ACTA IMEKO
ISSN: 2221-870X
December 2026, Volume 15, Number 4, 1 - 11

A raster-based workflow for classifying image-derived point clouds in cultural heritage GIS applications

Barbara Marana[1], Giorgio Ubbiali[2]

[1] *DISA, Department of Engineering and Applied Sciences, University of Bergamo, Viale G. Marconi 5, 24044 Dalmine (Bergamo), Italy*
[2] *Diemme Strumenti, Via Grandi 2B/c, 22063 Cantù (Como), Italy*

ABSTRACT
This work presents a raster-based workflow for classifying image-derived point clouds within a 3D GIS environment for cultural heritage applications. The case study is based on a dataset from UAV Structure-from-Motion (SfM) surveys and complemented by terrestrial SLAM-based data. Since automated point-level building extraction did not provide satisfactory results for the dataset, the proposed approach adopts a rule-based strategy using raster products to organise the point cloud into thematic classes in a transparent and reproducible manner.
The method operates on raster layers derived from the point cloud, including the Digital Surface Model (DSM), the Digital Terrain Model (DTM), the normalised Digital Surface Model (nDSM), and local surface roughness. These layers, together with automated ground classification, are used to construct thematic masks representing buildings, high vegetation, and ground, which are refined and combined into a classification scheme (LAS classes 6-5-2 for Building, High Vegetation, and Ground, respectively).
The classification is then transferred to the point cloud, allowing it to be organised and visualised within the same 3D GIS environment. The workflow follows a structured sequence of steps that takes into account the geometric characteristics of the dataset and provides a reproducible procedure for classifying image-derived point clouds within a GIS environment.

Section: RESEARCH PAPER

Keywords: cultural heritage; 3D GIS; image-derived point clouds; point cloud classification; raster-based workflow

Citation: B. Marana, G. Ubbiali, A raster-based workflow for classifying image-derived point clouds in cultural heritage GIS applications, Acta IMEKO, vol. 15 (2026) no. 4, pp. 1-11. DOI: 10.21014/actaimeko.v15i4.2356

Section Editor: Fabio Leccese, Università Degli Studi Roma Tre, Rome, Italy

Received March 31, 2026; **In final form** November 19, 2026; **Published** December 2026

Copyright: This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License.

Corresponding author: Barbara Marana, e-mail: barbara.marana@unibg.it

1. INTRODUCTION

Metric documentation of cultural heritage is increasingly based on tools capable of coherently integrating geometric representation, descriptive information, and GIS-based spatial processing [1]. In recent years, the combined use of different survey techniques and three-dimensional Geographic Information Systems (3D GIS) has been widely explored in the literature, highlighting the potential of 3D GIS to organise heterogeneous datasets and integrate geometric and semantic information for 3D data management in cultural heritage applications [1]–[3].

In parallel, UAV photogrammetry based on Structure-from-Motion (SfM) has become a widely adopted approach for generating high-density point clouds in cultural heritage applications. Numerous authors have investigated the use of SfM for three-dimensional documentation in these contexts [4]–[7]. However, image-derived point clouds differ structurally from multi-return LiDAR acquisitions. Multi-return LiDAR systems can record several returns along a single laser pulse, capturing distinct surfaces at different elevations (e.g. building roofs and the ground beneath). By contrast, SfM-based point clouds reconstruct only the surfaces visible from the available viewpoints and do not register multiple vertical returns. Consequently, although height-based visualisation may allow a visual separation of elements, vertical stratification is not intrinsically encoded in image-derived point clouds, as they represent only the visible surfaces.

These structural differences are especially relevant for the derivation of the Digital Terrain Model (DTM), as its computation requires the identification of ground points (e.g. points beneath buildings or vegetation), which is not intrinsically encoded in image-derived point clouds [4], [8]. Consequently, products derived from the DTM, such as the normalised Digital Surface Model (nDSM), are also affected. These layers, together with surface roughness derived from the Digital Surface Model (DSM), constitute the basis for the raster-based classification

workflow adopted in this work. The objective is to organise the point cloud into thematic classes through the generation and combination of raster-derived products.

In this study, the dataset combines UAV-based Structure-from-Motion (SfM) surveys and terrestrial laser scanning data acquired using SLAM technology. While this integration improves geometric completeness through complementary aerial and terrestrial acquisitions, the resulting point cloud does not provide multi-return information comparable to LiDAR data and, therefore, retains characteristics typical of image-derived surface reconstructions. Consequently, the adopted classification strategy is designed to be consistent with the characteristics of image-derived point clouds.

In addition to raster-based GIS workflows [9], various approaches for point cloud classification in cultural heritage contexts have been reported in the literature [10]. These approaches include learning-based techniques, fully automatic methods, and manual editing strategies, all of which operate directly on three-dimensional point clouds. In cultural heritage applications, learning-based techniques rely on training data derived from previously labelled point clouds [11]–[13]. While learning-based methods can achieve high classification performance, their applicability depends on the availability and diversity of training data, which must adequately reflect the variability of architectural forms typically encountered in cultural heritage contexts. Fully automatic classification approaches rely on predefined algorithmic rules with adjustable parameters, but they may struggle in complex environments characterised by articulated architectural geometries. Manual editing of three-dimensional point clouds can be effective in specific cases, but it remains inherently case-specific and dependent on operator judgement, which may limit reproducibility across different datasets.

In the analysed dataset, the automated building classification tool showed clear limitations. These limitations motivated the development of an alternative workflow that reduces reliance on direct point-based approaches, and operates primarily on two-dimensional raster products derived from the point cloud within a GIS environment. In this workflow, the automated ground classification tool is first applied at point level to identify ground points, which are used for the generation of the DTM and subsequently incorporated into the raster-based workflow. Above-ground elements are identified using a combination of height-above-ground information and surface roughness. Both ground and above-ground elements are then used for the construction and refinement of the thematic masks. This approach makes the decision logic explicit and reproducible through a structured sequence of operations. However, the adoption of a two-dimensional raster representation introduces aggregation effects inherent to grid-based elevation models [14]. When elements of different types are vertically projected onto the same raster cell, they are necessarily assigned to a single class, as each raster cell can store only one attribute value.

The present contribution concerns a site previously investigated and reported in [15], where a 3D GIS geodatabase was developed for selected interior portions of the Sanctuary of Madonna dell'Olmo (Verdellino, Province of Bergamo, Italy), based on terrestrial laser scanning and close-range photogrammetry data. The study presented here addresses a different research objective by introducing a raster-based classification workflow applied to a point cloud derived from UAV Structure-from-Motion (SfM) surveys, and complemented by terrestrial SLAM-based laser scanning data. The dataset describes the exterior areas of the complex, and is processed in order to organise and visualise the point cloud as thematic classes within a 3D GIS environment.

The remainder of the paper is structured as follows: Section 2 describes the acquisition and integration of point clouds; Section 3 presents the proposed raster-based methodology for point cloud classification; Section 4 illustrates the results and validation obtained on the test case; Section 5 discusses the advantages and limitations of the workflow; Section 6 concludes the paper and outlines future developments.

The novelty of the proposed approach does not lie in the introduction of new classification algorithms, but in the explicit formalisation of a raster-based, rule-driven workflow implemented within a GIS environment. The method provides a transparent and reproducible procedure, in which all processing steps and decision criteria are explicitly defined and can be directly modified by the user.

