Enhancing Transparent Visualization of Cultural Heritage Point Clouds through Adaptive Opacity Control

Hatsuki Nojiri[1], Satoshi Takatori[1], Liang Li[1], Motoaki Adachi[2], R. Michael Feener[3], Satoshi Tanaka[1]

[1] College of Information Science and Engineering, Ritsumeikan University, Japan
[2] Shrewd Design Co., Ltd., Japan
[3] Center for Southeast Asian Studies, Kyoto University, Japan
stanaka@is.ritsumei.ac.jp

Keywords: Cultural heritage, 3D scanned point clouds, Transparent visualization, Adaptive opacity control, Region highlighting.

Abstract

Three-dimensional (3D) scanning is widely used to preserve cultural heritage as large-scale point clouds. While these datasets contain rich geometric information, transparent visualization of such massive point clouds often suffers from visual clutter and reduced clarity, particularly when both external and internal structures are involved. Previous work resolved the problem of normal orientations, laying the foundation for robust shading in transparent visualization. Building on this foundation, this paper introduces a novel method of adaptive opacity control for region highlighting, which interprets shading as a distribution of opacity. By adjusting the lighting direction, effective opacity can be locally controlled without modifying the original point cloud data. This mechanism enables selective highlighting of user-specified regions, enhances the visibility of complex structures, while also allowing interactive dynamic shading by continuously changing the lighting direction. The effectiveness of the proposed method is demonstrated using culturally significant heritage point clouds, including UNESCO World Heritage sites, where intricate internal structures can be more clearly analyzed. Beyond cultural heritage, the proposed method is also applicable to modern architectural and other large-scale 3D scanned objects with similarly complex forms.

1. Introduction

Three-dimensional (3D) scanning is widely used to preserve tangible cultural heritage as digital data. The resulting large-scale 3D point clouds capture detailed geometric information, supporting restoration planning and research in archaeology, architecture, and art history. However, when the objective is analytical rather than photorealistic, conventional visualization has clear limitations. Understanding internal structures and local features requires non-photorealistic analytic visualization. Such point clouds often contain hundreds of millions of points, and effective analytic visualization remains challenging.

The objective of this study is to develop a method of transparent visualization for the analysis of 3D scanned point clouds. Such point clouds often record both the external and internal structures of the target cultural heritage object. Therefore, transparent visualization, which enables simultaneous observation and analysis of both, is highly effective. When visualizing complex 3D structures in a transparent manner, however, the rendered images often become cluttered due to overlapping elements, resulting in reduced clarity. One effective approach to alleviate this problem is to introduce shading effects by adjusting the position of the light source. Shading emphasizes the boundaries of structural elements, thereby improving the visibility of complex shapes.

However, a key difficulty is that raw 3D scanned point clouds usually do not contain normal vector information required for shading. A widely used approach is to estimate local normal vectors using Principal Component Analysis (PCA), but ambiguity in their orientation can cause rendering artifacts. Determining robust normal orientations remains challenging.

A recent study resolved the problem of normal orientations in transparent visualization (Nojiri et al., 2025), laying the foundation for further development. Building on this foundation, the novel contribution of this paper is the proposal of enhancing transparent visualization through adaptive opacity control for massive point clouds, especially those of cultural heritage. The key idea is that shading in transparent visualization is not only perceived as surface shading, but also as a distribution of opacity. By adjusting the direction and intensity of the light source, the shaded brightness changes, which effectively corresponds to a local change in opacity. This mechanism enables adaptive opacity control without modifying the original point cloud data, allowing users to selectively highlight regions of interest. As a result, specific structures within complex cultural heritage data can be more clearly revealed, providing analytical insights that were previously difficult to obtain.

To demonstrate the contributions described above, the proposed method is applied to tangible cultural heritage objects of high value, including UNESCO World Heritage sites. These case studies confirm that the method is particularly effective for analyzing cultural heritage with complex internal and external structures. Such characteristics are especially important for World Heritage sites and other culturally significant monuments, which often exhibit large-scale, structurally complex 3D forms. For their non-destructive analysis and long-term preservation, visualization techniques that selectively highlight regions of interest without modifying the original digital record are particularly valuable.

While the focus of this study is cultural heritage, the method is not limited to this domain; it can also be applied to modern architectural structures and other large-scale 3D scanned objects with similarly complex forms.

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-XI-4-2026-427-2026

2. Related Work

A common and practical approach for visualizing 3D scanned point clouds is point-based rendering, which directly renders the points themselves rather than reconstructing continuous or polygonal surfaces. When used for opaque visualization, that is, photorealistic visualization, such point-based rendering techniques provide highly effective and visually detailed representations, and they have been extensively applied to the digital documentation and visualization of cultural heritage objects (El-Hakim et al., 2005, Guidi et al., 2005, Ikeuchi et al., 2007, Yastikli, 2007, Laycock et al., 2008, Remondino et al., 2009, Koller et al., 2009, Dylla et al., 2009, Kersten et al., 2012, Aicardia et al., 2018, Discher et al., 2018, Thiel et al., 2018, Kawato et al., 2021).

Recently, image-based neural rendering techniques have attracted considerable attention as new approaches for photorealistic visualization of real-world scenes. Representative methods include Neural Radiance Fields (NeRF) (Mildenhall et al., 2020) and 3D Gaussian Splatting (3DGS) (Kerbl et al., 2023), which synthesize novel views from multi-view photographs. These approaches differ fundamentally from point-based rendering methods that explicitly handle and visualize geometric 3D models. While NeRF and 3DGS produce highly realistic images, they do not explicitly reconstruct underlying geometry and are therefore less suitable for applications requiring accurate visualization of internal or volumetric structures.

Regarding transparent visualization of massive 3D scanned point clouds, the effectiveness of stochastic algorithms has recently been demonstrated, and such methods have also been applied to large-scale cultural heritage sites (Tanaka et al., 2016, Uchida et al., 2020, Pan et al., 2021). However, these approaches have not examined the role of shading in transparent visualization, nor have they explored improvements in transparency through adaptive opacity control. Recent studies have also extended neural rendering methods to transparent or semi-transparent materials (Huang et al., 2025, Li et al., 2025), but their application to the transparent visualization of opaque objects, as pursued in this study, remains challenging.

