for flaw recognition, allowing this crucial judgment, currently performed by the human operator, to be progressively automated.

Nonetheless, classification remains retrospective at the current stage, emerging from corrected outputs. This feedback configuration is not structurally required by the method but constitutes a transitional strategy for formalising tacit knowledge. Only once a robust taxonomy is established, flaws may be directly tagged at the point of recognition, allowing the system to evolve from a descriptive tool into a diagnostic and predictive framework.

3.2. Dataset

The empirical basis from which this formalisation process was initiated consists of two separate datasets, both derived from recent digitisation campaigns involving heterogeneous Cultural Heritage collections. The first dataset includes naturalistic specimens and facial casts used for anthropological and medical studies, originally displayed in a recent temporary exhibition (reference omitted for double-blind review), held between (dates omitted for blind review) at (venue omitted for blind review) in (location omitted for blind review). The second consists mainly of paleontological and geological items belonging to the paleontological collection of (reference and location omitted for blind review), some of which have been described in previous work (Ammirati et al., 2025b)().

Both datasets were selected as pilot cases due to their diversity in scale, material, morphology, and acquisition context. Combined, they include taxidermied animals, fossils, wax models, numismatic items, decorated surfaces, printed volumes, xylographic plates, and archaeological artefacts, offering a broad representation of tangible heritage typologies and associated digitisation challenges.

The models were generated through standard SfM/MVS pipelines and manually processed by several operators, including the author. However, although the preliminary observation of defects was conducted on the extended dataset, the characterisations reported in the present study were verified exclusively on the digitisations conducted by the author. This methodological choice was made because algorithmic interventions depend both on the design of the reconstruction algorithm and on its parametrisation, as well as on the specific acquisition setup. While empirical observation allows the identification of hallucinated morphologies recurring across different workflows, restricting the formal analysis to a single pipeline makes it possible to control external variables and therefore attribute the variability of the observed defects solely to morphometric and densitometric conditions.

The acquisition setup varied depending on the possibility of manipulating the physical object and controlling the ambient illumination. Under optimal conditions, a turntable placed inside a portable diffuse box retroilluminated through diffusers was employed. As for the photogrammetric equipment, a Sony A7 I camera equipped with a 28–70 mm zoom lens and a 36 × 24 mm full-frame sensor was mounted

on a tripod. Image acquisition was performed with a fixed focal length, which was adjusted according to the overall size of each object in order to maximise the use of the available pixel resolution within the image frame, resulting in variable GSD values. According to the average reflectance of the object, exposure time was optimised around the histogram peak starting from the smallest aperture allowed by the selected focal length (from f/22 to f/36), while ISO sensitivity was fixed at 200 to minimise noise. For each camera position, or rotation of the turntable relative to the fixed camera position, between one and four images were captured with manually varied focus distances in order to compensate for the limited depth of field.

Each image (6000 × 4000 pixels) was recorded in RAW format (14-bit ARW), processed in Lightroom, and exported as 16-bit TIFF files for processing in Metashape. The SfM network was computed from these images through iterative error-driven refinement of the bundle adjustment during image alignment, while Surface Reconstruction was performed directly from the depth maps estimated from the aligned images at maximum quality, with interpolation enabled. The resulting 3D model was then exported for mesh processing in ZBrush and subsequently in Modo. Textures were also acquired under a controlled colour-correction pipeline; however, since the present study primarily concerns geometry, these aspects are not discussed in detail.

During this processing, particularly in the cleaning and optimisation phases, recurrent geometric issues were encountered, prompting the identification of flaw patterns that later informed the structure of the proposed taxonomy. A subset of assets managed directly by the author was retrospectively examined using this emerging interpretive framework. These models were not selected to "test" the method, but rather became part of the iterative process through which it was conceived. Other assets, although processed with different workflows and objectives, contributed to defining the broader operational context in which flaw types and frequency patterns were recognised.

The taxonomy, therefore, was not applied across the dataset as a finished diagnostic tool, but rather derived inductively from selected segments of practice. It remains non-validated in inter-operator terms and has not yet been used to annotate the dataset in a systematic way. Nonetheless, the corpus provided a grounded empirical substrate for articulating preliminary morphological classes and diagnostic frequencies. The classification should thus be interpreted as a provisional structuring of expert insight, not as the result of a completed validation protocol.

This initial formalisation, by explicitly encoding practitioner experience into a reproducible structure, now enables the possibility of systematic verification. A future intra-operator application, built upon the present conceptual groundwork, may serve to consolidate the taxonomy and flaw characterisation, while also potentially transforming them through extended and comparative implementation. Such a follow-up study would represent a necessary step towards assessing internal coherence, operator consistency, and diagnostic utility within controlled use cases.

3.3. Automatic filtering in RAWp

The present method integrates with the RAW - DCHOo pipeline, specifically at the transition from RAWp to DCHO, maintaining its structure and intent. The crucial distinction, preserved here, is that DCHO includes all critical, authorial corrections addressing complex flaws that cannot be resolved through automated filtering, of which RAWp is the final output.

Indeed, RAWp is not the raw mesh, but a pre-processed version derived through automated procedures aimed at removing well-characterised connectivity defects such as outliers, noise, non-manifold elements, degenerate polygons, self-intersections, anisotropies, or inconsistent normal orientation, as extensively discussed in (Sorgente et al., 2023). These issues are typically addressed through deterministic or AI-based approaches, many of which are embedded in standard mesh-cleaning tools commonly used in digitisation workflows.

3.4. Tracing interventions in DCHO

Having clarified the role of RAWp and anticipated that the operations leading to DCHO are critical and authorial in nature, it follows that correction is not a neutral refinement but a privileged moment of interpretive decision. The operator identifies surface anomalies and applies remedial actions, such as sculpting, re-modelling, deletion and reconstruction. These interventions have a primarily operational function, aimed at improving mesh quality, but at the same time, they carry epistemic value, as by establishing what requires correction, they implicitly define what is considered defective and help delineate the boundaries for classification.

To support this process, the method implements a variation of the active tracing mechanism embedded in the DCHO phase outlined in Bordignon et al. (2026), with the difference that the tracing deploys polygon groups instead of vertex colour maps (Fig. 4). As corrections are performed, the mesh is segmented using a standardised set of semantic categories, which describe the nature of the intervention applied rather than the specific tool or technique.

- *Hole-filling*: Address regions where photogrammetric interpolation is unreliable, typically contiguous zones with low confidence and anomalous edge length, but are surrounded by consistent geometry. These areas are manually removed to expose the underlying gap, which is then manually filled by inferring the missing geometry by adjacent topology and comparison with the raw photographic dataset.
- *Sculpting or morphological adjustment:* Corrects surface anomalies by altering the morphology without altering mesh connectivity.
- *Insertion or reconstruction of geometry:* Is reserved for regions where no reliable geometry can be recovered. In these cases, the photogrammetric model serves only as a general reference, and new geometry is directly modelled to restore the intended form. Examples include reflective or translucent surfaces, or fragmented structures such as pearl strings.
- *Surface smoothing:* Addresses high-frequency noise or residual microgeometry. It is applied over larger areas but does not lend itself to discrete segmentation, and is therefore treated as a diffuse, untagged intervention.

This segmentation is recorded manually, with the operator defining boundaries and assigning labels. While tools and techniques may vary, annotation should conform to the defined categories to ensure semantic consistency across interventions.

It should be noted that these intervention categories do not correspond deterministically to specific flaw classes. The same morphological anomaly may require different corrective strategies depending on its local geometric context, the surrounding topology, and the operator's interpretation of the reconstruction artefact. Conversely, a single intervention type may be applied to address morphologically different flaws. For this reason, the taxonomy of flaws presented in Section §4 and the classification of corrective interventions introduced here are treated as analytically distinct layers: the former describes the observable manifestations of reconstruction anomalies, while the latter records the interpretive actions undertaken to resolve them, even if the two segmentation are closely related.

