Journal of Computer Science

Research Article

MANAT: A Filtering-Based Method for Denoising Nonuniform Photogrammetric Point Clouds

Yun Sin Chong[1], Hui Hui Wang[2] and Yin Chai Wang[2]

[1]*Faculty of Computer Science and Information Technology, Universiti Malaysia Sarawak, Kota Samarahan, Sarawak, Malaysia*
[2]*Institute for Tourism Research and Innovation, Universiti Malaysia Sarawak, Kota Samarahan, Sarawak, Malaysia*

Article history
Received: 05-01-2026
Revised: 09-03-2026
Accepted: 01-04-2026

Corresponding Author:
Yun Sin Chong
Faculty of Computer Science and Information Technology, Universiti Malaysia Sarawak
Kota Samarahan, Sarawak, Malaysia
Email: yschong220@gmail.com

Abstract: Three-dimensional point clouds reconstructed from photogrammetry often exhibit noise and non-uniform sampling density, which challenges existing denoising methods that rely on precise normal estimation or extensive parameter tuning. This study presents the Multi Attribute Neighbour Attraction Technique (MANAT), a novel single-stage, density-adaptive filtering method that jointly leverages spatial position, surface normals, and color as inherent photogrammetric attributes for unified noise removal. MANAT assesses each point's consistency within its k-nearest neighbourhood using local geometric, orientation, and color statistics, enabling effective discrimination between valid surface points and noise in real-world photogrammetric data. On a large-scale heritage dataset of 141.7 million points, MANAT achieved 23.78% noise removal with improvements of 9.60, 6.91, and 4.40% in surface roughness, local and global normal standard deviations respectively. Comparison with DBSCAN confirms that spatial density alone is insufficient to characterise embedded photogrammetric noise, highlighting the necessity of multi-attribute denoising. These results demonstrate MANAT's practical effectiveness as a robust framework for enhancing the accuracy and reliability of photogrammetric 3D reconstructions under realistic acquisition conditions.

Keywords: Point Cloud Denoising, Photogrammetry, Density-Adaptive Filtering, 3D Reconstruction

Introduction

Three Dimensional (3D) scanning and reconstruction have revolutionized the way we capture and study the physical world, enabling applications ranging from heritage documentation (Fabris et al., 2023; Maset et al., 2022) to medical imaging (Mazzonetto et al., 2022; Quispe-Enriquez et al., 2023). These technologies generate point clouds as collections of 3D coordinates that represent surface geometry, commonly acquired through LiDAR or photogrammetry (Tang et al., 2024; Yang et al., 2023). While LiDAR delivers highly accurate measurements across diverse environments, photogrammetry has gained popularity due to its accessibility and lower cost, particularly when leveraging consumer-grade cameras and smartphones for budget-limited projects (Jasińska et al., 2023; Mengiste et al., 2022).

Despite these advantages, converting 2D images into 3D structures presents greater challenges than direct acquisition via LiDAR. Photogrammetric point clouds are frequently contaminated by noise and geometric artifacts stemming from lighting variations, texture inconsistencies, and motion blur (Gonizzi Barsanti et al., 2024; Liu et al., 2025). Noise suppression is likewise a critical challenge in other research domains, particularly in medical imaging, where denoising plays a central role in enhancing diagnostic reliability (Diwakar and Singh, 2020; Diwakar et al., 2024; 2022; Singh et al., 2023; Katta et al., 2024). This cross-domain emphasis on robust denoising further highlights the need for reliable noise mitigation strategies in photogrammetric reconstruction. Furthermore, existing camera calibration workflows are sensitive to device settings and require extensive manual effort, limiting their practicality for robust error correction (Ataiwe et al., 2021; Patonis, 2024; Rezaei et al., 2024; Van and Naprstkova, 2024). Similar to findings in other domains, where human-dependent frameworks introduce operational vulnerabilities (Aljumaiah et al., 2025), reliance on manual intervention increases the risk of inconsistency and reduces reproducibility. This further

motivates the development of automated post-reconstruction filtering strategies. Since these photogrammetric errors arise during reconstruction and are difficult to control during data capture, such filtering offers a practical and scalable solution. By targeting artifacts in the final model, post-processing avoids the need for additional hardware or controlled capture environments, making it a practical approach for improving 3D model reliability.

However, significant challenges remain in processing real photogrammetric data. A primary limitation is that many studies rely on synthetic noise models (e.g., Gaussian or random noise) that do not capture the complex variability of real-world photogrammetry (Ambrosino et al., 2024; Gao et al., 2018; Charles et al., 2017; Zhou et al., 2022). Moreover, existing filtering methods depend heavily on accurate normal estimation (Yang et al., 2025; Yuan et al., 2020) and carefully tuned parameters to guide denoising (Cheng et al., 2021; Ren et al., 2021). These requirements frequently prove inadequate when faced with the high noise levels and uneven sampling densities common in photogrammetric datasets. Many filters also rely on LiDAR-specific attributes or spatial coordinates (Le et al., 2022; Wang et al., 2024), thereby underutilizing to exploit the diagnostic potential of the chromatic (RGB) and orientation data inherent to photogrammetry. As existing methods evaluate visual and spatial signals independently (Jia et al., 2019; Kim and Yun, 2024), they struggle to distinguish valid geometric structures from stochastic noise in complex reconstructions. Additionally, the move toward multi-stage pipelines has increased algorithmic complexity and the demand for further tuning (Liu et al., 2025; Ren et al., 2021). These limitations motivate the development of denoising strategies that adapt to local density variations without relying on precise normal estimation or complex multi-stage workflows.

To address these challenges, we propose the Multi-Attribute Neighbour Attraction Technique (MANAT), a filtering-based method that advances prior work in several ways. MANAT is evaluated on real-world photogrammetric data rather than synthetic noise assumptions, and it uses reconstruction-provided normals as diagnostic indicators instead of relying on iterative normal estimation as a guide for point movement. Rather than depending on LiDAR-specific metadata, MANAT exploits attributes inherent to photogrammetric reconstruction, namely spatial position, normals, and color. These attributes are evaluated jointly, enabling MANAT to distinguish valid geometric structures from noise with greater attribute awareness than methods that evaluate visual and spatial signals in isolation. In addition, MANAT operates within a single-stage, density-adaptive framework, avoiding the extensive tuning and multi-stage pipelines common in prior approaches.

The major contributions of this work are as follows:

(1) A novel single-stage, density-adaptive filtering method that jointly exploits spatial position, surface normals, and color as inherent photogrammetric attributes for unified noise removal
(2) Evaluation on real-world photogrammetric data with non-uniform density distributions, in contrast to the synthetic noise models commonly adopted in prior works, demonstrating practical applicability for photogrammetric applications

The significance of this study lies in improving the reliability of photogrammetric 3D reconstruction for downstream processing, while enhancing accessibility for budget-limited projects and providing a practical reference for future denoising research under realistic photogrammetric scenarios. This work details the design of MANAT and evaluates its performance through quantitative metrics and visual inspection. The remainder of the paper reviews existing literature, describes the materials and experimental setups, shows the design of MANAT, and reports the findings and conclusions.

Related Work

Filtering-based methods represent one of the earliest and most intuitive approaches for point cloud denoising, originally adapted from classical image processing techniques (Digne and Franchis, 2017). These techniques evaluate each point within a locally defined neighborhood constructed using k-nearest neighbors or a fixed-radius search. Furthermore, they rely on attributes such as spatial positions, surface normals, density, and color information to identify noisy points or outliers.

Traditional Filtering

Traditional filters such as Statistical Outlier Removal (SOR) and Radius Outlier Removal (ROR) are effective for isolated outliers but struggle with the non-uniform densities of photogrammetric data. SOR uses standard deviation thresholds based on neighbor distances (Kim and Yun, 2024), which may erroneously remove valid points or retain clustered noise in complex scenes. Similarly, ROR relies on fixed-radius neighbor counts (Cao et al., 2022), showing limitations in adapting to the extreme sampling variances found in photogrammetry (Jia et al., 2019). While Conditional Outlier Removal (COR) allows for attribute-based filtering, its dependence on manual thresholds often leads to overfitting. Collectively, these methods remain analytically limited regarding surface orientation and color inconsistencies because they evaluate geometric attributes independently using non-adaptive constants rather than density-adaptive logic.