Unlike automated or learning-based approaches, the proposed workflow does not rely on training data or complex parameter tuning, but emphasises explicit classification rules, user control over the processing steps, and adaptability to different datasets, which are particularly relevant in cultural heritage contexts, where datasets are often heterogeneous and site-specific.

2. DATA ACQUISITION AND POINT CLOUD PREPARATION

The survey of the exterior areas of the Sanctuary of Madonna dell'Olmo was carried out through an integrated approach combining UAV-based photogrammetry and terrestrial SLAM-based laser scanning, supported by GNSS measurements for georeferencing purposes. An overview of the surveyed site is shown in Figure 1.

The surveyed area includes the main church building, characterised by roofs at different elevations, apsidal volumes, and a bell tower, as well as a small ancillary building and surrounding vegetation.

SLAM (Simultaneous Localization and Mapping) is a mobile laser scanning approach in which data are acquired while the sensor is moving, without the need for fixed scanning positions, allowing rapid coverage of complex environments. The terrestrial SLAM survey complements the UAV-based SfM acquisition, improving the geometric completeness of near-ground and vertical surfaces that are typically poorly reconstructed from aerial imagery alone.



Figure 1. UAV view of the Sanctuary of Madonna dell'Olmo, used as test site for the proposed raster-based classification workflow.

The objective of this phase was to generate a single three-dimensional point cloud by aligning and integrating SfM and SLAM data, suitable for subsequent raster-based processing and classification within a GIS environment, as described in Section 3.

2.1. UAV data acquisition

UAV imagery was acquired using a DJI Mini 3 Pro (DJI). The survey was conducted through two circular flight paths around the Sanctuary, ensuring high image overlap and redundant coverage of vertical and inclined surfaces. Additional UAV images were acquired to improve the coverage of areas not fully captured during the main flight paths, including façade details, the inner courtyard, and geometrically complex zones.

2.2. GNSS georeferencing

To ensure accurate metric positioning of the dataset, the terrestrial SLAM point cloud was georeferenced using Ground Control Points (GCPs) derived from a previously georeferenced survey. All datasets were referenced in the WGS84/UTM32N coordinate system (EPSG:32632), which was adopted throughout the entire workflow, including all subsequent raster-based analyses performed in ArcGIS Pro 3.5 (Esri).

2.3. SLAM-based terrestrial acquisition

The terrestrial survey was performed using the XGRIDS Kity K1 laser-based mobile mapping system (XGRIDS Technologies), operating using SLAM technology.

The acquisition produced a dense terrestrial point cloud intended to complement the UAV-based SfM reconstruction, particularly in near-ground and vertical areas.

Data processing was carried out using the proprietary Lixel Studio software (XGRIDS Technologies), which was used to generate the georeferenced SLAM point cloud for the subsequent integration phase.

2.4. Photogrammetric reconstruction and data integration

UAV images were processed using 3DF Zephyr (3DFlow) to generate a dense photogrammetric point cloud. This point cloud was not initially referenced to the selected spatial reference system. A set of tie points was manually identified between the SLAM point cloud and the dense photogrammetric point cloud. These points were used to reference and initially align the SfM reconstruction to the SLAM dataset. Subsequently, an Iterative Closest Point (ICP) algorithm was applied to refine the alignment, ensuring geometric consistency between the two datasets. The aligned SLAM and SfM point clouds were then integrated into a single dataset, exported in LAS format (standard point cloud file format) and used as input for the raster-based workflow described in Section 3.

3. METHODOLOGY

This section describes the methodology developed to classify the point cloud in LAS format, derived from UAV-based SfM and terrestrial SLAM surveys, entirely implemented within ArcGIS Pro.

The workflow is organised into sequential steps, from the preparation of raster products to the generation of thematic masks, leading to the classification of the point cloud into buildings, vegetation and ground (classes 6-5-2, according to the LAS Specification defined by the American Society for Photogrammetry and Remote Sensing (ASPRS) [16]), as summarised in Figure 2. Each phase addresses specific characteristics of the dataset, including the absence of multiple-return information and the effects introduced by the two-dimensional raster representation, namely the aggregation of vertically overlapping elements belonging to different classes when projected onto the same cell. The following subsections present the processing pipeline in detail.

In ArcGIS Pro, automated classification tools can be applied directly to the three-dimensional point cloud. In the present case, automated ground detection using the Classify LAS Ground tool provided satisfactory and stable results when applied to the dataset, as it relies on the morphology of the visible surface rather than on multi-return information, which is not available in image-derived point clouds, whereas automated building extraction using the Classify LAS Building tool did not produce reliable outcomes.

Although the dataset combines SfM and SLAM surveys, the resulting point cloud exhibits characteristics typical of image-derived point clouds and does not provide multi-return information comparable to LiDAR. The Classify LAS Building tool is primarily designed for LiDAR data, although it includes an option for photogrammetric datasets. In this study, different extraction parameters were tested and the photogrammetric data option was enabled; however, the tool detected only limited portions of roof surfaces and did not correctly delineate the architectural structures. For this reason, ground classification was successfully performed at point level and its results were integrated into a raster-based workflow, in which all thematic classes were created and refined through raster operations.

The workflow was implemented through a sequence of raster operations within ArcGIS Pro. These operations were used to create, refine, and combine thematic masks into a single classification raster. The resulting raster contains a unique class assignment for each cell prior to transferring the class codes to the LAS point cloud.

For clarity and reproducibility, a schematic overview of the raster-based classification workflow is provided in Appendix A.

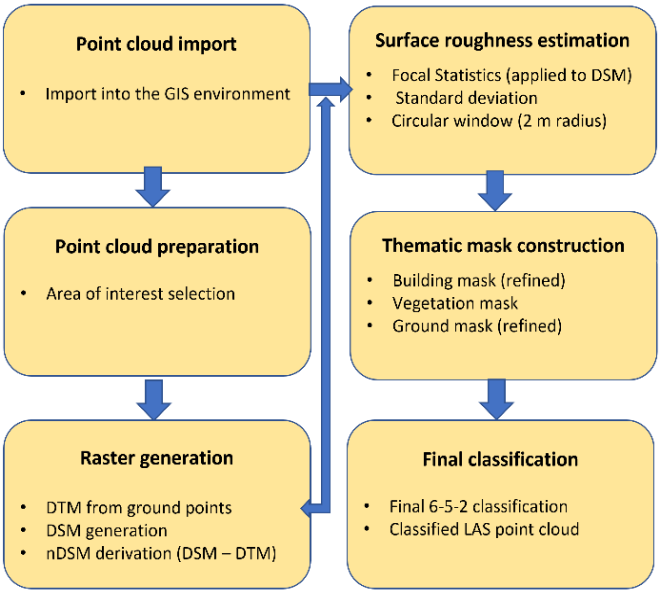

Figure 2. Raster-based workflow adopted for the classification of the point cloud, showing point cloud import and preparation within the GIS environment, DTM and DSM generation, derivation of the nDSM, surface roughness estimation, construction, refinement and fusion of thematic masks, and final point cloud classification (6-5-2).

3.1. LAS cloud processing inside the GIS environment

The point cloud was processed in ArcGIS Pro within the WGS84/UTM32N reference system. As previously discussed, the automated ground classification produced satisfactory results, whereas the automated building extraction did not yield outcomes suitable for the intended classification. Consequently, the classification approach was reorganised into a raster-based workflow, as described in the following section.