A more detailed discussion on the characterisation of mesh-processing interventions, developed along a different analytical axis from the operational one presented here, considers qualitatively the relationship between the consistent information available in the sampling and the information inferred during correction, as well as the role of the agent performing the inference. This complementary perspective, based on the same empirical dataset, is developed in a parallel previous study (Sullini, 2026b)().

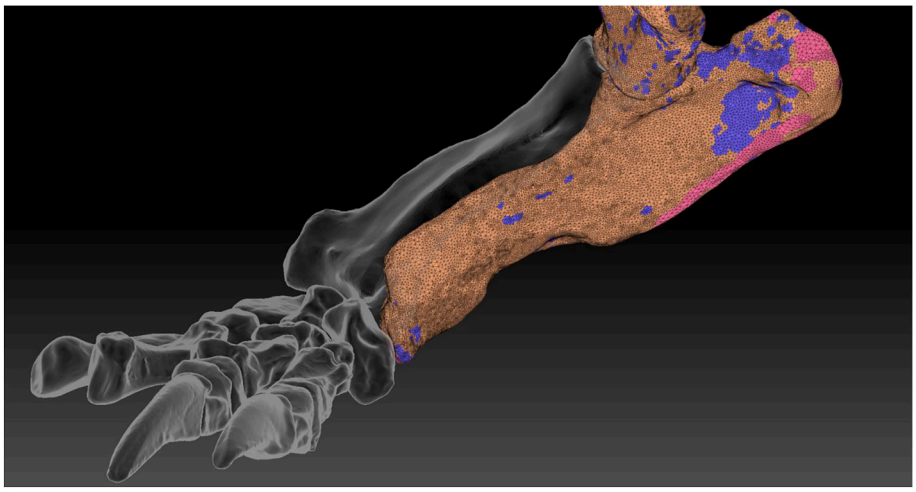

Fig. 4. Coloured segments differentiate operator interventions on the asset "Sid". Pink segments correspond to occluded areas whose geometry had been reconstructed with consistent geometry using topological guides; violet segments are related to hole-filling, after removal of inconsistent morphologies; the orange segment reflects only minor localised smoothing. The editing technique is reported here as an example of the corrections tracing, although its implementation falls outside the scope of the present work. Screencapture from ZBrush.

3.5. *Retro-projection of segments to RAWp*

From a technical standpoint, corrections modify the RAWp mesh, sometimes subtly, as in smoothing operations, and sometimes radically, through sculpting or replacement of geometry. Areas deemed reliable are typically left untouched or gently regularised to remove residual noise. As a result, DCHO is not topologically or morphologically identical to RAWp.

This difference can be exploited. The method described here applies a retro-projection process, in which segments defined in DCHO are spatially mapped back onto RAWp via raycasting modulated by geometric proximity (Ammirati et al., 2025a)(). This operation, referred to here as retro-projection, identifies on RAWp the regions that have been altered in DCHO, thereby exposing the areas that were implicitly recognised as flawed and subjected to correction (Fig. 5). The procedure is not intended to provide accurate pointwise correspondence between RAWp and DCHO, but to produce a workable approximation of corrected zones on the unedited mesh.

This enables RAWp to be geometrically segmented based on the local corrections applied in DCHO. However, this backward detection acquires a different meaning within RAWp: here, the segments do not reflect corrections, but rather identify the flaws themselves, established as such precisely because they were corrected, thus transforming RAWp into a diagnostic surface marked by operator-recognised anomalies.

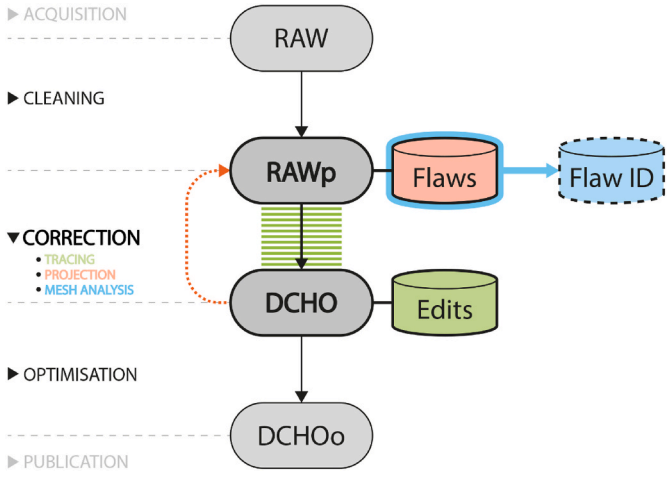

Fig. 5. Retro-projection from DCHO to RAWp in the heuristic iteration enables the segmentation of flaws based on past corrections, leading to the emergence of flaw types. This feedback loop functions as a device for extracting and structuring tacit knowledge.

These segments can then be analysed in conjunction with mesh indicators, curvature, edge length, confidence, to characterise each region in both morphological and procedural terms. Retro-projection establishes the knowledge loop that grounds classification in real-world operations, elevating RAWp from passive product of automated processing to epistemic layer: a representation of the object that includes not only artefacts, but the operator's recognition of them.

Although technically auxiliary, this step becomes methodologically central during the initial iterations of the workflow, which constitute a heuristic phase in which tacit recognition is progressively systematised, thus establishing a structured bridge between correction and classification. In later stages, once flaws are analytically defined and can be directly tagged, retro-projection may no longer be required. In the current configuration, however, it plays a foundational role in turning embodied knowledge into communicable structure.

Yet retro-projection introduces a structural vulnerability, as annotations are made during correction, but classification of flaws occurs afterwards, once the original interpretive act has already taken place. The retro-projection is therefore temporally and semantically detached from its source, and spatial correspondence may be imprecise, and categories inferred a posteriori may not fully reflect the rationale behind each correction. These shifts, both geometric and interpretive, may distort the intended meaning of the correction and lead to inconsistencies in flaw classification.

Crucially, this feedback mechanism is explicitly provisional. Its purpose is not to ensure perfect segment alignment, but to support the foundational, heuristic phase in which flaw types are defined, validated, and normalised. The immediate priority is to surface, name, and structure recurring defect patterns. Consequently, minor inconsistencies are tolerated, as they do not compromise the emergence of robust categories. Once the taxonomy is established and internalised, segmentation will occur directly on RAWp, enabling a consistent propagation of semantically meaningful regions throughout the workflow. The fragility of the current feedback process is thus a necessary condition of its foundational role, and its obsolescence is already embedded in the method's own trajectory.

3.6. Application of mesh analysis in the method

Within the proposed framework, mesh analysis refers to the interpretation of morphometric, sampling and photogrammetric indicators extracted from the RAWp model. Their function is not to certify accuracy, but to assist in the local inspection of surface behaviour, supporting flaw identification through patterns such as curvature and edge length anomalies, or low-confidence regions. During the heuristic iterations of the workflow, these indicators are explored within the segments identified through retro-projection downstream of corrective interventions, in order to recognise recurrent patterns. Once such patterns have been stabilised into explicit characterisations, they can be sought directly on the RAWp mesh.

With regard to the morphometric component of the analysis, two PCA-derived descriptors are systematically employed within this framework, selected for the specific diagnostic capacity they have demonstrated in relation to morphological hallucinations.

- *Omnivariance*: computed as the cubic root of the product of the three eigenvalues of the local covariance matrix, captures the overall magnitude of spatial variation and serves as a key indicator of morphological deviation.
- *Linearity*: computed from the relative dominance of the first eigenvalue, reflects the directional character of that variation. This distinction is particularly relevant for differentiating point-like features (e.g., spikes or pits) from elongated anomalies (e.g., ridges or bridging).