Enhanced Bilateral Filtering

Beyond simple statistical checks, bilateral filtering is widely used for point cloud denoising as it simultaneously considers both spatial distances and surface normal information. However, applying this to photogrammetry requires solving the dual problems of uneven density and unreliable normals. To mitigate density issues, researchers have implemented density-adaptive neighborhood selection via octree-based queries (Digne and Franchis, 2017) or feature-preserving downsampling (Zou et al., 2020) to regulate point distribution prior to filtering. Concurrently, to improve the accuracy of normal estimation, Yang et al. (2025) incorporated iterative, weighted PCA-based normal estimation, giving higher weights to nearby and reliable points. Similarly, Yuan et al. (2020) proposed anisotropic feature-aware neighborhoods to better adapt to local geometry for enhanced normal estimations. However, these methods increase computational demand and remain constrained by assumptions of local planarity.

Despite these advancements, a gap exists in leveraging the stable, pre-existing orientation attributes generated during the photogrammetry reconstruction process. Instead of attempting to recover surface orientation through iterative smoothing, there is an opportunity to use the original orientation data as a diagnostic indicator to identify structural inconsistencies. By shifting the role of the normal from a guide for point movement to a diagnostic tool for point removal, the need for complex normal estimations can be bypassed. This motivates the need for a multi-attribute mechanism that utilizes inherent reconstruction data to maintain accuracy without the overhead of an iterative process.

Hybrid Filtering

Hybrid methods combine neighborhood geometry with statistical analysis, projection, or reference information to improve noise removal. Density-adaptive statistical filtering was originally developed for autonomous driving LiDAR, where point density decreases with sensor distance. Methods such as DROR and DSOR introduced range-dependent scaling to adjust search radii and statistical thresholds dynamically (Charron et al., 2018; Kurup and Bos, 2021). AGDOR extends this idea by incorporating intensity values to distinguish low-intensity points under varying densities (Le et al., 2022), while AKA-LDGS integrates global statistics such as the mean and standard deviations of point coordinates (Wang et al., 2024). Although these frameworks demonstrate the value of density-adaptivity, their reliance on LiDAR-specific metadata and sequential or iterative workflows often introduces computational bottlenecks. Consequently, there is an opportunity to exploit inherent photogrammetric attributes as functional alternatives to sensor metadata, which supports the development of a single-stage, density-adaptive mechanism without sequential latency.

Projection-based hybrid methods evaluate point alignment with local planes or reference surfaces. Early approaches relied on plane projection or bilateral filtering for non-flat points (Yan and Jinlei, 2015), while later frameworks incorporated density-aware strategies to improve robustness under non-uniform sampling (Ambrosino et al., 2024; Zhao et al., 2021; Cheng et al., 2021) applied a statistical grid approach using PCA and kNN segmentation to separate dense structural points from sparse noise. While these methods demonstrate the benefits of combining geometric and projection guidance, a significant limitation of projection-based logic is the reliance on rigid mathematical primitives. This requirement risks discarding the fine, irregular details that characterize authentic historical surfaces (Cheng et al., 2021).

In some cases, approaches further rely on external reference data, such as surface roughness from high-quality terrestrial laser scans, to guide filtering (Di Filippo et al., 2024). Although effective at preserving surface detail, this strategy is unable to adapt to varying point densities independently and depends heavily on the availability of reference datasets. Overall, the reliance on rigid fitting or external benchmarks limits generalizability across photogrammetric projects and underscores the need for denoising strategies that avoid predefined surfaces and external geometric references.

Multiscale Filtering

Noise in point clouds often occurs at multiple scales, ranging from fine surface irregularities to dense, clustered artifacts. Single-scale filters may over-smooth fine details or prove inadequate in removing large, structured noise. Multiscale filtering addresses this by applying different methods across multiple levels to handle both large- and small-scale noise. For example, Ren et al. (2021) combined statistical filtering with bilateral refinement using PCA-based surface change factors, while Liu et al. (2025) employed DBSCAN clustering to remove large-scale platform noise before applying statistical and bilateral filtering for local denoising. Although these strategies improve robustness across scales, they necessitate complex multi-stage pipelines and extensive parameter tuning, which restrict automation and computational efficiency.

Color-Based Filtering

Color-based filtering offers an alternative to traditional, bilateral, hybrid, and multiscale methods by using image-derived color information for point cloud denoising. For example, Jia et al. (2019) transformed RGB data into HSV space and applied 2D segmentation

using Otsu thresholding to remove background points before mapping results back to 3D. Despite its efficiency, this approach depends on strong color contrast and precise RGB-D alignment, which may limit its use in uncontrolled environments. Similarly, Lin et al. (2023) proposed a graph-based color filtering to correct color noise without altering spatial positions, achieving notable computational efficiency. However, this approach assumes geometrically clean input and cannot address spatial outliers independently. Nevertheless, color remains underexplored as a diagnostic attribute when jointly combined with geometric and normal features in a unified single-stage filtering framework, helping to distinguish noise from valid points.

In summary, existing filtering methods reveal several critical limitations. Traditional approaches ignore orientation and color attributes, while bilateral methods remain constrained by unreliable normal estimation and local planarity assumptions. Hybrid frameworks demonstrate the value of density-adaptivity but remain tied to LiDAR-specific metadata or external reference datasets, and multiscale strategies improve noise coverage at the cost of complex multi-stage pipelines and extensive parameter tuning. Color-based methods, meanwhile, treat color as an isolated processing channel rather than a joint diagnostic signal alongside spatial and normal attributes. Critically, no existing single-stage filtering method has jointly exploited spatial position, surface normals, and color inherently available from photogrammetric reconstruction in a unified, density-adaptive framework a gap that directly motivates the development of MANAT.

Materials

To evaluate the proposed point cloud denoising strategy in real-world scenarios, experiments were conducted on photogrammetric datasets reconstructed from images featuring complex structures. This section describes the datasets, photogrammetric reconstruction workflow, implementation environment, and evaluation metrics used throughout the study.

Dataset Overview

The case study focuses on Fort Margherita, a prominent historical landmark in Kuching, Sarawak, where all data were captured directly by the authors through photogrammetric scanning. Its architectural complexity and combination of repetitive and unique features make it well-suited for evaluating reconstruction quality and noise filtering in photogrammetry. The fort includes a three-storey tower block with a battlement watchpoint, a courtyard enclosed by high walls with cannon platforms, and two smaller turret towers facing the Sarawak River. Several views of Fort Margherita are shown in Fig. 1.

Data Acquisition and Reconstruction

The study area was surveyed in April 2025 during the Northeast Monsoon, which provided slightly cloudy conditions favorable for photogrammetry. A consumer-grade smartphone, Xiaomi 14 T, equipped with a 50MP lens (f/1.7), was used to capture all images. All images were captured in manual mode to maintain consistent exposure and quality. ISO was set between 80-120 to reduce noise, and exposure times faster than 1/2000s were used to avoid motion blur. To minimize hidden spots, images were taken from multiple viewing positions, as recommended in previous studies (Caciora et al., 2021; Triviño-Tarradas et al., 2024). Terrain and building features, such as staircases and battlements, were used to capture high-angle shots for maximum coverage. In addition, a minimum of 70% overlap between consecutive images was maintained to enhance alignment quality (Lin et al., 2024; Sangmanee and Suwanwerakamtorn, 2023). Physical markers were also placed on the structure to further support image registration and geometric accuracy.

Finally, the images were reviewed to remove blurry or duplicate shots, leaving 1,514 usable photos covering the entire fort. Image alignment in RealityScan successfully registered 1,499 images, resulting in a high-detail dense point cloud with 148,695,469 points. The resulting raw point cloud for denoising experiments is shown in Fig. 2.