3.2. Alternative raster-based workflow

The implementation of the alternative raster-based workflow starts with the preparation of the point cloud and the generation of raster products derived from the dataset.

To enable efficient point cloud management and raster generation within the GIS environment, the LAS dataset was imported into ArcGIS Pro. Although the integrated survey covered the entire Sanctuary, the surrounding ancillary building, and the adjacent garden, the classification objective of this study focuses primarily on the built structures. The point cloud was therefore spatially extracted to concentrate the subsequent raster-based analysis on the architectural elements, while retaining portions of surrounding vegetation within the study area and reducing the computational load prior to raster generation (Figure 3).

From the extracted dataset, raster-based elevation models were generated to support the subsequent classification steps. The derivation and use of these products are described in the following subsections.

3.3. DTM, DSM, and nDSM generation

A Digital Terrain Model (DTM), representing a continuous terrain surface, was generated by applying the automated Classify LAS Ground tool available in ArcGIS Pro and exporting a raster using only points classified as ground (LAS Filter: Ground class selected). The LAS format includes a Classification attribute for each point, defined according to the ASPRS LAS specification [16], which is used for filtering and class-based operations throughout the workflow. The resulting surface was interpolated from ground points and includes reconstructed areas where direct observations were not available (e.g., beneath buildings or vegetation), due to the surface-based nature of image-derived point clouds.

A Digital Surface Model (DSM), representing the visible surface of the scene, was generated by exporting a raster from the complete point cloud (LAS Filter: All classes selected). The DSM is derived from the interpolation of all available points and therefore describes only the observed external surfaces.

A normalised Digital Surface Model (nDSM) was subsequently obtained as the difference between DSM and DTM using the Raster Calculator. This raster expresses the height of each element relative to the terrain surface and constitutes a key input for the subsequent mask construction steps, as it enables the identification of above-ground features with respect to the ground level (Figure 4 a). Height-above-ground information derived from DSM–DTM differencing is widely adopted in building detection and urban feature extraction approaches, where elevation thresholds are used to distinguish above-ground elements from terrain [17].

All raster-based processing steps in the workflow, including DTM and DSM generation, and the subsequent derivation of nDSM, surface roughness, and thematic masks, were carried out using a raster resolution of 0.5 m. The selected resolution reflects the thematic objective of the workflow, which aims to distinguish buildings, vegetation, and ground rather than to model fine architectural details. Consequently, the raster representation is not intended to provide a precise geometric delineation of object boundaries, but to support a thematic separation between the main classes, with expected boundary effects inherent to the selected resolution. The adopted value should therefore be regarded as an empirical parameter that may require adjustment when applying the workflow to datasets with different characteristics. For this reason, the selection of threshold values cannot be defined through fixed guidelines, but must be adapted to the specific characteristics of each dataset, particularly in cultural heritage contexts where architectural complexity and the presence of vegetation can significantly influence height and roughness patterns. A finer raster resolution may improve the geometric representation of object boundaries by reducing aggregation effects; however, it may also introduce increased variability in elevation and surface roughness values, leading to less stable threshold-based classification.

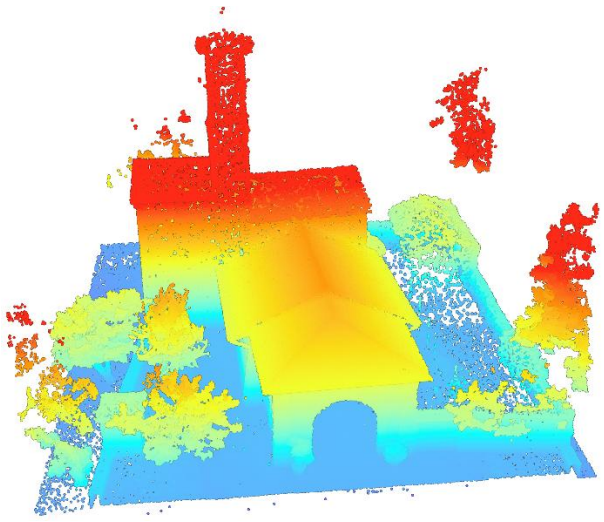

Figure 3. Spatially extracted portion of the point cloud, visualised in a 3D GIS environment using elevation-based rendering. This dataset is used as input for DTM, DSM and nDSM generation, and subsequent raster-based processing.

3.4. Surface roughness estimation

In this work, the term surface roughness refers to a raster-based measure of local elevation variability derived from the DSM and does not refer to a physical roughness parameter of the surface, but to a geometric descriptor commonly used in elevation model analysis [18]. To complement elevation-based information, surface roughness was computed from the DSM using the Focal Statistics tool in ArcGIS Pro. Specifically, the standard deviation of DSM elevation values was calculated within a circular neighbourhood with a radius of 2 m. This value was selected to account for elevation variability in the surrounding area of each cell (0.5 m resolution, consistent with the raster products used in the workflow), and may require adjustment depending on the characteristics of the analysed dataset. For each raster cell, the resulting value expresses the degree of elevation variability within the neighbourhood.

As the DSM is expressed in metres, the resulting standard deviation values are also expressed in metres and quantify the vertical variability of the surface. In the present dataset, values range approximately from 0.03 m in relatively flat or homogeneous areas to about 5.94 m in zones characterised by

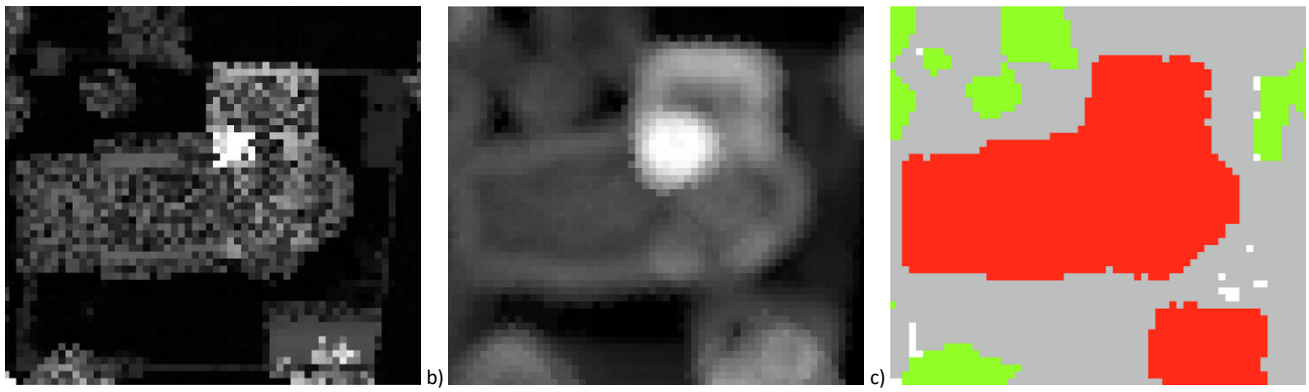

Figure 4. a) nDSM derived from the DSM and DTM through subtraction, b) Local surface roughness raster (black = low roughness; white = high roughness) computed from the DSM using the Focal Statistics tool, c) Final classification raster mask obtained through fusion of the Building (6, red), High Vegetation (5, green), and Ground (2, grey) thematic layers.

pronounced elevation variability. In the study area, elevated values (lighter tones in Figure 4 b) are observed in correspondence with building edges, vertical architectural elements such as the bell tower, and vegetation canopies.

