

Research article

Published
2026-03-02

Cite as
Stefanos Costopoulos and
Georgios Papaioannou (2026)
*The impact of normal mapping
on the appearance of
geometrically simplified
archaeological 3D models*, Peer
Community Journal, 6: e17.

Correspondence
ste.kostopoulos@aueb.gr

Peer-review
Peer reviewed and
recommended by
PCI Archaeology,
[https://doi.org/10.24072/pci.
archaeo.100627](https://doi.org/10.24072/pci.archaeo.100627)



This article is licensed
under the Creative Commons
Attribution 4.0 License.

The impact of normal mapping on the appearance of geometrically simplified archaeological 3D models

Stefanos Costopoulos¹ and Georgios Papaioannou ¹

Volume 6 (2026), article e17

<https://doi.org/10.24072/pcjournal.683>

Abstract

A great challenge in 3D digitization and modelling lies in striking a balance between surface detail and model size, while accommodating the geometric information representation requirements of diverse archaeological practices. 3D meshes that contain a high polygon count severely impact storage requirements, data transmission, and 3D rendering. It is therefore typical for high-resolution digitized objects and modelled assets alike to be sub-sampled and simplified into levels of detail (LODs) to match the requirements of the intended applications. By doing so, however, geometric detail is lost, negatively affecting the model's visual fidelity during 3D visualization. Normal mapping offers a reliable and practical solution to this challenge by transferring surface orientation data from the original, detailed models into simplified versions. This approach preserves the key visual characteristics present on the surface of a high-detail model, without reintroducing geometric complexity into the simplified structure. In archaeological contexts, where both accuracy and accessibility are important, normal mapping enables lightweight models to maintain the perceptual qualities of high-resolution originals, while supporting diverse visualization applications, from analytical inspection to public dissemination. In this paper, our proposed methodology explores the robustness of detail transfer through normal mapping on simplified archaeological models, highlighting its effectiveness as a tool for efficient cultural heritage visualization and immersive applications. The study claims that normal-mapped models can effectively stand in for the original high-detail scans in a variety of scenarios, thus offering a viable and efficient alternative to the demanding storage and transmission of high-fidelity 3D models.

¹Athens University of Economics and Business - Athens, Greece

Introduction

Advances in computer vision algorithms have significantly contributed to the emergence of 3D digitization as a transformative tool in various disciplines, particularly in archaeology, where the preservation, restoration, and documentation of cultural material are of paramount importance (Aicardi et al., 2018; Remondino & Rizzi, 2010). By leveraging continuing technological innovations in digitization and encouraging researchers and professionals in the field to adopt the wide range of methods available (Pavlidis et al., 2007; Graham et al., 2017), each suited to different object types and documentation needs, it has become feasible to recreate highly detailed and accurate digital surrogates.

Early archaeological applications of 3D technologies primarily focused on digital documentation, including 3D acquisition and manual modeling of archaeological sites and 3D recreations (reconstructions) of archaeological interpretations of archaeological sites and objects. Through these applications of 3D technologies, visualization became a key aspect of archaeological study, workflow management and interpretation. The first 3D acquisition systems that were commercially available and successfully applied to the archaeological context were primarily based on active scanning, including laser striping scanners and systems employing visible light patterns (structured light scanners). Time-of-flight distance measurement technology then became popular, with LiDARs dominating long-range sensing, while image-based approaches began to establish themselves as budget solutions, where detail was not critical. From this point on, image-based technologies such as structure-from-motion (SfM) photogrammetry (Hartley & Zisserman, 2003), neural radiance fields (NeRFs) (Mildenhall et al., 2020) and 3D Gaussian Splatting (3DGS) (Kerbl et al., 2023), have significantly evolved to become particularly prominent due to their flexibility, scale independence, portability, relative simplicity, and low cost. It is possible to acquire detailed recordings of artifacts and sites, enabling the capture of intricate shapes, registering minute textural details, and subtle color variations with remarkable precision (Adamopoulos et al., 2021; Manferdini & Russo, 2013). More recently, conversion from alternative 3D representations, such as NeRFs and 3DGS to meshes has become a significant and active research and practical application field, with many successful approaches substantially improving the geometric detail reconstruction of intricate, fuzzy, translucent and thin objects (Oechsle et al., 2021; Yu et al., 2024), thus extending the range of applicability of image-based techniques.

Today, digitized content is used beyond digital documentation, to build digital twins of the original objects and sites (ARTEMIS, 2025-2027) and complement often traditionally invasive archaeological investigations (Papaioannou et al., 2017), which are inherently destructive. They allow the exploration of multiple hypotheses, continuous study and accessibility through virtual means, all while avoiding the requirement of direct handling of the originals or risking anthropogenic damage and alteration. As such, the creation of widely and efficiently accessible digital surrogates enables the mass data collection and interactive analysis of cultural heritage objects, supporting scholarly research, scientific interpretation and dissemination to wider audiences (Storeide et al., 2023).

This study focuses on the evaluation of the robustness of normal mapping in the context of archaeological 3D visualization. It is primarily addressed to archaeologists and cultural heritage experts familiar with 3D digitization workflows and curation of digital assets. The aim of the work is to provide insights on best practices for efficient and perceptually faithful 3D model representation.

By assessing how well simplified models enhanced with baked normal maps preserve the perceptual qualities of their high-resolution counterparts, the study supports the suitability of normal-mapped models as effective stand-in representations for the original high-detail scans, offering a viable and efficient alternative to the demanding storage and transmission of high-fidelity 3D models in inspection, visualization and hypothesis testing tasks. The importance of surface shading as a tool

to capture relief details and study intricate patterns and especially inscriptions, was understood many years ago and specialized image-based techniques such as Reflectance Transformation Imaging (RTI) have been successfully used in this context to recover subtle details, without the generation of geometric 3D information (Looten et al., 2025).

Problem Statement

One of the significant challenges of 3D digitization and 3D modelling lies in striking a balance between surface detail and model size. When digitally reconstructing a high-resolution model, the resulting mesh often contains millions of triangles that capture fine geometric structure with high fidelity. While such dense geometry is valuable for research and archival purposes, it comes at the cost of heavy storage demands and high computational overhead (Niven & Richards, 2017; Özyeşil et al., 2017; Jiménez Fernández-Palacios et al., 2017). It is therefore typical for high-resolution digitized objects and modelled assets to be sub-sampled and simplified into levels of detail (LODs) to meet the requirements of intended applications (Luebke et al., 2003, pp. 3-17). By doing so, however, geometric detail is lost progressively, reducing overall visual fidelity during visualization (see Figure 1).