Fig. 1: Fort Margherita captured by Authors: (1) Courtyard, (2) Side view, (3) Front view

Fig. 2: Photogrammetry model from different perspectives in RealityScan

Since the dense cloud was too large for efficient experimentation, it was segmented into seven regions using Cloud Compare. During this process, peripheral non-target elements, such as surrounding vegetation and ground clutter, were manually removed to isolate the primary structures for analysis. It is important to note that this manual removal only excluded non-structural elements outside the target architecture and did not alter the intrinsic point distribution of the primary structures. The photogrammetric reconstruction artifacts within the architectural structures, such as surface noise and scattered outliers were retained and not manually corrected, ensuring that all reconstruction noise inherent to the target structures remained present for the algorithm to address. This cleaning process resulted in a final dataset of 141,691,598 points. The segmented dataset is shown in Fig. 3, with yellow boundaries highlighting the segmentation. This preprocessing allows the dataset to be more manageable for processing and enables experiments to be conducted within reasonable computational time without memory bottlenecks. As a result, the evaluation metrics are not artificially inflated by prior manual cleaning and will reflect the algorithm's performance specifically on photogrammetric noise rather than general scene segmentation, which aligns with the method's intended application domain.

Implementation Environment

The experiments of MANAT were conducted using a workstation equipped with an AMD Ryzen 7 7700X CPU, an NVIDIA RTX 4070 Super GPU, and 32GB Kingston DDR5 RAM. This configuration provided sufficient computational power and storage for processing large point cloud datasets.

Evaluation Metrics

A set of representative metrics is used to summarize the behavior of MANAT, including runtime, noise removal percentage, surface roughness, and dual-scale normal consistency metrics. Runtime indicates computational cost, while noise removal percentage reflects the proportion of points identified as inconsistent. Surface roughness measures local geometric irregularities by computing distances from each point to the best-fit plane of its neighborhood, with lower values indicating smoother surfaces. Dual-scale normal consistency metrics are employed to assess MANAT's output, evaluating surface coherence at two levels:

- Local Normal Standard Deviation (NSD) measures the angular spread of each point's normal relative to its k = 8 nearest neighbors, reflecting fine-scale surface jitter
- Global Normal Standard Deviation (NSD) measures the angular spread across all points, serving as a benchmark for overall geometric integrity and ensuring that large-scale architectural features are preserved

Together, these metrics provide a concise overview of MANAT's denoising effectiveness, capturing both fine-scale surface coherence and global structural preservation.

In addition, representative visualizations of the filtered point cloud are presented alongside the original dataset, rendered from identical viewpoints to ensure consistency. The inspection focuses on geometric feature preservation, the presence of isolated points, and the continuity of surfaces after filtering. These observations provide a clear illustration of the method's performance when applied to real-world point cloud data.

To contextualise MANAT's performance against a conventional approach, DBSCAN was applied as a density-based baseline. Due to memory constraints at the scale of the segmented dataset, DBSCAN was applied on voxel-down sampled versions of each segment at 0.005 m resolution. The eps parameter was determined via k-distance graph analysis on the down sampled point cloud, yielding eps = 0.10 m, with min_points = 8 set consistent with MANAT's neighbourhood size k = 8 to maximize methodological fairness.

Fig. 3: Model segmented into seven regions for easier processing, with boundaries highlighted in yellow

Methodology

This section presents a detailed description of the design of MANAT. To ensure a clear and structured presentation, the design is organized into five subsections, each focusing on a particular aspect of the method.

Point Cloud Neighborhood Selection

A fundamental component of MANAT is the definition of local neighborhoods, which provides the context for evaluating each point relative to its surrounding points. MANAT adopts the k-nearest neighbors (kNN) approach, where each point is associated

with a fixed number of neighbors based on Euclidean distance in 3D space. The neighborhood of a candidate point p_i, excluding the point itself, is denoted as $N(p_i)$.

kNN is preferred over radius-based neighborhoods because variable point densities in photogrammetry can lead to too few or too many neighbors within a fixed radius. A consistent number of neighbors provides a stable foundation for MANAT's multi-attribute evaluation. This approach enhances robustness near edges and corners, particularly when processing the complex geometries of large outdoor cultural heritage sites.

Point Cloud Representation and Feature Vector

Each point is represented by its spatial position, surface normal, and color, forming a multi-attribute feature vector. To formalize this representation, the point cloud is represented as a finite set of points:

$$P = \{ p_i \mid i = 1,2, \dots, N \}, p_i \in R^3 \quad (1)$$

Where N denotes the total number of points in a point cloud. The spatial position of each point p_i is defined as:

$$p_i = (x_i, y_i, z_i), p_i \in P \quad (2)$$

Where the Cartesian coordinates of the points are given in x, y and z directions in 3D Euclidean space. Additionally, the point normal for each point p_i is represented as:

$$n_i = \left(n_{xi}, n_{yi}, n_{zi}\right), n_i \in \mathbb{R}^3, \ \| n_i \| = 1 \quad (3)$$

where each point is assigned a unit normal vector. This ensures only the vector orientation is considered, avoiding potential bias caused by magnitude variations. In this study, these point normals are provided directly by RealityScan during reconstruction. Utilizing reconstruction-based normals ensures consistency with the photogrammetry-generated geometry and avoids the mismatches that often occur with post-estimated normals.

The color of each point p_i is defined as:

$$c_i = (L_i, a_i, b_i) \quad (4)$$

Where each RGB point is converted into the CIELab color space to provide a perceptually uniform representation. Unlike RGB or HSV, where numerically similar values may correspond to perceptually different colors, CIELab is designed such that numerical differences in each channel correspond more closely to human visual perception. This makes color variability in CIELab a more reliable indicator of visual inconsistency when evaluating neighborhood coherence. The color attribute thus complements spatial position and normal information. Combining these attributes, the complete multi-attribute feature vector for each point p_i is as follows:

$$f_i = (p_i, n_i, c_i) \quad (5)$$

Based on the feature vector f_i, MANAT evaluates inconsistency across multiple attributes including spatial position, normal, and color.

Neighborhood Inconsistency Determination

To determine neighborhood inconsistency with and without the candidate point p_i, three types of measures are computed: Geometry-based, normal-based, and color-based inconsistencies.

Among these, the geometry-based inconsistency measure evaluates how much adding a point increases the spread of its neighborhood. The centroid of its local neighborhood is computed to serve as a representative center for each point p_i. Let cen_{wo} denote the centroid of the k-nearest neighbors $\mathcal{N}(p_i)$ without p_i, and let cen_w denote the centroid including p_i. The centroids are represented as follows:

$$cen_{wo} = \frac{1}{\mathrm{k}} \sum_{j \in \mathcal{N}(p_i)} p_j \, , cen_w = \frac{1}{\mathrm{k+1}} \sum_{j \in \mathcal{N}(p_i) \cup \{p_i\}} p_j \quad (6)$$

The average spatial distance of all neighbors to their centroid is then computed as a metric of neighborhood "tightness". The difference between the mean Euclidean distance of neighbors to their centroids with and without p_i defines the geometric inconsistency score $d_{\mathrm{geo}}(p_i)$:

$$d_{geo}(p_i) = \frac{1}{k+1} \sum_{j \in N(p_i) \cup \{p_i\}} \left\| p_j - cen_w \right\| - \frac{1}{k} \sum_{j \in N(p_i)} \left\| p_j - cen_{wo} \right\| \quad (7)$$

A positive value in $d_{geo}(p_i)$ means that adding the candidate point p_i increases the spread of the neighborhood, suggesting that the point does not conform to the local geometric structure. This approach captures how the addition of p_i changes the overall spatial arrangement. If the point is inconsistent, adding it shifts the centroid cen_w toward itself, which then increases the average distance of its neighboring points from the centroid, resulting in a higher $d_{geo}(p_i)$. Conversely, a value near zero or negative implies the point is consistent with or tightens the existing distribution.