Surface roughness was computed on the DSM rather than on the nDSM, as it is intended to describe the local height variability of the surface directly represented by the point cloud. The nDSM, being derived from the DSM-DTM subtraction, expresses relative height above ground, and therefore does not directly represent the original elevation variability of the observed surface.

In the proposed workflow, surface roughness and nDSM are used as complementary descriptors in the subsequent raster-based analysis: height above ground identifies elements located above the terrain, while surface roughness characterises the local geometric variability of the observed surface.

3.5. Thematic masks construction: Building, High Vegetation, and Ground

The raster-based workflow presented here is implemented through the progressive construction of thematic masks, each designed to identify a specific class of elements (buildings, vegetation, and ground) within the point cloud using criteria specific to each class. Height above ground and surface roughness are used to describe above-ground elements, which are used to construct an initial Building Mask. Under the adopted thresholds, this mask includes both built structures and vegetation. Vegetation is subsequently edited within the set of above-ground elements, while ground areas are derived from the automated ground classification tool and used to generate the Ground Mask. The thematic masks are then progressively refined and combined within the raster-based workflow.

In the proposed approach, the LAS class High Vegetation (class 5) is adopted to represent above-ground vegetation. This choice reflects a deliberate methodological simplification, as the objective of the workflow is not to perform a detailed analysis of vegetation structure, but to distinguish vegetation from built elements and ground within a raster-based classification.

The criteria and processing steps used for each mask are detailed below.

Building Mask

The Building Mask was generated using the Con operator in the Raster Calculator by combining height-above-ground information derived from the nDSM with surface roughness values. Cells with nDSM values greater than 0.5 m and roughness values lower than 3.5 m were classified as built surfaces and assigned to class 6.

The height threshold of 0.5 m was selected as a practical separation between terrain and above-ground elements in the analysed dataset. The adopted roughness threshold was defined through iterative testing within the Raster Calculator, in order to balance class separability and preservation of architectural features. More restrictive values resulted in the partial exclusion of complex roof geometries and vertical architectural elements, whereas higher threshold values progressively increased the extent of areas classified as built surfaces, including vegetated regions.

Under the adopted thresholds, vegetated areas also satisfy the selection criteria and are fully included in the Building Mask. At this stage, the mask does not distinguish between built structures and vegetation.

The differentiation between built structures and vegetation is completed during the subsequent refinement phase, in which the Vegetation Mask is used to refine the initial Building Mask.

High Vegetation Mask

In the proposed approach, vegetation is not extracted directly through height-above-ground or surface roughness thresholds. Given the characteristics of the point cloud, vegetation and built structures satisfy similar height and roughness conditions; therefore, threshold-based selection alone does not allow a reliable separation at this stage. In the analysed case, vegetation elements exhibit values comparable to those of built structures, and are therefore fully included in the initial Building Mask.

Vegetation areas were subsequently edited through a manual editing phase within the GIS environment, based on visual inspection of the Building Mask and verification against the point cloud. Although the editing procedure involves operator judgement, it is performed within the set of above-ground elements identified by the raster-based processing and therefore remains consistent with the overall workflow structure. The resulting vector layer was converted into a raster layer and used as the High Vegetation Mask (class 5).

When vegetation overlaps roof surfaces, the overlapping raster cells are deliberately excluded from the Vegetation Mask and retained in the Building Mask, preserving architectural structures. This choice is made in view of the subsequent fusion of the thematic masks, where a single class must be assigned to each raster cell. In this way, the coexistence of vegetation and built structures, resulting from the projection of three-dimensional elements onto the same raster cell, is handled by

assigning a single class to each cell according to the adopted classification rules, while the underlying three-dimensional configuration of overlapping elements remains unchanged.

In the proposed workflow, vegetation is treated as a single above-ground class, and no further subdivision into vegetation types was considered.

Ground Mask

The Ground Mask was generated using the points classified as ground by the Classify LAS Ground tool. After selecting the ground class (LAS Filter: Ground class selected), these points were converted into a raster representation describing the spatial distribution of ground points using the LAS Point Statistics As Raster tool, which computes the number of points within each raster cell. This raster was subsequently transformed, using the Con operator in the Raster Calculator, into a thematic raster, in which cells containing at least one ground point were assigned the ground class code (class 2), while all other cells were set to NoData. The resulting raster constitutes the Ground Mask and was subsequently refined and combined with the other thematic masks during the final classification phase.

Because the workflow operates on two-dimensional raster representations, each raster cell is assigned a single class, even when elements belonging to different classes at different elevations project onto the same cell.

These cases are subsequently managed in a dedicated refinement step prior to the final fusion of the thematic masks.

3.6. Mask refinement and final classification

Before the final combination of the thematic layers, the Building Mask was locally refined using the High Vegetation Mask as a filtering criterion. Since the initial Building Mask includes both built structures and vegetation under the adopted thresholds, a conditional operation in the Raster Calculator was applied to remove vegetation-related values from the Building Mask based on the Vegetation Mask. This step is specifically designed in view of the subsequent fusion of the thematic masks, where a single class must be assigned to each raster cell. At raster level, a single class must be assigned to each cell; therefore, when vegetation and built structures occupy the same cell, their coexistence cannot be represented and is resolved through the adopted classification rules, as described in the High Vegetation Mask section. In particular, the Vegetation Mask is defined only in areas where vegetation is considered separable from built structures, while cells corresponding to vegetation overlapping roof surfaces are excluded from it and retained in the Building Mask. As a result, the refined Building Mask retains the full spatial extent of architectural elements, while vegetation is removed only where it is treated as a separate class.

The Ground Mask was then refined by removing ground values in raster cells where the refined Building Mask or the High Vegetation Mask were present. This step ensures that class 2 values are preserved only in terrain areas not occupied by above-ground elements and is consistent with the requirement to assign a single class to each cell during the final fusion of the thematic masks.

After refinement, the three thematic masks were combined into a single classified raster using a sequence of conditional operations in the Raster Calculator, in which class values were assigned based on the presence of the different masks within each raster cell. The resulting raster, encoding the final 6-5-2 classification scheme, is shown in Figure 4 c.

White gaps visible in the final classification raster correspond to NoData cells, mainly associated with areas not fully reconstructed in the original point cloud (e.g., vertical surfaces or occluded regions). These gaps do not represent misclassification errors but reflect intrinsic limitations of surface reconstruction, and are deliberately preserved to avoid introducing artificial or unsupported classifications.

3.7. Application to the LAS point cloud

Once the final 6-5-2 classification raster was generated, the corresponding class information was transferred to the point cloud using the Set LAS Class Codes Using Raster tool. Through this operation, the numeric values assigned in the raster were written into the LAS Classification field of each point according to the value of the raster cell in which it falls. During raster processing, class codes are handled by the GIS software as integer labels without intrinsic semantic meaning. Their standard interpretation is restored when the final classification is written back to the LAS dataset, where values 2, 5, and 6 correspond to the Ground, High Vegetation, and Building classes, according to the ASPRS LAS specification [16].

The resulting point cloud, classified according to the 6-5-2 scheme and organised according to standard LAS class codes, can be directly visualised within the 3D GIS environment (Figure 5).

The gaps visible in the 3D classified view (Figure 5) correspond to areas lacking point cloud data and reflect the same reconstruction incompleteness discussed for the raster representation in Figure 4 c. These gaps do not represent misclassification errors but are preserved to avoid introducing artificial or unsupported class assignments.

