Reports on Geodesy and Geoinformatics, 2026, Vol. 121, pp. 1–18

DOI: 10.2478/rgg-2026-0001
Received: 15 November 2025 / Accepted: 20 March 2026
Published online: 15 April 2026

ORIGINAL ARTICLE

Evaluation of open-source libraries and commercial software utilised for dense point cloud generation: a case study of cultural heritage objects

Anna Michałek [1*], Jakub Markiewicz [2], Justyna Wójcik-Leń [1], Adam Kostrzewa [1,3], Michał Kowalczyk [2], Sławomir Łapiński [2] and Adrian Macek [2]

[1]Faculty of Environmental Engineering, Geodesy and Renewable Energy, Kielce University of Technology, 7 Tysiąclecia Państwa Polskiego Avenue, 25-314 Kielce, Poland
[2]Faculty of Geodesy and Cartography, Warsaw University of Technology, 1 Politechniki Square, 00-661 Warsaw, Poland
[3]Institute of Civil Engineering, Warsaw University of Life Sciences, Nowoursynowska 166, 02-787 Warsaw, Poland

*amichalek@tu.kielce.pl

Abstract

The development of image-based approaches for 3D dense reconstruction has found wide application in the architectural documentation of cultural objects and sites. For this reason, it is crucial to understand the potential and limitations of Multi-View Stereo (MVS) algorithms used in 3D shape reconstruction. This article aims to evaluate the quality and accuracy of dense point generation for cultural heritage objects using open-source (MicMac, OpenMVS, RealityScan), commercial algorithms (Agisoft Metashape and Pix4D), and an Unmanned Aerial Vehicle (UAV) in parallel and skew configurations at two historic wooden structures. To achieve this, a workflow for data evaluation was proposed, incorporating eigenvalues and quality factors, such as planarity, roughness, normal vector variance, and cloud-to-TLS (Terrestrial Laser Scanning) cloud deviation analysis. The results demonstrate that the image-acquisition geometry strongly influences reconstruction accuracy. Parallel configurations consistently produced higher completeness, narrower deviation distributions, and improved geometric fidelity compared to skew configurations. Among commercial solutions, Agisoft Metashape and Pix4D achieved the best overall performance. Open-source algorithms yielded variable results: OpenMVS-SGM achieved competitive accuracy under optimal conditions, while the OpenMVS-Patch-Based Approach offered a balance between density and roughness.

Key words: cultural heritage, dense point clouds, Multi-View Stereo, quality assessment

1 Introduction

Cultural heritage plays a fundamental role in modern societies, preserving tangible and intangible evidence of the past. Therefore, it is essential not only to safeguard and protect it, but also to digitise and present cultural heritage assets using modern technologies in order to increase their accessibility to a broader audience (Bochenska et al., 2023).

Recent years have seen significant development and wider application of both active methods (also known as passive methods, such as close-range photogrammetry and multi-view stereo approaches) and range-based methods (e.g., Terrestrial Laser Scanning) for the

inventory of cultural heritage objects and the generation of architectural documentation (Abbate et al., 2019; Arif and Essa, 2017; Doroszuk and Markiewicz, 2022; Giżyńska et al., 2022; Muradov et al., 2022; Murtiyoso and Grussenmeyer, 2017; Remondino, 2011; Tobiasz et al., 2019).

Close-range photogrammetry (one of the image-based/passive techniques) has a long history of use in the inventory of cultural heritage. In recent years, significant advances have been made in the methods and algorithms used for image processing to digitise cultural heritage (Grilli and Remondino, 2019; Murtiyoso and Grussenmeyer, 2017; Murtiyoso et al., 2022; Stathopoulou et al., 2021; Stathopoulou and Remondino, 2023). 3D shape reconstruction from images (combined Structure-from-Motion and Multi-View Stereo) is widely used in heritage documentation due to its lower cost and high accuracy (Schönberger et al., 2016; Schonberger and Frahm, 2016). However, it requires more expertise in data acquisition and processing.

In close-range photogrammetry, point clouds and 3D meshes are among the primary outputs, requiring less user interaction. The emergence of Multi-View Stereo (MVS) methods has enabled the production of dense point clouds from multiple images. For this reason, it is essential to understand the capabilities and limitations of the MVS algorithms and methods used to generate these dense point clouds.

Open-source solutions often provide greater flexibility for users, allowing customisation of processing workflows and access to algorithmic parameters that are otherwise hidden in commercial software (Kostrzewa et al., 2025). Additionally, their lower cost and accessibility enable a wider range of researchers and practitioners to experiment with and adopt advanced MVS methods.

In parallel with the continued development of traditional MVS techniques, recent years have seen rapid progress in learning-based approaches to 3D reconstruction, particularly Neural Radiance Fields (NeRFs) (Mildenhall et al., 2021) and, more recently, 3D Gaussian Splatting (Kerbl et al., 2023). These methods represent a paradigm shift, enabling continuous scene representations learned directly from image collections and enabling high-quality novel-view synthesis and dense scene reconstruction (Müller et al., 2022; Tancik et al., 2023).

NeRF-based approaches have been successfully applied to cultural heritage documentation, demonstrating impressive visual realism and the ability to handle complex geometries and lighting conditions. However, their practical use in metric documentation and quantitative analysis is still limited by high computational demands, long training times, and challenges related to scale, georeferencing, and the extraction of explicit, metrically reliable point clouds or surfaces (Matsuki et al., 2024; Wang et al., 2025).

Recently proposed 3D Gaussian Splatting methods addressed some of these limitations by enabling faster training and real-time rendering, while still relying on implicit or semi-implicit scene representations. Although open-source implementations of both NeRF- and Gaussian-based methods are increasingly available, their integration into established photogrammetric workflows and their suitability for rigorous geometric evaluation remain active research topics.

For these reasons, this study focuses on well-established MVS algorithms that directly generate dense, explicit point clouds, which remain the primary data product in architectural documentation, conservation, and quantitative analyses of cultural heritage objects. Nevertheless, learning-based reconstruction methods are recognised as an important and rapidly evolving research direction that may complement or extend traditional MVS techniques in future studies.

This article aims to compare the Semi-Global Matching (SGM) and Patch-Based approaches (MVS algorithms) implemented in both open-source OpenMVS (2025), RealityScan (2025), MicMac (2025), and in commercial applications Agisoft Metashape (2025) and Pix4D (2025). To evaluate the applicability and analyse the capabilities and limitations of each solution, images acquired using a low-cost Unmanned Aerial Vehicle (DJI Mini 2) were used. The test sites included a historic granary and a wooden church in the Kielce Countryside Museum – Ethnographic Park in Tokarnia. To determine the most suitable solution, the following aspects were analysed: (1) point cloud density, (2) roughness, (3) planarity, and (4) vector normal variance. Additionally, the generated point clouds were compared with data obtained from a Leica BLK 360 laser scanner.

This paper is divided into five main sections. Section 2 outlines the primary principles of the MVS method, as well as the algorithms implemented in both commercial and open-source applications and libraries. In Section 3, the overview of the utilised sensors and the proposed methodology for data processing was presented. Section 4 summarises the results of the point cloud assessments generated using commercial and open-source solutions. In the conclusion (Section 5), future work is proposed, and the possibilities and limitations of utilised algorithms and approaches are summarised.

2 Multi-view stereo algorithms

The MVS is a computer vision technique that reconstructs three-dimensional scenes or objects from two-dimensional images captured from various viewpoints. This method applies the triangulation principle to determine the spatial coordinates of points in the scene, producing a dense and accurate representation of the geometry of the object.

The MVS process involves identifying corresponding points visible across multiple images and then triangulating them to compute their 3D coordinates. The effectiveness of this technique depends on advanced algorithms capable of handling challenges such as occlusions, lighting variations, and differences in image perspective. Progress in computational power and the availability of modern algorithms have significantly improved the efficiency and accuracy of MVS methods. As a result, with its ability to generate detailed 3D reconstructions from sets of 2D images, MVS has become a widely adopted tool for object modelling, scene reconstruction, and spatial analysis across various scientific and industrial fields.

Numerous commercial and open-source applications and solutions that utilise the MVS approach are available on the market. One of the key solutions is Agisoft Metashape, which is currently considered state-of-the-art and widely used for generating dense point clouds from image networks, among other things. It is commercial software for the photogrammetric processing of digital images and the generation of 3D spatial data, used in GIS applications, cultural heritage documentation, visual effects production, and indirect object measurements. Valued for its ability to process digital images and automatically generate high-quality 3D spatial data, Agisoft Metashape utilises the SGM algorithm to generate dense point clouds. This method computes depth maps from overlapping stereo image pairs by maximising pixel correspondence while minimising energy cost. The depth maps are then used to densify the sparse point cloud, resulting in a high-resolution 3D representation of the scene (Agisoft Photogrammetry Workflow Compendium, 2025).

Another commonly used software, especially for processing UAV imagery, is Pix4Dmapper, an advanced photogrammetry tool that generates accurate 3D models and maps. It can process data from traditional RGB cameras, LiDAR sensors, or a combination of both technologies. Based on the application behaviour and the available parameter settings, it can be inferred that the Pix4D software implements a patch-based approach to generate dense point clouds.

A separate category of solutions includes open-source applications and libraries. One of the main advantages of using such tools is that they are not treated as black boxes, and the algorithms used to generate dense point clouds are well-documented and transparent. A prominent example is COLMAP, which is commonly used for im-

age orientation using the Structure-from-Motion (SfM) approach. It relies on feature extraction and matching algorithms (including a modified SIFT algorithm), optimised Bundle Adjustment procedures, and a standardised camera parameter model. COLMAP ensures high accuracy in estimating projection matrices and the geometric structure of the 3D scene. The output data generated by COLMAP, namely internal camera parameters obtained through self-calibration during the bundle adjustment stage, external orientation parameters of the images, and the sparse point cloud, serve as the foundation for further data processing with MVS, Neural Radiance Fields (NeRF), and 3D Gaussian Splatting algorithms. For this reason, COLMAP serves not only as a reference SfM method but also as an essential component in modern reconstruction workflows.

Another commonly utilised method is the OpenMVS library, which is used for 3D shape reconstruction. This library contains implementations of both the Semi-Global Matching and patch-based approaches for generating dense point clouds. It should be noted that OpenMVS is among the most widely utilised open-source MVS workflows, offering transparent, reproducible, and modular processing steps that provide a robust framework for evaluating dense reconstruction quality across different workflows.

MicMac is an advanced open-source photogrammetric software developed by the Institut National de l'Information Géographique et Forestière and the École Nationale des Sciences Géographiques within the LASTIG laboratory. The project was initiated in 2005 and has been distributed under the CECILL-B license since 2007 (Rupnik et al., 2017). MicMac is widely used for processing satellite imagery, aerial photographs (both analogue and digital), and drone images. Dense point cloud generation in MicMac relies on several algorithms differing in methodology, scalability, and computational efficiency (MicMac, 2025): (1) Malt – a classical dense matching module based on the SGM method. It produces orthorectified images for orthomosaic creation, as well as binary .tif and associated .xml files representing dense point clouds. Malt offers three modes – Ortho, Urban, and GeoImage – optimised for diverse types of terrain and image content; (2) PIMs (Parallel Image Matching System) – a newer and more scalable algorithm employing PMVS. It supports parallel and distributed computation, significantly accelerating processing compared to Malt, and generates denser point clouds; (3) C3DC (Campagne 3D Cloud) – the latest module for dense reconstruction in MicMac, offering several operation modes: BigMac (full dense reconstruction), MicMac (faster, slightly less accurate reconstruction), and QuickMac (a rapid preview mode for assessing preliminary results). MicMac was selected for this study because it utilises a photogrammetric approach and is characterised by its multi-module dense matching framework. This framework provides a valuable point of comparison for evaluating how different methodological approaches influence reconstruction quality.

RealityScan (formerly RealityCapture) is a free photogrammetry application for individuals, small businesses, and educational users. It transforms real-world scenes into detailed digital models using a patch-based MVS algorithm to generate dense point clouds. Widely used across industries from game development to cultural heritage preservation, it supports scalable data processing with no limit on image count, offers Command Line Interface and API access, and includes an internal processing node, Node RealityScan, for automation. This solution was chosen for this comparative study because it represents a modern, widely adopted commercial MVS solution that provides high-quality, dense reconstructions with a streamlined workflow.

3 Methods and materials

3.1 Overview of the approach

This study evaluates the algorithms implemented in open-source and commercial software utilised for dense point cloud generation. The parameters listed in Table 1 were used to generate dense point clouds. Their selection ensured the generation of dense point clouds with comparable characteristics across all tested solutions. This consistency was essential for enabling a fair and objective evaluation of the different methods.

Various geometric features were calculated to assess the quality of the point clouds. These features are typically derived using Principal Component Analysis (PCA), which analyses the overall structure of the dataset to extract directional or gradient information indicating where most data points are concentrated. In this context, eigenvectors represent the directions of the principal gradients, while eigenvalues indicate the magnitude or strength of variation along those directions. This approach enables a detailed understanding of the spatial characteristics of the point cloud, facilitating further analysis of its accuracy and structural consistency. Eigenvalues were used to calculate the values of planarity and normal change rate. The eigenvalues λ_1, λ_2, and λ_3 of the structural tensor, where $\lambda_1 > \lambda_2 > \lambda_3 > 0$, provide a direct description of the local 3D structure (Weinmann et al., 2013, 2017). These geometric features are derived from the eigenvalues of the covariance matrix computed within a local neighbourhood around each point. In addition to the eigenvalues, the analysis also considers the distribution of points around the corresponding eigenvectors and features calculated within vertical columns centred on each point (Hackel et al., 2016). The covariance matrix reflects how much the data varies along each spatial dimension relative to the mean. The eigenvalues of this matrix quantify the degree of variation in the direction of their associated eigenvectors (Bazazian et al., 2015). Consequently, the largest and smallest eigenvalues correspond to strong and weak spatial variation directions, respectively, and are essential for describing surface characteristics such as flatness and curvature.