In addition, the normal-based inconsistency measure evaluates how much a candidate point p_i disturbs the orientation coherence of its local neighborhood. To compute a representative neighborhood direction, a mean normal is calculated within its neighborhood. For the neighborhood without and with the candidate point p_i, the

mean normal is defined by $\bar{n}_{wo}$ and $\bar{n}_w$ respectively, where each surface normal n_j is a unit vector:

$$\bar{n}_{wo} = \frac{1}{\mathrm{k}}\sum_{j\in\mathcal{N}(p_i)} \mathrm{n}_j \,, \bar{n}_w = \frac{1}{\mathrm{k}+1}\sum_{j\in\mathcal{N}(p_i)\cup\{p_i\}} \mathrm{n}_j \tag{8}$$

As these sums are not guaranteed to be of unit length, they are normalized by their magnitude to produce the unit mean normals. The unit mean normals without and with the candidate point p_i are denoted as $\hat{n}_{wo}$ and $\hat{n}_w$ respectively:

$$\hat{n}_{wo} = \frac{\bar{n}_{wo}}{\|\bar{n}_{wo}\|}, \|\bar{n}_{wo}\| \geq 10^{-8} \tag{9}$$

$$\hat{n}_w = \frac{\bar{n}_w}{\|\bar{n}_w\|}, \|\bar{n}_w\| \geq 10^{-8} \tag{10}$$

If the magnitude of the mean normal $\|\bar{n}\|$ is below 10^{-8}, the neighborhood normals are considered to nearly cancel each other, indicating no dominant orientation. In this case, the angular spread is set to the default maximum value of 90°, treating the neighborhood as maximally scattered. This 90° limit corresponds to the largest meaningful deviation between unoriented surface normals. This choice provides a bounded and numerically stable upper limit for orientation incoherence in neighborhoods with no well-defined normal direction.

In contrast, if a dominant orientation exists ($\|\bar{n}\| > 10^{-8}$), the angular differences of each unit normal n_j from the unit mean normal $\hat{n}$ is given by:

$$\theta_j = \cos^{-1}(\mathrm{n}_j \cdot \hat{n}) \,, \theta_j \in [0,\pi], \hat{n} = \begin{cases} \hat{n}_{wo}, & \text{without } p_i \\ \hat{n}_w, & \text{with } p_i \end{cases} \tag{11}$$

Where the dot product $\mathrm{n}_j \cdot \hat{n}$ gives the cosine of the angle between the unit vectors. The angular spread of these angular differences around the mean normal direction is measured using the standard deviation. This formulation evaluates how strongly the inclusion of a candidate point affects the overall dispersion of neighborhood orientations, ensuring that points with highly divergent normals are not masked by a simple averaging process. Finally, the normal inconsistency score for the candidate point p_i is defined as the difference in the angular spread with and without the point:

$$d_{\text{norm}}(p_i) = \text{std}\left(\{\theta_j\}_{j\in\mathcal{N}(p_i)\cup\{p_i\}}\right) - \text{std}\left(\{\theta_j\}_{j\in\mathcal{N}(p_i)}\right) \tag{12}$$

A large positive value of $d_{\text{norm}}(p_i)$ indicates that including the candidate point p_i increases the angular spread, meaning its orientation is inconsistent with the dominant neighborhood structure. Conversely, values near zero indicate that p_i aligns well with the neighboring normals. This measure identifies points whose surface orientation deviates from the local structure, enabling robust detection of noisy points in the point cloud.

For the color-based inconsistency measure, the goal is to evaluate how much a candidate point p_i disrupts the appearance coherence within its local neighborhood. It provides redundancy and strengthens outlier detection in regions where color carries complementary information alongside spatial positions and point normals.

Firstly, the standard deviation of each color channel across the neighboring points is calculated to quantify the local color variation in the neighborhood without and with the candidate point p_i. A single spread measure is then computed by averaging these standard deviations over the three channels as follows:

$$c_{wo}(p_i) = \frac{1}{3}\sum_{c\in\{L,a,b\}} \text{std}(c_j),\ j \in \mathcal{N}(p_i) \tag{13}$$

$$c_w(p_i) = \frac{1}{3}\sum_{c\in\{L,a,b\}} \text{std}(c_j),\ j \in \mathcal{N}(p_i)\cup\{p_i\} \tag{14}$$

The color inconsistency score is defined as the difference between the two averaged values:

$$d_{col}(p_i) = c_w - c_{wo} \tag{15}$$

A positive value of $d_{col}(p_i)$ indicates that including p_i increases the spread of colours in the local neighborhood, implying that the candidate point's appearance is inconsistent with its neighbors. This approach quantifies the color-based disruption to the local structure, aiding in the identification of noisy points that are spatially near but visually distinct from the surface.

Multi-Attribute Fusion

After computing the inconsistency scores for the three attributes, they are combined into a single scalar to represent the overall inconsistency for each point. To ensure they are scale-independent and comparable, each raw score of p_i is normalized using the median and median absolute deviation (MAD) of the corresponding attribute over the entire point cloud:

$$\hat{d}_j(p_i) = max\left(\frac{d_j(p_i) - median\,(d_j)}{MAD\,(d_j)}, 0\right),\ j \in \{geo, norm, col\} \tag{16}$$

The median represents the typical score in the dataset, while the MAD reflects the standard variability of inconsistency across the point cloud. By utilizing MAD, normalization remains robust against extreme outliers, ensuring that the filtering behavior is driven by the majority of clean data points. During this process, only positive deviations are retained. By clipping negative results to zero and weighing the attributes equally, this prevents a point that is highly consistent in one attribute

from offsetting a significant inconsistent deviation in another. The final weighted inconsistency score for the candidate point p_i is then computed as follows:

$$W(p_i) = \frac{1}{3}(\hat{d}_{\text{geo}}(p_i) + \hat{d}_{\text{norm}}(p_i) + \hat{d}_{\text{col}}(p_i)) \quad (17)$$

This score reflects the extent to which point p_i disrupts local structure across geometry, surface orientation, and appearance.

Base Thresholding

After computing the final weighted inconsistency scores, MANAT defines thresholds to classify points as inliers or outliers. Thresholding ensures that points with unusually high inconsistency are removed while preserving valid structures. To handle varying point densities, MANAT employs a global base threshold modified by a density-adaptive adjustment. This adjustment incorporates both a local density factor and a global factor to ensure consistency across different regions of the dataset. The first step is to establish the global base threshold that reflects the typical inconsistency level across the entire dataset. This threshold is defined as follows:

$$T_{base} = W_{med} + mult \cdot W_{MAD} \quad (18)$$

where W_{med} and W_{MAD} are the median and MAD of all weighted inconsistency scores in a point cloud. Using median-based statistics provides a statistical baseline that is robust to the presence of extreme outliers. The multiplier *mult* controls the pruning aggressiveness. A higher value produces a looser threshold, allowing more points to be retained, while a lower value increases the filtering intensity. This global threshold serves as the reference boundary for the density-adaptive adjustments, ensuring the filtering remains anchored to the overall statistical profile of the point cloud.

Adaptive Thresholding

To account for varying point densities, MANAT adjusts the threshold using a local density measure. For each point p_i, the average Euclidean distance to its *k*-nearest neighbors is computed as a local density measure ρ_i. This is then normalized relative to the entire cloud using a robust z-score based on the median and MAD of all density measures:

$$z_i^{\rho} = \frac{\rho_i - median(\rho)}{MAD(\rho)} \quad (19)$$

This density z-score z_i^{ρ} measures local sparsity relative to the global distribution. If $z_i^{\rho} \leq 0$, the point lies in a region that is denser than or equal to the median, and the base threshold is applied directly with no modification. If $z_i^{\rho} > 0$, the point lies in a sparser region, then the base threshold T_{base} is multiplied by a loosen factor L as follows:

$$L_i = 1 + \alpha \cdot tanh(\max(z_i^{\rho}, 0)), \alpha = mean\left(\max\left(z_i^{\rho}, 0\right)\right) \quad (20)$$

In this case, $\max(z_i^{\rho}, 0)$ ensures that only positive density deviations influence the adjustment, while negative deviations are set to 0. The global factor, density alpha α is defined as the mean of these deviations, representing the overall sparsity of the dataset. The hyperbolic tangent (*tanh*) function provides the local scaling. This nonlinear response allows the factor to increase smoothly for small z_i^{ρ} values before saturating for larger values, providing a self-correcting mechanism for non-uniform data. In practice, this adaptivity caps the maximum loosening at $1 + \alpha$, establishing an upper limit for threshold relaxation in extremely sparse regions. Conversely, for dense points, the $tanh$term remains small, so the resulting L stays close to 1, leaving the base threshold nearly unchanged.