In addition, small portions of roof surfaces, particularly along roof edges, are visible in Figure 5 as being classified as ground (class 2, shown in grey), despite belonging to roof surfaces. These effects occur at the boundaries of architectural elements, where raster discretisation and the adopted threshold-based classification logic may lead to raster cells that do not satisfy the adopted classification thresholds. As a result, portions of roof surfaces may be classified as ground, while in other areas the Building Mask may extend beyond the actual point cloud due to the aggregation of values within each raster cell.

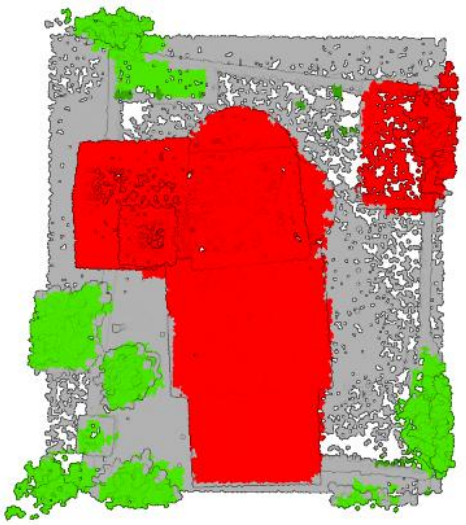

Figure 5. Classified point cloud shown in a top-down view, displaying Building (class 6, red), High Vegetation (class 5, green), and Ground (class 2, grey) after fusion of the raster masks and application of class codes to the LAS point cloud.

4. RESULTS

The proposed raster-based workflow, as described in Section 3, produced a classified point cloud organised according to the 6-5-2 scheme (Building, High Vegetation, and Ground) from the analysed dataset. The sequential generation of DTM, DSM, nDSM and surface roughness, together with the use of automated ground classification, enabled the construction and refinement of three thematic masks: Building, High Vegetation, and Ground, which were combined to obtain the final raster-based classification (Figure 4).

The spatial distribution of thematic classes, shown in Figure 4 c, was derived from the attribute table of the final classification raster mask. Considering the raster resolution of 0.5 m (cell area = 0.25 m²), the Ground class represents 51.5 % of the classified raster cells (428.3 m²), the Building class 37.4 % (310.8 m²), and High Vegetation 11.1 % (92.0 m²), corresponding to a total classified area of approximately 831.0 m². White gaps visible in the raster mask correspond to NoData cells associated with areas not fully reconstructed in the original point cloud.

When applied to the LAS point cloud, the resulting 6-5-2 classification scheme correctly identified the roof surfaces and vertical architectural elements of the Sanctuary, including the bell tower, as well as the surrounding vegetated areas and the terrain (Figure 5).

An overlay in 2D between the final classification raster mask and the original point cloud is shown in Figure 6. Visual inspection confirms that class assignments are geometrically consistent with the original dataset; minor discrepancies are primarily confined to class boundaries and reflect the aggregation of values within raster cells resulting from elements belonging to different classes, and the application of the adopted classification thresholds, rather than misclassification of the underlying point cloud. These discrepancies arise because class values are assigned at the raster-cell level and subsequently transferred to all points within each cell. As a result, cells located at the boundaries of architectural elements may include points with different geometric characteristics: in some cases, roof points are assigned to the ground class when the cell does not satisfy the building criteria (based on height-above-ground and surface roughness), while in other cases ground points are included within cells classified as building, producing local expansions of the building class. Similar effects can also be observed in the three-dimensional representation, including cases involving ground and vegetation, although these are not clearly visible in the two-dimensional raster representation, due to the aggregation of values within each cell. These effects are more clearly observable in the three-dimensional representation, as illustrated in Figure 7. The classified point cloud shows a clear spatial separation between built surfaces, vegetation, and ground, illustrating how the raster-based approach organises the dataset within a 3D GIS environment.

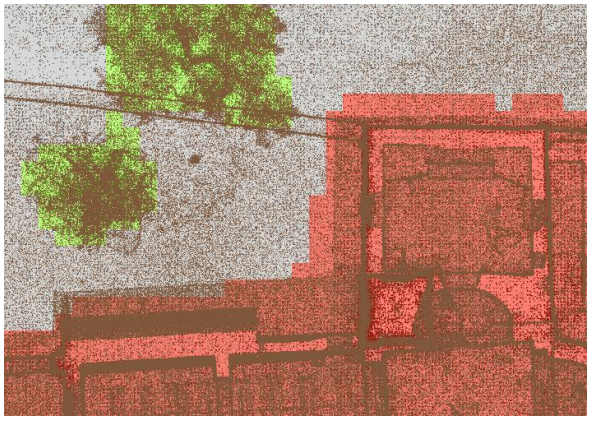

Figure 6. 2D overlay between the final classification raster mask and the original point cloud (displayed using intensity values). The raster mask encodes Building (red), High Vegetation (green), and Ground (grey). The figure shows geometric consistency between thematic classes and the underlying data, as well as local boundary effects related to raster discretisation and classification criteria.

The aim of the proposed workflow is not to optimise classification performance through statistical accuracy assessment, but to define a transparent and reproducible classification procedure within a GIS environment. The quantitative values previously reported refer to the spatial distribution of thematic classes derived from the final raster mask and do not constitute a statistical validation of classification performance. For this reason, the effectiveness of the method is assessed through both visual inspection of the resulting classification and sample-based quantitative validation, recognising that the observed limitations are intrinsically related to the adopted raster-based workflow.

A representative example of these limitations is observed where a tree crown partially overlaps the roof of a small ancillary building, as shown in Figure 7 a. This case represents a situation in which different three-dimensional elements project onto the same two-dimensional raster cell. In such cases, the raster assigns a single class value to each cell, which does not allow the individual contributions of vegetation and roof surfaces to be distinguished. This limitation becomes relevant when assigning a unique class value during the final combination of the thematic masks. In the proposed workflow, this situation is handled by

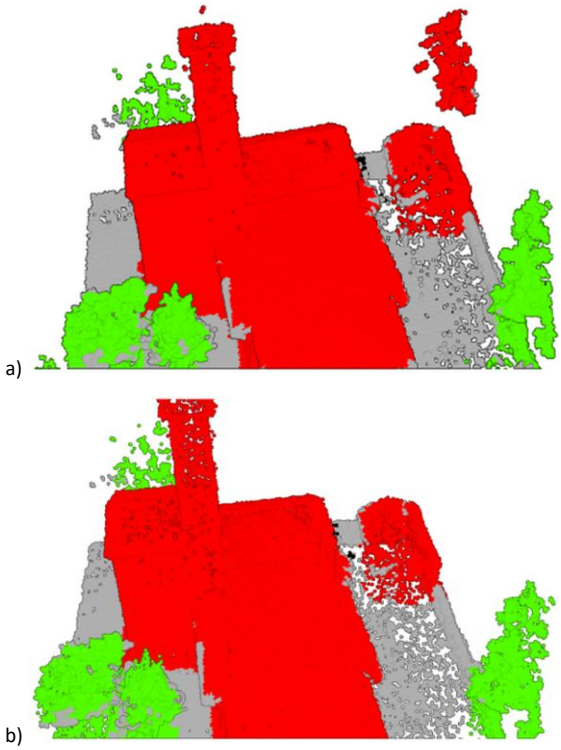

Figure 7. a) Example of classification effects in the point cloud, where a tree crown partially overlaps the roof of the ancillary building and both elements are assigned to the same raster cell. b) Resulting classification after manually removing canopy points from the point cloud for illustrative purposes, where the overlap between vegetation and roof surfaces is no longer present.

assigning the cell to the Building class in cases of overlap. When the raster classification is transferred back to the point cloud, this decision becomes visible in the three-dimensional representation: both the roof surface and the overlapping vegetation are assigned to the Building class and are therefore visualised in red. Vegetation is assigned only in areas where it occupies raster cells distinct from built structures. For illustrative purposes, the overlapping vegetation was manually removed within the same processing software used to generate and integrate the point cloud. This additional editing step was performed solely to provide a visual comparison and does not constitute a standard component of the proposed workflow. After this illustrative modification, the final 6-5-2 classification raster was re-applied to the corrected LAS dataset. The resulting classification, shown in Figure 7 b, no longer exhibits in the point cloud the aggregation effect between roof surfaces and vegetation previously visible in Figure 7 a.