To evaluate the quality of the point clouds, the following parameters were analysed:

(1) Point density

The first and most fundamental aspect of point cloud quality is density, which can be expressed using Equation (1).

$$\text{density} = \frac{(1 + \text{number\_of\_nn\_points\_in radius})}{\frac{4}{3}\pi r^3} \tag{1}$$

where r – radius used for determining nearest neighbour points, number_of_nn_points_in radius – number of nearest points.

It is typically measured in points per cubic meter (pts/m^3) and reflects the number of points distributed within a defined surface area of the point cloud. A higher surface density is generally associated with more reliable reconstruction results and improved accuracy correlating internal and external structural features, such as those found in wooden architectural elements.

(2) Roughness

This parameter reflects the distance between the point and the best-fit plane based on the nearest neighbours, providing insight into the spatial distribution and potential noise within the point cloud. In this study, elevated roughness values may indicate the presence of noise or artefacts in the data. Roughness is commonly quantified as the standard deviation of point elevations within a local neighbourhood around each point – that is, the deviation of points from the best-fit plane, which is typically associated with the smallest eigenvalue.

(3) Planarity

This geometric descriptor quantifies how well a set of points in a local neighbourhood aligns with a two-dimensional plane. It is calculated using the eigenvalues of the covariance matrix derived from the 3D coordinates of the points in the neighbourhood, as

Table 1. The parameters utilised for point cloud generation used in different software and algorithms

No.	Software/Algorithm	Type	Quality Parameter (Image resolution)	Number of Views/ Number of matches	Fusion mode	Number of iterations	Matching window size
1	Agisoft Metashape (version 2.1.2)	Commercial	1	x	x	x	x
2	Pix4D (version 4.8.4)	Commercial	1 (Full resolution)	3	x	x	7x7
3	RealityScan (version 2.0.1)	Commercial/ Open-source	High (full resolution)	x	x	x	x
4	OpenMVS – PM (version 2.0)	Open-source	1	0 – all possible	0	6	x
5	OpenMVS – SGM (version 2.0)	Open-source	1	0 – all possible	-1 -2	6	x
6	MicMac	Open-source	1	>2	best results	x	5x5

follows:

$$\text{planarity} = \frac{(\lambda_2 - \lambda_3)}{\lambda_1} \tag{2}$$

where:

- $\lambda_1, \lambda_2, \lambda_3$ are the eigenvalues of the covariance matrix, ordered such that $\lambda_1 > \lambda_2 > \lambda_3 > 0$,
- λ_1 represents the direction of the most significant variance (typically aligned with the surface normal),
- λ_2 andλ_3 represent the intermediate and smallest variances, respectively.

This ratio expresses the difference in spread between the second and third principal directions, normalised by the largest eigenvalue. A higher planarity value indicates that the points are more likely to lie on a flat, two-dimensional surface, as the variance in the third direction (perpendicular to the plane) is minimal. Conversely, lower values suggest a more isotropic or irregular distribution of points, potentially indicating surface roughness, curvature, or noise.

(4) Normal Vector Variance

Normal Vector Variance is a geometric descriptor that quantifies the angular variation of surface normals within a local neighbourhood. It indicates the smoothness or irregularity of a surface: (1) Low variance suggests a smooth, planar surface with consistent normal directions; (2) High variance may indicate surface roughness, noise, or complex geometry such as edges or corners. This parameter is calculated using the eigenvalues of the covariance matrix derived from the 3D coordinates of points in a local neighbourhood, as follows:

$$\text{normal vector variance} = \frac{\lambda_3}{(\lambda_1 + \lambda_2 + \lambda_3)} \tag{3}$$

where:

- $\lambda_1, \lambda_2, \lambda_3$ are the eigenvalues of the covariance matrix, ordered such that $\lambda_1 > \lambda_2 > \lambda_3 > 0$,
- λ_3 corresponds to the smallest eigenvalue, representing the variance in the direction orthogonal to the local surface.

This ratio expresses the relative contribution of the smallest eigenvalue to the neighbourhood total variance. A higher value indicates greater deviation from planarity, suggesting more irregular or noisy surfaces, whereas a lower value implies a more uniform, smoother surface.

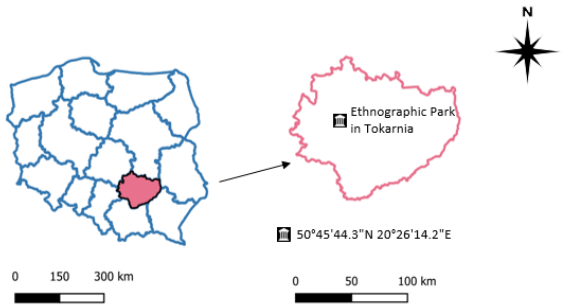

Figure 1. Localisation of the Ethnographic Park in Tokarnia – an open-air museum of the Kielce Countryside [1]

3.2 Test site description

To verify the impact of selecting an algorithm or application for dense point cloud generation, two types of Test Sites located in the Museum of the Kielce Countryside (Poland) were used (see Figure 1). Ethnographic Park in Tokarnia – an open-air museum of the Kielce Countryside, located in the Czarna Nida Valley, in the village of Tokarnia, Chęciny commune, Kielce County, Świętokrzyskie voivodeship. The museum, covering 65 hectares, comprises 38 groups of buildings arranged into sectors: small-town, upland, manor and farmstead, Świętokrzyskie architecture, Vistula River region, and loess terrain.

Two Test Sites were selected for the study (see Figure 2). The selection of two Test Site types was driven by the characteristics of the acquired data, which exhibit heterogeneous structural attributes and surface geometries, allowing for the assessment of the efficacy and quality of generating dense point clouds using different algorithms. The sites were selected by the staff of the Kielce Countryside Museum, who chose the most valuable and representative buildings within the Ethnographic Park.

The Test Site I is a granary from Rogów, built in 1684. This date is confirmed by the carving on two original copper weathervanes mounted on the roof ridge, one of which bears the painted coats of arms of the former owners – the Firlej and Wodzicki families. The log walls of the building were made of wood and rein-

1 Source:https://www.geoportal.gov.pl/pl/dane/panstwowy-rejestr-granic-prg

(a) (b)

Figure 2. Localisation of the Test Sites in the open-air museum of the Kielce Countryside: (a) granary from Rogów, (b) wooden church from Rogów [2]

forced with a post-and-plank framework. The foundation and the ground floor are built of red brick. These red brick foundations and the ground floor were added around 1870, when the estate owner moved the granary from its original location on the Vistula River to the outskirts of the village near the seventeenth-century Rogów manor house, which no longer exists. The building measures approximately 26 × 14 × 5 metres, forming a compact rectangular structure. The walls of the third storey were constructed using the corner-post-and-filler technique. The post-and-plank structure features massive corner posts with a 50 × 50 centimetres cross-section, each made from a single piece of larch. Originally, the corner-notched log walls of the lower storeys were made of firwood and were additionally reinforced with a Lusatian-type structure, of which the larch corner posts are a characteristic element. Among the surviving original architectural details are the oak trick doors on the first and second floors, set on wrought belt-strap hinges.

The entire building is covered with a hipped roof, double-layered with wooden shingles. The granary window openings are semicircular and secured with forged iron grilles (Granary from Rogów, 2026).

The structure examined as Test Site II is a wooden church originating from Rogów, donated by the Diocesan Curia in Kielce. The hospital church, dedicated to Our Lady of Consolation, was founded in Rogów in 1763 by Michał Wodzicki, heir to the Rogów estate, Grand Crown Under Chancellor, and Bishop of Przemyśl. The church walls were constructed using a log-building technique, with larch logs set on oak foundations. The rafter-and-purlin roof structure is covered with wooden shingles. Inside the church, particular attention is drawn to the richly ornamented pulpit and the original Rococo side altars. This building constitutes an exemplary representation of Baroque wooden ecclesiastical architecture. Baroque wooden churches were characterised by a tall, steeply pitched shingle-covered roof that visually dominated the low wall structure beneath. At the roof ridge stands an openwork bell turret (sygnaturka) crowned with a Baroque helmet-shaped cupola. Externally, the walls are clad with vertical boarding, giving the façades a smooth, uniform appearance. The church measures approximately 16 × 7 × 11 metres, forming a characteristic, compact Baroque wooden structure. In the open-air museum, the church continues to serve its original function, remaining an active sacred site where masses are regularly celebrated (Baroque church from Rogów, 2026).

To evaluate the quality of algorithms used to generate dense point clouds, images acquired with the DJI Mini 2 UAV were selected. The sensor parameters used are presented in Table 2. The images were captured in two scenarios – parallel and oblique to the studied object. An example of image distribution is shown in Figure 3. The selection of the DJI Mini 2 drone was motivated by its low cost, which may be associated with higher measurement uncertainties and lower image quality compared to platforms equipped with medium-format cameras. Yet, it is widely utilised in the documentation of cultural heritage objects. The images were acquired with a nominal overlap of approximately 80%. For Test Site I, lighting conditions varied due to shadows, penumbra, and variable illumination across the object. In contrast, Test Site II was surveyed under full shading conditions, resulting in uniform illumination of the object. These circumstances provided optimal conditions for consistent image acquisition.

2 Ethnographic Park of the Kielce Village Museum, Source: https://fotopolska.eu – Poland in Photography, 2025

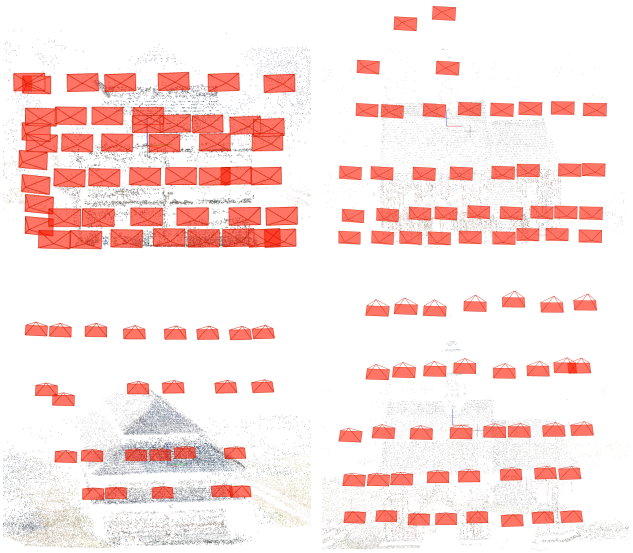

(a) (b)

Figure 3. Example of (a) parallel and (b) skew images obtained for the utilised Test Sites

As reference data, point clouds obtained from the terrestrial laser scanner Leica BLK 360 were used. Table 2 presents the basic parameters of the equipment used, the accuracy of point cloud orientation, and key information regarding the reference data.

4 Results and discussion

4.1 Quality assessment of Structure-from-Motion

The results of data processing using the Structure-from-Motion (SfM) method were evaluated in the first stage to determine the possibility of utilising open-source and commercial applications and algorithms for generating dense point clouds. For this purpose, Agisoft Metashape, Pix4D, Colmap, RealityScan, and MicMac were examined. The investigation was based on the assessment of image orientation reprojection error values. Additionally, the number and distribution of tie points were analysed, which is crucial for the MVS. These points are used to densify the point cloud (Shen, 2013). In Figures 4a and 4b, the values of the reprojection error (Root Mean Square Error – RMSE) in pixels are presented, while Figures 4c and 4d show the number of detected tie points. Figure 5 presents the spatial distribution of tie points.

Analysis of RMSE values revealed consistent patterns across both test sites, with similar trends observed in all software packages. For Test Site I (Figure 4a), the lowest RMSE values were obtained using Pix4D (0.18 and 0.17 pixels). For Test Site II (Figure 4b), Pix4D again produced the most accurate results (0.20 and 0.19 pixels), confirming the high reconstruction reliability of this software. The highest RMSE values were recorded for Agisoft Metashape and MicMac – for Test Site I, 0.68 and 0.78 pixels in Agisoft Metashape and 0.82 and 0.73 pixels in MicMac. In contrast, for Test Site II, the corresponding errors reached 0.82 and 0.89 pixels (Agisoft Metashape)

Table 2. Parameters of the image-based (DJI Mini 2) and range-based (Leica BLK360) approaches used in this investigation.