Final Thresholding

Finally, the global base threshold and the density-adaptive adjustments are combined to define a per-point threshold. The final threshold T is obtained by applying the loosen factor L to the baseline threshold T_{base}:

$$T_i = T_{base} \cdot L_i \quad (21)$$

Points whose weighted inconsistency score $W(p_i)$ exceeds the final threshold T_i are removed. This formulation dynamically relaxes the decision boundary in sparse regions while maintaining stricter filtering in dense areas, allowing isolated but valid structures such as thin edges or small features to be retained. By combining global statistics with local density scaling, MANAT establishes a unified thresholding criterion capable of handling non-uniform sampling within a single execution pass.

Computational Complexity

MANAT operates in a single-pass, non-iterative framework, and its computational complexity can be analyzed both per point and globally across a point cloud of N points. To construct local neighborhoods, MANAT employs a KD-tree for spatial indexing, which requires O (N log N) time to build. Querying the k-nearest neighbors for all N points incurs O (k log N) time per point, resulting in a total neighborhood construction cost of O (Nk log N).

For each point, all three inconsistency scores are computed through fixed-size operations over its k neighbors:

- Geometric inconsistency (centroid and mean distance): O(k)
- Normal inconsistency (mean normal, angular deviations, standard deviation): O(k)
- Color inconsistency (per-channel standard deviation over k neighbors): O(k)

Each point therefore incurs O(k) cost across all three measures, giving a global complexity of O(Nk) for the full inconsistency computation stage. The subsequent normalization, including computation of the median and MAD for both the inconsistency scores and the density z-score, requires O (N log N) due to sorting. Score fusion and thresholding are both O(N).

The overall complexity is thus dominated by the KD-tree query stage, yielding O (Nk log N). Since k is a fixed, user-defined constant independent of N, the algorithm effectively scales as O (N log N), confirming that MANAT grows near-linearly with dataset size and remains practical for large-scale photogrammetric point clouds.

Results and Discussion

Based on the experimental framework detailed previously, this section presents a comprehensive evaluation of the MANAT algorithm alongside DBSCAN for contextual comparison. By employing both quantitative metrics and qualitative visual assessments, we analyze the performance of the proposed method in handling the complex noise profiles of real-world photogrammetric data.

Quantitative Evaluation

For quantitative evaluation, results from the seven segments were aggregated for each method, representing combined performance across the dataset (Table 1). Under the representative MANAT configuration (k = 8, *mult* = 1), the dense 141.7 million-point cloud was processed in approximately 204 minutes (~1.44 min per million points), removing 23.78% of points identified as inconsistent. Surface roughness decreased from 0.000427 ± 0.000495 m to 0.000386 ± 0.000457 m, reflecting reduced local irregularities and smoother surface geometry. It is worth noting that this metric is computed only in well-sampled neighborhoods, and sparse or missing regions are excluded from the calculation. This ensures that roughness accurately represents dense, reliably reconstructed areas without being skewed by gaps or outliers.

Table 1: Performance comparison of MANAT and the DBSCAN baseline. DBSCAN results are based on voxel-downsampled data due to memory constraints and are provided for contextual reference. Metrics include the number of points retained, removed noise percentage, runtime, global and local NSD, and surface roughness

Method	Number of points	Removed Noise (%)	Runtime (mm:ss)	Global NSD (∘)	Local NSD (∘)	Roughness (m)
Original	141,691,598	-	-	34.11	3.91	0.000427± 0.000495
MANAT	108,000,076	23.78	204:17	32.61	3.64	0.000386± 0.000457
Downsampled	55,347,191	-	-	37.26	4.56	0.000491± 0.000582
DBSCAN	55,343,796	0.006	128:09	37.27	4.56	0.000491± 0.000582

In addition, local NSD decreased from 3.91° to 3.64°, indicating improved coherence within each point's immediate neighborhood. The local metric is computed by calculating the angular difference between each point's normal and the mean normal of its k = 8 nearest neighbors, providing a measure of small-scale surface jitter. This demonstrated that fine-scale surface features were preserved while noise was removed from the surface. In contrast, the global NSD decreased from 34.11° to 32.61°. This single-value metric represents the angular spread of all surface normals across the entire point cloud. The relative decrease in local NSD (~6.9%) is larger than the global reduction (~4.4%), suggesting that MANAT primarily targets localized jitter rather than altering the overall geometric distribution of the model. This indicates that while the cloud becomes more internally coherent, while the broad orientation of the architectural surfaces remains largely unchanged.

These dual-scale normal metrics therefore provide complementary insights: Local normals quantify precision within immediate neighborhoods, revealing the suppression of fine jitter, while the global standard deviation monitors the aggregate orientation of the dataset, ensuring that the overall structural alignment is preserved. Together with runtime, noise removal percentage, and surface roughness, these results demonstrate that MANAT balances local accuracy with global geometric fidelity.

To contextualise MANAT's performance, DBSCAN was evaluated as a widely adopted point cloud processing baseline, commonly applied as a noise removal step. DBSCAN removed only 0.006% of points, with surface quality metrics remaining virtually unchanged before and after processing. This outcome reveals that the noise present in the Fort Margherita dataset is not predominantly composed of spatially isolated outliers

alone but rather includes embedded surface noise inherent to photogrammetric reconstruction. Operating solely on spatial density, DBSCAN lacks the attribute awareness necessary to discriminate between noise and valid surface points in complex photogrammetric data. MANAT addresses this broader noise profile through multi-attribute evaluation of geometric position, surface normal, and colour, achieving 23.78% noise removal with measurable improvements across all surface quality metrics. Although DBSCAN processed a downsampled dataset containing 2.56 times fewer points, its runtime remained relatively high at 128 minutes compared to MANAT's 204 minutes on the full-resolution dataset, indicating more favourable scalability for MANAT relative to the number of processed points.

Fig. 4 presents a radar chart comparing MANAT and DBSCAN across five performance dimensions, with all quality metrics expressed as percentage improvement relative to their respective baselines and processing efficiency normalised by the maximum observed throughput. MANAT consistently dominates across all dimensions, achieving 23.78% noise removal and improvements of 4.40%, 6.91%, and 9.60% in global NSD, local NSD, and roughness respectively, alongside processing efficiency of 693,603 pts/min. In contrast, DBSCAN recorded negligible improvement across all quality metrics (≤0.006%), with processing efficiency of 431,894 pts/min despite operating on a significantly reduced point count. The global NSD for DBSCAN registered a marginal change of -0.027%, indicating negligible degradation. As the radar chart displays improvement metrics, negative values are floored to 0.00% for visualization purposes. The visual disparity between the two polygons confirms that multi-attribute filtering is more effective than DBSCAN's density-based approach for photogrammetric noise removal.