More generally, whenever vegetation, built structures, or terrain elements occupy the same raster cell at different elevations, the raster representation assigns a single class value to that cell, as required during the final classification stage. This value is subsequently transferred to the corresponding three-dimensional points, so that vertical separability is not preserved in the classified 3D visualisation, even though its geometry remains distinct. This behaviour reflects an intrinsic limitation of the raster representation and does not originate from point-level misclassification, but from the projection of three-dimensional geometry onto a two-dimensional raster representation.

Overall, the results confirm that the proposed raster-based workflow provides a coherent organisation of the point cloud, while explicitly managing the inherent limitations of raster-based representation.

4.1. Validation of the classification raster mask

A sample-based validation was carried out to provide a quantitative check of the classification raster mask. The validation was performed on selected reference areas manually edited on the original DSM of the dataset, and was organised into two complementary tests. The reference polygons were defined on the DSM in order to maintain consistency with the raster-based nature of the proposed workflow.

The two validation tests were designed to evaluate the behaviour of the proposed workflow under different conditions: the first focused on areas where Building, Ground, and High Vegetation could be unambiguously identified on the DSM, while the second focused on the Building–Ground interface of the church in order to assess local classification behaviour at class boundaries.

First, the original DSM was visually inspected, with continuous reference to the point cloud, in order to identify areas where the three target classes could be recognised without ambiguity. Five reference polygons were defined for each class (Building, High Vegetation, and Ground), focusing on clearly identifiable regions. Building polygons were selected on roof surfaces, Ground polygons on open terrain areas, and High Vegetation polygons on clearly recognisable vegetation crowns, providing a controlled evaluation of the classification logic in clearly separable areas. Transition zones (i.e., class boundaries) and the area affected by the tree canopy overlapping the small ancillary building were deliberately avoided.

Each reference polygon was manually edited within a single vector layer and assigned the corresponding LAS class code (6 for Building, 5 for High Vegetation, and 2 for Ground) through a dedicated attribute field. The dataset was then rasterised using the same spatial resolution (0.5 m) and grid alignment as the classification raster mask. The resulting validation raster consisted of 50 cells, distributed across the three classes (24 Ground, 13 High Vegetation, and 13 Building). The number of validation cells exceeds the number of reference polygons, because these polygons covered areas larger than the adopted raster cell size and therefore produced multiple raster cells after rasterisation.

The validation raster was then compared with the classification raster mask on a cell-by-cell basis using the Con function in the Raster Calculator, which assigns a value of 1 to matching cells and 0 to non-matching cells. As a result, all 50 validation cells were assigned a value of 1, indicating full agreement between the two rasters within the selected validation sample. The validation is intended as a sample-based quantitative verification of selected classified areas, rather than a comprehensive accuracy assessment of the entire dataset. This result confirms the coherence of the classification rules in areas where the target classes are clearly separable.

To complement the previous validation, an additional building-focused verification was carried out along the Building–Ground interface of the church, with reference polygons distributed along the entire perimeter (Figure 8). Small reference polygons were manually edited close to the boundary between the two classes in order to locally sample the behaviour of the classification along the Building–Ground interface. The polygons were positioned immediately adjacent to the class boundary while remaining within areas that could be clearly identified as Building or Ground on the DSM. Separate polygon layers were defined for the Building and Ground classes. A total of 14 polygons were defined for each class and assigned the corresponding class codes.

The reference polygons were then rasterised using the same spatial resolution and grid alignment as the classification raster mask, resulting in two validation rasters consisting of 16 Ground cells and 15 Building cells, respectively. Due to the relationship between the size of the reference polygons and the adopted raster resolution, the number of validation cells produced by the rasterisation process differed from the number of polygons. The validation rasters were subsequently compared with the classification raster mask using the Con function in the Raster Calculator.

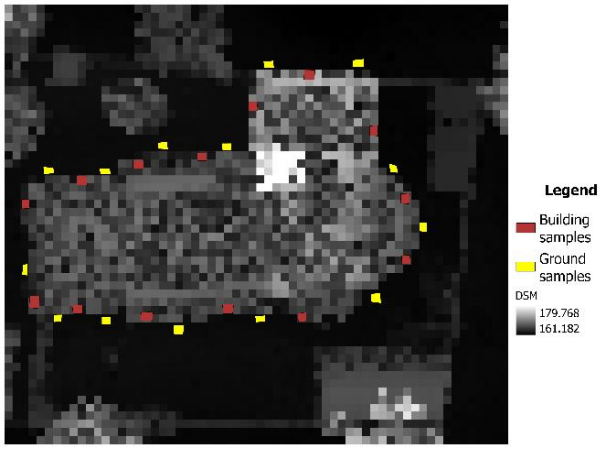

Figure 8. Reference polygons used for the building-focused validation test, distributed along the Building–Ground interface of the church. Building samples are shown in red and Ground samples in yellow over the DSM representation.

The comparison showed 27 correctly classified cells out of 31 total validation cells and 4 discrepancies, corresponding to a cell-based agreement of 87.1 %. The observed discrepancies occurred in both directions at the Building–Ground interface: two reference Building cells were classified as Ground in the raster mask, while two reference Ground cells were classified as Building.

These discrepancies reflect local effects related to raster discretisation and the adopted threshold-based classification logic, rather than systematic misclassification of the underlying point cloud.

5. DISCUSSION

The contribution of this work lies primarily in the structured integration of well-established raster-derived products within a GIS-based classification procedure. While individual components such as DSM, DTM, nDSM, and surface roughness are widely used, their explicit combination into a transparent and reproducible workflow tailored for image-derived point clouds represents a practical contribution for GIS-oriented applications in cultural heritage.

The proposed approach prioritises interpretability and reproducibility over algorithmic complexity, offering a method that can be directly implemented and adapted within standard GIS environments.

The results illustrate how the proposed workflow operates and clarify the limitations inherent to the raster-based approach. Although the processing is performed within a two-dimensional raster representation, the final outcome is a classified three-dimensional point cloud, as the class information is transferred back to the LAS dataset.

When distinct three-dimensional elements project onto the same raster cell, the raster representation assigns a single class value to the cell, and their individual contributions cannot be distinguished. These cases do not indicate errors in the three-dimensional point cloud itself but result from the reduction of three-dimensional geometry to a two-dimensional raster representation.

Since the workflow operates on raster products derived from the elevation models, the resulting classification raster mask remains consistent with the two-dimensional surface representation of the dataset. The main limitations arise in situations involving vertical overlap between different elements, where the raster representation cannot preserve three-dimensional separability.