		Test Site I	Test Site II
		Image-based approach	
Sensor parameters		Camera Model: DJI FC7303 (DJI mini2) Sensor Dimensions: 4000x2550 (px) Focal length: 4 mm Sensor Pixel Size: 1.79 μm	
No. of images	**Parallel**	22	38
	Skew	17	43
		Range-based approach (TLS)	
Sensor parameters		Leica BLK360 TLS (830 nm wavelength) Scanning resolution: 12 mm/10 m 3D point accuracy: 4 mm/ 10 m	
No. of scans		33	21
Registration accuracy		Strength: 66 % Overlap: 40 % RMS Error: 0.011 m	Strength: 43 % Overlap: 46 % RMS Error: 0.009 m

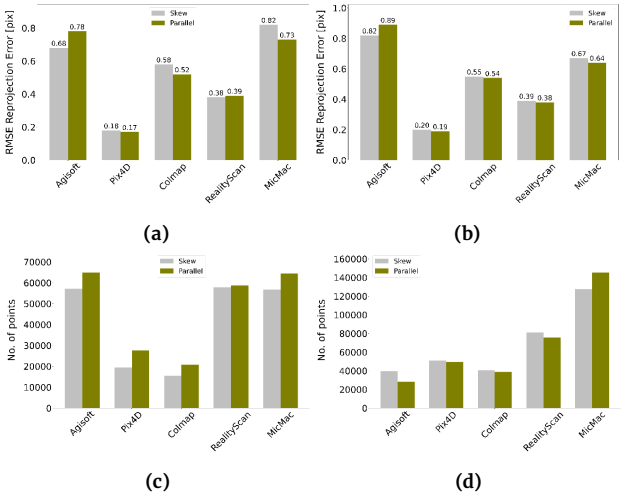

Figure 4. Plots of RMSE reprojection error for (a) Test Site I and (b) Test Site II, with the number of tie points for (c) Test Site I and (d) Test Site II

and 0.67 and 0.64 pixels (MicMac). Colmap and RealityScan yielded intermediate results, positioned between the highest and lowest RMSE values. No clear correlation was observed between the RMSE values and the image acquisition configuration (Parallel vs. Skew) or the processing software used. Similarly, the differences in RMSE between Test Site I and Test Site II were minimal, indicating that image configuration, coverage, and overall image quality did not significantly affect reconstruction accuracy. In terms of the number of Tie Points (Figures 4c and 4d), MicMac produced the densest point clouds – approximately 65,000 points for Test Site I and 140,000 for Test Site II – whereas Pix4D generated the smallest datasets (about 20,000 points for Test Site I and 50,000 for Test Site II). This suggests that Pix4D employs a better-implemented algorithm for tie point filtration. Agisoft Metashape exhibited more variable behaviour: for Test Site I, it produced one of the densest clouds, comparable to MicMac; however, for Test Site II, the number of Tie Points dropped significantly to around 30,000, which is fewer than those generated by Pix4D. This discrepancy may be attributed to differences in the distribution of textures and lighting conditions on the façade, both of which can affect feature-matching performance in the Agisoft Metashape tie-point extraction algorithm.

Another crucial aspect related to generating dense point clouds, in addition to image orientation accuracy and the number of detected tie points, is their spatial distribution. This factor influences the correctness of the self-calibration process and the point cloud densification procedure (Figure 5). Uniform, well-distributed tie points across the image set are essential for robust self-calibration and minimising systematic errors. Poor spatial distribution can lead to localised distortions and reduced geometric stability, negatively affecting the densification stage of point cloud generation.

The distribution of Tie Points (Figure 5) reflects the feature-matching strategies implemented in each software package. Since all Test Sites cover the same spatial extent, differences in the number of Tie Points primarily indicate variations in feature detection, matching thresholds, and the robustness of the respective algorithms. MicMac produced the highest number of Tie Points, consistent with its tendency to extract a broad set of features when sufficient texture is available. Pix4D generated noticeably fewer Tie Points, reflecting its more conservative feature-matching strategy, yet these points were well distributed across the façades. Agisoft Metashape, Colmap, and RealityScan yielded intermediate results, each showing characteristic patterns associated with their internal detection and matching pipelines. The distribution primarily provides insight into the matching behaviour of each software. As a result, the outputs from RelityScan, including both image orientation parameters and a sparse point cloud, were used in the next steps of data processing with the OpenMVS 2.0 library.

4.2 Visual point quality assessment

The subsequent investigation is based on the visual assessment of the generated dense point cloud. In this investigation, the completeness, point density, and quality of geometrical shape determination were evaluated (Appendix 1). As mentioned previously, for each Test Site, two scenarios were analysed: Parallel, where images were taken parallel to the façade, and Skew, where images were captured at an angle. The columns show the results of six methods: Agisoft Metashape, Pix4D, RealityScan, OpenMVS-PB, OpenMVS-SGM, and MicMac. Large white areas in the point cloud indicate missing data, which is a direct indicator of lower reconstruction quality.

For Test Site I in the Parallel configuration, commercial applications such as Agisoft Metashape, Pix4D, and RealityScan, as well as the open-source OpenMVS-PB, generate point clouds with high

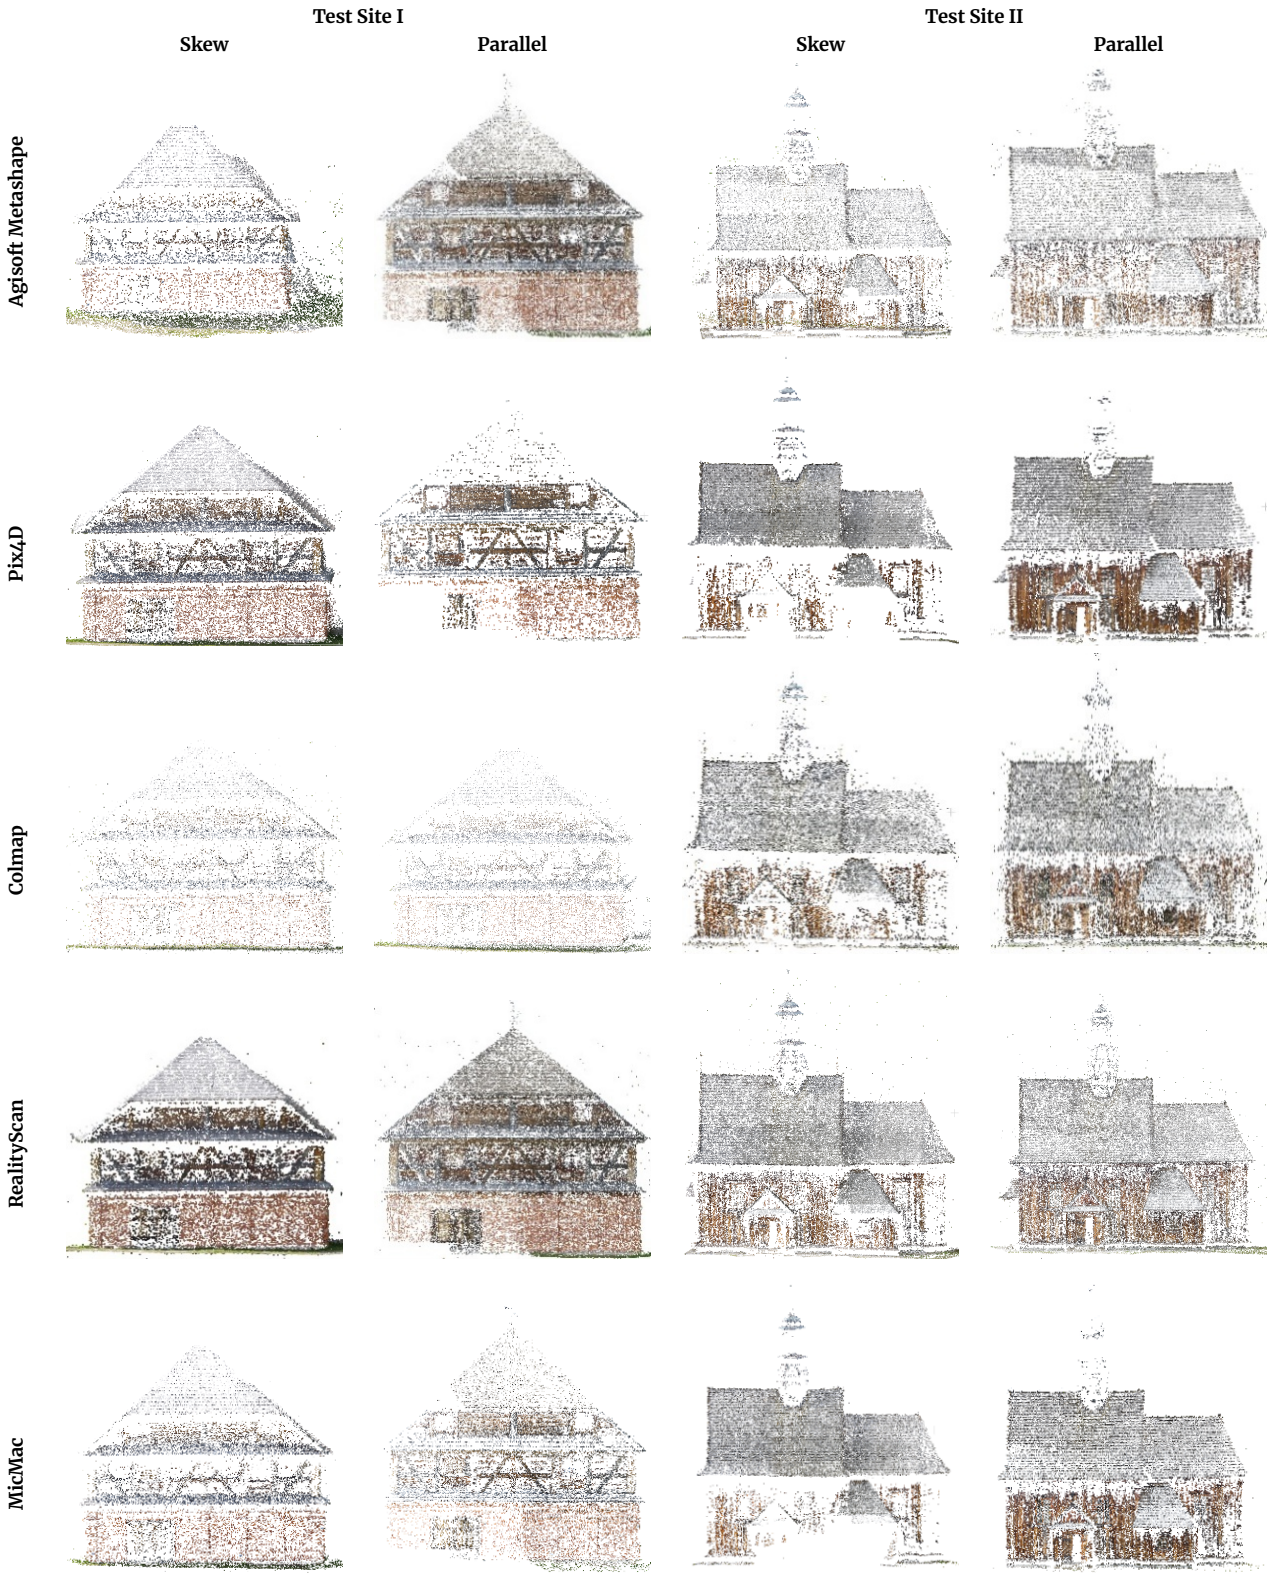

Figure 5. Tie Points distribution for Test Sites I and II

completeness, albeit with only minor gaps in the upper roof sections. In contrast, open-source methods, such as OpenMVS-SGM and MicMac, exhibit significant gaps on walls and roofs, which indicates limited reconstruction capability under this configuration. In the Skew configuration, quality degradation is evident across all methods, but most pronounced in MicMac and OpenMVS-SGM, where white areas dominate the central parts of the object. These algorithms are sensitive to changes in image acquisition angles.

Similar patterns are observed for Test Site II. In the Parallel configuration, commercial methods ensure good completeness, although gaps appear in the upper parts of the tower. Open-source methods, however, struggle with reconstructing tower details and wall surfaces. In the Skew configuration, the most significant gaps occur in OpenMVS-SGM and MicMac, where reconstruction is highly fragmented. Even commercial methods exhibit a decline in quality, yet they still preserve the basic structure of the object.

Visual analysis leads to several conclusions. Commercial solutions consistently deliver higher completeness and robustness, whereas open-source algorithms exhibit significant limitations, particularly under non-optimal image acquisition conditions. The presence of large white areas in open-source results suggests that these methods are susceptible to issues with image geometry and coverage.

Based on these observations, several recommendations can be made. First, commercial software such as Agisoft Metashape, Pix4D, and RealityScan should be preferred for projects requiring high completeness and reliability. Second, when using open-source methods, image acquisition should prioritise parallel configurations and include more images to improve coverage. Third, additional images should be captured for complex architectural details such as towers and roof edges, particularly when using open-source tools. Fourth, hybrid approaches combining image-based and range-based techniques (e.g., TLS) should be considered to enhance completeness in areas with problematic data. Finally, visual inspection of point clouds should be performed before further processing to prevent error propagation in subsequent analysis stages.