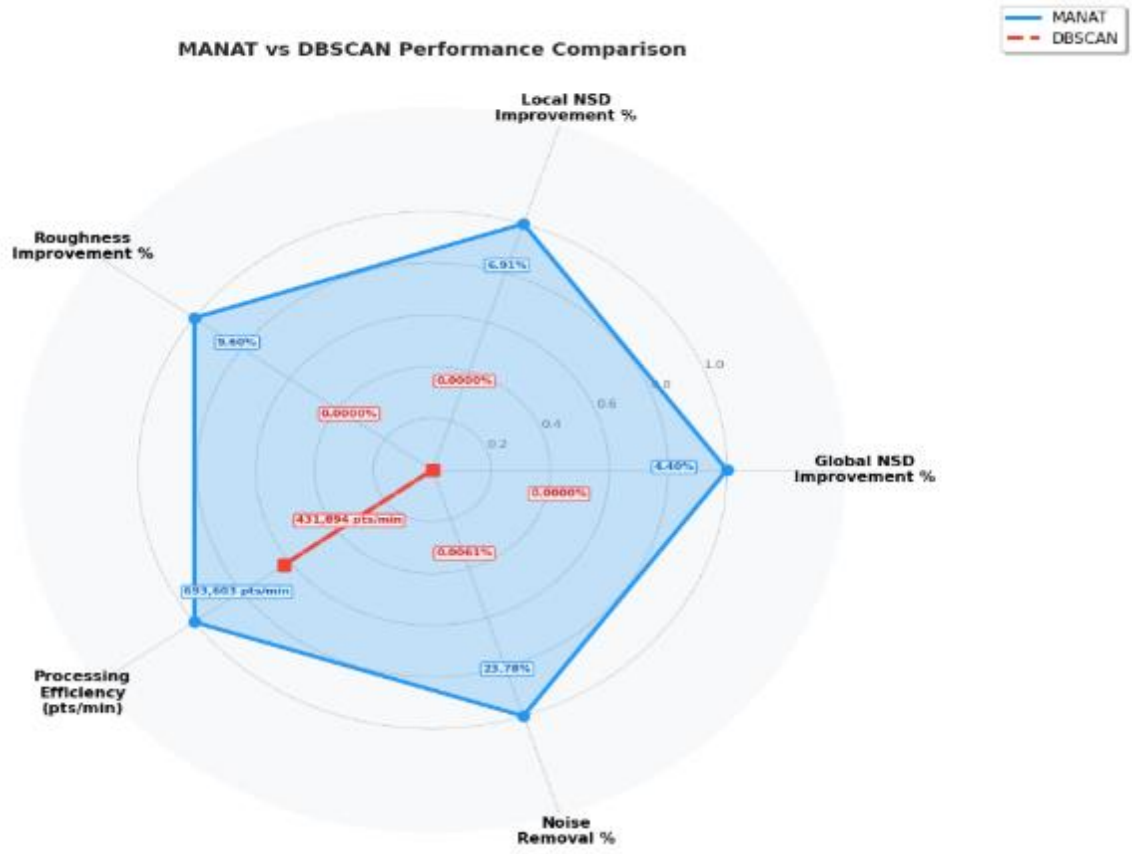

Fig. 4: Radar chart comparing MANAT and DBSCAN across noise removal percentage, global and local NSD improvement, roughness improvement, and processing efficiency

Qualitative Evaluation

To visually assess MANAT's performance, filtered point clouds were compared with the originals from identical viewpoints. The inspection focused on geometric feature preservation, the presence of isolated points, and the continuity of surfaces, providing an intuitive understanding of the method's behavior in real-world data. Visual

comparison for DBSCAN was excluded as it was applied on voxel-downsampled data. Direct visual comparison with the full-resolution original would be confounded by density differences rather than denoising effectiveness. Furthermore, no visually discernible difference was observed between the downsampled original and DBSCAN-filtered outputs, consistent with the quantitative findings reported in Table 1.

Fig. 5 presents an overview of the original and MANAT-filtered point clouds of Fort Margherita. At this scale, the overall structural layout remained visually consistent, demonstrating that MANAT preserves large-scale architectural geometry. Scattered and isolated points, which often represent measurement noise or poorly reconstructed regions, were substantially reduced, improving overall spatial coherence.

Regions of Interest (ROIs) are also presented to highlight common structural elements encountered in architectural point clouds, including planar surfaces, edges, and locally irregular areas.

The visual examples in Fig. 6 illustrate the filtering behavior of MANAT across regions with different structural characteristics and point densities. In the window region (a), the algorithm removed inconsistent points along the wall and window surfaces while maintaining crisp edges and planar continuity. This shows that MANAT can distinguish between true geometric features and isolated noise in densely sampled areas. In the top-down view of the battlements (b), regions with lower visibility and fewer observations during photogrammetric acquisition initially contained noisy scattered points. MANAT selectively removed these outliers, resulting in a more uniform spatial distribution without affecting the overall layout of the roof and battlements.

Fig. 5: Overview of the original and the MANAT-filtered point cloud

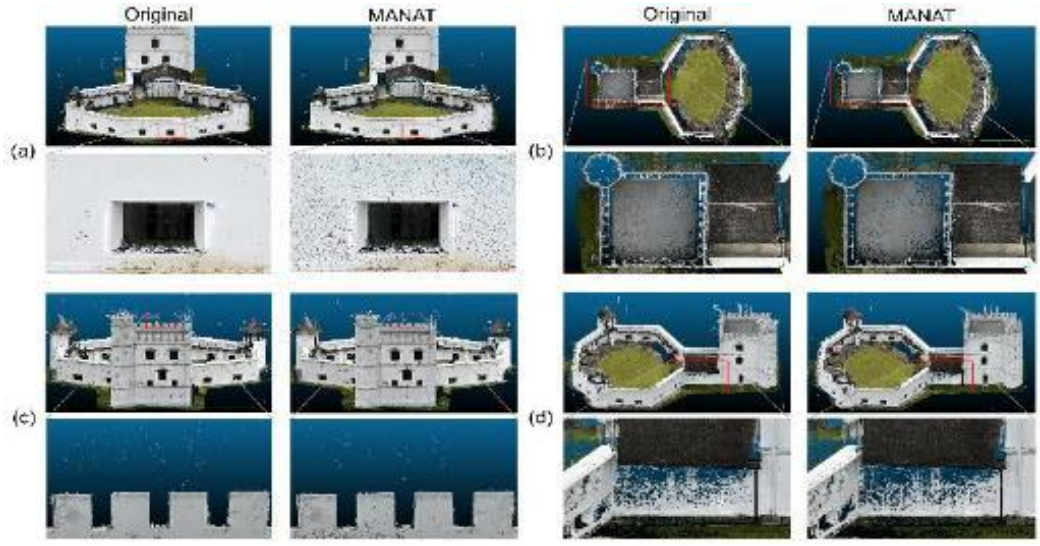

Fig. 6: A visual comparison between the original and MANAT-filtered point clouds across four ROIs. The ROIs include (a) an overview of the fort façade with a zoomed-in view of a window on the perimeter wall, (b) a top-down view of the entire structure with a magnified battlement region, (c) a side view of the fort focusing on the battlements with detailed views of the merlons and crenels, and (d) an oblique view highlighting a wall section beneath the roof

The side-view battlements (c) demonstrate MANAT's ability to handle repetitive structural patterns. Here, surface irregularities and isolated points were reduced, while the merlons and crenels were preserved, indicating that the method adapts its filtering to maintain periodic architectural features. In the sparsely sampled occluded wall region (d), some points were removed, but the dominant wall surface remained identifiable. This highlights MANAT's capacity to balance filtering strength with local point density, avoiding over-pruning in under-sampled areas.

Overall, these qualitative observations reveal that MANAT dynamically adjusts its filtering based on local density and structural context. Fine-scale noise is effectively suppressed, scattered outliers are removed, and planar or repetitive features are preserved. When

considered alongside quantitative metrics such as surface roughness and dual-scale NSD, the results indicate that MANAT achieves both local precision and global geometric integrity. Importantly, this adaptive behavior ensures that the point cloud remains visually coherent and that architectural details are retained, even in regions with sparse coverage or complex geometry.