Transparent visualization is a valuable tool for the visual analysis of complex 3D scanned objects that contain internal structures. However, conventional visualization methods for 3D scanned point clouds have not provided the ability to selectively highlight user-specified regions within the data. The novelty of the present work lies in using shading for highlighting user-specified regions, which requires estimating normal vectors at each point.

To realize shading-based highlighting, a normal vector must be estimated for each point. Since raw point clouds do not inherently contain surface orientation information, various normal estimation methods have been proposed. Among them, PCA-based approaches are widely used (Hoppe et al., 1992, Pauly et al., 2003, Mitra and Nguyen, 2003). However, ambiguity in normal orientation remains a challenge for complex 3D scanned point clouds. Recently, Nojiri et al. (Nojiri et al., 2025) addressed this ambiguity in transparent visualization, enabling stable shading.

In contrast to these conventional approaches, the proposed method leverages shading not only for surface perception but as a mechanism for adaptive opacity control, enabling selective region highlighting without modifying point density or the original data representation.

3. Methods Related to Transparent Visualization Used in this Study

This section reviews the techniques underlying the proposed method.

3.1 Stochastic Point-Based Rendering (SPBR) for Opacity Control

Stochastic point-based rendering (SPBR) is a precise transparent visualization technique widely used for massive 3D scanned point clouds (Tanaka et al., 2016, Uchida et al., 2020). Unlike traditional approaches that require depth sorting of all points, SPBR avoids expensive sorting by randomly dividing the input point cloud into multiple subsets. Each subset is rendered independently, and the final transparent image is obtained by averaging the results, which statistically reproduces correct transparency.

An important feature of SPBR is that it enables control of local surface opacity by adjusting point density. For a local surface region S, the opacity α is defined as:

$$\alpha = 1 - \left(1 - \frac{s}{S}\right)^{n/L_{\mathrm{R}}}, \tag{1}$$

where s is the cross-sectional area of a point (adjusted so that one point overlaps one pixel in the image plane), n is the number of points in region S after random downsampling (thinning) or upsampling (duplication), and L_{R} is the number of subsets used for averaging. This formulation enables local opacity control by regulating point density.

SPBR has been successfully applied to large-scale cultural heritage sites, including Borobudur Temple and the Gion Festival floats (Pan et al., 2021, Tanaka et al., 2016), providing the technical foundation for the present work.

3.2 Normal Orientation Determination with PCA

Since raw 3D scanned point clouds do not include normal vectors, they must be estimated for shading. A widely used approach applies PCA to the local structure of the point cloud. Specifically, for each point, a spherical neighborhood is extracted, and the covariance matrix of the point coordinates within this neighborhood is computed. The eigenvalues and eigenvectors of this covariance matrix are then obtained, and the eigenvector corresponding to the smallest eigenvalue is normalized and adopted as the normal orientation of the point.

However, the sign of the obtained normal vector is ambiguous. When such normals are directly used for shading, artifacts occur due to inconsistent normal signs.

A recently proposed solution (Nojiri et al., 2025) is to exploit the characteristics of transparent visualization. In transparent rendering, both the front and back surfaces are simultaneously visible, regardless of their geometric orientation or occlusion. Thus, all visible surfaces can be treated as front-facing and consistently shaded. Based on this idea, every PCA-derived normal vector is dynamically aligned toward the viewpoint. Let $\mathbf{n}$ denote the surface normal vector at a point $\mathbf{p}$, and $\mathbf{e}$ the position vector of the viewpoint (see Fig.1). The sign of the normal vector is dynamically determined such that

$$\mathbf{n} \cdot (\mathbf{e} - \mathbf{p}) > 0 \tag{2}$$

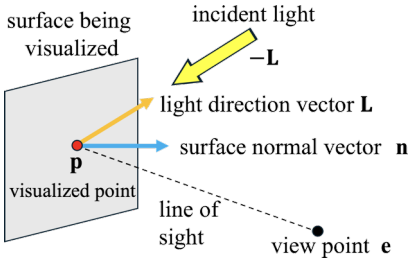

Figure 1. Vector notations used in this paper, showing the surface normal vector $\mathbf{n}$, the light direction vector $\mathbf{L}$, the viewpoint $\mathbf{e}$, and the visualized point $\mathbf{p}$.

is satisfied.

In other words, the normal vector is flipped, if necessary, so that it always faces the viewpoint. This viewpoint-based alignment resolves the sign ambiguity of normal vectors and enables robust shading in transparent visualization of point clouds.

3.3 Effect of Shading in Transparent Visualization

Here, we explain the shading model adopted in this study. The shaded color C_{shaded} is derived from the original point color C_{org}. Each color component (R, G, B) is normalized within the range $[0, 1]$.

In our study, we adopt the following simplified Lambert shading model:

$$C_{\mathrm{shaded}} = k_{\mathrm{d}}\ \max(0, \mathbf{L}\cdot\mathbf{n})\, C_{\mathrm{org}} + k_{\mathrm{a}}\, C_{\mathrm{org}}\ . \tag{3}$$

In this formula, $\mathbf{L}$ denotes the normalized light direction vector, defined as a unit vector pointing toward the light source, and $\mathbf{n}$ denotes the unit surface normal vector (see Fig.1). Note that the direction of the incident light is given by $-\mathbf{L}$, as illustrated in Fig.1. The parameters k_{d} and k_{a} denote the diffuse and ambient reflection coefficients, respectively, and both are assumed to take values within the range $[0, 1]$. If any color component of the right-hand side of Eq.(3) exceeds 1, it is truncated to 1. The factor $\max(0, \mathbf{L}\cdot\mathbf{n})$ ensures that when $\mathbf{L}\cdot\mathbf{n} \geq 0$ (acute angle between $\mathbf{L}$ and $\mathbf{n}$), the value is $\mathbf{L}\cdot\mathbf{n}$, and when $\mathbf{L}\cdot\mathbf{n} < 0$ (obtuse angle), the value is set to zero.