4.3 Density analysis

The point cloud density is a crucial metric in 3D reconstruction because it directly affects the likelihood of successful reconstruction (level of detail), completeness, and the accuracy of the final dense point cloud. Higher density means more measurement points per unit area and the potential for better recognition and representation of architectural details, edges, and complex structures, which are significant for the documentation of cultural heritage wooden objects. However, analysing density alone is insufficient, as it is also essential to consider the distribution of the point density. Uniform point distribution ensures consistent geometric fidelity, supports reliable surface fitting, and improves the performance of subsequent processes such as mesh generation or semantic segmentation. Poor distribution can lead to localised errors, especially in self-calibration and bundle adjustment, and may compromise the structural integrity of the digital model. For this reason, an analysis of point cloud density and its distribution for each method was conducted separately (Appendix 2):

- **Agisoft Metashape:** The point cloud density obtained from this algorithm is typically high and uniform across the entire object in both Test Sites. For Test Site I (both Parallel and Skew configurations), Agisoft Metashape effectively densified the points, even on less textured, smooth wall surfaces. For Test Site II (the wooden church), the algorithm maintained a high density, which ensures subsequent shape fidelity (low roughness) on the vertical boards. Nevertheless, the density on thin elements (e.g., crosses, railings) may be lower than in OpenMVS-PB.
- **Pix4D:** For the test site in parallel configuration, the density is not uniform and depends on the type of architectural details. However, in the Skew configuration, a drop in density was observed in areas with a large angle of light incidence, leading to "holes" and reduced façade completeness. For Test Site II, the density is uneven, with high density on the roof and low density on the vertical boards, which negatively impacts the reconstruction of details.
- **RealityScan:** Point cloud density is characterised by the density consistency in both Test Sites and configurations. For Test Site II (Skew configuration), RealityScan achieved remarkably low yet uniform density on both vertical elements and the roof, enabling the best preservation of edge recognisability.
- **OpenMVS-PB:** This algorithm is often optimised for density. For Test Site I, it yielded the highest point cloud density among all methods. Still, this result directly depends on the type of architectural details and the number of utilised images. In the case of Test Site I in the Parallel configuration, there was a lower number of overlapping images on the upper part, which resulted in a lower density. However, in the skew configuration (higher overlap), the density was more uniform and higher. Particularly high density was visible on the horizontal and slanted wooden boards for both configurations in Tets Site I. For Test Site II, the density was high and showed a uniform distribution, but remained strongly dependent on the image configuration (similar to Test Site I).
- **OpenMVS-SGM:** The point cloud density in OpenMVS-SGM is often higher than in OpenMVS-PB; however, the distribution is also uniform for both Test Sites and similar to that of OpenMVS-PB.
- **MicMac:** This algorithm generated the lowest point cloud density despite a higher number of tie points among all tested methods, particularly in areas with complex geometry (such as roofs) and low texture on Test Site I. In the Parallel configuration of Test Site II, the highest and most uniform density distribution was observed on the roof.

In summary, the highest density was achieved with OpenMVS-SGM, but it was accompanied by increased noise and roughness. The most optimal balance between high density was observed for RealityScan and OpenMVS-PB. In contrast, MicMac and Pix4D in the Skew configuration experienced the greatest difficulty in achieving sufficient model density and completeness.

4.4 Roughness analysis

The subsequent analysis aimed to evaluate surface roughness, defined as the distance of a point from a locally fitted plane constructed using its nearest neighbours within a 5 cm radius (Appendix 2). To illustrate the spatial distribution of roughness, a colour scale was used, with blue representing smooth, coherent, and well-reconstructed regions characterised by low roughness. These areas typically correspond to extensive flat surfaces or regularly curved geometries. Conversely, shades of green and yellow indicate areas where the surface exhibits greater variability, with numerous small-scale details or rough textures, such as those found on roof tiles or bricks. It is important to note that roughness is inversely proportional to surface flatness.

Analysing the results presented in Appendix 3 for each algorithm/method individually, it can be stated that for:

- **Agisoft Metashape:** The roughness values obtained place this algorithm just behind those of point clouds generated by OpenMVS-PB and RealityScan. For Test Site I (both Parallel and Skew configurations), shape reconstruction is satisfactory only for the wooden structural elements located in the upper part of the object. The 3D reconstruction of bricks was less accurate than that of RealityScan and OpenMVS-PB, as the roughness distribution appeared random. For Test Site II, vertical

wooden boards were reconstructed correctly, although the roof representation was inferior to that achieved by RealityScan and OpenMVS-PB.

- **Pix4D:** For Test Site I in the Parallel configuration, results are comparable to Agisoft Metashape for both wooden and brick details, but performance deteriorates under the Skew configuration. For Test Site II, the roughness distribution is less accurate than Agisoft Metashape, resulting in lower fidelity for both vertical boards and roof surfaces.
- **RealityScan:** For Test Site I in the Parallel configuration, the results are similar to those of the two preceding methods, while for Skew, they resemble those of Pix4D. In Test Site II, RealityScan achieved the highest consistency in roughness for both roof and vertical boards, while preserving edge recognizability. Consequently, this commercial software can be considered the most effective for this type of object.
- **OpenMVS-PB:** For Test Site I, roughness is accurately represented on horizontal and slanted wooden boards. Lower performance was observed for bricks in the lower part of the object under the Skew configuration, as gaps between bricks were flattened. For Test Site II, both vertical boards and roof structures were reconstructed correctly, with increased roughness visible along edges.
- **OpenMVS-SGM:** For Test Site I in the Parallel configuration (incomplete point cloud), results indicate that, similar to OpenMVS-PB, this algorithm accurately reconstructs wooden structural details. Under the Skew configuration, roughness is higher compared to OpenMVS-PB for both the roof and brick wall segments. This suggests that OpenMVS-SGM provides the most accurate brick reconstruction among the tested algorithms. For Test Site II, the results are inferior to those of OpenMVS-PB and RealityScan; however, roof reconstruction surpasses Agisoft Metashape and Pix4D in both Parallel and Skew configurations. Regarding vertical boards, the point cloud exhibits lower roughness (better shape fidelity) than Pix4D but higher roughness (lower quality) than Agisoft Metashape.
- **MicMac:** This algorithm produced the least accurate results among all tested methods for both brick-wood composite structures and the wooden church.

In summary, the examination of the results presented in Appendix 2 reveals that the most favourable roughness values (i.e., the lowest surface roughness) were obtained for point clouds generated using the OpenMVS-PB algorithm across both test sites, in both Skew and Parallel configurations. It is noteworthy that these point clouds maintain a balanced representation of individual details. In contrast, point clouds produced with MicMac (in both tests) and Pix4D (in Test Site I) exhibit the highest roughness values. These reconstructions accurately preserve material roughness and texture, albeit at the expense of increased local noise. Point clouds generated with Agisoft Metashape and RealityScan enable the creation of smooth wall surfaces while retaining roughness and fine details on smaller structural elements.

4.5 Planarity assessment

Planarity is the metric that determines the degree to which a set of points within a local neighbourhood conforms to a plane. This measure is computed using the eigenvalues of the covariance matrix derived from the 3D coordinates of the points in the neighbourhood. The metric ranges from 0 to 1, where 0 indicates a complete lack of planarity and 1 represents perfect planarity. The results of the planarity assessment are presented in Appendix 4. To visualise the spatial distribution of this factor, a colour palette was adopted in which red and orange correspond to areas where points form flat, coherent surfaces, typically associated with large, smooth regions. Conversely, green and yellow indicate areas where points deviate from the plane or exhibit high shape variance, which generally corresponds to edges, corners, and architectural details.

Evaluating the spatial distribution of planarity for each point presented in Appendix 4, it can be stated that for:

- **Agisoft Metashape:** The point clouds generated by the MVS algorithm implemented in this application exhibit extremely high planarity in both Test Sites. For Test Site I (both configurations), the model is almost entirely red, suggesting that Agisoft Metashape employs aggressive surface smoothing, achieving remarkably high planarity at the potential expense of preserving fine geometric details (e.g., gaps between bricks). This may reduce the accuracy of architectural documentation derived from point clouds and may also lead to misinterpretations by specialists in cultural heritage. A similar result, dominated by red, is visible for Test Site II.
- **Pix4D:** For Test Site I (Parallel configuration), planarity is also extremely high (predominantly red), comparable to Agisoft Metashape. However, in the Skew configuration, larger green/yellow areas are visible on the roof and lower sections, indicating a decrease in the ability to smooth and maintain surface flatness under more challenging data-acquisition conditions. For Test Site II (Skew configuration), lower planarity (more yellow/green) is observed compared to Agisoft Metashape and RealityScan.
- **RealityScan:** Demonstrated the most balanced planarity relative to Agisoft Metashape, especially for Test Site I. Although the model is mainly red, green and yellow zones are visible on the bricks and boards (particularly in the Parallel configuration), suggesting that this algorithm retains some degree of texture roughness/variance while maintaining the overall flatness of the main surfaces. For Test Site II, it achieves planarity comparable to Agisoft Metashape.
- **OpenMVS-PB:** This algorithm provides remarkably high planarity in both Test Sites, with a prevalence of the colour red. For Test Site I, however, there is a tendency to maintain green/yellow bands at the board junctions and in the brick section, which may indicate attempts to better represent gaps and details. For Test Site II (both configurations), the model is nearly uniformly red, confirming high smoothing.
- **OpenMVS-SGM:** For Test Site I (both configurations), planarity is slightly lower (more green/yellow details) than for Agisoft Metashape and OpenMVS-PB, especially in the brick section. This is consistent with earlier analysis. For Test Site II, high planarity (red) is visible, but with distinct green/yellow edges, indicating the preservation of some roughness.
- **MicMac:** This algorithm exhibits the lowest planarity among all tested methods. Significant areas of green and yellow are visible in both Test Sites (e.g., the lower part of Test Site I and a large part of Test Site II), indicating that MicMac is least effective at smoothing the reconstructed surfaces, resulting in high variability and a low degree of fit to a plane.

In summary, the highest planarity was observed for Agisoft Metashape, Pix4D, RealityScan, and OpenMVS-PB/SGM in the Skew configuration. These programs most effectively represent large planar surfaces as smooth and coherent. The poorest performance was recorded for MicMac and OpenMVS-SGM in the Parallel configuration. A high presence of green noise indicates that the point cloud in these areas is less smooth or more affected by noise, complicating the identification of a consistent surface.

4.6 Normal vector variance

The following analysis aims to evaluate the normal vector variances, which are a key indicator of the local quality and consistency of surface orientation in the reconstructed models. Low variance (represented by the colour blue in the visualisation) indicates consistent, well-defined surfaces. In contrast, high variance (represented by

the colours green/red) suggests reconstruction errors, noise, or extremely rough textures.

Analysing the results presented in the visualisation for each algorithm/method individually (Appendix 5), it can be stated that for:

- **Agisoft Metashape:** The normal vector variance is relatively low (blue/green), suggesting a consistent but overly smoothed surface. For Test Site I (both configurations), the surface appears to be excessively smoothed (predominantly blue), which may hide fine details but ensures overall vector consistency. For Test Site II in the Skew configuration, areas of higher variance (green) are visible on the roof, indicating issues with accurately determining the surface orientation in those areas.
- **Pix4D:** This algorithm demonstrates one of the highest variances (predominantly red/green) in most cases. For Test Site I (Parallel configuration), the walls are predominantly green, and the roof is red, indicating a high level of inconsistency in the normal vectors and significant geometric noise. For Test Site II (both configurations), the model exhibits the worst normal vector consistency (mostly red), showing serious errors in defining the surface orientation.
- **RealityScan:** Achieved the best consistency of normal vectors, especially for Test Site II. In both configurations (Parallel and Skew), the predominantly blue colour on the church model indicates that the wall and roof surfaces are exceptionally well-defined with low variance. In Test Site I, low variance is also maintained. However, green and red details are visible in the brickwork section, suggesting attempts to preserve the brick texture and gaps (where high variance is expected).
- **OpenMVS-PB:** This algorithm achieved the lowest variance among all methods for Test Site II (both configurations), where the model is almost entirely blue. This means OpenMVS-PB was most successful at obtaining an exceptionally consistent and smooth surface orientation for the wooden church. For Test Site I, blue is also dominant, but green/red zones are visible on the lower brick part, which may result from an attempt to preserve the gaps between bricks.
- **OpenMVS-SGM:** For Test Site I (large blue areas), consistency is high; however, green/yellow details are visible on the bricks, suggesting a more detailed but noisy representation of their orientation. For Test Site II in the Parallel configuration, the point cloud is less complete. Still, in the reconstructed areas (including the roof), it maintains good consistency (blue), surpassing Agisoft Metashape and Pix4D.
- **MicMac:** MicMac shows the greatest problem with orientation consistency, evident as large red areas across both Test Sites. The red patches (e.g., on the roof of Test Site I) show the highest variance, a direct indicator of geometric chaos and the unsuitability of these areas for subsequent surface analysis.

The examination of the visualisation reveals that the most consistent and lowest normal vector variance (predominantly blue colour) – indicating the best surface quality and smoothness – was achieved by the OpenMVS-PB and RealityScan algorithms, especially for Test Site II (the wooden church). In contrast, the Pix4D and MicMac algorithms consistently generated the highest variance (large red areas), confirming surface consistency issues and a high level of geometric noise.

4.7 Comparison between MVS and TLS point clouds – cloud-to-cloud distance

CloudCompare software was employed to evaluate the quality of the generated point clouds and to identify the most effective solution for 3D reconstruction. For benchmarking purposes, a point cloud utilised by Leica BLK TLS was selected as the reference dataset. In the figures, a histogram of deviations was presented, with the X-axis ranging from 0 to 0.05 m and the Y-axis showing the frequency of occurrence. The interpretation focuses on three aspects: distribution (shape and position of the peak), spread, and density (concentration of counts in the lowest bins).