Ablation Study

An ablation study was conducted to evaluate the individual contribution of each attribute within MANAT. In each single-attribute configuration, only the respective inconsistency score was computed and applied directly without fusion, while all other parameters were kept identical to the full MANAT configuration. Results are summarised in Table 2. The geometry-only configuration produced the lowest removal rate (0.889%), with NSD and roughness values nearly identical to the original scan, indicating that spatial proximity alone is insufficient to distinguish surface jitter from structural detail. Conversely, the normals-only configuration produced the highest removal rate (35.626%), with local NSD increasing to 3.68° compared to full MANAT's 3.64°, suggesting over-aggressive pruning that disrupts local coherence despite improving global orientation. The colour only configuration achieved a removal rate of 24.507%, however local NSD increased to 4.12°, exceeding the original scan's 3.91°. This highlights that colour driven pruning disrupts local surface continuity when used in isolation, as photometric cues lack geometric context to distinguish lighting variation from structural noise. Full MANAT achieves the best balance across all metrics, indicating that multi-attribute fusion is essential. Each attribute contributes complementary information: Geometry provides structural constraints, normals enable orientation refinement, and colour captures photometric inconsistency. Their integration enables selective noise removal that no single-attribute configuration achieves. The increased runtime of full MANAT reflects this additional computational load, which is justified by its more effective denoising performance.

Table 2: Ablation study evaluating the individual contribution of geometry, normal, and color attributes in MANAT

Method	Number of points	Removed Noise (%)	Runtime (mm:ss)	Global NSD (∘)	Local NSD (∘)	Roughness (m)
Full MANAT	108,000,076	23.78	204:17	32.61	3.64	0.000386± 0.000457
Geometry only	140,432,133	0.889	55:20	33.96	3.87	0.000422± 0.000485
Normals only	91,212,076	35.626	96:29	32.29	3.68	0.000370± 0.000460
Colors only	106,967,692	24.507	71:11	32.80	4.12	0.000425± 0.000487

Conclusion

This study presented MANAT, a novel single-stage, density-adaptive point cloud denoising algorithm designed for passive scanning of photogrammetry, which jointly exploits spatial position, surface normals, and color as inherent photogrammetric attributes for unified noise removal. Unlike prior approaches that rely on synthetic noise assumptions or LiDAR-specific metadata, MANAT was evaluated on real-world photogrammetric data with non-uniform density distributions, demonstrating practical applicability under realistic acquisition conditions. On a 141.7 million-point heritage dataset, MANAT achieved 23.78% noise removal with measurable improvements across surface roughness, local NSD, and global NSD of 9.60%, 6.91%, and 4.40%, respectively. Comparison against DBSCAN confirmed that spatial density alone is insufficient for photogrammetric noise removal, highlighting the necessity of multi-attribute filtering that jointly exploits geometric, normal, and color attributes. Future work will focus on validating MANAT across additional real-world photogrammetric datasets of varying complexity and scale, as well as conducting systematic parameter sensitivity analysis to further establish the robustness of the method without extensive parameter tuning. Overall, this work demonstrates the practical effectiveness of MANAT and establishes a robust multi-attribute framework for improving the reliability of photogrammetric 3D reconstruction.

Acknowledgement

The authors would like to thank the publisher for the opportunity to share our work and contribute to the field. We also extend our gratitude to the editorial team for their careful review and thorough editing of the article.

Funding Information

The authors acknowledged the financial support from the Ministry of Higher Education Malaysia through Fundamental Research Grant Scheme (FRGS) FRGS/1/2024/ICT09/UNIMAS/02/1.

Author's Contributions

Yun Sin Chong: Designed the methodology, performed the experiments, and drafted the original manuscript.

Hui Hui Wang and Yin Chai Wang: Supervised the research, reviewed and edited the manuscript, and approved the final version.

Ethics

This manuscript describes original research by the authors. All co-authors have reviewed and approved the manuscript, and no ethical concerns are associated with this work.

References

Aljumaiah, O., Jiang, W., Reddy Addula, S., & Amin Almaiah, M. (2025). Analyzing Cybersecurity Risks and Threats in IT Infrastructure based on NIST Framework. *Journal of Cyber Security and Risk Auditing*, *2025*(2), 12–26. https://doi.org/10.63180/jcsra.thestap.2025.2.2

Ambrosino, A., Di Benedetto, A., & Fiani, M. (2024). Hybrid Denoising Algorithm for Architectural Point Clouds Acquired with SLAM Systems. *Remote Sensing*, *16*(23), 4559. https://doi.org/10.3390/rs16234559

Ataiwe, T. N., Hatem, I., & Al Sharaa, H. M. J. (2021). Digital Model in Close-Range Photogrammetry Using a Smartphone Camera. *E3S Web of Conferences*, *318*, 04005. https://doi.org/10.1051/e3sconf/202131804005

Caciora, T., Herman, G. V., Ilieș, A., Baias, Ștefan, Ilieș, D. C., Josan, I., & Hodor, N. (2021). The Use of Virtual Reality to Promote Sustainable Tourism: A Case Study of Wooden Churches Historical Monuments from Romania. *Remote Sensing*, *13*(9), 1758–1789. https://doi.org/10.3390/rs13091758

Cao, Y., Huang, H., & Yu, D. (2022). Filter Methods for Removing Falling Snow from Light Detection and Ranging Point Clouds in Snowy Weather. *Sensors and Materials*, *34*(12), 4507. https://doi.org/10.18494/sam4047

Charron, N., Phillips, S., & Waslander, S. L. (2018). De-noising of Lidar Point Clouds Corrupted by Snowfall. *Proceedings of the 2018 15th Conference on Computer and Robot Vision (CRV)*, 254–261. https://doi.org/10.1109/crv.2018.00043

Cheng, D., Zhao, D., Zhang, J., Wei, C., & Tian, D. (2021). PCA-Based Denoising Algorithm for Outdoor Lidar Point Cloud Data. *Sensors*, *21*(11), 3703. https://doi.org/10.3390/s21113703

Charles, R. Q., Su, H., Kaichun, M., & Guibas, L. J. (2017). PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. *Proceedings of The2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 77–85. https://doi.org/10.1109/cvpr.2017.16

Di Filippo, A., Antinozzi, S., Limongiello, M., & Messina, B. (2024). An Effective Approach for Point Cloud Denoising In Integrated Surveys. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, *XLVIII-2/W4-2024*, 181–187. https://doi.org/10.5194/isprs-archives-xlviii-2-w4-2024-181-2024

Digne, J., & de Franchis, C. (2017). The Bilateral Filter for Point Clouds. *Image Processing On Line*, *7*, 278–287. https://doi.org/10.5201/ipol.2017.179

Diwakar, M., & Singh, P. (2020). CT image denoising using multivariate model and its method noise thresholding in non-subsampled shearlet domain. *Biomedical Signal Processing and Control*, *57*, 101754. https://doi.org/10.1016/j.bspc.2019.101754

Diwakar, M., Singh, P., & Garg, D. (2024). Edge-guided filtering based CT image denoising using fractional order total variation. *Biomedical Signal Processing and Control*, *92*, 106072. https://doi.org/10.1016/j.bspc.2024.106072

Diwakar, M., Singh, P., Swarup, C., Bajal, E., Jindal, M., Ravi, V., Singh, K. U., & Singh, T. (2022). Noise Suppression and Edge Preservation for Low-Dose COVID-19 CT Images Using NLM and Method Noise Thresholding in Shearlet Domain. *Diagnostics*, *12*(11), 2766. https://doi.org/10.3390/diagnostics12112766

Fabris, M., Fontana Granotto, P., & Monego, M. (2023). Expeditious Low-Cost SfM Photogrammetry and a TLS Survey for the Structural Analysis of Illasi Castle (Italy). *Drones*, *7*(2), 101. https://doi.org/10.3390/drones7020101

Gao, X., Hu, W., & Guo, Z. (2018). Graph-Based Point Cloud Denoising. *Proceedings of the 2018 IEEE Fourth International Conference on Multimedia Big Data (BigMM)*, 1–6. https://doi.org/10.1109/bigmm.2018.8499090

Gonizzi Barsanti, S., Marini, M. R., Malatesta, S. G., & Rossi, A. (2024). Evaluation of Denoising and Voxelization Algorithms on 3D Point Clouds. *Remote Sensing*, *16*(14), 2632. https://doi.org/10.3390/rs16142632