The Lambert shading model (3) indicates that the R, G, and B components of the shaded color C_{shaded} increase as $\mathbf{n}$ becomes more parallel to $\mathbf{L}$, and decrease as $\mathbf{n}$ becomes more perpendicular to $\mathbf{L}$. Consequently, a non-uniform brightness distribution emerges on the point-based surface, which can be controlled by adjusting the light direction vector $\mathbf{L}$. This control of brightness becomes most apparent when k_{a}, that is, the ambient term $k_{\mathrm{a}}C_{\mathrm{org}}$ in Eq.(3), is set to zero. In practice, however, k_{a} should remain nonzero to ensure a minimum level of brightness in the visualized image. Nevertheless, by making it sufficiently small, the intended control effect can be preserved.

The Lambert shading model (3) effectively enhances the perception of surface undulations in transparent visualization, as in conventional opaque rendering. Figure 2 illustrates this effect using the point cloud of a wall relief from Borobudur Temple, a UNESCO World Heritage site in Indonesia, visualized by SPBR (Pan et al., 2021). For shapes that are nearly planar, such as this dataset, normal vectors can be assigned in a straightforward manner. The upper image presents a photorealistic rendering of the entire relief panel, while the two lower images show

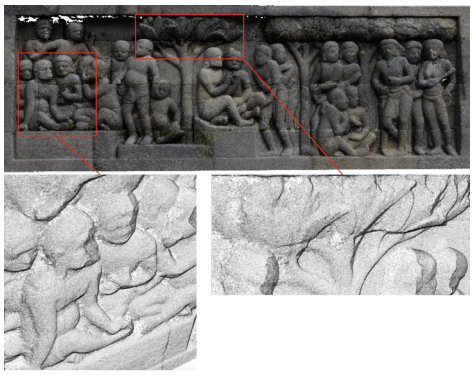

Figure 2. Photorealistic visualization (top) and transparent visualization (bottom, rendered by SPBR) of a wall relief from Borobudur Temple, a UNESCO World Heritage site in Indonesia, showing how shading enhances the perception of surface undulations.

close-up transparent visualizations from an oblique viewpoint. With shading, the subtle undulations of the relief are emphasized, and the analytical visibility is clearly improved.

4. Enhancing Transparent Visualization with Adaptive Opacity Control

Building on the methods reviewed in Section 3, we develop a theoretical framework for adaptive opacity control in transparent visualization of 3D scanned point clouds. The key idea is that the brightness distribution generated by shading can be interpreted as a distribution of effective opacity. This concept provides the theoretical basis for the adaptive opacity control method presented in Section 5.

4.1 Formulation of Adaptive Opacity Control (Black Background Case)

When shading is introduced into visualization, a brightness distribution appears on the point-based surface being visualized. In the rendering of opaque objects, this brightness distribution is simply perceived as surface shading. In transparent visualization, however, the situation differs: in the visualized image, regions of higher brightness appear more opaque, whereas regions of lower brightness appear more transparent. Thus, the brightness distribution is perceived by the observer as a distribution of effective opacity. In other words, by manipulating the lighting direction, shading effectively functions as a mechanism for controlling this perceived opacity. We call this concept *effective opacity*.

For simplicity, we consider the transparent visualization of a single isolated surface. Let C_{transp} denote the rendered color in the transparent visualization result, and C_{bg} the background color. Then C_{transp} can be expressed as

$$C_{\mathrm{transp}} = \alpha C_{\mathrm{shaded}} + (1-\alpha) C_{\mathrm{bg}}\ , \tag{4}$$

where α represents the user-defined surface opacity, as defined by the opacity formula (1).

Here, we focus on the case of a black background ($C_{\mathrm{bg}} = \mathbf{0} = (0, 0, 0)$), which represents the basic form of the proposed method. Under this condition, Eq.(4) reduces to

$$C_{\mathrm{transp}} = \alpha C_{\mathrm{shaded}}\ . \tag{5}$$

This simplified formulation reveals an essential property of the proposed transparent visualization method: the final transparent color C_{transp} depends solely on the product of the opacity α and the shaded color C_{shaded}. Consequently, a change in C_{shaded} and a change in α have equivalent effects on the resulting brightness. This equivalence implies that modifying the lighting direction effectively results in a change of effective opacity, explaining why shading control can be regarded as opacity control in transparent visualization.

As an illustrative example, Table 1 summarizes the results of the formulation under a black background ($C_{\mathrm{bg}} = \mathbf{0}$), using representative parameter values ($\alpha = 0.5$, $k_{\mathrm{d}} = 1$, and $k_{\mathrm{a}} = 0$). The table lists the corresponding values of C_{shaded} and C_{transp} for several typical configurations of the light direction vector $\mathbf{L}$ relative to the surface normal vector $\mathbf{n}$. When $\mathbf{L}$ is parallel to $\mathbf{n}$, the original surface color C_{org} is fully preserved in C_{shaded}, and C_{transp} becomes half of C_{org}, representing the highest effective opacity. As $\mathbf{L}$ rotates away from $\mathbf{n}$, both C_{shaded} and C_{transp} decrease proportionally, and they eventually converge to the background color ($C_{\mathrm{bg}} = \mathbf{0}$) when $\mathbf{L}$ is orthogonal to $\mathbf{n}$. This simple example clearly demonstrates that, under a black background, changing the light direction vector $\mathbf{L}$ directly controls the effective opacity in transparent visualization.

Table 1. Representative values of C_{shaded} and C_{transp} under typical lighting directions $\mathbf{L}$,with $C_{\mathrm{bg}} = \mathbf{0}$ (black background), assuming $\alpha = 0.5$, $k_{\mathrm{d}} = 1$, and $k_{\mathrm{a}} = 0$.