This study clearly shows that the quality of 3D shape reconstruction depends on the image acquisition method and the dense reconstruction algorithm used. The main relevant findings are presented below:

- **Agisoft Metashape**: Under the Parallel configuration at both test sites, Agisoft Metashape exhibits an extremely narrow error distribution with a pronounced peak within the 0 – 0.005 m range and minimal spread, resulting in high density in the lowest bins. Under the Skew configuration for Test Site I, the distribution remains significantly narrower. In contrast, at Test Site II, the peak shifts toward 0.01 – 0.015 m and the spread increases, indicating greater sensitivity to flight geometry.
- **Pix4D:** The generated dense point cloud exhibits a deviation distribution similar to that of Agisoft Metashape under the Parallel configuration, with a pronounced peak near 0.005 m and minimal spread. Under the Skew configuration at Test Site II, the peak shifts to approximately 0.015 – 0.02 m, increasing the proportion of points with deviations exceeding 2 cm. At Test Site I, particularly under the Parallel configuration, the results demonstrate high consistency, confirming the robustness of the method under favourable geometric conditions.
- **RealityScan:** This software exhibits the widest and flattest distributions among all methods. There is no pronounced peak near zero; instead, the profile is flattened, and density within the 0 – 0.01 m bins remains low. This indicates a large spread and poor uniformity of the point cloud. Even under the Parallel configuration, improvement is minimal, indicating that the method would require aggressive outlier filtering.
- **MicMac:** At Test Site I, the distribution of MicMac is similar to that observed in commercial algorithms, with a clear peak at minor deviations and a short tail, indicating moderate spread and good density near zero. At Test Site II, especially under the Skew configuration, the peak shifts toward 0.01 – 0.02 m and the tail lengthens, indicating increased sensitivity to scene complexity and flight geometry. The Parallel configuration yields improved results but still falls short of Agisoft Metashape and Pix4D.
- **OpenMVS-PB:** At Test Site I, the distribution peaks approximately at 0.005 m and has a relatively short tail, suggesting lower spread. The Parallel configuration narrows the profile, but density still concentrates slightly offset from zero. At Test Site II, the results are similar to those at Test Site I: a steep peak near zero and a slightly longer tail, indicating limited variability in the method across data and acquisition settings.
- **OpenMVS-SGM:** At Test Site II, OpenMVS-SGM exhibits a very steep peak below 0.003 m under the Skew configuration and a short tail, along with high density in the first bins. Even under the Skew configuration, the profile remains compact, making this method among the most metrologically accurate in the entire dataset.

In summary, the tested algorithms provide the highest metrological quality, with a sharp peak at 0 – 0.005 m and a rapidly decreasing tail, as observed in Agisoft Metashape and Pix4D under the Parallel configuration, and in OpenMVS-SGM at Test Site II. MicMac and OpenMVS-PB deliver moderate results, strongly influenced by scene complexity and acquisition geometry. Conversely, the poorest quality is observed in the RealityScan-generated point cloud for Test Site II. Across all tested methods, point clouds obtained under the Parallel configuration are significantly better, due to reduced spread and increased concentration of counts in the lowest bins. This contrasts with the Skew configuration, highlighting its critical role in achieving high reconstruction accuracy.

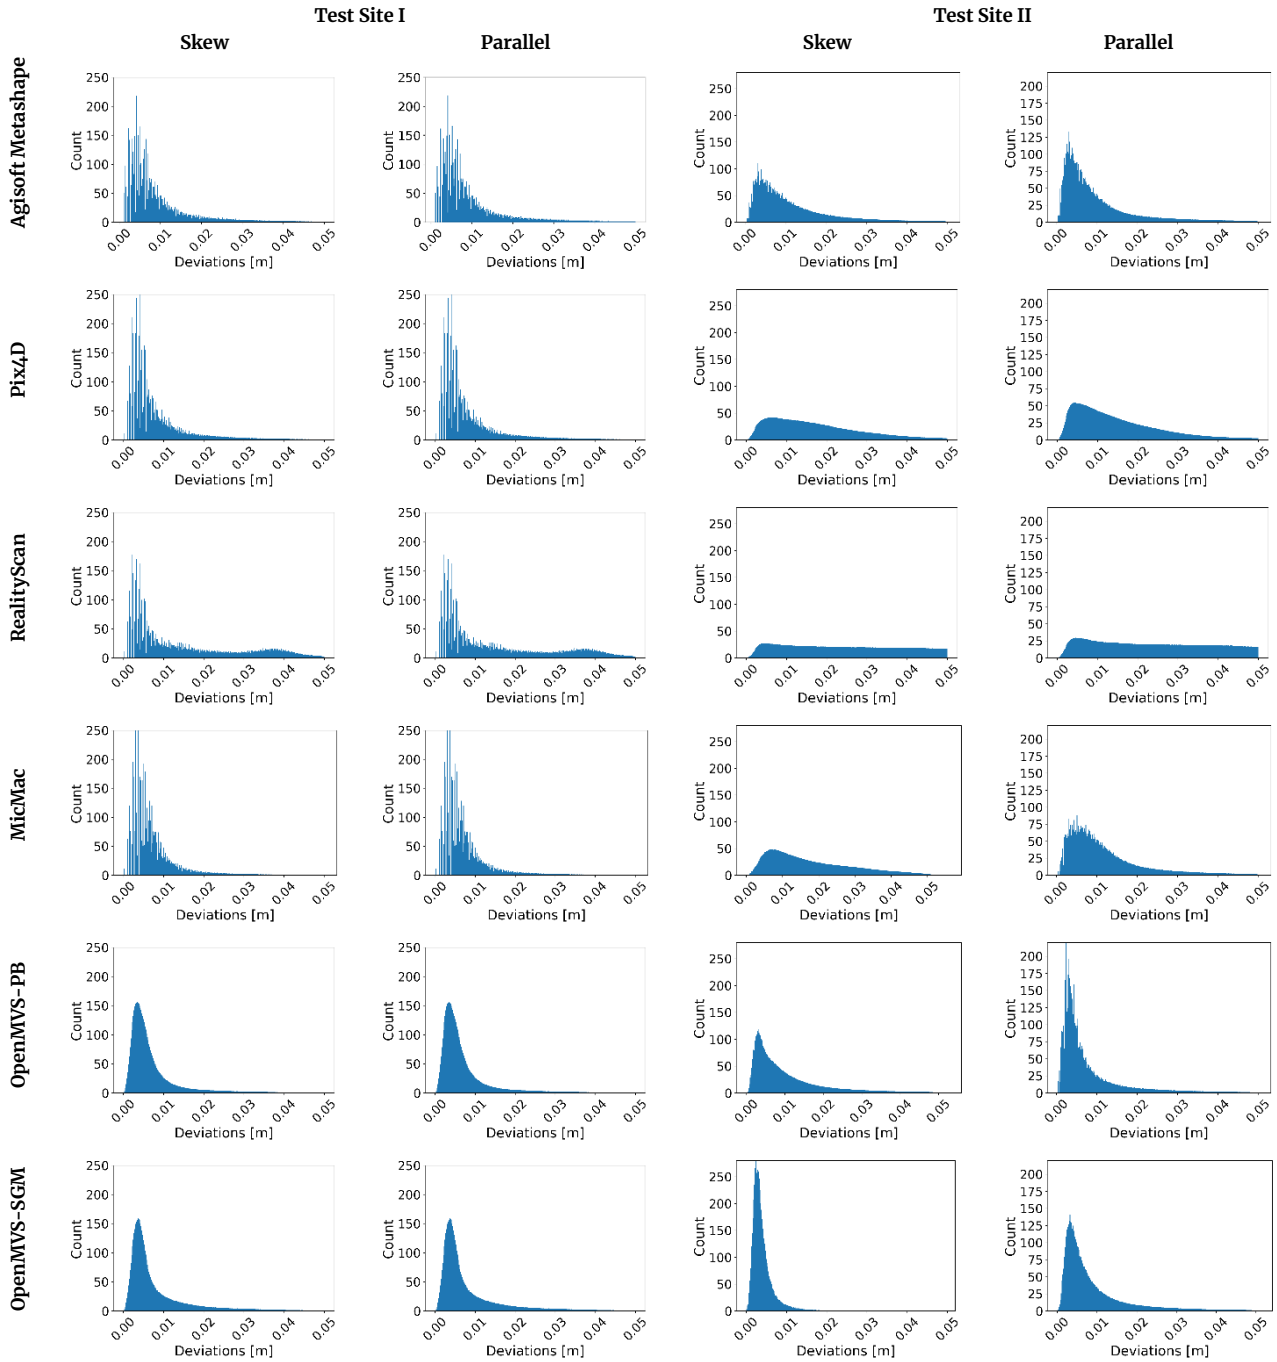

Figure 6. Probability density histogram of linear deviations between the TLS point cloud and the point cloud generated by the analysed algorithms for Test Sites I and II

5 Conclusion and summary

This study evaluated the performance of commercial and open-source MVS algorithms for generating dense point clouds in cultural heritage documentation of historic wooden structures. Open-source (RealityScan, MicMac, and OpenMVS: Patch-Based and Semi-Global Matching) and commercial applications (Agisoft Metashape and Pix4D) were assessed on UAV images acquired in two configurations (Parallel and Skew). To verify the metric quality of point clouds, TLS point clouds were used as reference data. The qualitative analysis focused on point density, roughness, planarity, normal vector variance, and cloud-to-cloud deviation.

The results clearly demonstrate that image acquisition geometry significantly influences reconstruction quality. Parallel configurations consistently produced higher accuracy, narrower deviation distributions, and greater point density than Skew configurations, resulting in increased roughness, reduced completeness, and longer deviation tails. This finding underscores the importance of planning flight paths to optimise image geometry for photogrammetric processing.

In terms of point density, OpenMVS-SGM achieved the highest values, followed by OpenMVS-PB and RealityScan. Agisoft Metashape and Pix4D provided balanced density with good completeness, whereas MicMac consistently generated the lowest density, resulting in gaps and poor representation of fine details. Roughness analysis showed that OpenMVS-PB delivered the smoothest surfaces while preserving structural details, with RealityScan also performing well for complex roof geometries. Conversely, MicMac

and Pix4D under Skew conditions exhibited the highest roughness, indicating noise and reduced fidelity.

Planarity assessment indicated that Agisoft Metashape, Pix4D, RealityScan, and OpenMVS-PB/SGM most effectively reconstructed large, flat surfaces. In contrast, MicMac and OpenMVS-SGM in Parallel configurations struggled with noise in planar regions. Normal vector variance analysis confirmed that OpenMVS-PB produced the most consistent surface orientation, whereas Pix4D and MicMac displayed high variance, suggesting geometric instability.

Cloud-to-cloud deviation analysis against TLS benchmarks highlighted Agisoft Metashape and Pix4D under Parallel configurations as the most accurate, exhibiting sharp peaks at 0–0.005 m and minimal tails. OpenMVS-SGM at Test Site II also achieved excellent results. RealityScan exhibited the widest distributions and longest tails, necessitating aggressive outlier filtering, whereas MicMac and OpenMVS-PB demonstrated moderate performance, strongly influenced by scene complexity and image geometry.

Overall, the commercial solution, Agisoft Metashape, consistently delivered robust and accurate results, particularly for parallel acquisition. Among open-source tools, OpenMVS-PB demonstrated a favourable trade-off between point density and surface roughness. RealityScan and MicMac proved less dependable under non-optimal conditions, with RealityScan requiring additional filtering and MicMac struggling with completeness and noise. For projects demanding high geometric fidelity, Parallel image acquisition should be prioritised, and hybrid workflows combining image-based and TLS data are recommended for capturing complex architectural details. Future research should focus on the integration of semantic enrichment, automated quality control, and adaptive strategies for dense point cloud fusion.

Beyond purely quantitative performance, the choice of software for dense point cloud generation should also account for user experience, workflow complexity, and the level of expertise required. Commercial solutions such as Agisoft Metashape and Pix4D provide highly automated, user-friendly environments that enable reliable results even for non-expert users, particularly under well-planned image acquisition conditions. These tools are well-suited to routine cultural heritage documentation tasks where robustness, reproducibility, and time efficiency are critical. In contrast, open-source libraries such as OpenMVS and MicMac require greater technical expertise, careful parameter tuning, and greater user involvement, but offer substantially higher flexibility and transparency of the processing pipeline.

The presented results demonstrate that, when properly configured by experienced users, open-source solutions can match or even surpass commercial software in terms of point density, surface quality, and metrological accuracy. This is particularly relevant for cultural heritage documentation, where object-specific characteristics and non-standard acquisition scenarios often require customised processing strategies. The maturity of open-source photogrammetric tools has reached a level at which they represent viable and powerful alternatives for advanced users and research-oriented applications. At the same time, commercial software remains preferable for standardised workflows and less experienced operators. These findings highlight that software selection should be guided not only by performance metrics, but also by user expertise, project requirements, and the desired level of methodological control.

References

Abbate, E., Sammartano, G., and Spanò, A. (2019). Prospective upon multi-source urban scale data for 3D documentation and monitoring of urban legacies. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLII-2/W11:11–19, doi:10.5194/isprs-archives-xlii-2-w11-11-2019.

Agisoft Metashape (2025). Available online: http://www.agisoft.com/. (accessed on 10 November 2025).

Agisoft Photogrammetry Workflow Compendium (2025). Available online: https://github.com/tamer017/Agisoft-Photogrammetry-Workflow-Compendium/blob/master/docs/dense_cloud_generation.md. (accessed on 9 November 2025).

Arif, R. and Essa, K. (2017). Evolving techniques of documentation of a world heritage site in Lahore. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLII-2/W5:33–40, doi:10.5194/isprs-archives-xlii-2-w5-33-2017.