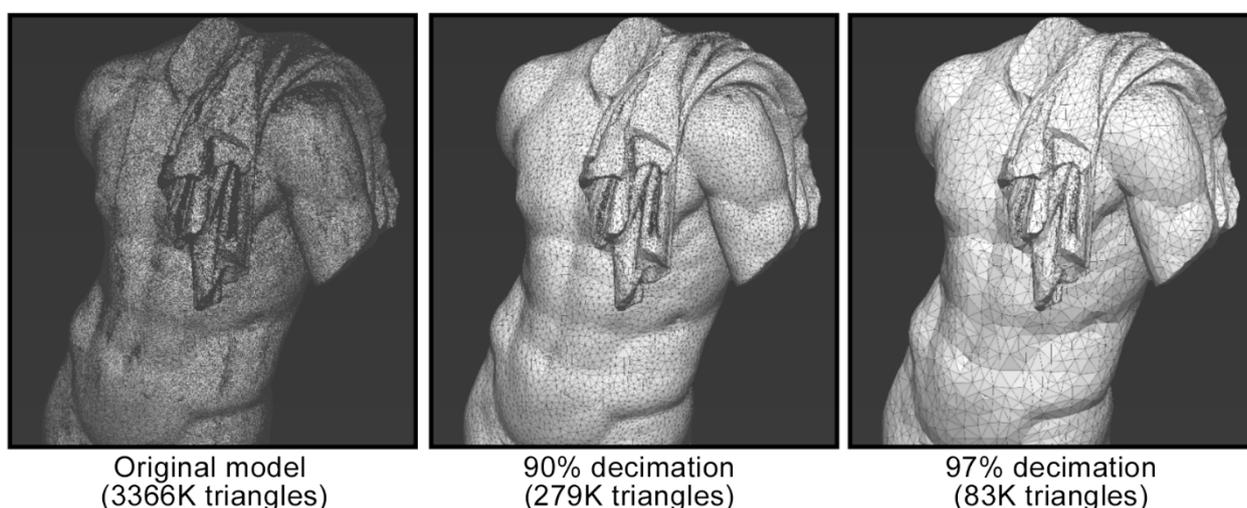


Figure 1 - Progressive mesh simplification. Left: Original high-resolution model (3,366K triangles). Middle: 90% decimation (279K triangles). Right: 97% decimation (83K triangles).

An established computer graphics method for preserving visual fidelity in low-detail versions is the estimation and transfer of the local surface orientation (the “normal” direction) from the high-resolution mesh onto a texture image. The “baked” local surface detail is then applied during rendering to the low-resolution version, simulating fine detail through illumination computation (Theoharis et al., 2013). This technique, known as normal mapping, provides a lightweight way to maintain perceptual detail even after significant simplification. It is therefore very important to investigate the extent to which normal mapping can be applied to simplified geometry, while preserving visual fidelity for visual exploration of the 3D models, since it can significantly reduce data storage, transmission and display requirements.

Background

Normal mapping, briefly introduced in the Problem Statement section, is a technique for the visual enhancement of 3D surfaces, during rendering (Halladay, 2019). The key observation is that when rendering detailed geometry, detail is visually communicated via the interaction of the light with the subtle geometric variations of the surface relief. In turn, the dominant factor affecting local shading is the surface orientation, as determined by the vector pointing directly outwards from the surface (the normal vector), which is directly used in lighting computations and therefore, the apparent coloring of the surface. When rendering any three-dimensional surface, we can artificially substitute the true (or interpolated) normal of the surface with one computed or supplied at runtime. The replacement normal vectors are usually defined in local tangent space of the surface and provided in a pre-computed texture image. Most of the time, the indexing of the texture is performed via a precomputed bijective mapping of surface locations to the (2D) parametric domain of the image (Theoharis et al., 2008, pp. 383-389).

Starting from a high-resolution surface, it is possible to produce simplified versions and transfer the differential geometric detail that is lost to the low-resolution surface via texture mapping. In essence, the lost detail manifests as a deviation in the elevation of the resulting model, registered in a *displacement map*. However, adapting the relief of the simplified mesh during rendering is rather expensive. On the other hand, using only the locally bent normal of the displaced high-resolution original surface without displacing it, we are able to mimic the correct illumination and preserve the appearance of the detailed surface, to a certain extent. It is known that the effect breaks at the silhouettes of objects, where the outline is controlled only by the simplified geometry, or at very oblique viewing angles, where the absence of true geometric detail prevents accurate self-occlusion (the shading and shadowing effects caused when protruding details occlude other nearby parts). Once texture parameterization of the surface of the geometry is established, computing the normal map is straightforward. Rays are cast from a low-resolution object perpendicular to the surface, both outward and inward. These are intersected with the high-definition surface and the normal vector at the closest intersection is registered, expressed in local tangent-space coordinates of the simplified surface and stored as an image, properly encoding Cartesian coordinates as color values (see Figure 2).

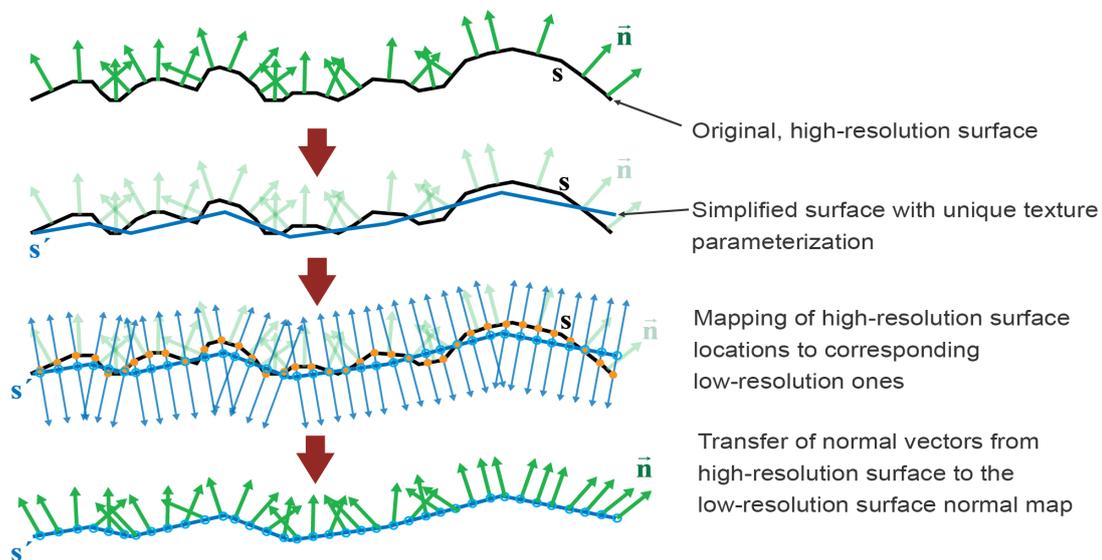


Figure 2 - Transfer of geometric detail from a high-resolution mesh to its low-resolution counterpart. The output is a normal map representing the surface normals of the original

detailed model. In the first step (second row), a simplified version of the geometry is generated. Next, for locations on the simplified mesh, the closest points on the high-resolution mesh are identified and finally, their normal vectors are mapped as replacement direction vectors on a texture to perform normal mapping on the low-resolution model.