Both chosen morphometric indicators are computed across a small range of neighbourhood scales proportional to the known local edge length, enhancing robustness and supporting the frequency-based reading of surface flaws introduced in the taxonomy. In the present implementation this evaluation is performed in CloudCompare. Specifically, since edge length is known for each mesh, the neighbourhood scale is estimated empirically. For each test asset, three reference areas are selected to represent positive and negative edge-length anomalies and average sampling density. A Euclidean neighbourhood with radius approximately twice the local edge length is then considered for each area, providing a consistent scale for qualitative inspection. In this initial heuristic stage of the workflow, the multiscale evaluation remains qualitative, aimed at recognising recurrent patterns, while more structured scale-dependent characterisations may emerge in later iterations of the method.

The associated scalar fields can then be visualised as false-colour maps (Fig. 6), enabling qualitative inspection of local morphometric behaviour. Sampling structure, in turn, is inspected visually through the mesh wireframe, which provides an immediate cue about local edge-length distribution without requiring the explicit computation of a dedicated scalar field.

Although their values are not metrologically absolute, these descriptors provide a reproducible morphometric substrate that supports interpretation and documentation. Their role is not to replace expert judgement but to reinforce it through localised and reproducible metrics.

Though approximate, morphometric estimation via Omnivariance and Linearity offers a consistent and operationally viable substrate grounded in mesh geometry. None of these descriptors is diagnostic in isolation: a sharp curvature anomaly may equally reflect a legitimate feature or a reconstruction artefact, and mesh analysis therefore relies on the combined interpretation of morphometric and sampling properties, evaluated in relation to the surrounding geometry and acquisition context. Rather than replacing expert judgement, these indicators reinforce it through localised and reproducible metrics, and may serve as anchors for future workflows including machine-learning-based flaw detection and multiscale segmentation.

To reduce this ambiguity, morphometric descriptors are

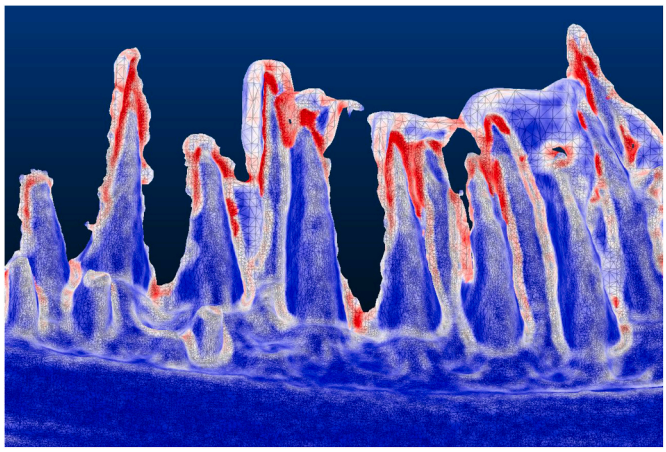

Fig. 6. Zoomed colour-mapped visualisation of the PCA-derived Omnivariance scalar field in CloudCompare, computed for the RAWp asset Lophius Piscatorius. The black wireframe overlay enables simultaneous visual inspection of edge length. At the selected analysis scale, the hallucinated regions are correlated with the combination of low omnivariance (white to blue) and relatively long edges, while high omnivariance (red) is concentrated along their margins, where it correlates with the algorithmic "stitching" between hallucinated clusters and consistent surface geometry.

complemented by sampling indicators, which provide contextual information about the acquisition process and help qualify the plausibility of surface patterns. Within this framework, PCA-derived indicators are not treated as standalone classifiers but as morphological guides, whose interpretation relies on the combined reading of morphometric, sampling, and photogrammetric descriptors. This integrated reading allows similar geometric configurations to be disambiguated through structural context rather than through threshold-based classification. Specifically, the contextual indicators adopted to complement the morphometric descriptors are.

- *Edge Length*: assessed through wireframe inspection to identify locally undersampled or degenerate regions associated with the interpolation of insufficiently sampled areas.
- *Confidence:* used as an indicator highlighting areas where reconstruction plausibility is low, signalling regions where the reliability of the reconstructed surface may be reduced.

3.7. Semantic structuring and paradata integration

The combined reading of these indicators, however, is not merely a technical operation: the decisions it informs are inherently interpretive, and the segments it produces carry semantic weight that extends beyond their geometric definition. Segmentation applied to 3D models is never purely geometric: it encodes interpretive choices made during correction on DCHO, namely paradata, and provides the foundation for plausibility assessment of reconstructed areas on RAWp, in line with the London Charter (2009). This semantic layer permeates the entire methodological framework and lies at the core of the problems outlined in Section 1, which this work aims to address through a structured approach.

This raises a key question: how can the embedded semantic information in DCHO and RAWp be interfaced with external metadata ontologies? Unlike external annotation models based on separate property files, the approach proposed here embeds semantic content directly within the 3D model. Two hypothetical strategies can be envisioned, both easily implementable without altering the method. The first, more lightweight, encodes meaning directly into the segment name. Each layer is assigned a UID (Unique Identifier) that includes the flaw or correction type and an alphanumeric code, allowing for both human

readability and automated parsing. The second, more articulated, involves associating structured metadata with each segment, for instance, typological class, frequency band, confidence level, or editing rationale, using custom property containers available in environments such as Blender. These attributes could then be retrieved by custom parsers or upon import into external systems. This would allow more expressive metadata to be embedded within the file itself, without relying on external descriptors, and would facilitate future alignment with ontologies or machine-readable interfaces.

These hypothetical strategies aim to explore the feasibility of embedding semantically structured information within 3D assets. While their full implementation exceeds the scope of this article, they are conceptually aligned with current efforts to encode geometric semantics through CIDOC CRM extensions (Giovannini, 2018; Amico and Felicetti, 2021), even if such approaches are typically tailored to architectural-scale models rather than object-scale artefacts.

What is particularly relevant to the present work is that both the cleaned mesh (DCHO) and the flaw-tagged version (RAWp) are already understood as dual-purpose artefacts: not only geometric representations, but also carriers of structured, interpretable paradata. Together, these layers define a stable and differentiated information system within the RAW → DCHOo workflow, where each mesh state is not merely a procedural step but a semantically meaningful unit of knowledge.

4. Taxonomy

4.1. General rationale

The taxonomy presented here stems from a first operational implementation of the method, directly applied by the author within real-world digitisation workflows. It reflects qualitative knowledge consolidated through repeated mesh correction tasks and should be understood as a preliminary articulation of classification logic, internally coherent but not yet externally validated. As defined in the methodology, flaw types have been derived retrospectively from expert-guided correction, then annotated and grouped based on recurring morphological patterns, spatial scale, and indicative mesh metrics.

Although grounded in general geometric principles, the taxonomy is inevitably influenced by the specific surface reconstruction algorithms used. Different strategies, such as Poisson reconstruction, Delaunay triangulation, or Marching Cubes, tend to produce distinct families of artefacts and modulate their appearance. For instance, Poisson reconstruction generates watertight surfaces by design and never leaves open holes, while Delaunay or Marching Cubes may result in incomplete or disconnected regions. A flaw may therefore appear differently or even not at all, depending on the reconstruction pipeline. In this implementation, the taxonomy is primarily informed by Depth Map fusion workflows, as consistently adopted across all test assets. However, it is deliberately structured to accommodate outputs from other reconstruction paradigms, ensuring adaptability across heterogeneous datasets and use cases.

While this version is based on a single-operator corpus, its goal is not to impose a fixed classification, but to demonstrate how semantic segmentation can emerge from procedural traceability and domain-specific judgement. It serves as a proof of concept: a foundational set of flaw types derived through structured intervention tracing and retro-projection. Although sufficiently mature to support analytical reflection, the taxonomy remains provisional. Broader adoption will require systematic application across diverse digitisation contexts by independent operators, within a future validation programme.