Baroque church from Rogów (2026). Available online: https://mwk.com.pl/aktualnosci/151-obiekty-qr/6905-baroque-church-from-rogow. (accessed on 2 February 2026).

Bazazian, D., Casas, J. R., and Ruiz-Hidalgo, J. (2015). Fast and Robust Edge Extraction in Unorganized Point Clouds. In *2015 International Conference on Digital Image Computing: Techniques and Applications (DICTA)*, page 1–8. IEEE, doi:10.1109/dicta.2015.7371262.

Bochenska, A., Kot, P., Muradov, M., Markiewicz, J., Zawieska, D., Hess, M., and Antoniou, A. (2023). Critical evaluation of cultural heritage architectual standard documentation methods across different European countries. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-M-2–2023:251–258, doi:10.5194/isprs-archives-xlviii-m-2-2023-251-2023.

Doroszuk, K. and Markiewicz, J. (2022). The possibility of using close-range photogrammetry in the inventory of historic complex basements – Case study. *Sensors and Machine Learning Applications*, 1(2), doi:10.55627/smla.001.02.0014.

Ethnographic Park in Tokarnia (2026). Available online: https://www.geoportal.gov.pl/pl/dane/panstwowy-rejestr-granic-prg. (accessed on 3 February 2026).

Giżyńska, J., Komorowska, E., and Kowalczyk, M. (2022). The comparison of photogrammetric and terrestrial laser scanning methods in the documentation of small cultural heritage object – case study. *Journal of Modern Technologies for Cultural Heritage Preservation*, 1(1).

Granary from Rogów (2026). Available online: https://mwk.com.pl/en/museum/ethnographic-park-in-tokarnia/151-obiekty-qr/6892-granary-from-rogow. (accessed on 2 February 2026).

Grilli, E. and Remondino, F. (2019). Classification of 3D Digital Heritage. *Remote Sensing*, 11(7):847, doi:10.3390/rs11070847.

Hackel, T., Wegner, J. D., and Schindler, K. (2016). Fast semantic segmentation of 3D point clouds with strongly varying density. *ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences*, III–3:177–184, doi:10.5194/isprsannals-iii-3-177-2016.

Kerbl, B., Kopanas, G., Leimkuehler, T., and Drettakis, G. (2023). 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics*, 42:1–14, doi:10.1145/3592433.

Kostrzewa, A., Płatek-Żak, A., Banat, P., and Wilk, Ł. (2025). Open-Source vs. Commercial Photogrammetry: Comparing Accuracy and Efficiency of OpenDroneMap and Agisoft Metashape. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-1/W4-2025:65–72, doi:10.5194/isprs-archives-xlviii-1-w4-2025-65-2025.

Matsuki, H., Murai, R., Kelly, P. H., and Davison, A. J. (2024). Gaussian splatting slam. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 18039–18048.

MicMac (2025). Available online: https://github.com/micmacIGN/micmac. (accessed on 10 November 2025).

Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., and Ng, R. (2021). NeRF: representing scenes as

neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, doi:10.1145/3503250.

Müller, T., Evans, A., Schied, C., and Keller, A. (2022). Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics*, 41(4):1–15, doi:10.1145/3528223.3530127.

Muradov, M., Kot, P., Markiewicz, J., Łapiński, S., Tobiasz, A., Onisk, K., Shaw, A., Hashim, K., Zawieska, D., and Mohi-Ud-Din, G. (2022). Non-destructive system for in-wall moisture assessment of cultural heritage buildings. *Measurement*, 203:111930, doi:10.1016/j.measurement.2022.111930.

Murtiyoso, A. and Grussenmeyer, P. (2017). Documentation of heritage buildings using close-range UAV images: dense matching issues, comparison and case studies. *The Photogrammetric Record*, 32(159):206–229, doi:10.1111/phor.12197.

Murtiyoso, A., Pellis, E., Grussenmeyer, P., Landes, T., and Masiero, A. (2022). Towards Semantic Photogrammetry: Generating Semantically Rich Point Clouds from Architectural Close-Range Photogrammetry. *Sensors*, 22(3):966, doi:10.3390/s22030966.

OpenMVS (2025). Available online: https://github.com/cdcseacave/openMVS. (accessed on 10 November 2025).

Pix4D (2025). Available online: https://pix4d.com/. (accessed on 10 November 2025).

RealityScan (2025). Available online: https://www.realityscan.com/. (accessed on 10 November 2025).

Remondino, F. (2011). Heritage Recording and 3D Modeling with Photogrammetry and 3D Scanning. *Remote Sensing*, 3(6):1104–1138, doi:10.3390/rs3061104.

Rupnik, E., Daakir, M., and Pierrot Deseilligny, M. (2017). MicMac – a free, open-source solution for photogrammetry. *Open Geospatial Data, Software and Standards*, 2(1), doi:10.1186/s40965-017-0027-2.

Schonberger, J. L. and Frahm, J.-M. (2016). Structure-from-motion revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4104–4113.

Schönberger, J. L., Zheng, E., Frahm, J.-M., and Pollefeys, M. (2016). Pixelwise view selection for unstructured multi-view stereo. In *European conference on computer vision*, pages 501–518. Springer, doi:10.1007/978-3-319-46487-9_31.

Shen, S. (2013). Accurate Multiple View 3D Reconstruction Using Patch-Based Stereo for Large-Scale Scenes. *IEEE Transactions on Image Processing*, 22(5):1901–1914, doi:10.1109/tip.2013.2237921.

Stathopoulou, E. K., Battisti, R., Cernea, D., Remondino, F., and Georgopoulos, A. (2021). Semantically Derived Geometric Constraints for MVS Reconstruction of Textureless Areas. *Remote Sensing*, 13(6):1053, doi:10.3390/rs13061053.

Stathopoulou, E. K. and Remondino, F. (2023). A survey on conventional and learning-based methods for multi-view stereo. *The Photogrammetric Record*, 38(183):374–407, doi:10.1111/phor.12456.

Tancik, M., Weber, E., Ng, E., Li, R., Yi, B., Wang, T., Kristoffersen, A., Austin, J., Salahi, K., Ahuja, A., Mcallister, D., Kerr, J., and Kanazawa, A. (2023). Nerfstudio: A Modular Framework for Neural Radiance Field Development. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Proceedings*, SIGGRAPH '23, page 1–12. ACM, doi:10.1145/3588432.3591516.

Tobiasz, A., Markiewicz, J., Łapiński, S., Nikel, J., Kot, P., and Muradov, M. (2019). Review of Methods for Documentation, Management, and Sustainability of Cultural Heritage. Case Study: Museum of King Jan III's Palace at Wilanów. *Sustainability*, 11(24):7046, doi:10.3390/su11247046.

Wang, G., Pan, L., Peng, S., Liu, S., Xu, C., Miao, Y., Zhan, W., Tomizuka, M., Pollefeys, M., and Wang, H. (2025). NeRFs in robotics: A survey. *The International Journal of Robotics Research*, doi:10.1177/02783649251374246.

Weinmann, M., Jutzi, B., and Mallet, C. (2013). Feature relevance assessment for the semantic interpretation of 3D point cloud data. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, II-5/W2:313–318, doi:10.5194/isprsannals-ii-5-w2-313-2013.

Weinmann, M., Jutzi, B., Mallet, C., and Weinmann, M. (2017). Geometric features and their relevance for 3D point cloud classification. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, IV-1/W1:157–164, doi:10.5194/isprs-annals-iv-1-w1-157-2017.

Appendix 1. Probability density histogram of linear deviations between the TLS point cloud and the point cloud generated by the analysed algorithms for Test Sites I and II

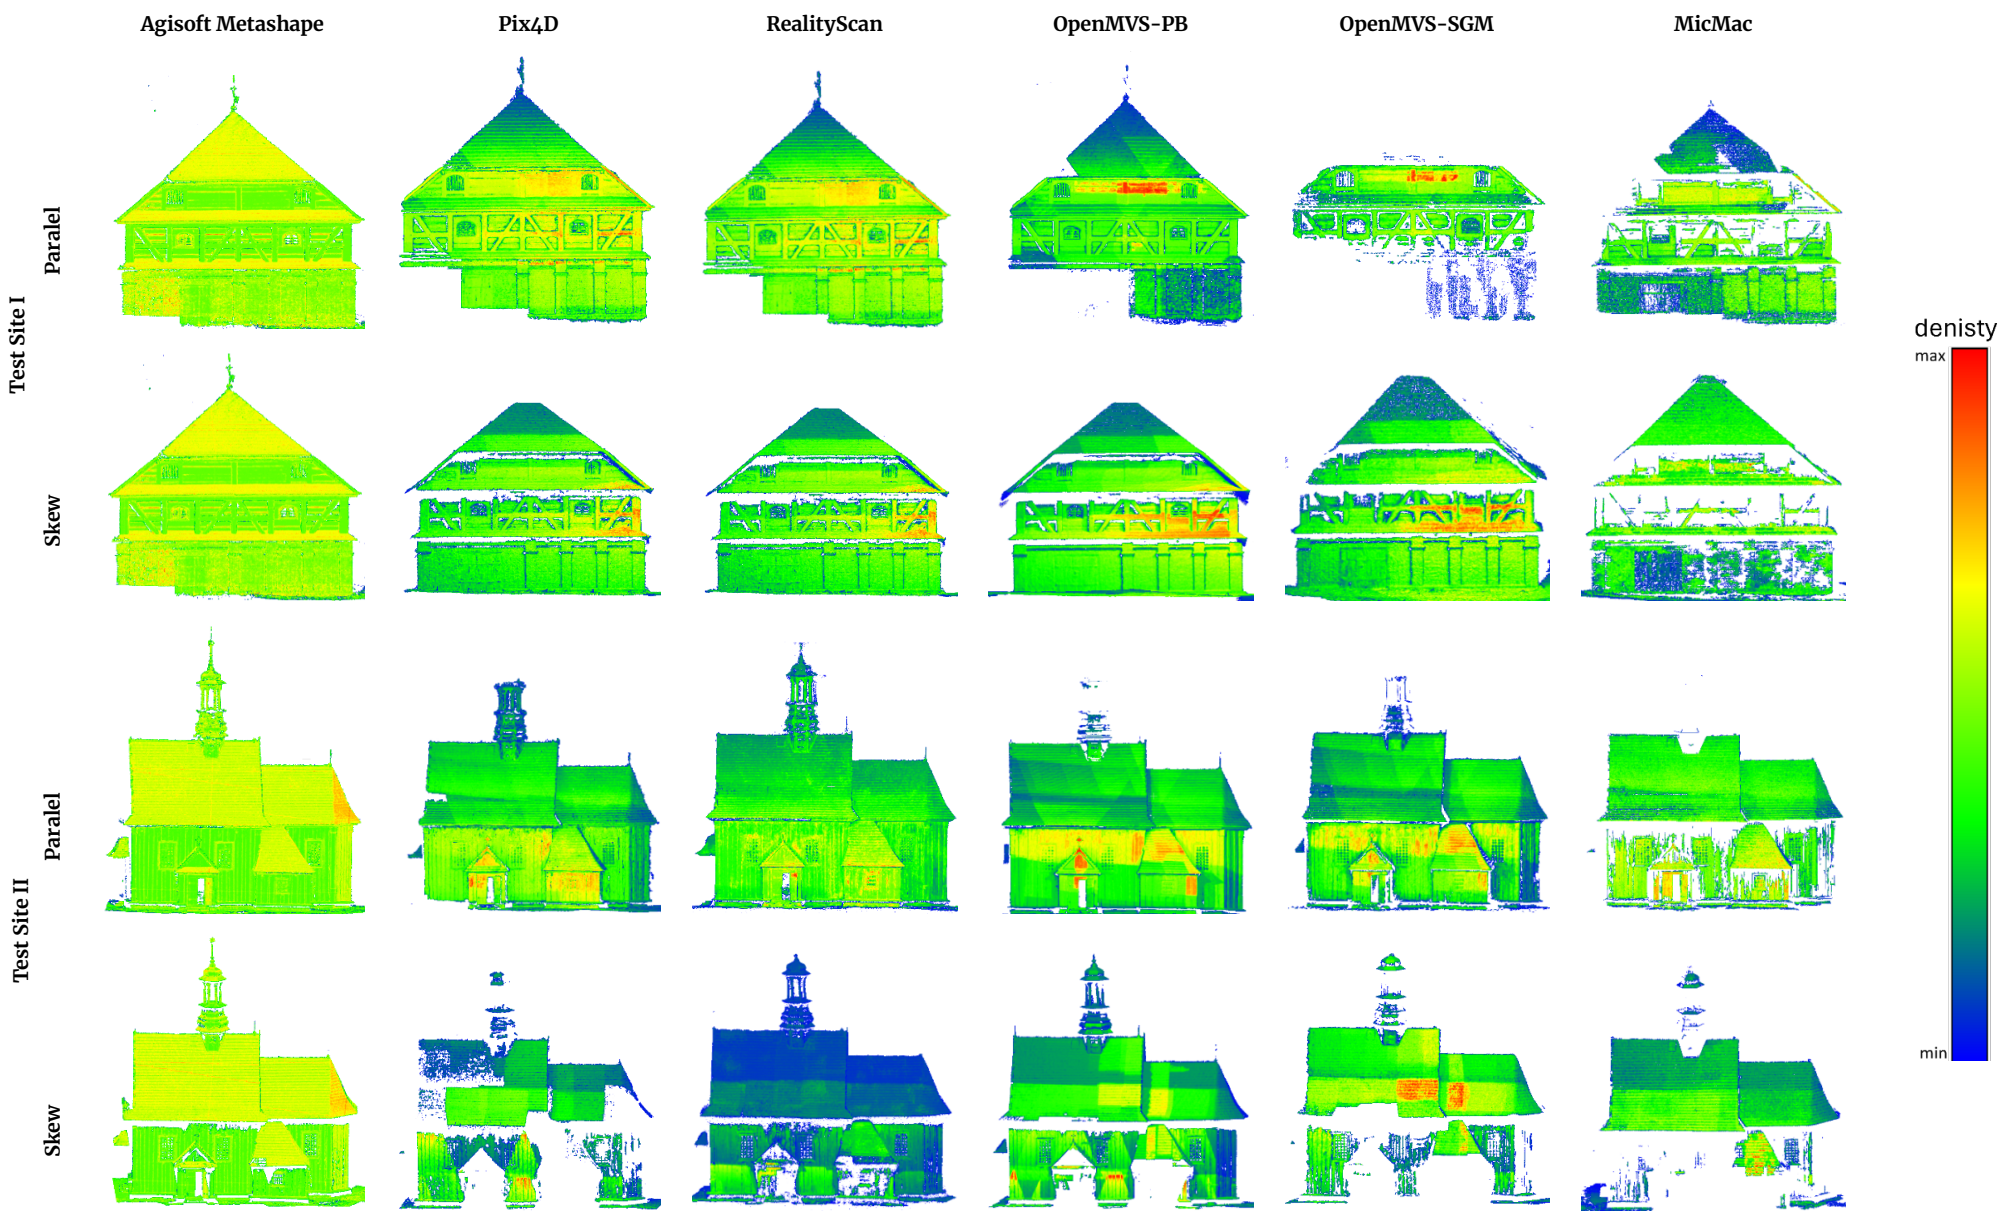

Appendix 2. The example of density for point clouds from parallel and skew images for utilised Test Sites