In contemporary visualization workflows, normal baking and mapping constitute essential methods for realistic appearance rendering, along with physically-based material modelling and shading. For archaeological visualization, they allow simplified models to retain features such as inscriptions, relief patterns, and subtle material variations that are essential for scholarly interpretation and study. Despite its strengths, normal mapping has limitations. While it convincingly simulates surface relief in most conditions, it cannot reproduce true geometric variation at silhouettes or extreme viewing angles. Furthermore, the absence of true geometric detail means that computational analysis of the surfaces that targets detailed information must also take into account the normal map information, instead of relying only on the raw geometric data. Even so, the method remains an effective compromise between efficiency and fidelity, which makes it particularly valuable for archaeological 3D visualization.

Normal-mapped Model Generation

3D Model Generation

Our proposed workflow begins with 3D acquisition using Structure from Motion (SfM) photogrammetry, employing either ground-based or aerial photography, depending on the scale and accessibility of the subject. The 3D model generation was conducted using the Agisoft Metashape¹ Standard Edition software (Agisoft Metashape, 2024). Textural information was encoded at an 8192² resolution to ensure a detailed representation of the surface's distinct properties and features. The acquisition of the image dataset was performed with a Nikon D3300 DSLR camera (24.2 MP, 6000×4000 resolution) equipped with an 18–55 mm zoom lens. For the test subjects located indoors the acquisition was further complicated by constantly shifting natural lighting, as illumination depended solely on sunlight. The camera settings, including aperture, shutter speed, and ISO, were frequently modified during capture to compensate for these fluctuations. For aerial acquisition of the large monuments, a DJI Mavic 2 Zoom drone equipped with a ½.3" CMOS camera sensor was employed. UAV flights were conducted to provide elevated viewpoints and full coverage of structures that could not be captured effectively from the ground. Due to lighting inconsistencies all captured images were batch post-processed in Adobe Lightroom to correct the variations in exposure and to standardize color balance before photogrammetric reconstruction.

Simplification

Following the creation of a high-polygon mesh and the pose and scale rectification, the original mesh is progressively simplified and smoothed, using the MeshLab software (Callieri et al., 2011), to drastically lower the original model resolution to 3-10% of its original polygon count, to make it compatible with typical requirements of mid-range rendering hardware, browser-based rendering solutions (browser executed JavaScript) and download speeds.

As extensively discussed in standard computer graphics literature, different algorithms for simplification exist, with the choice ultimately depending on the trade-off between accuracy of appearance and processing cost (Akenine-Möller et al., 2018, pp. 706–715). A popular approach is Quadric Edge Collapse decimation (Heckbert & Garland, 1999), a vertex-collapse method that provides controlled reduction of polygon count, optionally followed by Laplacian smoothing (Sorkine

¹ <https://www.agisoft.com>

et al., 2004) to reduce sharp transitions caused by the removal of triangles. It is important in this process to maintain the original texture mapping, if present, to ensure that the UV mapping remains stable during the polygon collapse operations. However, under drastic decimation, warping the texture map correspondences is unavoidable, erroneously merging textured patches, which can cause visible artifacts during the rendering, a risk also associated with topological and geometric degeneracies introduced by aggressive simplification (Attene et al., 2013; Sander et al., 2001). In such cases, a new texture parameterization computation of the simplified object version is advisable. In our test cases, we did not come across such a situation, yet the 3D models must be thoroughly inspected after simplification. For reproducibility, we note that the decimation was performed using MeshLab's default parameter settings, including the default quality threshold (0.3), which we found sufficient for stable behavior across all models. Furthermore, we opted for percentile polygon reduction instead of absolute polygon count (3–10% of the original triangle count), performed iteratively, allowing us to observe the behavior of collapses and evaluate the resulting geometry and storage footprint at each stage. Decimation was stopped when the target percentage was reached and no topological and volumetric distortions were introduced. The stopping criterion was therefore primarily related to the requirement that the model remained compact without introducing collapse-related flaws. The selected simplification ratios were adapted to the scale and surface complexity of each object; whereas large architectural remains tolerate a higher level of geometric reduction, smaller artifacts, where fine detail is important to preserve, require a more conservative approach to decimation. In this process, small-scale details, such as inscriptions and decorative or deterioration details are lost and this is exactly where normal mapping contributes to recover the visual aspect of such intricate geometric structures. With the simplification stage concluded, a careful inspection of the mesh is advised in order to regulate spikes, irregular facets, or any other artifacts that may have been previously introduced. If such issues appear, a single iteration of Laplacian smoothing is advised and is typically sufficient to correct minor irregularities.

Normal map baking

In our 3D model processing pipeline, the resulting low-poly meshes are subsequently imported into Adobe Substance 3D Painter (Adobe, 2024), where a normal-map baking procedure is used to transfer high-resolution surface detail onto the simplified geometry. In the current version of the software (10.1.2), the workflow begins by importing the low-resolution mesh and assigning the corresponding high-resolution model as the “High-Definition Mesh” during the Bake Mesh Maps setup. It is recommended to adjust the ray-casting cage via the Max Frontal Distance and Max Rear Distance parameters in the baking settings, as these values control how far rays are cast along the low-poly surface normal direction to capture details from the high-resolution mesh. Proper configuration of these distances minimizes the risk of over-projection, which may otherwise introduce shading artifacts². Normal map baking assumes a sufficiently smooth and consistent correspondence between the low-resolution and high-resolution meshes; in regions where aggressive simplification alters local topology, ray projection may become ambiguous, potentially introducing shading inconsistencies.

Since the normal information is encoded in tangent space, the baked normal maps remain consistent across all LODs. This step also includes the generation of supporting textures such as albedo, roughness, and metallicity, which collectively preserve the fine visual qualities of the object. Recent work has shown that careful treatment of albedo, roughness, metalness, and normal maps within a PBR framework substantially improves the realism and appearance of generated 3D objects under different lighting conditions, ensuring that textures and geometry remain well aligned (Wang et al., 2024). Care must be exercised when producing relatively low-resolution normal maps (e.g. 1024²

² <https://www.youtube.com/watch?v=hYtHp4IXvsM>

or 2048^2), since geometric detail variations are underrepresented in them, yet their presence affects the statistical distribution of the surface (micro-) facets. In other words, when intentionally averaging normal information, the variance must be carried over to the roughness information of the surface finish, increasing its value (Olano & Baker, 2010). Figure 3 shows the resulting albedo texture and the corresponding baked normal map for a mosaic floor segment.