Condition	C_{shaded}	C_{transp}
$\mathbf{L} \mathbin{/\!/} \mathbf{n}$	C_{org}	$0.5C_{\mathrm{org}}$
$\mathbf{L} \cdot \mathbf{n} = 0.5$	$0.5C_{\mathrm{org}}$	$0.25C_{\mathrm{org}}$
$\mathbf{L} \perp \mathbf{n}$	$\mathbf{0}$ $(= C_{\mathrm{bg}})$	$\mathbf{0}$ $(= C_{\mathrm{bg}})$

4.2 Extension to General Background Case

We now extend the above formulation to accommodate a general, non-black background color. The essential idea of this extension is to modify the Lambert shading model in Eq.(3) so that the background color C_{bg} contributes to the shaded result:

$$C_{\mathrm{shaded}} = KC_{\mathrm{org}} + (1 - K)C_{\mathrm{bg}} \,, \tag{6}$$

where K is defined as

$$K = k_{\mathrm{d}} \max(0, \mathbf{L} \cdot \mathbf{n}) + k_{\mathrm{a}} \,. \tag{7}$$

In practice, K is clamped to the range $[0, 1]$. The extended shading model (6) naturally reduces to the basic form of Eq.(3) when the background is black, i.e., $C_{\mathrm{bg}} = \mathbf{0}$.

Unlike the black background case, the background color C_{bg} always contributes to the resulting color, so that the visualized surface appears as a blend of the original surface color C_{org} and the selected background color C_{bg}. Nevertheless, the fundamental principle remains unchanged: the effective opacity can still be controlled by adjusting the light direction vector $\mathbf{L}$.

When $\mathbf{L}$ and $\mathbf{n}$ are parallel, the color components of C_{shaded} reach their maximum values according to the extended shading model in Eq.(6). In this case, the color recorded in the 3D scan, C_{org}, is preserved with minimal loss of brightness. Consequently, C_{transp}, rendered by SPBR according to Eq.(4), receives the maximum contribution from C_{org}. This condition corresponds to the highest effective opacity applied to the visualized surface, so that the surface region appears most strongly highlighted.

On the other hand, when $\mathbf{L}$ and $\mathbf{n}$ are orthogonal, K in Eq.(7) reaches its minimum (near-zero) value. In this case, Eq.(6) approximately reduces to $C_{\mathrm{shaded}} \simeq C_{\mathrm{bg}}$, effectively corresponding to the lowest opacity (i.e., the highest transparency), where the background becomes visible through the visualized surface, resulting in the weakest highlighting.

We summarize representative examples in Table 2 for a parameter setting of $\alpha = 0.5$, $k_{\mathrm{d}} = 1$, and $k_{\mathrm{a}} = 0$. This parameter configuration represents an extreme setting in which the effect of adaptive opacity control is most clearly manifested. When $\mathbf{L}$ and $\mathbf{n}$ are parallel, $K = 1$ holds, and the original surface color C_{org} is perfectly preserved as C_{shaded}, that is, $C_{\mathrm{shaded}} = C_{\mathrm{org}}$. Consequently, C_{transp} receives the maximum contribution from C_{org}, which corresponds to the highest effective opacity applied to the visualized surface. As $\mathbf{L}$ rotates toward orthogonality with $\mathbf{n}$, K decreases accordingly. In this case, Eq.(6) indicates that the contribution of C_{org} to both C_{shaded} and C_{transp} gradually decreases, while the influence of C_{bg} correspondingly increases. In the limiting case where $\mathbf{L}$ and $\mathbf{n}$ are perfectly orthogonal, $K = 0$ holds, and C_{transp} converges to C_{bg}. This represents a state in which the effective opacity vanishes and complete transparency is achieved.

Table 2. Representative values of C_{shaded} and C_{transp} under typical lighting directions $\mathbf{L}$,with the general background color C_{bg}, assuming $\alpha = 0.5$, $k_{\mathrm{d}} = 1$, and $k_{\mathrm{a}} = 0$.

Condition	C_{shaded}	C_{transp}
$\mathbf{L} \mathbin{/\!/} \mathbf{n}$	C_{org}	$(C_{\mathrm{org}} + C_{\mathrm{bg}})/2$
$\mathbf{L} \cdot \mathbf{n} = 0.5$	$(C_{\mathrm{org}} + C_{\mathrm{bg}})/2$	$(C_{\mathrm{org}} + 3C_{\mathrm{bg}})/4$
$\mathbf{L} \perp \mathbf{n}$	C_{bg}	C_{bg}

5. Proposed Method: Adaptive Opacity Control for Highlighting User-Specified Regions

Building on the formulation presented in Section 4, this section summarizes the concrete steps of the proposed method of adaptive opacity control designed to highlight user-specified regions. At the end of this section, we also demonstrate the effectiveness of the method through transparent visualization of a cube model containing an internal structure.

Given an input 3D scanned point cloud, the proposed visualization procedure proceeds as follows:

Step 1 (pre-processing): Normal Orientation A normal orientation is assigned to each point. Following the PCA-based approach explained in Section 3.2, a local spherical neighborhood is defined for each point, and the covariance matrix of the neighboring point coordinates is computed. The eigenvectors of this covariance matrix are then obtained, and the orientation of the eigenvector corresponding to the smallest eigenvalue is adopted as the local normal orientation. Note that only the normal orientation is determined here; the sign of the normal vector remains undetermined at this stage.

Step 2: Rendering with SPBR Render the point cloud using SPBR with the specified background color. For each point, the sign of the normal vector is dynamically determined so that it faces the viewpoint, satisfying Eq.(2). During rendering, adjust the light direction vector $\mathbf{L}$ to become parallel to the normal vectors of the regions of interest.

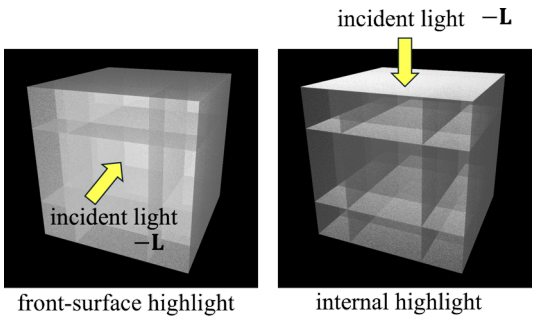

Figure 3. Transparent visualization of a point-based cube model with an internal structure. The arrows indicate the direction of incident light, i.e., $-\mathbf{L}$.