Jasińska, A., Pyka, K., Pastucha, E., & Midtiby, H. S. (2023). A Simple Way to Reduce 3D Model Deformation in Smartphone Photogrammetry. *Sensors*, *23*(2), 728. https://doi.org/10.3390/s23020728

Jia, C., Yang, T., Wang, C., Fan, B., & He, F. (2019). A new fast filtering algorithm for a 3D point cloud based on RGB-D information. *PLOS ONE*, *14*(8), e0220253. https://doi.org/10.1371/journal.pone.0220253

Katta, S., Singh, P., Garg, D., & Diwakar, M. (2024). A Hybrid Approach for CT Image Noise Reduction Combining Method Noise-CNN and Shearlet Transform. *Biomedical and Pharmacology Journal*, *17*(3), 1875–1898. https://doi.org/10.13005/bpj/2991

Kim, Y.-K., & Yun, S.-H. (2024). Performance analysis of point cloud preprocessing algorithms suitable for construction progress analysis. *Library Progress International*, *44*(2), 631–638. https://doi.org/10.48165/bapas.2024.44.2.1

Kurup, A., & Bos, J. (2021). DSOR: A scalable statistical filter for removing falling snow from LiDAR point clouds in severe winter weather. *Computer Vision and Pattern Recognition*. https://doi.org/10.48550/arXiv.2109.07078

Le, M.-H., Cheng, C.-H., Liu, D.-G., & Nguyen, T.-T. (2022). An Adaptive Group of Density Outlier Removal Filter: Snow Particle Removal from LiDAR Data. *Electronics*, *11*(19), 2993. https://doi.org/10.3390/electronics11192993

Lin, G., Li, G., Giordano, A., Sang, K., Stendardo, L., & Yang, X. (2024). Three-Dimensional Documentation and Reconversion of Architectural Heritage by UAV and HBIM: A Study of Santo Stefano Church in Italy. *Drones*, *8*(6), 250. https://doi.org/10.3390/drones8060250

Lin, W.-C., Lee, M.-Z., Chou, H.-S., Lin, Y.-J., Li, K.-C., Lin, T.-L., & Chen, S.-L. (2023). 3D Point Cloud Denoising Based on Color Attribute. *Proceedings of the 2023 Asia Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, 1512–1516. https://doi.org/10.1109/apsipaasc58517.2023.10317301

Liu, S., Mohd Jaki, B. M., Huang, Y., Dong, M., & Liu, Q. (2025). Multiscale hierarchy denoising method for heritage building point cloud model noise removal. *Npj Heritage Science*, *13*(1), 199. https://doi.org/10.1038/s40494-025-01639-5

Maset, E., Valente, R., Iamoni, M., Haider, M., & Fusiello, A. (2022). Integration of photogrammetry and portable mobile mapping technology for 3D modeling of cultural heritage sites: The case study of the Bziza Temple. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, *XLIII-B2-2022*, 831–837. https://doi.org/10.5194/isprs-archives-xliii-b2-2022-831-2022

Mazzonetto, I., Castellaro, M., Cooper, R. J., & Brigadoi, S. (2022). Smartphone-based photogrammetry provides improved localization and registration of scalp-mounted neuroimaging sensors. *Scientific Reports*, *12*(1), 10862. https://doi.org/10.1038/s41598-022-14458-6

Mengiste, E., Prieto, S. A., & García de Soto, B. (2022). Comparison of TLS and Photogrammetric 3D Data Acquisition Techniques: Considerations for Developing Countries. *Proceedings of the International Symposium on Automation and Robotics in Construction (IAARC)*, 491–494. https://doi.org/10.22260/isarc2022/0067

Patonis, P. (2024). A Comparative Study on the Use of Smartphone Cameras in Photogrammetry Applications. *Sensors*, *24*(22), 7311. https://doi.org/10.3390/s24227311

Quispe-Enriquez, O. C., Valero-Lanzuela, J. J., & Lerma, J. L. (2023). Smartphone Photogrammetric Assessment for Head Measurements. *Sensors*, *23*(21), 9008. https://doi.org/10.3390/s23219008

Ren, Y., Li, T., Xu, J., Hong, W., Zheng, Y., & Fu, B. (2021). Overall Filtering Algorithm for Multiscale Noise Removal from Point Cloud Data. *IEEE Access*, *9*, 110723–110734. https://doi.org/10.1109/access.2021.3097185

Rezaei, S., Maier, A., & Arefi, H. (2024). Quality Analysis of 3D Point Cloud Using Low-Cost Spherical Camera for Underpass Mapping. *Sensors*, *24*(11), 3534. https://doi.org/10.3390/s24113534

Sangmanee, W., & Suwanwerakamtorn, R. (2023). Integrated Framework for Virtual Tours and 3D Visualization of Cultural Tourism in Pattani, Thailand Based on WebGIS Platform. *International Journal of Geoinformatics*, *19*(9), 44–60. https://doi.org/10.52939/ijg.v19i9.2835

Singh, P., Diwakar, M., Pandey, N. K., Singh, R., Sisodia, D., Arya, C., & Chakraborty, C. (2023). Low-dose COVID-19 CT Image Denoising Using CNN and its Method Noise Thresholding. *Current Medical Imaging Reviews*, *19*(2), 182–193. https://doi.org/10.2174/1573405618666220404162241

Triviño-Tarradas, P., Mohedo-Gatón, A., Carranza-Cañadas, P., & Hidalgo-Fernandez, R. E. (2024). The Application of Metaverse in the Tourism Sector as a Tool for Enhancing Sustainability—Case Study: A Medieval 'Perfume Burner' of the Local Historical Museum of Montilla (Cordoba, Spain). *Sustainability*, *16*(16), 6966. https://doi.org/10.3390/su16166966

Tang, W. T., Siu, T., Villupuram, S. S., Lee, T., & Ji, M. (2024). Digitalisation of asset management in Chinese temples in Hong Kong. *HKIE Transactions*, *30*(3), 34–43. https://doi.org/10.33430/v30n3thie-2022-0053

Van, T. N., & Naprstkova, N. (2024). Accuracy of Photogrammetric Models for 3D printed Wrist-hand Orthoses. *Manufacturing Technology*, *24*(3), 458–464. https://doi.org/10.21062/mft.2024.048

Wang, L., Chen, Y., Song, W., & Xu, H. (2024). Point Cloud Denoising and Feature Preservation: An Adaptive Kernel Approach Based on Local Density and Global Statistics. *Sensors*, *24*(6), 1718. https://doi.org/10.3390/s24061718

Yang, H., Wang, W., Wang, Y., & Wang, P. (2025). Novel method for robust bilateral filtering point cloud denoising. *Alexandria Engineering Journal*, *127*, 573–585. https://doi.org/10.1016/j.aej.2025.04.099

Yang, S., Hou, M., & Li, S. (2023). Three-Dimensional Point Cloud Semantic Segmentation for Cultural Heritage: A Comprehensive Review. *Remote Sensing*, *15*(3), 548. https://doi.org/10.3390/rs15030548

Yuan, X., Liu, B., & Ma, Y. (2020). Anisotropic neighborhood searching for point cloud with sharp feature. *Measurement and Control*, *53*(9–10), 1943–1953. https://doi.org/10.1177/0020294020964245

Yan, F., & Jinlei, Z. (2015). Research on scattered points cloud denoising algorithm. *Proceedings of the 2015 IEEE International Conference on Signal Processing, Communications and Computing (ICSPCC)*, 1–5. https://doi.org/10.1109/icspcc.2015.7338887

Zhao, Q., Gao, X., Li, J., & Luo, L. (2021). Optimization Algorithm for Point Cloud Quality Enhancement Based on Statistical Filtering. *Journal of Sensors*, *2021*(1), 7325600. https://doi.org/10.1155/2021/7325600

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

Zou, B., Qiu, H., & Lu, Y. (2020). Point Cloud Reduction and Denoising Based on Optimized Downsampling and Bilateral Filtering. *IEEE Access*, *8*, 136316–136326. https://doi.org/10.1109/access.2020.3011989