4.2. Characterisation of flaws

As flaw classification in the current framework follows the intervention based on the retro-projected mapping of corrected areas onto RAWp, there is a structural risk that categories may be shaped by the techniques used to correct them. To prevent this, the taxonomy has been deliberately constructed to remain independent of specific remediation strategies. Rather than deriving flaw types from the nature of the intervention, it defines them through observable morphological traits and recurring causal patterns. Correction strategies may align with these categories, but they do not determine them.

This separation is essential to avoid any circular reasoning in which the classification of a flaw would merely reflect the technique used to correct it. By defining flaws independently of remediation strategies, the taxonomy promotes an objective understanding of mesh defects and preserves its applicability across different tools, workflows, and software ecosystems. It also prevents premature assumptions about optimal correction, which may vary by context or evolve with technological progress.

Each flaw entry in the taxonomy follows a consistent template structured into three sections.

- *Visual example*: a representative image drawn from the annotated dataset, showing a real-world occurrence of the flaw. For conciseness and to avoid redundancy, only three illustrative examples are included in this publication.
- *Morphological description*: a concise qualitative account of the flaw's appearance, typical geometric configuration, and recurrent locations on the mesh (e.g., tips, folds, recesses). Terminology is chosen to remain scientifically grounded while accessible to non-specialist users. Where no established label existed in the literature, new names were introduced following the same communicative intent.
- *Analysis*: a qualitative reading of how the flaw correlates with mesh indicators, Photogrammetric *Confidence*, visually assessed *Edge Length*, and two PCA-derived descriptors (*Omnivariance* and *Linearity*). Consistently with the role assigned to them in the methodology, these indicators support operator recognition but are not used as classification criteria. Their interpretation relies on visualisation through false-colour maps and wireframes.

It is important to point out that PCA is known to be unstable when applied to sparsely sampled areas, due to the limited number of vertices within the analysis kernel (Pauly et al., 2003). This issue is particularly relevant for photogrammetric meshes, which are typically non-isotropic. To mitigate this weakness at an operational level, the mesh can be densified (Jaiswal et al., 2024); in the current implementation, this is achieved by resampling the entire mesh so that large polygons are populated with additional points, thereby improving the stability of the PCA computation.

Moreover, PCA-derived indicators provide stable and interpretable proxies for local surface variation, but lack directional information, as the eigenvalues of a covariance matrix are inherently positive. As a result, these descriptors cannot distinguish between convex and concave regions. Yet for a human operator, this distinction is both immediate and decisive: a cusp and a pit, or a ridge and a fold, are perceived as fundamentally different.

This gap highlights the descriptive limits of scalar indicators when decoupled from embodied recognition. Hence, such distinctions are consistently reflected in the flaw categories and differentiated mainly via morphological description. In future iterations, they may be formalised analytically through more expressive mesh analysis and integrated into semi-automated pre-assessment workflows. The proposed classification should therefore be understood as a first operational implementation, structured to accommodate future refinements.

4.3. Differentiation in LOW, MID, HIGH-frequency

The classification of flaws in this taxonomy is primarily organised by the spatial scale at which they manifest. Coherently, PCA descriptors are computed at three neighbourhood scales using kernels of appropriate size. Rather than linking each indicator to a specific frequency band, the

analysis evaluates local surface behaviour across all scales. This multi-scale comparison supports the identification of flaws by detecting patterns both at individual frequencies and across the full spectrum.

To enable comparison across objects of different sizes, a relative measure of spatial frequency is adopted, defined as $\lambda^{-1} = D/d$, where D is the characteristic size of the object (e.g., the bounding box diagonal), and d is the estimated extent of the flaw, understood as the region over which the geometric irregularity is visibly expressed.

This formulation does not rely on absolute thresholds, but instead relates flaw size to object scale. For example, a spike of fixed dimension may appear negligible on a large object or as a major deformation on a small one. The parameter λ^{-1} captures this proportionality, allowing classification based on perceptual and morphological salience rather than absolute metrics.

In practice, d is not derived from mesh primitives such as triangle width or edge length, but corresponds to a qualitative footprint estimated by the operator—e.g., the diameter of a bulge, the radius of a void, or the length of a fold. While based on expert judgment, this estimate is grounded in visual and structural cues. Future refinement using curvature fields, confidence maps, or error metrics is anticipated but not implemented here.

Based on λ^{-1}, three empirical bands are defined: LOW frequency ($\lambda^{-1} < 5$), MID frequency ($5 \leq \lambda^{-1} \leq 10$), and HIGH frequency ($\lambda^{-1} > 20$). Intermediate values ($10 < \lambda^{-1} \leq 20$) are treated as transitional and assigned based on dominant morphology. These bands are heuristic and non-predictive; they serve to support operator-level recognition and enable comparative classification. Some flaws may span multiple bands, but each is assigned to the frequency where it most typically occurs in the analysed corpus.

To reduce redundancy and improve readability, the bands are abbreviated as LF (Low Frequency), MF (Mid Frequency), and HF (High Frequency) throughout flaw descriptions. These designations indicate the scale at which each indicator, particularly Omnivariance, is evaluated within the multiscale PCA analysis.

Finally, many high-frequency flaws commonly recognised in the literature, such as noise, spikes, and mesh connectivity errors, are not included in this taxonomy, as they are systematically filtered out during the automatic processing leading to RAWp. These artefacts are typically morphologically degenerate rather than ambiguous and thus require no further differentiation within this framework.

4.4. Flaw taxonomy

4.4.1. Low-frequency flaws

4.4.1.1. F_M_BLO; Glancing Angle Blobs. Irregular swollen growths, often asymmetrical and detached from the underlying morphological logic of the item and extend outward from the expected surface envelope, typically along edges or tips of the object. Correlated to dataset images showing the item against a contrasting background (Fig. 7a).

- *Confidence:* Low or uneven across the affected protrusion.
- *Edge Length:* Longer than the surrounding area, isotropic over the anomaly.
- *Omnivariance:* Low at LF and MF, moderate HF near the protrusion's border.
- Linearity: High, aligned with the edge tangent.

4.4.1.2. F_L_CEM; Connection Membranes. These artefacts appear smooth and continuous, often planar or twisted, and typically resemble thin membranes stretched between two structural anchor curves. They occur between folds, cavities, or opposed surfaces, and are often located in narrow voids or connecting opposed low-sampled surfaces (Fig. 7b).

- *Confidence:* Low, with short gradient at the boundary
- *Edge Length:* Low.
- *Omnivariance:* Generally low, with localised HF peaks along boundaries.
- *Linearity:* Low, with directional spikes near boundaries.

4.4.1.3. F_L_PAT; Low-confidence Patch. Smooth areas that appear geometrically plausible but exhibit a visibly higher frequency than their surroundings, often through edge interpolation. While such flaws correspond to actual holes in some reconstruction pipelines, in DM workflows they manifest as subtly inflated closures or false continuities, frequently aligned with low-visibility or photometrically challenging zones (Fig. 7c).

- *Confidence:* Low with a smooth gradient at the boundary.
- *Edge Length:* High, increasing with the size of the anomaly.
- *Omnivariance:* Low and locally smooth, with limited internal variation.
- *Linearity:* Weak and diffuse, with no dominant directionality.

4.4.2. Mid-frequency flaws

4.4.2.1. F_M_ATT; Crests Attenuation. Linear features such as ridges, folds, or sharp edges appear softened or rounded, losing their sharp definition. The crest becomes thicker and less distinct, occasionally bloated, with reduced contrast across adjacent surface orientations.