The objective of this study is not to optimise three-dimensional point-level classification, but to define a transparent and reproducible procedure for organising point clouds within a raster-based GIS workflow, and subsequently transferring the classification to the three-dimensional dataset.

In this context, the proposed workflow is intended as a complementary GIS-oriented alternative to learning-based and point-level classification approaches, prioritising transparency and explicit control of the classification logic over automated optimisation.

In the proposed approach, elevation above ground (nDSM) and surface roughness are used as complementary geometric descriptors to characterise above-ground elements. However, when applied through threshold criteria, these descriptors do not provide a stable separation between vegetation and built structures in the analysed dataset, since vegetation elements satisfy the same height and roughness conditions as architectural features. For this reason, the Building Mask includes both built structures and vegetation, ensuring that built structures are not excluded, even though vegetation is consequently included within the same mask.

Vegetation is subsequently edited within the set of above-ground elements defined by the Building Mask. In cases where vegetation can be clearly separated from built surfaces, it is explicitly identified and assigned to the corresponding class. When vegetation and roof surfaces project onto the same raster cell, priority is assigned to the Building class in order to preserve architectural structures. In this way, each raster cell receives a single class assignment, ensuring internal consistency within the raster representation and supporting the subsequent fusion of the thematic masks.

However, the geometric configuration of overlapping elements in the three-dimensional space remains unchanged, although their class assignment in the visualised point cloud may differ due to the raster-based classification. When vegetation physically overlaps roof surfaces, the raster representation assigns a single class to the corresponding cell. The loss of vertical separability therefore reflects an intrinsic limitation of representing three-dimensional geometry within a two-dimensional raster representation, rather than a misclassification of the underlying point cloud.

This behaviour is illustrated by the case shown in Figure 7, and applies more generally whenever buildings, vegetation, or terrain elements project onto the same raster cell at different elevations. In such cases, the raster representation assigns a single class value to the cell, preventing the explicit preservation of vertical separability.

Within the proposed workflow, these cases are addressed through explicit classification decisions embedded in the construction and refinement of the thematic masks. Vegetation is removed from the Building Mask in cases where it can be clearly distinguished, while ground values are excluded wherever above-ground classes are present. In this way, overlaps are resolved at raster level in view of the final fusion of the thematic masks, where a single class value must be assigned to each raster cell, ensuring internal consistency of the classification.

In addition, given the adopted raster resolution (0.5 m), the objective of the workflow is not to obtain a geometrically precise identification of vegetation boundaries, but to support the thematic separation between the main classes within the raster-based classification. This limitation applies more generally to the geometric definition of all classes, including built structures, due to the inherent discretisation of the raster representation.

The classification behaviour is also influenced by the adopted threshold values for elevation and surface roughness. These effects are consistent with the sample-based quantitative validation, which highlighted local classification discrepancies at class boundaries. In these areas, raster discretisation and local variations in elevation and surface roughness may lead to cells being assigned to adjacent classes. The adopted threshold values depend on the geometric characteristics of the analysed dataset and must be selected according to site-specific conditions. In contexts characterised by complex architectural geometries or close spatial interaction between built structures and vegetation, threshold selection becomes particularly sensitive and may lead to local misclassification or incomplete identification of both built and ground elements. Nevertheless, the modular structure of the workflow allows specific thematic masks to be adapted or omitted depending on the characteristics of the analysed site (for example in sites characterised by a single building or simple

architectural structures, or in the absence of vegetation). While the threshold values may vary across datasets, the sequence of operations and the underlying classification logic remain transferable to other image-derived point clouds.

Gaps visible in both the point cloud and the derived raster products reflect incomplete reconstruction mainly associated with the image-based component of the dataset (SfM), particularly in areas such as vertical surfaces or occlusions. The workflow deliberately preserves these gaps, avoiding interpolation or forced classification in areas where no reliable geometric information is available.

Despite the above-mentioned limitations, the resulting classified point cloud can be effectively organised and visualised within a 3D GIS according to the defined thematic classes. The workflow therefore provides a transparent and reproducible procedure for organising image-derived point clouds in cultural heritage applications while explicitly acknowledging the implications of raster-based representation.

The proposed approach is particularly suited for datasets where transparency and reproducibility are prioritised over fully automated classification performance.

6. CONCLUSIONS

This work presents a raster-based workflow for classifying image-derived point clouds within a 3D GIS environment for cultural heritage documentation. The proposed approach combines elevation-derived raster products, surface roughness information, and automated ground classification through the construction, refinement, and fusion of thematic masks, providing a transparent and reproducible procedure for organising point clouds in cultural heritage sites, including those characterised by articulated architectural structures and surrounding vegetation. Although the workflow is reproducible in its structure, the threshold values, the vegetation editing and the configuration of the thematic masks can be adapted to the specific characteristics of each dataset.

Application to the case study shows that the fusion of the Building, High Vegetation, and Ground masks (6-5-2) produces a raster-based classification layer. In this representation, each raster cell receives a single class assignment derived from the combined thematic masks. The resulting classification can be transferred to the three-dimensional point cloud, enabling class-based visualisation and organisation within a 3D GIS environment.

When three-dimensional elements are represented within the same raster cell at different elevations, a single class value is assigned and vertical separability is therefore not preserved. Consequently, all points falling within that cell are visualised as belonging to the same class and are displayed using the same class-based symbology in the three-dimensional representation. This behaviour reflects an intrinsic property of grid-based modelling, rather than a point-level misclassification.

Where vertical overlaps are absent, the proposed workflow provides a coherent classification of built structures, vegetation, and ground. The sample-based quantitative validation confirmed the internal consistency of the workflow in clearly separable areas, while also highlighting local classification effects at class boundaries related to raster discretisation and threshold-based classification.

In areas where different elements project onto the same raster cell at different elevations, the raster representation assigns a single class value to each cell; in accordance with the adopted classification logic, built structures are preserved, vegetation is assigned only where it can be clearly separated, and ground is assigned only where no above-ground elements (i.e., buildings and vegetation) are present. This results in class attribution effects in areas affected by vertical overlap, which become evident in the three-dimensional representation and arise from the requirement to assign a single class value to each raster cell during the final combination of the thematic masks. These effects include spatial discontinuities of the ground class, the attribution of overlapping vegetation to the Building class, and local classification effects at element boundaries.

In the analysed dataset, the Classify LAS Building tool did not provide satisfactory results for identifying built structures, whereas the Classify LAS Ground tool produced stable and reliable outcomes. These limitations motivated the adoption of the proposed raster-based workflow, which addresses these issues through explicit rule-based decisions implemented on raster products within a GIS environment.

Although the workflow has been tested on a single cultural heritage case study, the proposed methodology is based on transferable GIS-based operations and explicit rule-based procedures that can be applied and adapted to other image-derived point clouds characterised by different architectural and environmental conditions. The selected case study is therefore intended not as a statistically representative validation dataset, but as a representative operational scenario used to illustrate the applicability, behaviour and limitations of the proposed workflow in a complex cultural heritage context. Future developments may focus on applying the proposed workflow to additional datasets and cultural heritage scenarios in order to further assess its applicability across different architectural configurations and site conditions, as well as to investigate the effects of raster discretisation and vertical overlaps. Further work may address the distinction between vegetation and built surfaces in image-derived point clouds, particularly in cases where they produce similar geometric responses. Additional tests may also explore the use of finer raster resolutions in order to evaluate their influence on boundary aggregation effects and classification stability.