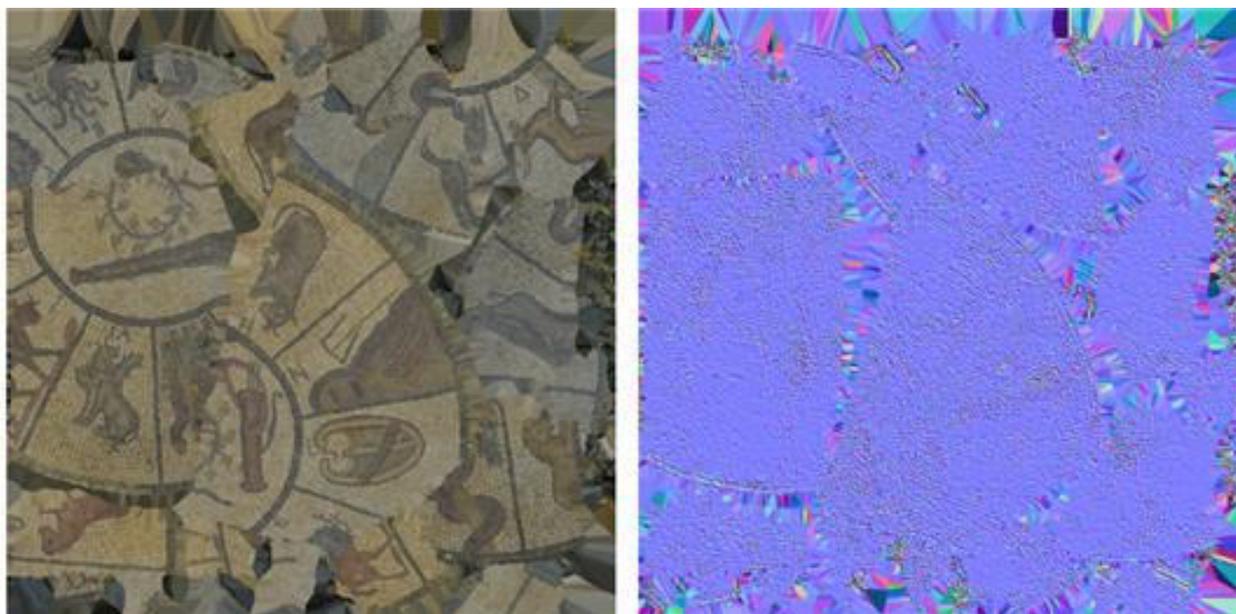


Figure 3 - Diffuse reflectance (albedo) and normal map of a floor mosaic segment.

Illumination pre-processing

Since the final models are intended for inspection focusing on the representation of the details, all baked illumination variations registered on the texture map during acquisition must be decoupled from the albedo and removed. Otherwise, they would a) interfere with the experiment, significantly boosting the impression of local surface relief and b) hinder the application of new lighting conditions by potentially contradicting the intended illumination direction(s). For this purpose a de-lighting pass was conducted using Unity Engine's delighting plugin (available in Unity version 6). By decoupling acquisition-dependent illumination from the albedo, comparability across the selected test cases is further improved, supporting consistent visual evaluation.

Additionally, since roughness information already occupied a separate color channel of an additional texture map, we could afford to also pre-compute and store ambient occlusion information in the same texture map. Ambient occlusion is a statistical measure of the local "openness" of a point on the surface, measuring its ability to gather light from a distant environment (Pharr & Green, 2004, pp. 279-292). During rendering, this information can be used to suppress environmental illumination, e.g. coming from a sky illumination model or environment map, to further accentuate surface details in model cavities. This complementary effect was deemed important in digitized models with significant cavities or deep surface relief patterns (e.g. brickwork), where variations in shadowing strongly influence perception.

Parameters affecting the presence (strength) of normal mapping and ambient occlusion were carefully adjusted to ensure consistency across test cases and approximate the appearance of the

real objects as best as possible. These adjustments were necessarily subjective, based on direct visual comparison between the physical artifacts and their digital counterparts.

Evaluation

Test Cases

The focus of our study was the digitization of material from the Archaeological Museum of Elis, Greece, along with monuments located in the surrounding area. The dataset included both small to medium sized objects and large-scale architectural remains, which were deliberately chosen due to their distinctive characteristics. The aim was to assess how different surface types, from polished marble to rough weathered textures, as well as the presence of cavities and heavy relief geometry, would perform under representation by normal mapping.

The five distinctive test cases used for the evaluation are shown in Figure 4. The first subject is a Roman era marble funerary stela, whose front face is heavily inscribed and bears a protruding sculpted figure dressed in a himation with multiple folds. The inscriptions are shallow and easily eliminated during simplification. The sides and back of the stela have a rough and irregular shape due to weathering and embedding into the soil. The second subject is a Roman- era ruin named “Octagon”, a site that remains largely absent from published scholarship due to the fact that recent archaeological surveys carried out by the local ephorate have yet to be released. The structure consists of surviving masonry with collapsed walls, uneven fragments, and deep crevices across several sections. The construction combines stone and brick, producing surfaces that range from smooth weathered edges to coarse rubble. These irregularities are easily reduced or flattened when the mesh is simplified, particularly in the case of cavities and fractured wall sections. The third subject is a marble statue of the god Hermes (Papathanasopoulos, 1968, p. 133-134). The surface presents smooth muscle transitions and finely carved draped cloth. These delicate features are difficult to preserve during simplification, making the statue a representative case for assessing the treatment of fine, close-range geometry on a medium-sized artifact. The fourth subject of interest is the ancient theater of Elis (Gialouris, 1969, pp. 70-72), which is a monument that consists of partially surviving masonry and architectural remains, spanning a large area. At this large scale, fine geometric details are less critical than the overall form and layout of the structure. The fifth subject is a Roman mosaic floor decorated with scenes from the Twelve Labors of Herakles (Papathanasopoulos, 1968, p.133). The surface is composed of tesserae arranged in intricate patterns with recessed joints that outline the portrayed topics. The combination of vivid coloration, detailed motifs, and shallow relief makes the mosaic a distinctive case for assessing how small patterns are represented.



Figure 4 - The five characteristic case studies used for our evaluation. From left to right: Funerary Stelae of Aurelia, Roman “Octagon” ruin, Hermes of Elis, Ancient Theatre of Elis, mosaic floor segment.

Evaluation Methodology

The simplified meshes served as the basis for testing, while the full-detail versions were retained as ground-truth references. For each test case three model versions were prepared and presented during evaluation: a low-resolution model with only the base color texture present, the same model but with normal maps present and the original high-resolution reconstruction, serving as reference. The rendering of all models was done in Unity Engine (Unity Technologies, 2024) - version 2022.3.29f1. In terms of visualization, we installed consistent illumination setups for indoor and outdoor sets. For the 3D models located inside the museum, local lighting and lightmaps with baked global illumination from off-screen virtual reflectors were used to provide consistent indoor illumination.

The perceptual evaluation was carried out through a questionnaire-based study involving 80 participants with varying levels of familiarity with 3D visualization. The questionnaire was conducted via Google Forms and was divided into two sections. The first section served as a compilation of demographic and experiential data to interpret the responses of the participants and to subdivide them into meaningful categories based on their level of familiarity with 3D visualization, from level 1: "I have no experience evaluating or interacting with visual content" to 5 "I have formal professional training or educational experience in visual or digital media".

In the second section, each participant was presented with the three versions of the test cases, where the order of appearance was randomized to reduce bias. These renderings showcased a simplified mesh with only albedo textures applied, a simplified mesh with normal mapping, and a high-resolution mesh. The process was repeated three times for each model, using a different viewing angle each time, to highlight areas of interest and essential features that effectively showcased each model's uniqueness and visual characteristics. The participants were requested to select the version they perceived as the most realistic and of the highest visual quality. They were also asked to indicate one or more other visual criteria that influenced their choice: texture clarity, color accuracy, detail preservation, shadow and lighting effects, and silhouette accuracy.