- *Confidence:* Comparable to the surrounding surface.
- *Edge Length:* Slightly longer across the smoothed transition.
- *Omnivariance:* Low at LF and MF; moderate HF response near inflection points.
- *Linearity:* Moderate along the crest, decreasing in smoothed regions

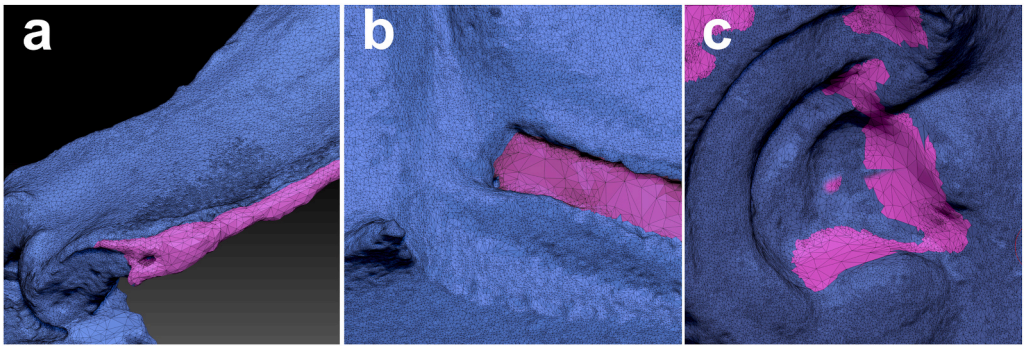

Fig. 7. Zbrush postprocessed screenshots showing examples of Low-frequency flaws: *Glancing Angle Blobs F_L_BLO* (a) *in the "Sid" asset; Connection Membranes F_L_CEM (b) in the "Limulo" asset; Low-Confidence Patch F_L_PAT (c) in the "Sardo" asset.*

4.4.2.2. F_ M_BLU; TTip Blunting. Rounded or truncated terminations of pointed features. These areas lose their sharp angular profile, replaced by domed or flattened ends and are shortened in respect to the real item. The entity of blunting is inversely proportional to sampling.

- *Confidence:* High and comparable to the neighbourhood.
- *Edge Length:* Slightly longer than the neighbourhood.
- *Omnivariance:* Low, with smoothed HF variation near the tip.
- *Linearity:* Locally variable, moderate along the preserved axis but low at the blunted tip.

4.4.2.3. F_M_COM; Collapsed Membranes. Localised collapse of very thin structures. The do not resolve the expected separation between layers, resulting in missing geometry and creating additional non-existing holes, in the topological sense, whose perimeter is formed by high curvature junctions between opposite sides of the structure (Fig. 8a).

- *Confidence:* Even with neighbourhood, suddenly increased for the junctions.
- *Edge Length:* Generally short and even with surroundings, also in the junctions due to thinness.
- *Omnivariance:* Low at LF, with irregular spikes at MF and HF near topological disruptions.
- *Linearity:* Locally high along the junction at HF.

4.4.2.4. F_M_FIL; Pit Filling. Depressions or cavities that appear shallower than expected or completely filled in, resulting in the loss or flattening of concave features. The mesh forms a smooth convex or planar patch, obscuring the original indentation.

- *Confidence:* Comparable to the surrounding surface.
- *Edge Length:* Slightly longer within the filled patch.
- *Omnivariance:* Low at LF and MF; locally increased HF near the edge of the filled area.
- *Linearity:* Low and incoherent across the flattened surface.

4.4.2.5. F_M_GHO; Reflection Ghosts. Detached or floating geometries, often semi-coherent but visibly disconnected. These artefacts may mimic nearby features in distorted form or appear as duplicated, misaligned volumes hovering over the surface. Correlated to image sets with flat reflective surfaces.

- *Confidence:* Variable, but generally low.
- *Edge Length:* Mostly short, but uneven.
- *Omnivariance:* Variable and non-characteristic, ranging from low to high depending on local structure.
- *Linearity:* Variable, with localised directionality emerging in elongated fragments.

4.4.2.6. F_M_ISL; High Sampling Islands. Localised areas of detailed geometry embedded in otherwise smooth, low-curvature surfaces. These patches exhibit coherent polarity with their surroundings but disrupt visual continuity, appearing as dense meso-structures within homogeneous regions. Often correlated with photometric anomalies, such as high-contrast zones within otherwise uniform textures.

- *Confidence:* High, consistent with the surrounding surface.
- *Edge Length:* Significantly shorter than neighbouring regions.
- *Omnivariance:* Low at LF and MF; moderate HF within the island.
- *Linearity:* Nearly null with no consistent directional structure.

4.4.2.7. F_M_SHA; Crevices Shallowing. Narrow grooves or fissures that appear partially filled, resulting in an often irregular but generally shallower geometry on the bottom of the crevice. The original separation between adjacent features is lost, and the geometry is smoothed in areas where a clear incision should be present (Fig. 8b).

- *Confidence:* High, with irregular drop near the former groove.
- *Edge Length:* Short, but inconsistent in directionality and depth.
- *Omnivariance:* Low at LF and MF, moderate HF where local shape breaks occur.
- *Linearity:* Weak and unstable along the expected groove axis.

4.4.2.8. F_M_SPC; Specular Cavities. Isolated or clustered depressions appearing like irregular corrosions, misaligned with actual geometry and located in areas affected by strong specular highlights, especially if homogeneously coloured. They have smooth edges and shallow depth, disrupting otherwise continuous, flat or low-curvature surfaces (Fig. 8c).

- *Confidence:* Moderate, often inconsistent with surrounding surface.
- *Edge Length:* Regular but elongated in the cavity.
- *Omnivariance:* Low at LF, locally increased at MF around contour.
- *Linearity:* Low, with low directionality at the border.

4.4.3. High-frequency flaws

4.4.3.1. F_H_BUB; Offset Bubbles. Localised, rounded swellings that protrude slightly from smooth surfaces. Typically appear in small clusters or isolated spots, creating vesicle-like outliers not coherent with the object's geometry. Often observed on otherwise regular regions, and likely caused by misaligned point clouds or depth outliers.

- *Confidence:* High at core, lower near perimeter.
- *Edge Length:* Slightly longer than surroundings, locally irregular.
- *Omnivariance:* Low at LF, Moderate at MF, with sharp rise at HF around the edges.
- *Linearity:* Low throughout, no dominant orientation.

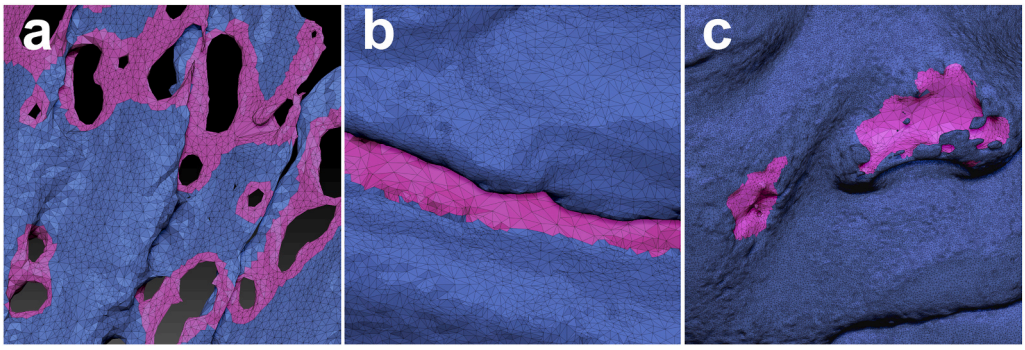

Fig. 8. Zbrush postprocessed screenshots showing examples of Mid-frequency flaws: Collapsed Membranes F_M_COM (a) in the "Lofoforo" asset; *Crevices Shallowing F_M_SHA (b) in the "Lophius" asset;* Specular Cavities *F_M_SPC (c) in the "Kung" asset.*

4.4.3.2. F_H_CRU; Offset Crust. Plausible surface patches slightly offset from the true geometry, often hard to detect centrally but producing visible steps along their boundary and presenting irregular corroded margins, then small islands emerging with constant offset around the proper margin, typically localised on low curvature areas. This type may subsume seam-like artefacts in active-sensor workflows (Fig. 9a).