By following these steps, the user-specified region of interest attains higher effective opacity according to Eq.(5) or Eq.(4), in conjunction with the shading model defined in Eq.(3) or Eq.(6). Consequently, the region of interest is highlighted in the transparent visualization with enhanced visibility.

Note that Step 1 serves as a pre-processing stage performed only once for a given dataset; the computed normal orientations can then be reused in subsequent visualizations, whereas Step 2, which performs rendering based on SPBR, can be executed at an interactive speed (Uchida et al., 2020). Therefore, the core mechanism of the proposed method, adaptive highlighting of user-specified regions through adjustment of the lighting direction, can be achieved interactively.

As a simple example, we applied the proposed method to the transparent visualization of a cube model with an internal structure, represented as a point-based model. The result is shown in Fig.3. In the left image, the lighting direction is parallel to the front-face normal, so the front surface is highlighted and the internal structure becomes less visible. In the right image, the lighting direction is parallel to the top-face normal, so the horizontal surfaces are highlighted and the internal structure becomes more visible. Thus, by controlling the lighting direction, the user can adaptively increase the opacity of a surface of interest to highlight it, or decrease the opacity of an obstructing surface to reveal the structures behind it.

This lighting-based control mechanism distinguishes the proposed method from global opacity reduction and density-based highlighting, which are conventional approaches to opacity manipulation. Unlike these methods, no modification or resampling of the original point cloud data is required. Region highlighting is achieved solely by adjusting the lighting direction, enabling interactive performance for large-scale datasets. Moreover, the approach improves the visibility of regions of interest. Global opacity reduction uniformly decreases brightness, often reducing distinguishability between regions of interest and surrounding structures. Density-based methods alter point density to emphasize selected regions, but do not provide intuitive control over geometrically meaningful planar regions such as walls or floors. Therefore, the method is particularly effective for visual analysis of World Cultural Heritage sites, which typically exhibit large-scale, structurally complex 3D architectures, both externally and internally.

6. Case Studies

In this section, we present case studies demonstrating the effectiveness of the proposed region-highlighting method based on adaptive opacity control. The method is applied to large-scale 3D scanned point clouds of culturally significant heritage structures, including a mosque with complex internal geometry, an ancient shrine office building designated as an Important Cultural Property of Japan, and a World Heritage palace complex.

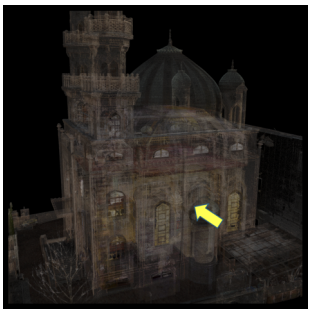

(a) Front-surface highlight

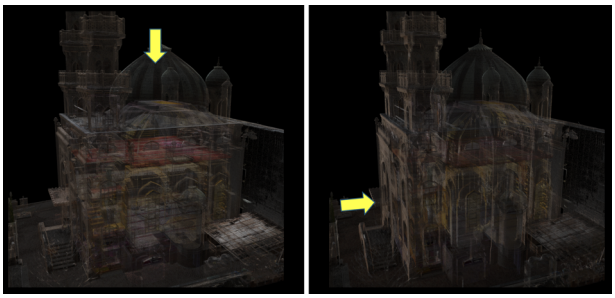

(b) Internal highlight (focus on horizontal planes)

(c) Internal highlight (focus on vertical planes)

Figure 4. Transparent visualizations with front-surface and internal highlighting of the 3D scanned point cloud of the Kobe Muslim Mosque (Hyogo, Japan). The arrows indicate the direction of incident light. Figure (a) shows front-surface highlighting, while Figures (b) and (c) show internal highlighting, focusing on horizontal and vertical planes, respectively.

Each case was selected to demonstrate a distinct aspect of our method: internal structural highlighting, orientation-based wall extraction and eye-position–based roof visibility control in a multi-room building, and large-scale dynamic shading in an architectural heritage complex. Sections 6.1 and 6.2 present examples with black and white backgrounds, respectively.

All point-cloud datasets presented here were acquired by the authors through 3D scanning.

6.1 Black-Background Transparent Visualization

Figures 4 and 5 show transparent visualizations of the Kobe Muslim Mosque dataset (45 million points; Hyogo, Japan) on a black background. The arrows indicate the direction of incident light. The rendering parameters are: $k_{\mathrm{d}} = 0.7$, $k_{\mathrm{a}} = 0.3$, and $L_{\mathrm{R}} = 75$.

In Fig.4, the highlighted regions change adaptively according to the selected lighting direction, similarly to Fig.3. Figure 4(a) demonstrates front-surface highlighting, obtained by setting $\mathbf{L}$ parallel to the surface normal of the mosque’s front face. Figures 4(b) and 4(c) demonstrate internal highlighting, focusing on horizontal planes (such as floors) and vertical planes (such as walls), respectively, by setting $\mathbf{L}$ perpendicular to the corresponding planes.

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-XI-4-2026-427-2026 | © Author(s) 2026. CC BY 4.0 License.

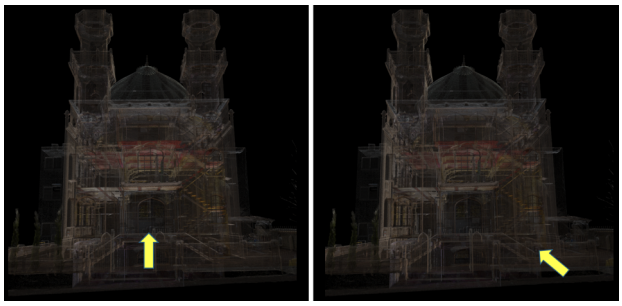

(a) Focus on upper floor (b) Focus on upper-left region

(c) Focus on right wall (d) Focus on left wall

Figure 5. Transparent visualizations with *internal highlighting* of the same data as Fig.4, using various lighting directions perpendicular to the viewing direction. Surfaces closer to the light source become more transparent, whereas those farther appear more opaque and are highlighted.