Quantitative degradation evaluation

Before conducting the user evaluation, we performed a quantitative evaluation of the error introduced by extreme simplification (removal of 97% of the original geometry). The geometric evaluation of the decimated meshes was performed in MeshLab using the *distance-from-mesh* tool. This method calculates the deviation between a sample set on one mesh and the corresponding closest points on the surface of another model and vice-versa. In Figure 5, we present indicative results for the degradation of two distinctive test cases, the funerary stelae and the Roman "Octagon" building. In the case of the funerary stelae, localized deviations appeared on the garment and frame edges, while inscriptions were particularly vulnerable, collapsing into indistinct surfaces. Maximum surface deviation from the original scan was 0.07mm. In the case of the Roman "Octagon" ruin, cavity interiors and collapsed sections began to accumulate differences. Significant deviations appeared across fractured masonry and rubble zones, flattening complex relief onto indistinct, smooth planar features. The maximum surface deviation from the original scan was 1.14cm. From an archaeological perspective, these geometric deviations may indeed affect the representation of important features such as carved details, surface contours, and recessed or eroded areas in the context of computational geometric analysis, since the relevant algorithms and tools operate on true geometric features (e.g. surface-based measurements and mesh alignment). In the present work, where we evaluate how the perceptual aspects of the digitized geometry are retained, the geometric error acts as a measure of comparison between what is *truly lost* in terms of geometric definition and what is *perceived* as lost; such geometric distortions can be acceptable when the overall surface appearance is preserved, thus motivating the perceptual evaluation that follows, especially considering the relative scale of the object or site.

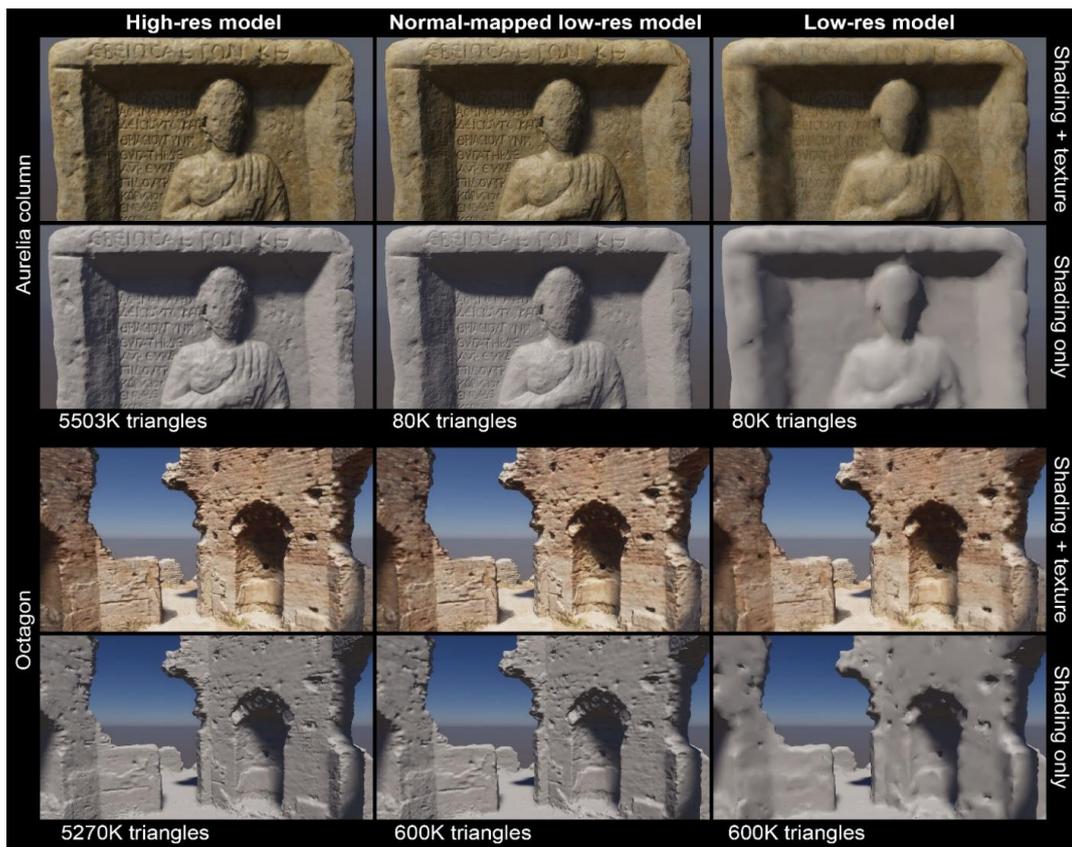


Figure 5 - Visual comparison of the flattening and over-smoothing effects of decimation. When normal maps computed from the hi-resolution meshes are applied, most visual detail is rectified in the rendered 3D model.

Evaluation Results

For the funerary stelae, the normal-mapped version was considered the most convincing by 60.0% of respondents, compared to 36.7% for the high-resolution model and only 3.3% for the plain low-resolution mesh. This strong preference highlights how normal maps were highly effective in reproducing both the inscribed letters and the carved surface details. A similar outcome was observed in the Octagon ruin. Here, 52.5% of participants favored the normal-mapped version, while 30.3% selected the high-resolution mesh and 17.2% chose the simplified mesh without enhancement. This strong preference indicates that normal mapping successfully conveyed surface texture and depth, making it the most efficient yet convincing option for this type of architectural structure. The visual evaluation of the Hermes statue model showed that the high-resolution version was preferred by 43.2% of the participants, closely followed by the version with normal mapping at 39.4%. The low-resolution model received 17.4% of the votes. The high number of votes for the low-resolution version can be justified by the overall smooth appearance of the object, which provides few clues of the decimation. For the fourth case study, the ancient theater site, the results revealed that the simplified model enhanced with normal maps was preferred by 50.0% of participants, while the original high-resolution model received 37.5% of the votes, and the low-resolution model a 12.5%. This preference can be attributed to the large architectural scale of the structure, where geometric detail is less critical and normal maps convincingly convey surface depth and form. The last case is the mosaic floor segment for which the simplified model containing normal maps was once more selected by most participants with 50.0%, while the high-resolution model received 33.8% and the low-resolution

simplified model 16.2%. This preference suggests that normal maps were effective in reproducing the recessed textures and visual depth of the floor variations.

The results are summarized in figures 6 and 7. It is evident that a normal-mapped 3D model with carefully tuned presence of the normal map and the addition of an ambient occlusion contrast enhancement rendering pass can be quite convincing and effective at conveying a sense of detail. The surprising fact is the dominance of the choice of normal-mapped models over the high-resolution ones in the expert users category (see Figure 7). We attribute this to the fact that expert users tend to focus their evaluation on the presence of detail and correct illumination/shadowing effects that stem from a detailed relief, as indicated by the factors presented in Figure 8. Furthermore, across all expertise levels, the most frequently cited factors for the decision to choose a particular version included preservation of detail, texture clarity, and shading, aspects in which normal mapping is particularly effective.