- *Confidence:* High on the crust, drop-offs near its borders might be present.
- *Edge Length:* Even with surroundings.
- *Omnivariance:* Low at LF and MF, moderate HF along the rim.
- *Linearity:* Weak or inconsistent, with localised alignment where the crust follows a geometric feature.

4.4.3.3. F_H_PEA; String of Pearls. Elongated features appear fragmented into rounded, bead-like segments separated by gaps or depressions. The artefact disrupts linear continuity, especially in thin ridges or filaments, producing an alternation of swollen nodes and voids (Fig. 9b).

- *Confidence:* High, homogeneous along the structure.
- *Edge Length:* Locally short, alternating with slightly longer gaps.
- *Omnivariance:* Low at LF, locally moderate at MF and high at HF around bead edges.
- Linearity: Alternating; low within each bead, higher in segments connecting them.

4.4.3.4. F_M_SKI; Double Skin. Duplicated, sheet-like surface layers running near and roughly parallel to the actual geometry, forming thin offset shells on low curvature geometry. These artefacts may be continuous or fragmented and tend to follow the real surface with slight misalignment. The corresponding area of the true geometry is often less resolved.

- *Confidence:* Moderate to high, decreasing towards the duplicated layer's boundaries.
- *Edge Length:* Slightly shorter within the duplicated region; longer on the underlying surface.
- *Omnivariance:* Low at LF and MF, with scattered HF anomalies along overlaps.
- *Linearity:* Moderate, aligned with the curvature direction of the true surface.

4.4.3.5. F_H_TEX; Texture-induced Microgeometry. Fine-scale surface artefacts that replicate texture patterns rather than true geometry. They appear as low-relief ridges, dots, or grain aligned with photometric structures—e.g., brush strokes, veins, or prints—over otherwise smooth or weakly curved regions (Fig. 9c).

- *Confidence:* High, consistent with surrounding surface.
- *Edge Length:* Shorter than average, uniformly distributed.
- *Omnivariance:* Low at LF and MF, moderate HF where false details emerge.
- *Linearity:* Weak, scattered and locally directional if texture is structured.

5. Conclusions

This work introduces a methodological framework for the semantic annotation of flaws in 3D reconstructions, rooted in expert-guided correction and retroactive flaw recognition. Rather than relying on pre-established taxonomies or automated classification, the method is grounded in procedural traceability and localised mesh inspection, supported by geometric and metrological indicators. Its function is not to offer definitive correction strategies, but to systematise and encode the tacit knowledge routinely applied during expert intervention.

The core contribution lies in the operationalisation of this logic: by annotating retro-projected corrections and interpreting them through scalar indicators such as PCA-derived descriptors, edge length, and photogrammetric confidence, the method provides a reproducible substrate for flaw characterisation and its translation into an explicitly formalised form, hence accessible, interoperable and reusable. Furthermore, it enables the emergence of semantically meaningful segmentations and consolidates a structured record of operator judgement without imposing rigid analytical thresholds.

This positioning entails a shift from validation paradigms based on external ground-truth references towards an intrinsic assessment of the reconstructed geometry itself. Rather than evaluating the performance of specific acquisition instruments or reconstruction pipelines under controlled conditions, the method operates directly on the produced mesh, where artefacts and inconsistencies are identified through their geometric manifestation. In this sense, quality assessment becomes instance-specific and independent of the technical configuration that generated the model, making the approach applicable even in constrained or suboptimal acquisition contexts, such as those frequently encountered in small and medium GLAM institutions.

As a first implementation, the method is applied to a single-operator dataset. The flaw taxonomy derived from this process, while internally coherent and grounded in recurring morphologies, should be regarded as provisional. It demonstrates the viability of the proposed framework and offers a functional basis for mesh diagnosis, but does not claim exhaustiveness or universal applicability. The taxonomy serves both as an outcome and as a validation mechanism for the method: it assesses whether structured, repeatable patterns can be extracted from the correction process, and whether they can support informed classification of geometric anomalies.

The proposed method is conceived within the RAW–DCHO framework as a refinement of the automatic quality control of Cultural

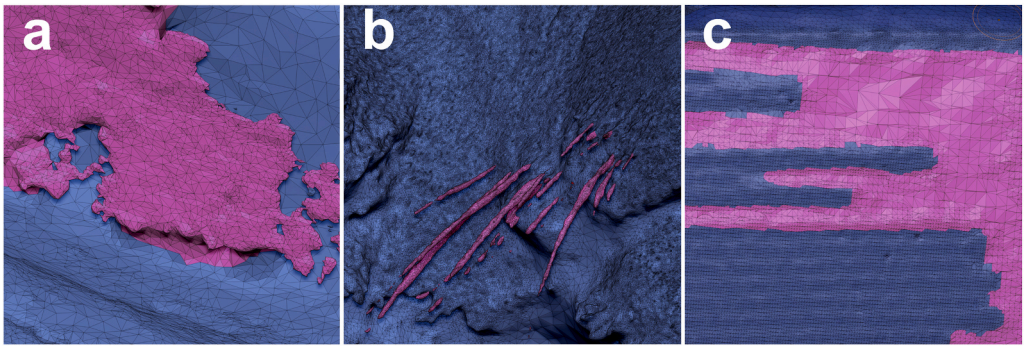

Fig. 9. Zbrush postprocessed screenshots shwing examples of High-frequency flaws: Double Skin F_H_SKI (a) in the "Sid" asset; *String of Pearls F_H_PEA (b) in the "Lions" asset; Texture-Induced Microgeometry F_H_TEX (c) in the "Spino" asset;.*

Heritage 3D models leading to the RAWp state. A further operational constraint should be noted, as aggressive smoothing or remeshing, typically applied within Mesh Quality Assessment workflows to adapt meshes to applications requiring pseudo-isotropic sampling (e.g., finite element analysis), is incompatible with this method, as such operations alter precisely the microstructural evidence leveraged as diagnostically significant for hallucination detection. This constraint does not, however, preclude its application directly on unprocessed mesh outputs of Surface Reconstruction, where it may operate in synergy with compatible quality assessment approaches.

In its current form, the method operates as a device for tacit knowledge collection. Its use is retrospective: it gathers expert decisions, transforms them into structured annotations, and distils them into a semantic taxonomy. This mode will remain valid as the taxonomy evolves, potentially accommodating new reconstruction algorithms, acquisition tools, or object classes. While this retrospective mode remains central, the structured formalisation of recurring flaw patterns also enables a controlled transition towards a more operative use of the method. Once stabilised, the taxonomy entries may be translated into configurations of numerical indicators, expressed through consistent combinations of scalar descriptors and relative thresholds. In this form, the method can support the pre-assessment of suspicious morphologies directly on raw reconstructed surfaces, functioning as a feedback mechanism that complements expert inspection without replacing it.

The concrete outcome of this process, namely the segmentation of hallucinated regions on the raw mesh output from acquisition, is already available from the heuristic iterations, and its intelligibility and interoperability are already being addressed. A dedicated line of research underway concerns the visualisation of annotations as interactive semantic textures (Sullini, 2026b), applied using the same parametrisation employed by canonical PBR maps. This enables direct visualisation in a web browser without specific installations, through compatibility with the WEB3D technological ecosystem, either as an alternative albedo map in generic viewers, or through dedicated tools such as TRACE (Ammirati et al., 2025a)(), which allow the user to modulate the relative visibility of the albedo and semantic layers.

Furthermore, the glTF delivery format further extends this capacity: its *EXT_mesh_features* and *EXT_structural_metadata extensions* allow the semantic layer to be embedded directly into the model geometry and made queryable in already compatible viewers such as Cesium.js, as well as prospectively in any viewer capable of retrieving the information when structured according to the specifications of these extensions. These implementations directly support the intelligibility of flaw documentation across diverse operator profiles, enabling transparent and reproducible communication of diagnostic outcomes.