Figure 5 demonstrates *internal highlighting*, which enhances the perceptual clarity of complex structures. The visualized data are the same 3D scanned point cloud as in Fig. 4. Figures 5(a)–(d) show the results for lighting directions from below, lower right, left, and right. In all cases, surfaces closer to the light source become more transparent, whereas those farther appear more opaque and are highlighted.

The contrast between the surfaces closer to and farther from the light source, which is demonstrated in Fig.5, originates from the sign of $\mathbf{L}\cdot\mathbf{n}$. Note that the sign of the surface normal vector $\mathbf{n}$ is determined according to Eq.(2). Then, $\max(0, \mathbf{L}\cdot\mathbf{n})$ vanishes in Eq.(3) for surfaces closer to the light source, while it retains a finite value for surfaces farther from the light source, thereby contributing to increased brightness. For example, in Fig. 5(a), the normal vector $\mathbf{n}$ of the lower floor points upward, which is parallel to the incident lighting direction $-\mathbf{L}$, that is, anti-parallel to $\mathbf{L}$, and therefore $\max(0, \mathbf{L}\cdot\mathbf{n})$ vanishes. On the other hand, for the upper floor, $\mathbf{n}$ points downward, which is parallel to $\mathbf{L}$, and therefore $\max(0, \mathbf{L}\cdot\mathbf{n})$ remains finite.

Next, we present another dataset (100 million points) of the ancient office building at Tamaki Shinto Shrine (Nara Prefecture, Japan), which forms part of the UNESCO World Heritage property “Sacred Sites and Pilgrimage Routes in the Kii Mountain Range”, and is designated as an Important Cultural Property of Japan (see the photos taken during our scanning activity in Fig.6). The building consists of many small rooms, making its 3D structure complex. Transparent visualization is effective for observing and understanding its overall configuration, and the rendering was performed with a black background. The rendering parameters are $k_{\mathrm{d}} = 0.7$, $k_{\mathrm{a}} = 0.3$, and $L_{\mathrm{R}} = 200$.

Figure 7 demonstrates transparent visualization focusing on the vertical walls extending from back to front in the ancient office building, making the other walls and horizontal floors less visible. This focusing is achieved by setting the lighting direction normal to those vertical walls, that is, parallel to their surface normal vectors. The walls of the building are decorated with many culturally significant paintings, and accurate extraction of the wall regions is important for further academic analysis.

Figure 6. Photos of the ancient office building at the Tamaki Shinto Shrine (Nara Prefecture, Japan), which is designated as an Important Cultural Property of Japan.

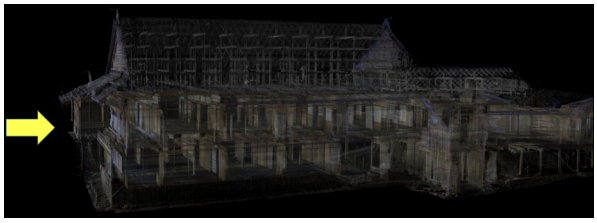

Figure 7. Transparent visualization focusing on the vertical walls extending from back to front in the ancient office building shown in the photos of Fig. 6, achieved by setting the lighting direction normal to those walls.

Figure 8 demonstrates transparent visualization focusing on the horizontal floors of the office building, making the vertical walls less visible. In Fig.8(a), the roof region remains visible, whereas in Fig. 8(b), it becomes almost invisible. This difference originates from the change in eye position in relation to the condition given in Eq.(2). In Fig.8(a), the building is viewed from a position slightly above the roof; thus, most of the surface normal vectors $\mathbf{n}$ in the roof region point upward, forming an acute angle with $\mathbf{L}$ (which is directed upward). As a result, $\max(0, \mathbf{L}\cdot\mathbf{n})$ remains finite. Conversely, in Fig.8(b), the building is viewed from a position slightly below the roof, and the majority of the surface normal vectors $\mathbf{n}$ in the roof region point downward, forming an obtuse angle with $\mathbf{L}$ (which is directed upward). Consequently, $\max(0, \mathbf{L}\cdot\mathbf{n})$ tends to vanish.

By effectively utilizing the property demonstrated in this subsection, users can flexibly and interactively control which surfaces are highlighted in the visualization. This capability enhances the freedom of focus and leads to a deeper, more intuitive understanding of the object’s overall structure.

6.2 White-Background Transparent Visualization

In this subsection, we present visualization examples with a white background, which generalize the black-background visualizations demonstrated in Section 6.1. The white background enables clearer perception of external contours and allows more versatile applications in documentation and presentation.

We demonstrate transparent visualization with a white background using the 3D scanned point cloud (96 million points) of the Ninomaru-goten Palace of Nijo-jo Castle (Kyoto, Japan),

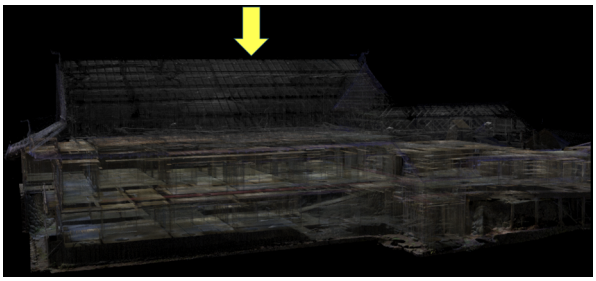

(a) Focus on floors with the roof region visible

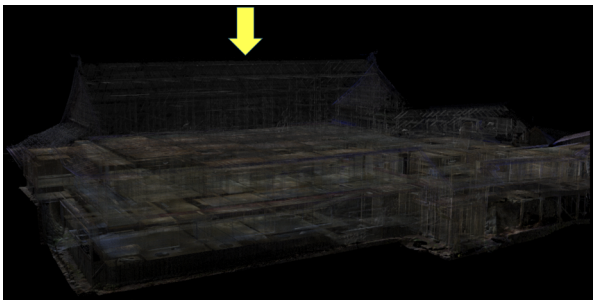

(b) Focus on floors with the roof region less visible

Figure 8. Transparent visualization focusing on horizontal floors of the ancient office building shown in Fig. 6. Figure (a) shows the result when viewed from a position slightly above the roof, where the roof region remains visible. Figure (b) shows the result when viewed from a position slightly below the roof, where the roof region becomes almost invisible.