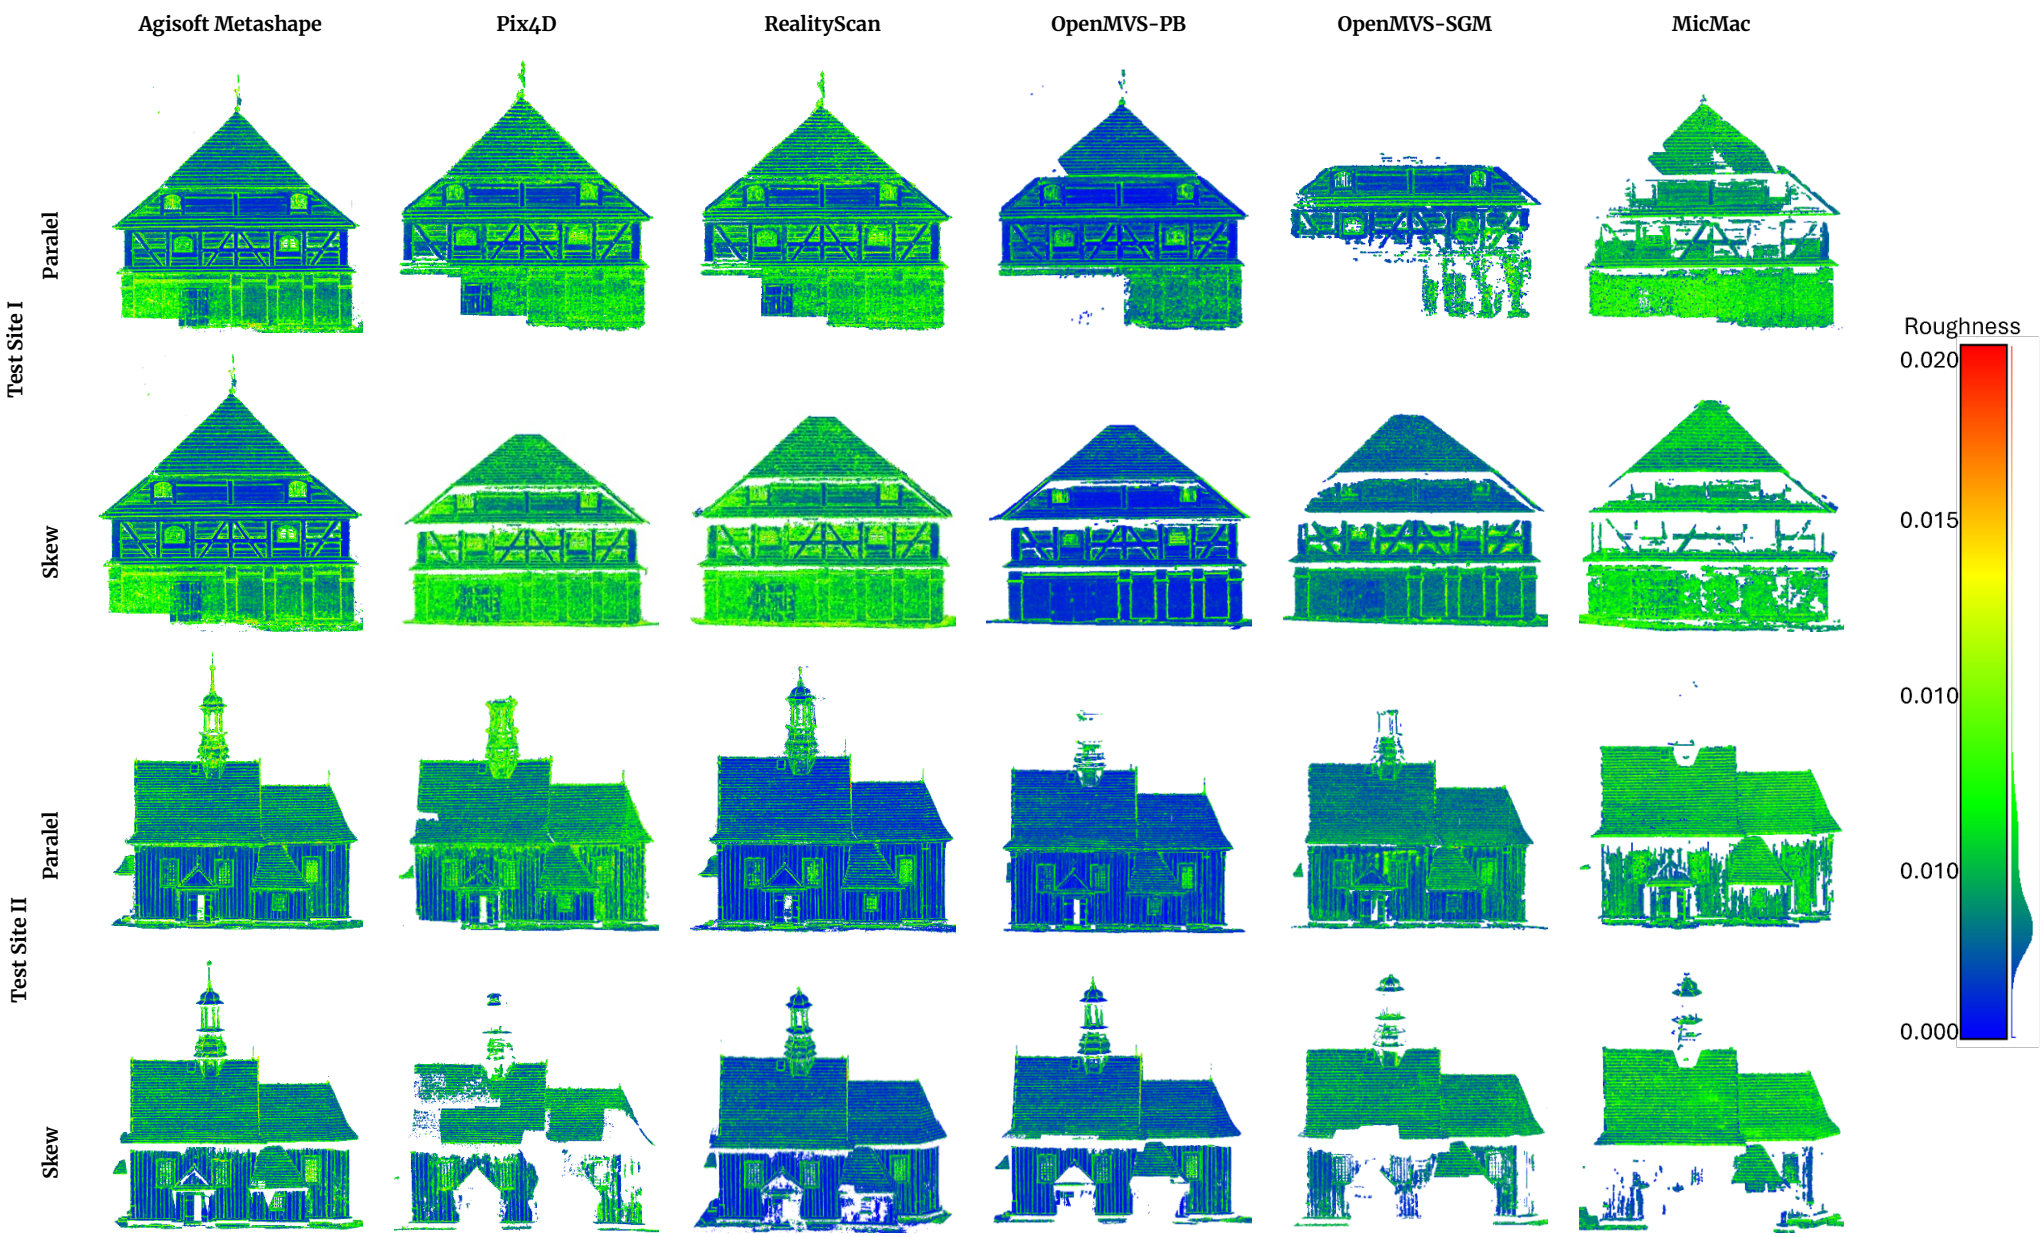

Appendix 3. The example of roughness for point clouds from parallel and skew images for utilised Test Sites