Within this framework, the proposed approach establishes a traceable and interpretable link between reconstructed geometry and expert-driven diagnostic knowledge, positioning semantic annotation as a foundational component for both analytical assessment and future developments in automated validation of 3D cultural heritage assets.

6. Future works

Three main developments are envisioned.

The first concerns the extension of the method as a knowledge collection infrastructure across broader empirical contexts. To strengthen its generalisability, the method should be evaluated across different operators, datasets, and domains, enabling multi-author validation and comparative flaw profiling. Such a multi-operator extension can only be meaningfully undertaken now that the methodological framework has been defined and the preliminary hypotheses explored in the present work have been verified. The results obtained here provide a preliminary foundation for such an extension.

The second concerns the progressive consolidation of the taxonomy in an analytical direction and its potential integration into advanced quality control systems. As the heuristic phase stabilises, further refinements are expected to improve the precision of segment delineation, the robustness of segmentation granularity, and the balance between recall and false negatives. In this phase, the previously identified configurations of numerical indicators may be further stabilised and systematised as structured criteria for anomaly classification, enabling their direct application to the detection and categorisation of suspicious morphologies on RAWp models (Fig. 10).

Only once such analytical stability has been achieved may these configurations support predictive approaches. In this scenario, classification could be implemented either through rule-based systems grounded in indicator thresholds or through trained segmentation models, where the identified patterns function as features or attention mechanisms within machine learning frameworks. However, this perspective remains conditional upon the consolidation of the heuristic phase and the analytical validation of the taxonomy.

The third concerns the extension of the method beyond mesh analysis. With appropriate adaptations, most notably the replacement of confidence and edge-length indicators with density-based measures, the framework could be extended to the qualification of point clouds themselves, using Surface Reconstruction as an analytical tool to identify spatial regions where available samples induce hallucinated surfaces. This aspect is currently the subject of ongoing research.

Lastly, the taxonomy itself is conceived as an open-ended structure. While its current form is grounded in Depth Map–based reconstruction workflows, its design explicitly accommodates extension to other acquisition methods, surface generation algorithms, and diagnostic practices. This generalisability is not claimed as demonstrated here, but as a deliberate feature of the framework, one that will require future empirical verification across heterogeneous datasets, acquisition conditions, and processing pipelines.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this manuscript, the author used AI-assisted tools (ChatGPT) solely for the purpose of improving the readability and linguistic clarity of the text, which was originally drafted in another language. The use of such tools was conducted under strict human oversight, and the final version was critically reviewed and approved by the author, who takes full responsibility for its accuracy and integrity.

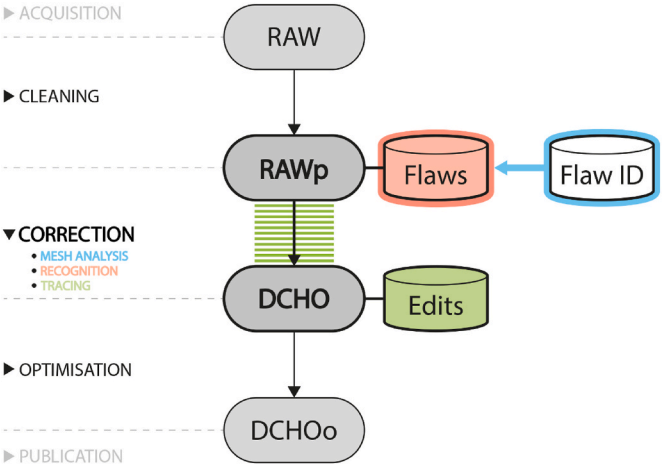

Fig. 10. Once the analytical characterisation of flaws within the taxonomy is established, mesh analysis on RAWp can support expert segmentation and tagging of flaws prior to correction, without altering the RAW → DCHOo workflow.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

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.

References

Amico, N., Felicetti, A., 2021. *Ontological entities for planning and describing cultural heritage 3D models creation* (arXiv:2106.07277). arXiv. https://doi.org/10.48550/arXiv.2106.07277.

Apollonio, F.I., Fantini, F., Garagnani, S., Garagnani, S., Gaiani, M., 2021. A photogrammetry-based workflow for the accurate 3D construction and visualization of museums assets. Remote Sens. https://doi.org/10.3390/rs13030486.

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., 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, e00309. https://doi.org/10.1016/j.daach.2023.e00309.

Barazzetti, L., 2020. Solid model generation for digitized organic bodies via T-Splines. Heritage 3 (3), 3. https://doi.org/10.3390/heritage3030035.

Barzaghi, S., Bordignon, A., Collina, F., Fabbri, F., Fanini, B., Ferdani, D., Gualandi, B., Heibi, I., Mariniello, N., Massari, A., Massidda, M., Moretti, A., Peroni, S., Pescarin, S., Rega, M.F., Renda, G., Sullini, M., 2025. A reproducible workflow for the creation of digital twins in the cultural heritage domain. Transformations: A DARIAH Journal, Workflows (Image-based workflows), 14773. https://doi.org/10.46298/transformations.14773.

Bentkowska-Kafel, A., Denard, H. (Eds.), 2016. Paradata and Transparency in Virtual Heritage. Routledge. https://doi.org/10.4324/9781315599366.

Beraldin, J.-A., Mackinnon, D., Cournoyer, L., 2015. Metrological characterization of 3D imaging systems: progress report on standards developments. In: Larquier, B. (Ed.), 17th International Congress of Metrology. EDP Sciences, p. 13003. https://doi.org/10.1051/metrology/20150013003.

Berger, M., Tagliasacchi, A., Seversky, L., Alliez, P., Guennebaud, G., Levine, J., Sharf, A., Silva, C., 2016. A Survey of surface reconstruction from point clouds. Comput. Graph. Forum 36. https://doi.org/10.1111/cgf.12802 n/a-n/a.

Bordignon, A., Collina, F., Fabbri, F., Fanini, B., Ferdani, D., Rega, M.F., Sullini, M., 2026. Balancing visual credibility and transparency: a FAIR approach to 3D digitisation for cultural heritage documentation and dissemination. Peer Community J. 6. https://doi.org/10.24072/pcjournal.696.

Botsch, M., Kobbelt, L., Pauly, M., Alliez, P., Levy, B., 2010. Polygon Mesh Processing. CRC Press. https://doi.org/10.1145/1057432.1057457.

Do Carmo, M.P., 2016. Differential Geometry of Curves and Surfaces: Revised and Updated Second Edition. Courier Dover Publications.

Cazals, F., Pouget, M., 2005. Estimating differential quantities using polynomial fitting of osculating jets. Comput. Aided Geomet. Des. 22 (2), 121–146. https://doi.org/10.1016/j.cagd.2004.09.004.

Cipriani, L., Fantini, F., 2017. Digitalization culture vs archaeological visualization: integration of pipelines and open issues. Int. Arch. Photogram. Rem. Sens. Spatial Inf. Sci. 195–202. https://doi.org/10.5194/isprs-archives-XLII-2-W3-195-2017.

Cipriano, Gregory, Phillips George, N., Gleicher, Michael, 2009. Multi-Scale surface descriptors. IEEE Trans. Visual. Comput. Graph. 15. https://doi.org/10.1109/TVCG.2009.168.

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: Final Study Report. Publications Office of the European Union. https://data.europa.eu/doi/10.2759/471776.

Farella, E., Morelli, L., Rigon, S., Grilli, E., Remondino, F., 2022. Analysing key steps of the photogrammetric pipeline for Museum artefacts 3D digitisation. https://www.mdpi.com/2071-1050/14/9/5740.

Frost, A., Mirashrafi, S., Sánchez, C.M., Vacas-Madrid, D., Millan, E.R., Wilson, L., 2023. Digital documentation of reflective objects: a cross-polarised photogrammetry workflow. 3D Research Challenges in Cultural Heritage III: Complexity and Quality in Digitisation 13125, 131.