which is part of the UNESCO World Heritage site “Historic Monuments of Ancient Kyoto” and is designated as a National Treasure of Japan. The rendering parameters are: $k_{\mathrm{d}} = 0.7$, $k_{\mathrm{a}} = 0.3$, and $L_{\mathrm{R}} = 70$.

Figure 9 shows the transparent visualizations, where the lighting direction was varied sequentially to highlight different wall orientations. In Fig.9(a), the light is cast from the front-left, highlighting walls extending from back to front. In Fig.9(b), the light comes from the front, highlighting both back–front and left–right walls. In Fig.9(c), the light is cast from the front-right, highlighting walls extending from left to right. By continuously changing the lighting direction, the proposed method enables *dynamic shading*, which provides an intuitive understanding of the architectural structure. Since this shading control does not require any modification of the point cloud data itself, structural analysis based on the dynamic shading can be performed at interactive rendering speed.

7. Conclusion

This paper presented a method for adaptive transparent visualization of large-scale point clouds to enhance the visibility of user-specified regions. The approach interprets shading brightness as a distribution of effective opacity. By adjusting the lighting direction, it locally controls this opacity, highlighting selected regions without modifying the original point cloud data.

The method was first formulated under a black background, where the equivalence between shading and opacity control is most clearly demonstrated. We then extended it to general background colors, broadening its applicability.

Case studies using large-scale cultural heritage datasets demonstrated the effectiveness of the approach. For both black and white backgrounds, we achieved selective region highlighting based on adaptive opacity control. This significantly improves visibility during internal observation, enhancing the usefulness of transparent visualization. Because the lighting direction can be changed interactively, dynamic shading is realized by varying it continuously. This is highly useful for analyzing complex 3D structures of scanned cultural heritage objects.

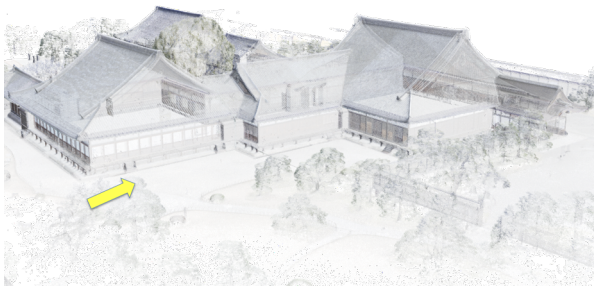

(a) Focus on walls extending from back to front

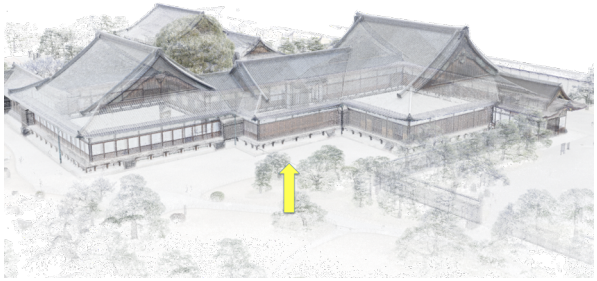

(b) Focus on both back–front and left–right walls

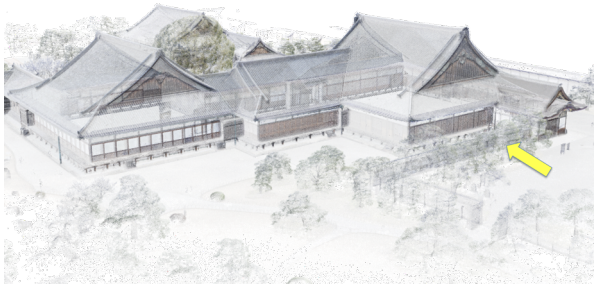

(c) Focus on walls extending from left to right

Figure 9. Transparent visualizations of the 3D scanned point cloud of the Ninomaru-goten Palace of Nijo-jo Castle (Kyoto, Japan). The visualizations are rendered with a white background, and the lighting direction is varied to highlight different wall orientations. The images illustrate how varying the lighting direction enables *dynamic shading*, allowing intuitive and interactive understanding of architectural structures.

Beyond cultural heritage, the method is also applicable to modern architectural structures and other large-scale 3D scanned objects with complex forms. Future work includes exploring these applications and extending the method to virtual reality for digital twins constructed from 3D scans. In particular, integration with immersive virtual reality supporting six-degree-of-freedom exploration, where lighting direction is controlled by a head-mounted device, could enable intuitive spatial analysis through natural viewpoint and light manipulation.

Acknowledgment

The point-cloud datasets used in this study were acquired by the authors through 3D scanning. We thank the Kobe Muslim

Mosque Community, the Maritime Asia Heritage Survey (MAHS) Japan field team, and Prof. Keiji Yano; the Chief Priest of Tamaki Shinto Shrine; the Kyoto City Government and the Nijojo Castle Office; Dr. Fadjar I. Thufail (BRIN) and Brahmantara (Borobudur Conservation Office) for their cooperation in data acquisition. This research was supported by JSPS KAKENHI Grant Number 21H04903.

References

Aicardia, I., Chiabrandob, F., Linguaa, A., Noardo, F., 2018. Recent trends in cultural heritage 3D survey: The photogrammetric computer vision approach. *Journal of Cultural Heritage*, 32, 257-266.