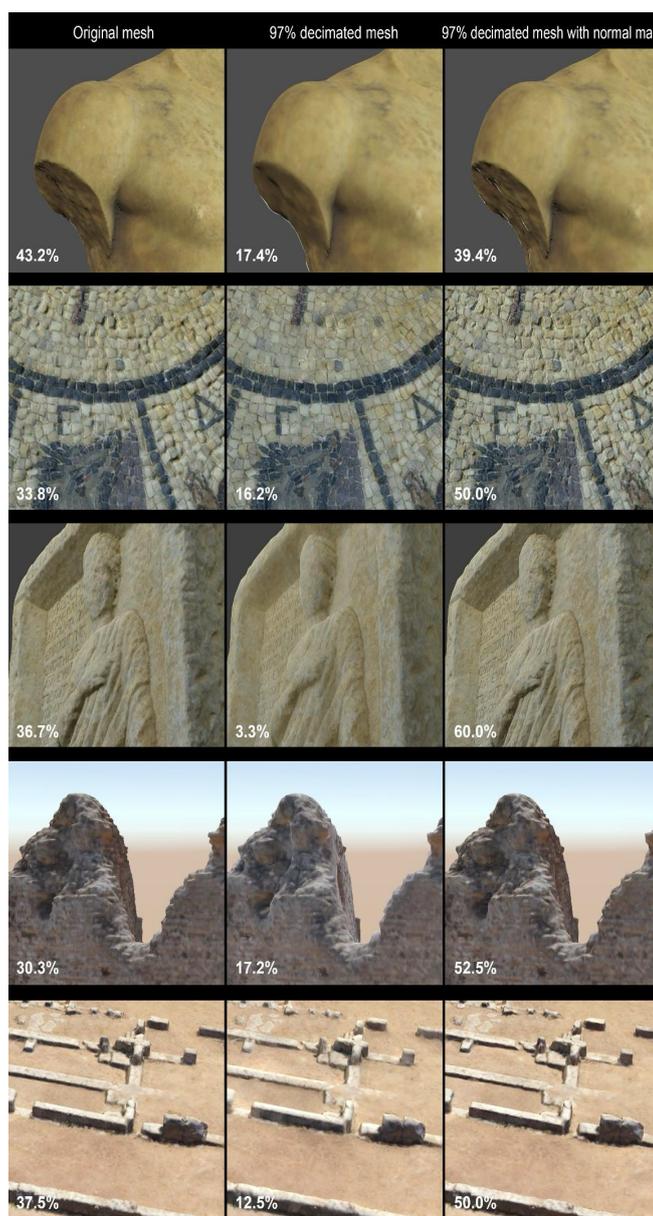


Figure 6 – Average user evaluation results.

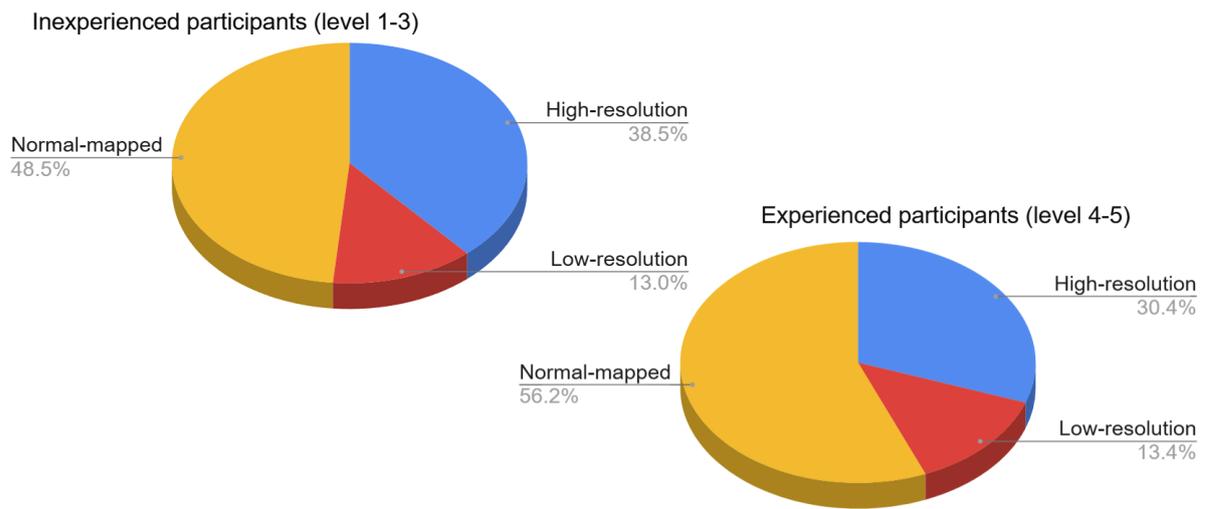


Figure 7 - Average user evaluation results per experience category.

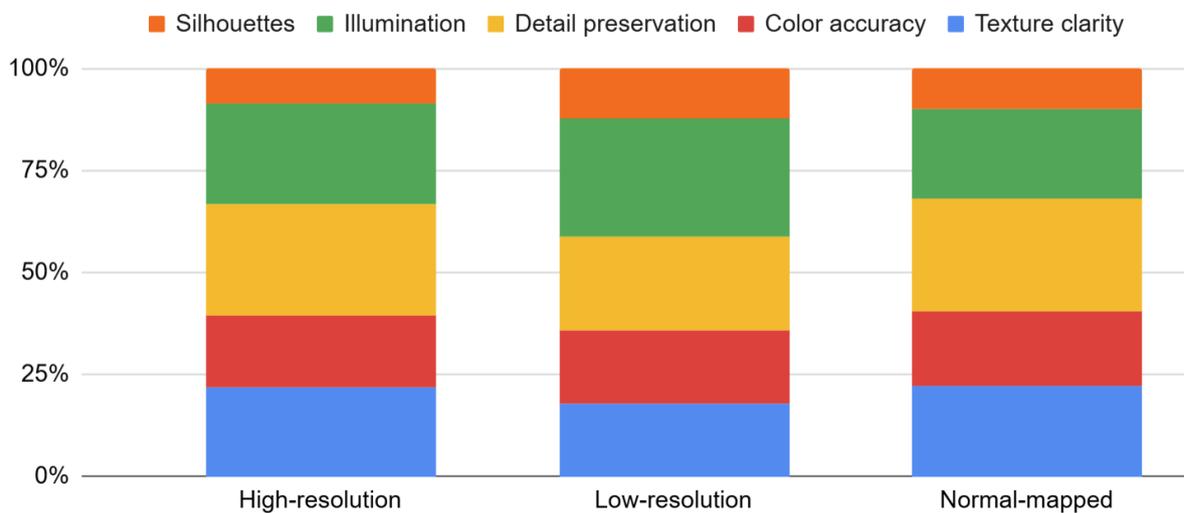


Figure 8 – Factors that affected the choice of the models.