Giovannini, E.C., 2018. Virtual Reconstruction Information Management. A Scientific Method and 3D Visualization of Virtual Reconstruction Processes. Doctoral Thesis. Alma Mater Studiorum - Università di Bologna. https://doi.org/10.6092/unibo/amsdottorato/8330.

Jaiswal, M., Corpuz, A.M., Hsu, M.-C., 2024. Mesh-driven resampling and regularization for robust point cloud-based flow analysis directly on scanned objects. Comput. Methods Appl. Mech. Eng. 432, 117426. https://doi.org/10.1016/j.cma.2024.117426.

Jolliffe, I.T., Cadima, J., 2016. Principal component analysis: a review and recent developments. Philos. Trans. R. Soc. A Math. Phys. Eng. Sci. 374 (2065), 20150202. https://doi.org/10.1098/rsta.2015.0202.

Kirchhöfer, M., Chandler, J., Wackrow, R., 2011. Cultural Heritage recording utilising low-cost closerange photogrammetry. Geoinformatics FCE CTU 6, 185–192. https://doi.org/10.14311/gi.6.24.

Koller, D., Frischer, B., Humphreys, G., 2010. Research challenges for digital archives of 3D cultural heritage models. Journal on Computing and Cultural Heritage 2 (3), 7:1–7:17. https://doi.org/10.1145/1658346.1658347.

Leksell, T., 2018. Can market specialization reduce the entry cost of automated largescale 3D-scanning of movable artifacts for culture and heritage preservation? https://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-234308.

Menna, F., Nocerino, E., Remondino, F., Dellepiane, M., Callieri, M., Scopigno, R., 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, 675–683. https://doi.org/10.5194/isprsarchives-XLI-B5-675-2016.

Meyer, M., Desbrun, M., Schröder, P., Barr, A.H., 2003. Discrete differential-geometry operators for triangulated 2-Manifolds. In: Hege, H.-C., Polthier, K. (Eds.), Visualization and Mathematics III. Springer, pp. 35–57. https://doi.org/10.1007/978-3-662-05105-4_2.

Moyano, J., Cabrera-Revuelta, E., Nieto-Julián, J.E., Fernández-Alconchel, M., Fernández-Valderrama, P., 2023. Evaluation of geometric data registration of small objects from non-invasive techniques: applicability to the HBIM field. Sensors 23 (3), 1730. https://doi.org/10.3390/s23031730.

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

Nocerino, E., Menna, F., Remondino, F., 2014. Accuracy of typical photogrammetric networks in cultural heritage 3D modeling projects. *The International Archives of the Photogrammetry*. Remote Sensing and Spatial Information Sciences XL–5, 465–472. https://doi.org/10.5194/isprsarchives-XL-5-465-2014.

Nonaka, I., 1994. A dynamic theory of organizational knowledge creation. Organ. Sci. 5 (1), 14–37. https://doi.org/10.1287/orsc.5.1.14.

Pacheco, A., Bolivar-Baron, H., Gonzalez-Crespo, R., Pascual-Espada, J., 2014. Reconstruction of high resolution 3D objects from incomplete images and 3D information. International Journal of Interactive Multimedia and Artificial Intelligence 2 (6), 7. https://doi.org/10.9781/ijimai.2014.261.

ParteciPa, 2025. Linee guida per la digitalizzazione del patrimonio culturale, 2.4. Come. (s.d.). Docs Italia. Recovered July 21th 2025, from. https://docs.italia.it/italia/icdp/icdp-pnd-digitalizzazione-docs/it/v1.0-giugno-2022/il-progetto-di-digitalizzazione/come.html.

Pauly, M., Keiser, R., Gross, M., 2003. Multi-scale feature extraction on point-sampled surfaces. Comput. Graph. Forum 22 (3), 281–289. https://doi.org/10.1111/1467-8659.00675.

Pérez, E., Salamanca, S., Merchàn, P., Adán, A., 2016. A comparison of hole-filling methods in 3D. Int. J. Appl. Math. Comput. Sci. https://doi.org/10.1515/amcs-2016-0063.

Rahaman, H., Champion, E., 2019. To 3D or not 3D: choosing a photogrammetry workflow for cultural Heritage groups. Heritage 2 (3), 1835–1851. https://doi.org/10.3390/heritage2030112.

Remondino, F., Del Pizzo, S., Kersten, T.P., Troisi, S., 2012. Low-Cost and open-source solutions for automated image orientation – a critical overview. In: Ioannides, M., Fritsch, D., Leissner, J., Davies, R., Remondino, F., Caffo, R. (Eds.), Progress in Cultural Heritage Preservation. Springer, pp. 40–54. https://doi.org/10.1007/978-3-642-34234-9_5.

Remondino, F., Nocerino, E., Toschi, I., Menna, F., 2017. A critical review of automated photogrammetric processing of large datasets. Int. Arch. Photogram. Rem. Sens. Spatial Inf. Sci. XLII-2-W5, 591–599. https://doi.org/10.5194/isprs-archives-XLII-2-W5-591-2017.

Sebar, L.E., Grassini, S., Parvis, M., Lombardo, L.. A low-cost automatic acquisition system for photogrammetry. https://doi.org/10.1109/I2MTC50364.2021.9459991.

Sevilla Principles, 2017. International Council on Monuments and Sites (ICOMOS). https://icomos.es/wp-content/uploads/2020/06/Seville-Principles-IN-ES-FR.pdf.

Sorgente, T., Biasotti, S., Manzini, G., Spagnuolo, M., 2023. A Survey of indicators for mesh quality assessment. Comput. Graph. Forum 42 (2), 461–483. https://doi.org/10.1111/cgf.14779.

Sullini, M., 2026. GLTF Suitability for Digital Heritage 3D Assets. A c. Di), The Future of Heritage Science and Technologies II, Switzerland, pp. 447–462. https://doi.org/10.1007/978-3-031-98379-5_34.

Sullini, M., 2026. From Measured to Trusted: Documenting Geometric Reliability in 3D Heritage Assets. In: session. Springer Nature, forthcoming. Zenodo, 30. Technologies (CHNT, New, p. UHDL2025. https://doi.org/10.5281/zenodo.18858970.

Wachowiak, M.J., Karas, B.V., 2009. 3d scanning and replication for Museum and cultural Heritage applications. J. Am. Inst. Conserv. 48 (2), 141–158. https://doi.org/10.1179/019713609804516992.

Zhang, H., Kaick, O. van, Dyer, R., 2007. Spectral methods for mesh processing and analysis. https://doi.org/10.2312/egst.20071052.

Zhou, L., Sun, G., Li, Y., Li, W., Su, Z., 2022. Point cloud denoising review: from classical to deep learning-based approaches. Graph. Models 121, 101140. https://doi.org/10.1016/j.gmod.2022.101140.

Denard, H. (2009) "The London Charter. For the Computer-Based Visualization of Cultural Heritage, Version 2.1." (https://www.londoncharter.org, accessed on 25.7.2025).

Ammirati, L., Bordignon, A., Collina, F., Fabbri, F., Fanini, B., Ferdani, D., Rega, M. F., & Sullini, M. (2025). TRACE-ing the Gaps: Mapping Interventions on Incomplete 3D Meshes. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, *X-M-2–2025*, 3–11. 30th CIPA Symposium "Heritage Conservation from Bits: From Digital Documentation to Data-driven Heritage Conservation" - 25–29 August 2025, Seoul, Republic of Korea. https://doi.org/10.5194/isprs-annals-X-M-2-2025-3-2025.

Ammirati, L., Bordignon, A., Collina, F., Fabbri, F., Ferdani, D., Rega, M. F., & Sullini, M. (2025). 3D Digitisation for Geological and Paleontological Specimens: Challenges and Solutions. *Digital Heritage*. Digital Heritage 2025. https://doi.org/10.2312/dh.20253255.