Discher, S., Masopust, L., Schulz, S., Richter, R., Döllner, J., 2018. A Point-Based and Image-Based Multi-Pass Rendering Technique for Visualizing Massive 3D Point Clouds in VR Environments. *Journal of WSCG*, 26(2), 76-84.

Dylla, K., Frischer, B., Mueller, P., Ulmer, A., Haegler, S., 2009. Rome reborn 2.0: A case study of virtual city reconstruction using procedural modeling techniques. *Proc. CAA 2009*, 62–66.

El-Hakim, S. F., Beraldin, J.-A., Gonzo, L., Whiting, E., Jemtrud, M., Valzano, V., 2005. A hierarchical 3D reconstruction approach for documenting complex heritage sites. *Proc. CIPA XX International Symposium*, 790–795.

Guidi, G., Frischer, B., Simone, M. D., Cioci, A., Spinetti, A., Carosso, L., Micoli, L. L., Russo, M., Grasso, T., 2005. Virtualizing ancient rome: 3D acquisition and modeling of a large plaster-of-paris model of imperial rome. *Proc. SPIE 5665*, Videometrics VIII, 119–133.

Hoppe, H., DeRose, T., Duchamp, T., McDonald, J., Stuetzle, W., 1992. Surface Reconstruction from Unorganized Points. *Computer Graphics (Proc. ACM Siggraph 1992)*, 26(2), 71-78.

Huang, L., Ye, D., Dan, J., Tao, C., Liu, H., Zhou, K., Ren, B., Li, Y., Guo, Y., Guo, J., 2025. TransparentGS: Fast Inverse Rendering of Transparent Objects with Gaussians. *arXiv preprint arXiv:2504.18768 (accepted by SIGGRAPH 2025)*. https://arxiv.org/abs/2504.18768.

Ikeuchi, K., Oishi, T., Takamatsu, J., 2007. Digital Bayon temple — e-monumentalization of large-scale cultural-heritage objects —. *Proc. ASIAGRAPH 2007*, 1(2), 99–106.

Kawato, M., Li, L., Hasegawa, K., Adachi, M., Yamaguchi, H., Thufail, F. I., Riyanto, S., Brahmantara, Tanaka, S., 2021. A digital archive of Borobudur based on 3D point clouds. *ISPRS Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci. (Proc XXIV ISPRS Congress)*, XLIII-B2(2021), 577-582.

Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G., 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics (Proc. SIGGRAPH 2023)*, 42(4), 1–14.

Kersten, T. P., Keller, F., Saenger, J., Schiewe, J., 2012. Automated generation of an historic 4D city model of hamburg and its visualisation with the ge engine. *Progress in Cultural Heritage Preservation (Lecture Notes in Computer Science 7616)*, 55–65.

Koller, D., Frischer, B., Humphreys, G., 2009. Research Challenges for Digital Archives of 3D Cultural Heritage Models. *ACM Journal on Computing and Cultural Heritage*, 2(3). Article No.7.

Laycock, R. G., Drinkwater, D., Day, A. M., 2008. Exploring Cultural Heritage Sites through Space and Time. *ACM Journal on Computing and Cultural Heritage*, 1(2). Article No.11.

Li, M., Pang, P., Fan, H., Huang, H., Yang, Y., 2025. TSGS: Improving Gaussian Splatting for Transparent Surface Reconstruction via Normal and De-lighting Priors. *arXiv preprint arXiv:2504.12799 (accepted by ACM MM 2025)*. https://arxiv.org/abs/2504.12799.

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. *Computer Vision – ECCV 2020*, Lecture Notes in Computer Science, 12346, Springer, 405–421.

Mitra, N. J., Nguyen, A., 2003. Estimating surface normals in noisy point cloud data. *Proc. SCG ’03 (the 19th Annual Symposium on Computational Geometry*, 322–328.

Nojiri, H., Takatori, S., Hasegawa, K., Li, L., Feener, R. M., Yano, K., Tanaka, S., 2025. Transparent visualization with dynamic shading for 3D scanned point clouds of cultural heritage objects. *2025 IEEE 18th Pacific Visualization Conference (PacificVis)*, 375–380.

Pan, J., Li, L., Yamaguchi, H., Hasegawa, K., Thufail, F. I., Brahmantara, Tanaka, S., 2021. Integrated High-Definition Visualization of Digital Archives for Borobudur Temple. *Remote Sensing*, 13(24), 5024.

Pauly, M., Keiser, R., Gross, M., 2003. Multi-scale Feature Extraction on Point-Sampled Surfaces. *Computer Graphics Forum*, 22(3), 281-289.

Remondino, F., Girardi, S., Rizzi, A., Gonzo, L., 2009. 3D Modeling of Complex and Detailed Cultural Heritage Using Multi-Resolution Data. *ACM Journal on Computing and Cultural Heritage*, 2(1). Article No.2.

Tanaka, S., Hasegawa, K., Okamoto, N., Umegaki, R., Wang, S., Uemura, M., Okamoto, A., Koyamada, K., 2016. See-Through Imaging of Laser-scanned 3D Cultural Heritage Objects based on Stochastic Rendering of Large-Scale Point Clouds. *ISPRS Ann. Photogramm. Remote Sens. Spatial Inf. Sci (Proc. XXIII ISPRS Congress, Prague, 2016)*, III-5, 73-80.

Thiel, F., Discher, S., Richter, R., Döllner, J., 2018. Interaction and Locomotion Techniques for the Exploration of Massive 3D Point Clouds in VR Environments. *ISPRS Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci. (Proc. 3D Spatial Information Science – The Engine of Change)*, XLII(4), 623-630.

Uchida, T., Hasegawa, K., Li, L., Adachi, M., Yamaguchi, H., Thufail, F. I., Riyanto, S., Okamoto, A., Tanaka, S., 2020. Noise-robust transparent visualization of large-scale point clouds acquired by laser scanning. *ISPRS Journal of Photogrammetry and Remote Sensing*, 161, 124-134.

Yastikli, N., 2007. Documentation of cultural heritage using digital photogrammetry and laser scanning. *Journal of Cultural Heritage*, 8, 423-427.