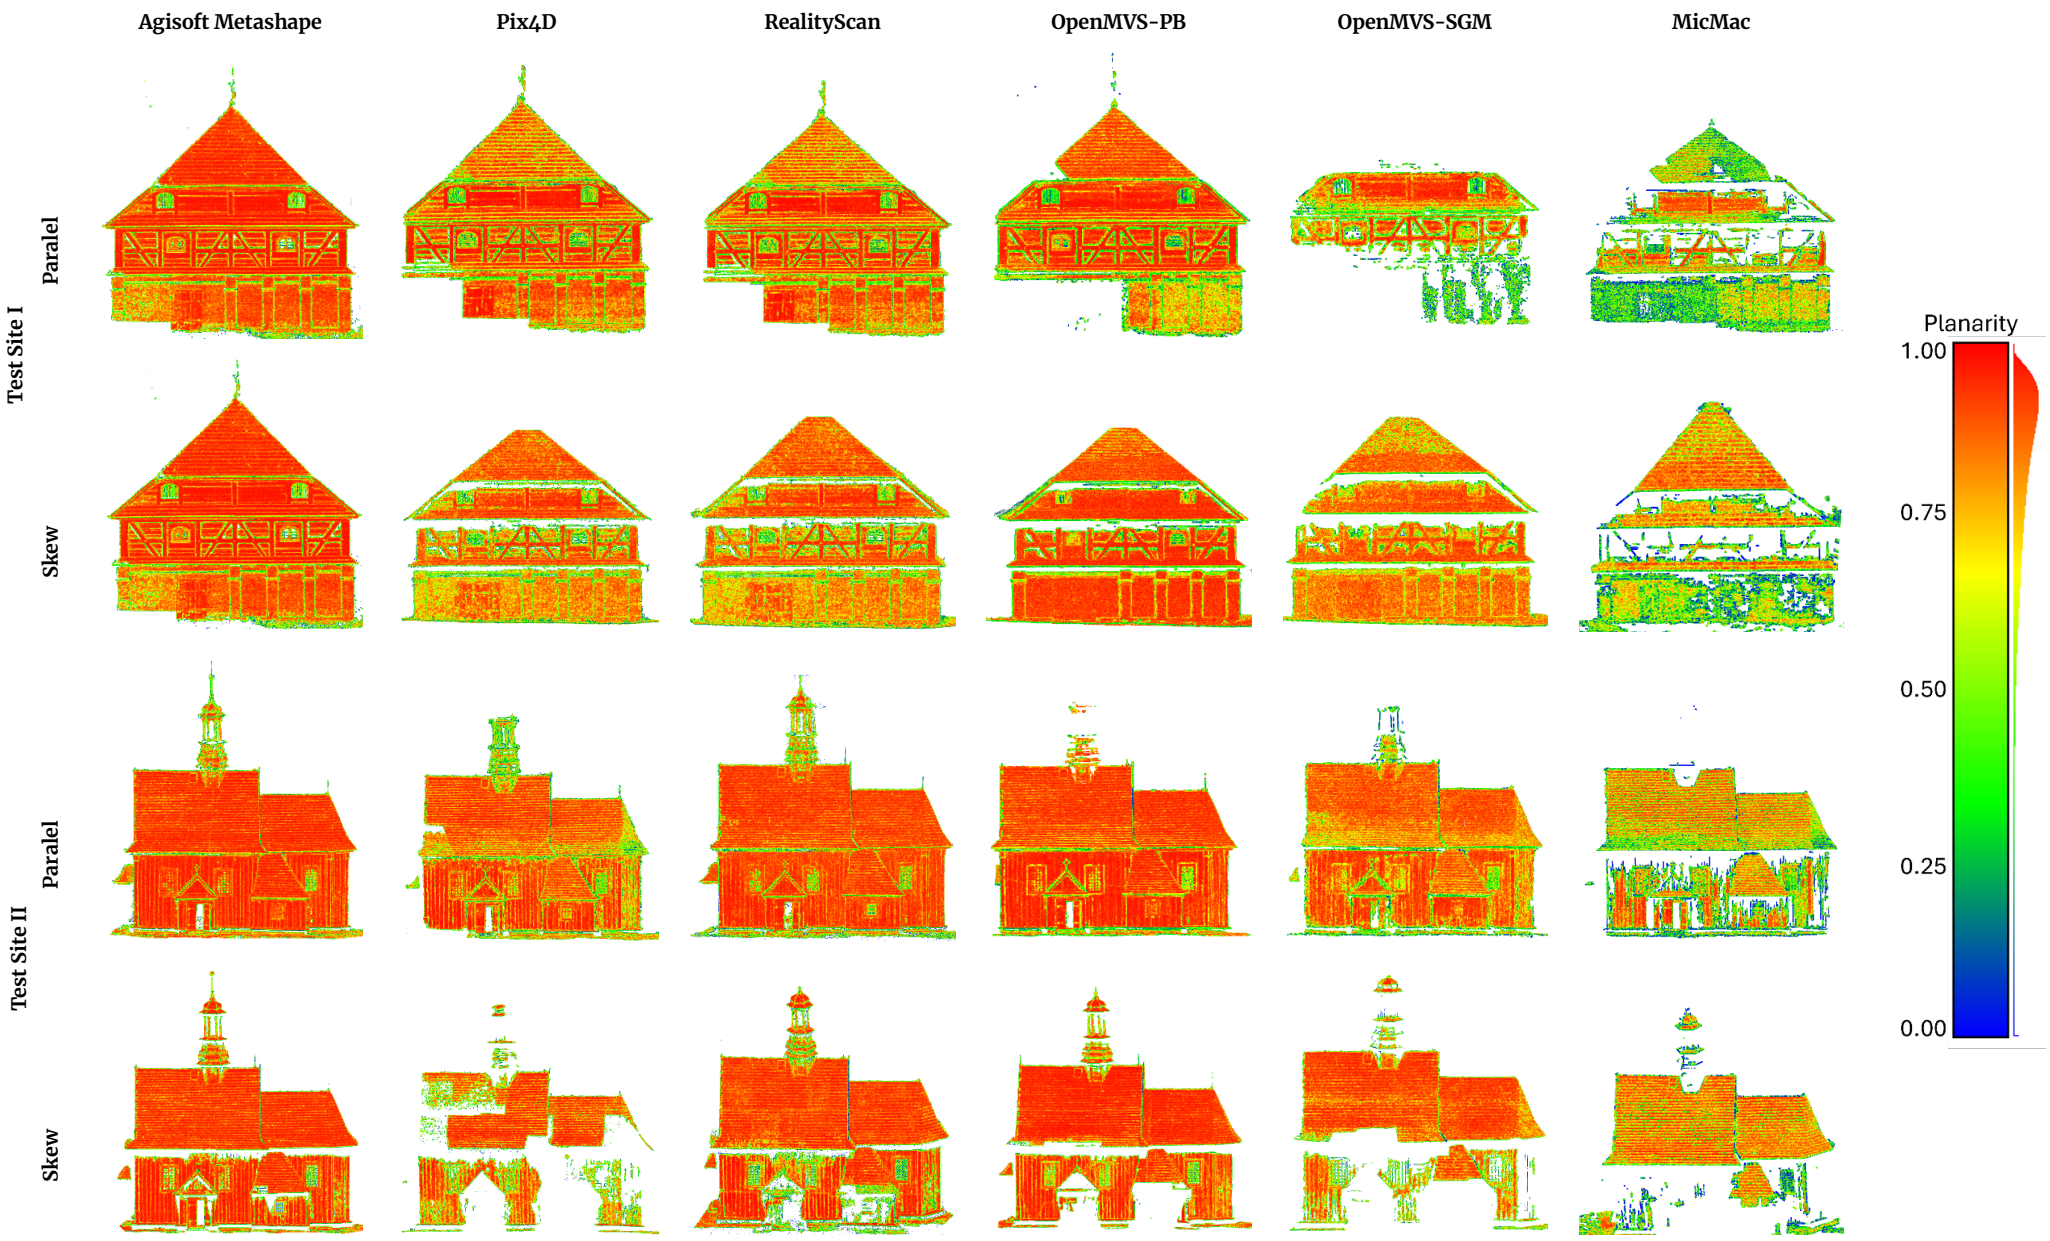

Appendix 4. The example of planarity vector variances for point clouds from parallel and skew images for utilised Test Sites

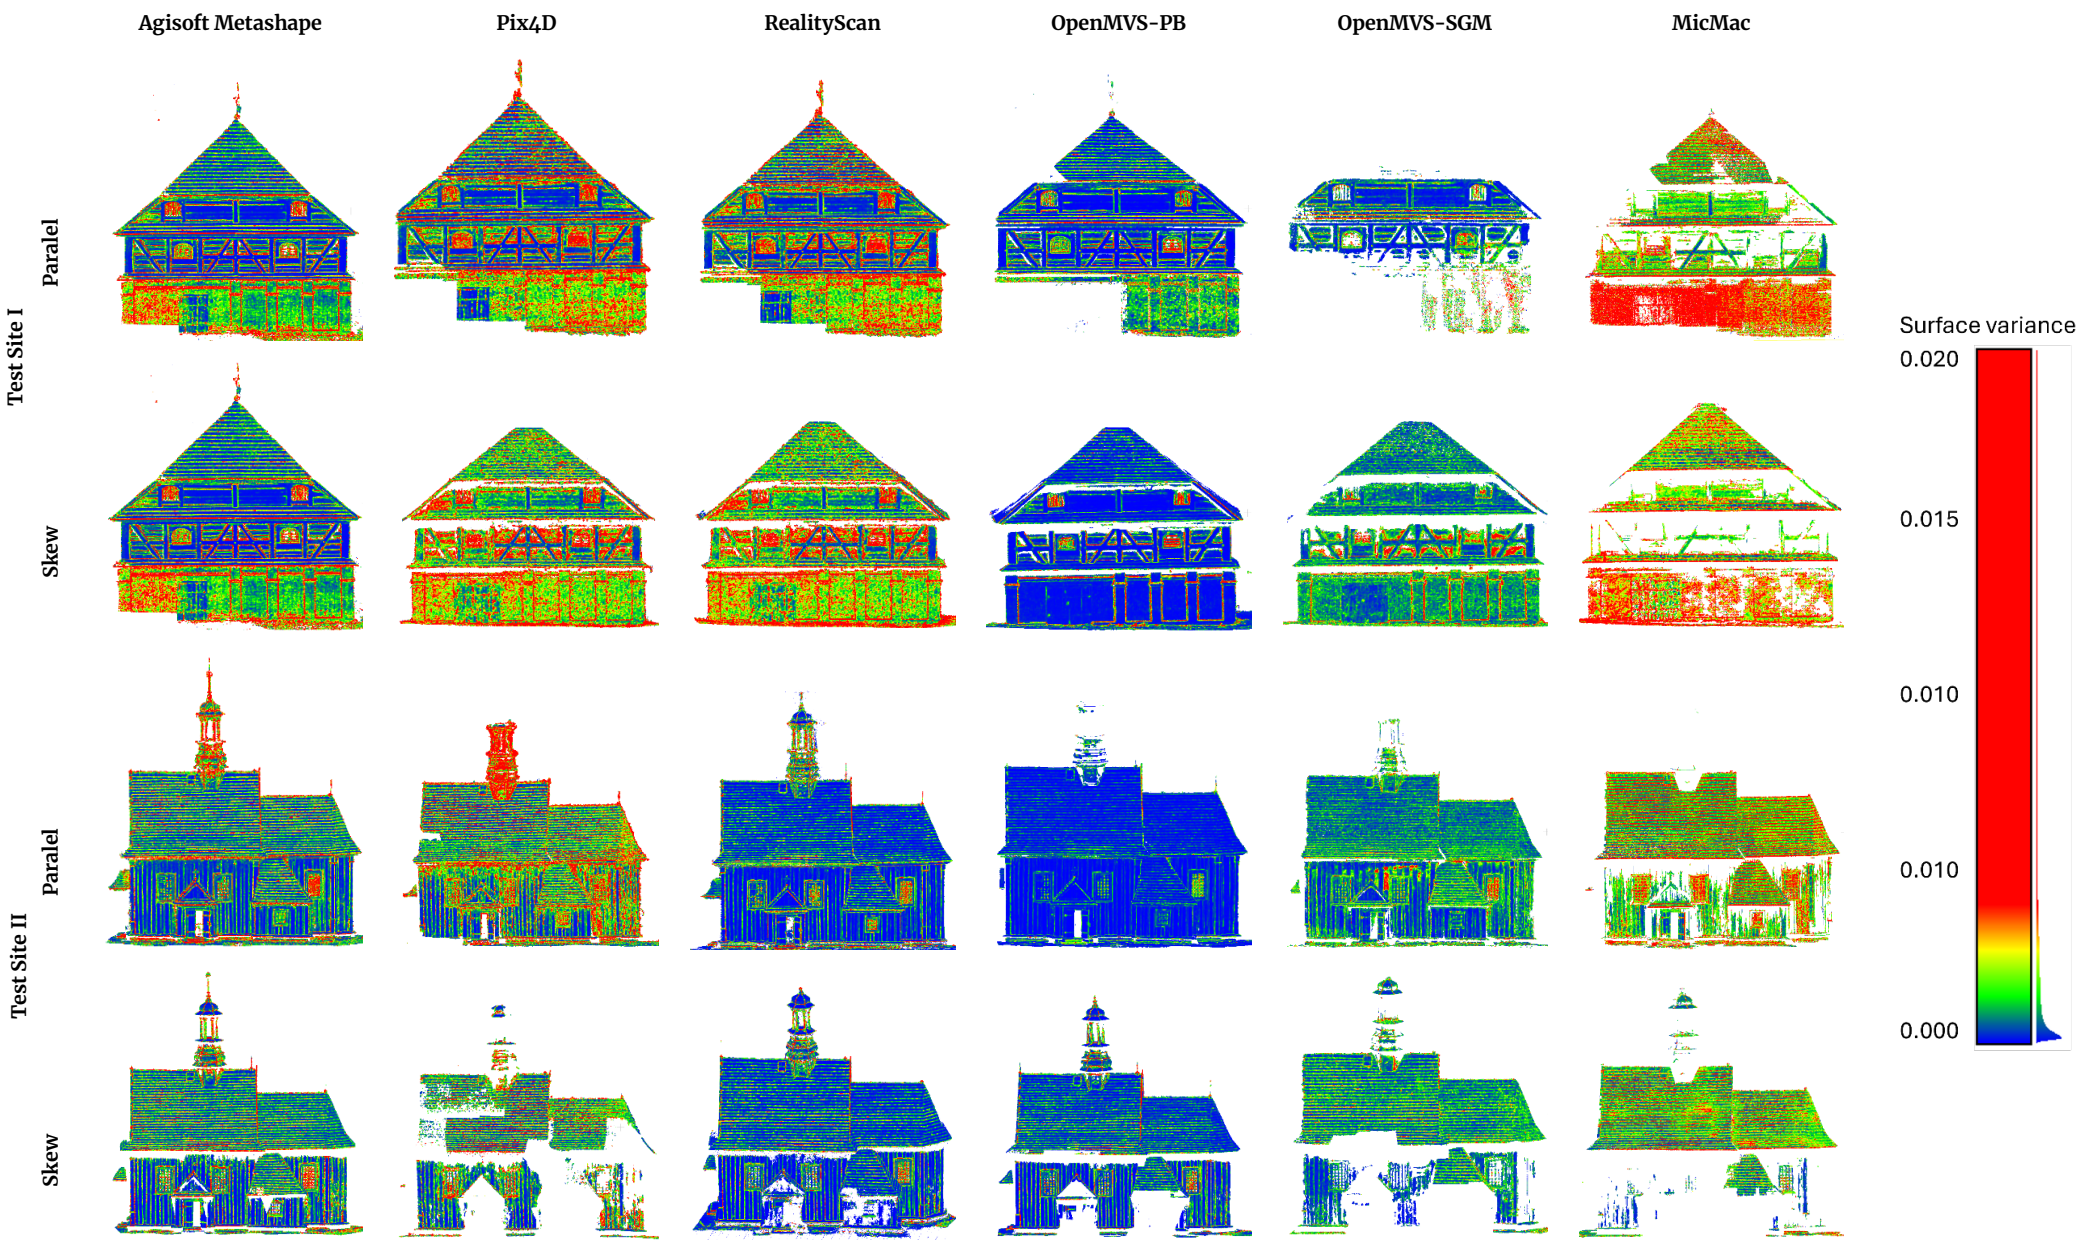

Appendix 5. The example of normal vector variances for point clouds from parallel and skew images for utilised Test Sites