AUTHORS' CONTRIBUTION

Barbara Marana: Conceptualization, Methodology, Formal analysis, Writing – original draft, Writing – review and editing.

Giorgio Ubbiali: Data curation, Investigation, Writing – review and editing.

REFERENCES

[1] I. M. E. Zaragoza, G. Caroti, A. Piemonte, The use of image and laser scanner survey archives for cultural heritage 3D modelling and change analysis, Acta IMEKO vol. 10 (2021) no. 1, pp. 1–8. DOI: 10.21014/acta_imeko.v10i1.847

[2] L. Centarti, C. Ferreyra, C. G. Guida, M. Limongiello, B. Messina, Documentation and management of complex 3D morphologies through digital technology, Acta IMEKO vol. 13 (2024) no. 2, pp. 1–7. DOI: 10.21014/actaimeko.v13i2.1845

[3] N. Dell'Unto, G. Landeschi, Archaeological 3D GIS, Routledge, London, pp. 176, ISBN 9781003034131, 2022. DOI: 10.4324/9781003034131

[4] A. Capolupo, Accuracy assessment of cultural heritage models extracting 3D point cloud geometric features with RPAS SfM-MVS and TLS techniques, Drones vol. 5 (2021) no. 4, art. 145. DOI: 10.3390/drones5040145

[5] C. Achille, A. Adami, S. Chiarini, S. Cremonesi, F. Fassi, L. Fregonese, L. Taffurelli, UAV-based photogrammetry and integrated technologies for architectural applications-Methodological strategies for the after-quake survey of vertical structures in Mantua (Italy), Sensors vol. 15 (2015) no. 7, pp. 15520–15539.
DOI: 10.3390/s150715520

[6] A. Ulvi, Documentation, three-dimensional (3D) modelling and visualization of cultural heritage by using unmanned aerial vehicle (UAV) photogrammetry and terrestrial laser scanners, International Journal of Remote Sensing vol. 42 (2021) no. 6, pp. 1994–2021.
DOI: 10.1080/01431161.2020.1834164

[7] F. Remondino, Heritage recording and 3D modeling with photogrammetry and 3D scanning, Remote Sens. vol. 3 (2011) no. 6, pp. 1104–1138.
DOI: 10.3390/rs3061104

[8] M. R. James, S. Robson, Straightforward reconstruction of 3D surfaces and topography with a camera: accuracy and geoscience application, Journal of Geophysical Research: Earth Surface vol. 117 (2012) no. F3, art. F03017.
DOI: 10.1029/2011JF002289

[9] A. Jochem, B. Höfle, V. Wichmann, M. Rutzinger, A. Zipf, Area-wide roof plane segmentation in airborne LiDAR point clouds, Computers, Environment and Urban Systems vol. 36 (2012) no. 1, pp. 54–64.
DOI: 10.1016/j.compenvurbsys.2011.05.001

[10] S. Yang, S. Xu, W. Huang, 3D point cloud for cultural heritage: A scientometric survey, Remote Sensing vol. 14 (2022) no. 21, art. 5542.
DOI: 10.3390/rs14215542

[11] M. Weinmann, B. Jutzi, C. Mallet, Semantic 3D scene interpretation: a framework combining optimal neighborhood size selection with relevant features, ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences vol. II-3 (2014), pp. 181–188.
DOI: 10.5194/isprsannals-II-3-181-2014

[12] M. Weinmann, B. Jutzi, S. Hinz, C. Mallet, Semantic point cloud interpretation based on optimal neighborhoods, relevant features and efficient classifiers, ISPRS Journal of Photogrammetry and Remote Sensing vol. 105 (2015), pp. 286–304.
DOI: 10.1016/j.isprsjprs.2015.01.016

[13] T. Hackel, J. D. Wegner, K. Schindler, Joint classification and contour extraction of large 3D point clouds, ISPRS Journal of Photogrammetry and Remote Sensing vol. 130 (2017), pp. 231–245.
DOI: 10.1016/j.isprsjprs.2017.05.012

[14] T. Hengl, Finding the right pixel size, Computers & Geosciences vol. 32 (2006) no. 9, pp. 1283–1298.
DOI: 10.1016/j.cageo.2005.11.008

[15] B. Marana, G. Ubbiali, From surveying techniques to a 3D GIS geodatabase, Proc. of the 2025 IMEKO TC26 Int. Conf. on Metrology for Archaeology and Cultural Heritage (MetroArchaeo), Bergamo, Italy, 15–17 October 2025, pp. 527–532.
DOI: 10.21014/tc26-2025.098

[16] ASPRS, LAS Specification Version 1.4 – Revision 15, American Society for Photogrammetry and Remote Sensing, 2019.

[17] K. Zhang, S. C. Chen, D. Whitman, M. L. Shyu, J. Yan, C. Zhang, A progressive morphological filter for removing nonground measurements from airborne LiDAR data, IEEE Transactions on Geoscience and Remote Sensing vol. 41 (2003) no. 4, pp. 872–882.
DOI: 10.1109/TGRS.2003.810682

[18] C. H. Grohmann, M. J. Smith, C. Riccomini, Multiscale analysis of topographic surface roughness in the Midland Valley, Scotland, IEEE Transactions on Geoscience and Remote Sensing vol. 49 (2011) no. 4, pp. 1200–1213.
DOI: 10.1109/TGRS.2010.2053546

APPENDIX A

Raster-based classification workflow

This appendix provides a schematic overview of the raster-based classification procedure adopted in the proposed workflow.

Table A1 summarises the sequence of processing steps, the input data, the GIS tools employed and the resulting outputs, with the aim of improving clarity and reproducibility.

Table A1. Overview of the raster-based classification workflow (Input → Tool/Process → Output).

Step	Input data	Tool/Process	Output
1	Point cloud (LAS)	Import into GIS environment	Point cloud available for raster-based processing
2	Point cloud (all points)	Automated ground classification	Point cloud with ground class
3	Point cloud (ground class)	Export Raster	Digital Terrain Model (DTM)
4	Point cloud (all points)	Export Raster	Digital Surface Model (DSM)
5	DSM; DTM	Raster Calculator (DSM – DTM)	Normalised Digital Surface Model (nDSM)
6	DSM	Focal Statistics (standard deviation, circular neighbourhood, radius 2 m)	Surface roughness raster
7	nDSM, Surface roughness	Raster Calculator (Con operator)	Building Mask (class 6)
8	Building Mask, including both buildings and vegetation	Vector-based editing and vector-to-raster conversion	High Vegetation Mask (class 5)
9	Point cloud (ground class)	LAS Point Statistics as Raster (point count)	Ground point distribution raster
10	Ground point distribution raster	Raster Calculator (conversion to thematic raster)	Ground Mask (class 2)
11	Building and High Vegetation masks	Raster Calculator (remove building where high vegetation is present in clearly separable areas)	Refined Building Mask
12	Ground, Building, High Vegetation masks	Raster Calculator (remove ground where building or high vegetation are present)	Refined Ground Mask
13	Refined Building and Ground masks, High Vegetation Mask	Raster Calculator (Con operations)	Final classification raster (6-5-2)
14	Final classification raster, point cloud (LAS)	Set LAS Class Codes Using Raster	Classified LAS point cloud