Conclusion

In this work we applied normal mapping to simplified 3D models of archaeological material in order to evaluate its effectiveness in preserving quality and visual similarity, with a particular focus on visualization-oriented archaeological workflows. The objective analysis, carried out on objects of different scales and surface materials, confirmed that geometric reduction introduces measurable deviations, particularly in areas with carvings, inscriptions, or localized surface complexity such as folds and cavities. We further explored how human perception responds in these cases, showing that normal mapping strongly reinforces the impression of preserved detail, even when geometry has been significantly reduced, thereby supporting its use as a practical visualization aid for archaeologists and cultural heritage professionals, particularly in inspection and dissemination contexts rather than in tasks requiring high-precision geometric analysis.

An important finding, which is consistent with the relevant theoretical foundations of reflectance models, is that normal mapping is more impactful on relatively glossy surfaces (funerary stelae, hermes statue, mosaic), since shading produces more high-contrast highlights that are accentuated by shifts of surface orientation. For surfaces with high roughness, normal mapping can be complemented by non-physically-based shading approaches, to accentuate relief and orientation transitions. Similar to normal mapping, the related quantities can be pre-computed on the original, high-resolution meshes and baked on textures mapped to the low-resolution model versions. These shading options include ambient occlusion (Zhukov et al., 1998) and radiance scaling (Vergne et al., 2010). Ambient occlusion can be also efficiently rendered as a screen-space effect, using the normal vectors and depth information, but radiance scaling depends on curvature estimation, which may require careful preprocessing, especially in the case of uneven mesh topology (Andreadis et al., 2015).

Despite its effectiveness for visualization, the proposed approach has inherent limitations. Normal mapping implies the existence of an injective mapping between the mesh surface and the parametric plane, where the normal texture gets stored, so that all points are uniquely mapped to texture locations. However, digitized objects may have no texture parameterization, either because they are devoid of color information or the latter is stored on the point samples. In such cases, an additional texture parameterization (UV unwrapping) step is required, which can be performed by both commercial and open-source modelling and mesh processing tools. Finally, normal maps encode surface orientation rather than true geometric displacement, and therefore simplified and normal-mapped models may lack the descriptiveness required to directly perform precise measurements on them, despite being uncannily visually similar to their high-resolution counterparts.

Acknowledgements

The authors gratefully acknowledge the Ephorate of Antiquities of Elis for granting permission to conduct the scanning campaign that forms the basis of this research.

«©Υπουργείο Πολιτισμού, Εφορεία Αρχαιοτήτων Ηλείας, Αρχαιολογικός Χώρος Ήλιδας/Αρχαιολογικό Μουσείο Ήλιδας - Οργανισμός Διαχείρισης και Ανάπτυξης Πολιτιστικών Πόρων»,

This work was partially supported by the ARTEMIS EU project - Applying Reactive Twins to Enhance Monument Information Systems. (2025-2027). HORIZON-INFRA-2024-TECH-01-04, Project No. 101188009.

Preprint version 4 of this article has been peer-reviewed and recommended by Peer Community In Archaeology (<https://doi.org/10.24072/pci.archaeo.100627>; Pavlidis, 2026)

Funding

The authors received no financial support for this research.

Conflict of interest disclosure

The authors declare that they comply with the PCI rule of having no financial conflicts of interest in relation to the content of the article.

References

Adamopoulos, E., Rinaudo, F., & Ardissono, L. (2021). A Critical Comparison of 3D Digitization Techniques for Heritage Objects. *ISPRS International Journal of Geo Information*, 10(1), p. 10. <https://doi.org/10.3390/ijgi10010010>.

- Adobe. (2024). *Substance 3D Painter* (10.1.2) [Software]. Adobe Systems.
- Agisoft Metashape. (2024). Agisoft Standard Edition Version 2.1 [Software]. Agisoft LLC
- Aicardi, I., Chiabrando, F., Lingua, A. M., & Noardo, F. (2018). Recent trends in cultural heritage 3D survey: The photogrammetric computer vision approach. *Journal of Cultural Heritage*, 32, pp. 257–266. <https://doi.org/10.1016/j.culher.2017.11.006>.
- Akenine-Möller, T., Haines, E., Hoffman, N., Pesce, A., Iwanicki, M., & Hillaire, S. (2018). *Real-Time Rendering* (4th ed.). A K Peters/CRC Press. <https://doi.org/10.1201/b22086>.
- Andreadis, A., Papaioannou, G., & Mavridis, P. (2015). A Parametric Space Approach to the Computation of Multi-scale Geometric Features. In *Proceedings of the 10th International Conference on Computer Graphics Theory and Applications*, pp. 5–15. <https://doi.org/10.5220/0005225700050015>.
- ARTEMIS EU project (2025-2027) Applying Reactive Twins to Enhance Monument Information Systems. HORIZON-INFRA-2024-TECH-01-04, Project No. 101188009, <https://www.artemis-twin.eu/>.
- Attene, M., Campen, M., & Kobbelt, L. (2013). Polygon mesh repairing: An application perspective. *ACM Computing Surveys*, 45(2), pp. 1–33. <https://doi.org/10.1145/2431211.2431214>.
- Callieri, M., Dellepiane, M., Ranzuglia, G., Cignoni, P., & Scopigno, R. (2011). *MeshLab as a complete open tool for the integration of photos and color with high-resolution 3D geometry data*. <https://api.semanticscholar.org/CorpusID:14519487>.
- Gialouris, N. (1969). Anaskaphi Ilios hypo Nikolaou Gialouri [Excavation of Elis by Nikolaos Gialouris. *Praktika tes en Athenais Archaologikes Etaireias*, 124, pp. 70–72. <https://ir.lib.uth.gr/xmlui/handle/11615/17932>.
- Graham, C. A., Akoglu, K. G., Lassen, A. W., & Simon, S. (2017). Epic dimensions: A comparative analysis of 3D acquisition methods. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLII-2/W5, pp. 287–293. <https://doi.org/10.5194/isprs-archives-XLII-2-W5-287-2017>.
- Halladay, K. (2019). Normal Mapping. In *Practical Shader Development*, pp. 181–200. Apress. https://doi.org/10.1007/978-1-4842-4457-9_10.
- Hartley, R., & Zisserman, A. (2003). *Multiple view geometry in computer vision* (2nd ed.). Cambridge University Press. <https://doi.org/10.1017/CBO9780511811685>.
- Heckbert, P. S., & Garland, M. (1999). Optimal triangulation and quadric-based surface simplification. *Computational Geometry: Theory and Application*, 14(1-3), pp. 49-65. [https://doi.org/10.1016/S0925-7721\(99\)00030-9](https://doi.org/10.1016/S0925-7721(99)00030-9).
- Jiménez Fernández-Palacios, B., Morabito, D., & Remondino, F. (2017). Access to complex reality-based 3D models using virtual reality solutions. *Journal of Cultural Heritage*, 23, pp. 40–48. <https://doi.org/10.1016/j.culher.2016.09.003>.
- Kerbl, B., Kopanas, G., Leimkuehler, T., & Drettakis, G. (2023). 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics*, 42(4), pp. 1–14. <https://doi.org/10.1145/3592433>.
- Looten, J., Gravina, B., Muth, X., Villaeys, M., & Bordes, J.-G. (2025). Towards a more robust representation of lithic industries in archaeology: A critical review of traditional approaches and modern techniques. *Peer Community Journal*, 5, e52. <https://doi.org/10.24072/pcjournal.562>.
- Luebke, D., Reddy, M., Cohen, J. D., Varshney, A., Watson, B., & Huebner, R. (2003). Chapter 10—Temporal Detail. In D. Luebke, M. Reddy, J. D. Cohen, A. Varshney, B. Watson, & R. Huebner (Eds.), *Level of Detail for 3D Graphics*, pp. 301–xvii. Morgan Kaufmann. <https://doi.org/10.1016/B978-155860838-2/50012-1>.
- Manferdini, A. M., & Russo, M. (2013). Multi-scalar 3D digitization of Cultural Heritage using a low-cost integrated approach. In *Proceedings of the 2013 Digital Heritage International Congress (Digital Heritage)*, 1, pp. 153-160. <http://dx.doi.org/10.1109/DigitalHeritage.2013.6743725>.

- Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., & Ng, R. (2020). NeRF: Representing scenes as neural radiance fields for view synthesis. In *Proceedings of the European Conference on Computer Vision (ECCV 2020)*, pp. 405–421. https://doi.org/10.1007/978-3-030-58452-8_24.
- Niven, K., & Richards, J. D. (2017). Chapter 12—The Storage and Long-Term Preservation of 3D Data. In D. Errickson & T. Thompson (Eds.), *Human Remains: Another Dimension*, pp. 175–184. Academic Press. <https://doi.org/10.1016/B978-0-12-804602-9.00013-8>.
- Oechsle, M., Peng, S., & Geiger, A. (2021). UNISURF: Unifying Neural Implicit Surfaces and Radiance Fields for Multi-View Reconstruction. In *Proceedings of the 2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 5569–5579. <https://doi.org/10.1109/ICCV48922.2021.00554>.
- Olano, M., & Baker, D. (2010). LEAN mapping. In *Proceedings of the 2010 ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, pp. 181–188. <https://doi.org/10.1145/1730804.1730834>.
- Özyeşil, O., Voroninski, V., Basri, R., & Singer, A. (2017). A survey of structure from motion. *Acta Numerica*, 26, pp. 305–364. <https://doi.org/10.1017/S096249291700006X>.
- Papaioannou, G., Schreck, T., Andreadis, A., Mavridis, P., Gregor, R., Sipiran, I., & Vardis, K. (2017). From Reassembly to Object Completion: A Complete Systems Pipeline. *Journal on Computing and Cultural Heritage*, 10(2), pp. 1–22. <https://doi.org/10.1145/3009905>.
- Papathanasopoulos, G. (1968). Anaskaphai eis Archaian Elida [Excavations at ancient Elis]. *Archaiologika Analekta Ex Athenon [Athens Annals of Archaeology]*, 1(2), pp. 128–135. https://www.aegeussociety.org/wp-content/uploads/2022/04/AAA_001_02.pdf.
- Pavlidis, G., Koutsoudis, A., Arnaoutoglou, F., Tsioukas, V., & Chamzas, C. (2007). Methods for 3D digitization of Cultural Heritage. *Journal of Cultural Heritage*, 8(1), pp. 93–98. <https://doi.org/10.1016/j.culher.2006.10.007>.
- Pavlidis, G. (2026). A robust method for simplifying 3D models while keeping high fidelity visualization of archaeological artifacts. *Peer Community in Archaeology*, 100627. <https://doi.org/10.24072/pci.archaeo.100627>.
- Pharr, M., & Green, S. (2004). *Chapter 17. Ambient Occlusion*. NVIDIA Developer. <https://developer.nvidia.com/gpugems/gpugems/part-iii-materials/chapter-17-ambient-occlusion>.
- Remondino, F., & Rizzi, A. (2010). Reality-based 3D documentation of natural and cultural heritage sites—techniques, problems, and examples. *Applied Geomatics*, 2, pp. 85–100. <https://doi.org/10.1007/s12518-010-0025-x>.
- Sander, P. V., Snyder, J., Gortler, S. J., & Hoppe, H. (2001). Texture mapping progressive meshes. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, pp. 409–416. <https://doi.org/10.1145/383259.383307>.
- Sorkine, O., Cohen-Or, D., Lipman, Y., Alexa, M., Rössl, C., & Seidel, H.-P. (2004). Laplacian surface editing. In *Proceedings of the 2004 Eurographics/ACM SIGGRAPH Symposium on Geometry Processing*, pp. 175–184. <https://doi.org/10.1145/1057432.1057456>.
- Storeide, M. S. B., George, S., Sole, A., & Hardeberg, J. Y. (2023). Standardization of digitized heritage: A review of implementations of 3D in cultural heritage. *Heritage Science*, 11(1), p. 249. <https://doi.org/10.1186/s40494-023-01079-z>.
- Theoharis, T., Papaioannou, G., Bjørlykke, K., Ekroll, Ø., Rieke-Zapp, D., Lein, A., Andreadis, A. (2013). Multi-scale 3D Digitization at Nidaros Cathedral: from archiving to large-scale visualization. In *Proceedings of the 18th Conference on Cultural Heritage and New Technologies*, pp. 1–16. https://archiv.chnt.at/wp-content/uploads/Theoharis_etal_final_2014.pdf.
- Theoharis, T., Papaioannou, G., Platis, N., & Patrikalakis, N. M. (2008). *Graphics and Visualization: Principles & Algorithms (1st ed.)*. A K Peters/CRC Press. <https://doi.org/10.1201/b10676>.
- Unity Technologies. (2024). Unity Engine version 2022.3.29f1 [Software]. Unity Technologies.

- Vergne, R., Pacanowski, R., Barla, P., Granier, X., & Schlick, C. (2010). Radiance Scaling for versatile surface enhancement. In *Proceedings of the 2010 ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, pp. 143–150. <https://doi.org/10.1145/1730804.1730827>.
- Wang, Y., Xu, X., Ma, L., Wang, H., & Dai, B. (2024). Boosting 3D object generation through PBR materials. In *SIGGRAPH Asia 2024 Conference Papers*, pp. 1–11. <https://doi.org/10.1145/3680528.3687676>.
- Yu, Z., Sattler, T., & Geiger, A. (2024). Gaussian Opacity Fields: Efficient Adaptive Surface Reconstruction in Unbounded Scenes. *ACM Transactions on Graphics*, 43(6), pp. 1–13. <https://doi.org/10.1145/3687937>.
- Zhukov, S., Iones, A., & Kronin, G. (1998). An ambient light illumination model. *Rendering Techniques '98*, pp. 45–55. https://doi.org/10.1007/978-3-7091-6453-2_5.