Multi-Layered 3D Survey Process to Represent Visible and Invisible Data

Michele Russo [1], Valeria Cera [2], Martina Casciola [1], Maurizio Porcu [3]

[1] History, Representation, and Restoration of Architecture, Sapienza University of Rome - (m.russo; martina.casciola)@uniroma1.it
[2] Dept. of Architecture, University of Naples Federico II, Naples - valeria.cera@unina.it
[3] Codevintec Italiana srl, Milan - maurizio.porcu@codevintec.it

Keywords: TLS survey; GPR; Data integration; Multi-layer process; Architecture stratification.

Abstract

Over the last 25 years, active and passive 3D survey techniques applied in the Cultural Heritage domain have achieved a high degree of reliability. The degree of maturity in the integrated use of these methodologies is well exemplified by their broad application, across the public and private sectors. Nevertheless, surveying techniques able to acquire non-visible data (thermography, sonar, radar, etc.) related to wall packages, floors, foundations, etc.) remain little known in terms of data integration and representation. The latter in particular is crucial for the interpretation of signals at different frequencies that pass through the material, yielding responses that provide information on the type of material or on transitions between different dielectric properties. Often the experience has a basic role in understanding the significance of the maps and radargram, because it lacks the comparison and validation phase with the visible information. It leads to the paradox that to fully understand a complex architecture, analysing both its external and internal consistency, must be integrated data that presents different levels of uncertainty in its interpretation and representation, combining understandable information with others that require complex and non-immediate reading. The research work presented here focuses precisely on this dichotomy. Starting from heterogeneous data obtained from 3D laser scanners, photogrammetry, and ground-penetrating radar, this contribution offers a reflection on the management of this information and its possible integrated representation to make its interpretation more accessible, promoting its integration for a comprehensive understanding of the building.

1. Introduction

In the field of historic Heritage, 3D integrated layered models able to respond to the data stratification of architectural systems remain a challenging aspect. The superimposition of different buildings, restoration stages, transformation processes, and variations in social function defines a level of complexity that a single-instrument survey approach cannot address. The possibility of connecting heterogeneous data, such as historical documentation, geometry, materiality, voids, subsoil, and construction traces that are no longer visible, is essential to define an exhaustive knowledge about the artefacts.

Today, therefore, the epistemological framework of heritage documentation has undergone a profound renewal. The data availability from different sources makes their integration a necessary condition, rather than an optional possibility, for producing knowledge that is reliable, stratified, and scientifically verifiable.

The key issue is no longer merely the acquisition of information, but the correlation, verification, and cross-analysis of data that, by their nature, resolution, and scale of observation, belong to different disciplinary fields.

The integration of range-based and image-based techniques is a consolidated reality, based on well-known processes. On the other hand, the application of instruments as geo-radar in the architecture domain still represents a challenge for the possible misleadings in data interpretation. The possibility to merge the representation of data coming from survey techniques which present different levels of readability define a research area faced by this paper. Through the application of these techniques to a particularly complex case study, the Baptistery in Nocera, the research aims to reflect on the heterogeneity of the data and provide a simplified solution that promotes integration between the techniques and the overall interpretation of all the data entered into a single reference system.

2. State of Art

The growing demand for both quantitative and qualitative knowledge of the built environment requires an integrated approach in which differentiated surveying techniques, operating across distinct physical domains, contribute synergistically to the production of morphological and volumetric information. Within this context, an increasingly widespread approach combines advances in three-dimensional surveying with non-invasive geophysical prospection. Their integration makes it possible to overcome the rigid distinction between what is “above-ground” and what is “buried”, providing a unified and continuous representation of the built heritage and its context (Pieraccini et al., 2023; Pansini et al., 2024).In particular, the integration of Terrestrial Laser Scanning (TLS) and Ground Penetrating Radar (GPR) represents one of the most mature paradigms in this field, as it grounds knowledge in the systematic comparison between high-accuracy metric data and high-resolution stratigraphic–diagnostic data (Mele et al., 2025). The state of the art on TLS-GPR integration for stratified architectures highlights three main trajectories: (i) co-registration and 3D visualisation workflows; (ii) TLS-based physical constraint methods for GPR interpretation; (iii) information formalisation towards HBIM/XR and queryable information models.

Along the first axis, studies document an evolution from a posteriori integrations (overlay of profiles on plans/orthophotos) towards fully 3D integrations, in which the TLS point cloud constitutes the reference space for georeferencing radar traces (either in a local or absolute coordinate system, when both datasets are acquired through integration with GNSS surveying)

and for generating volumetric representations that are readable in relation to the actual geometry of the artefact (Barrile et al., 2017; Adamopoulos et al., 2021). Within this line of research, GPR data are treated either as 2D slices positioned within 3D space (Forte et al., 2021) or as blobs (iso-surfaces) processed using dedicated radar analysis software. Further developing this approach, a second line of investigation has emerged that formalises the role of TLS as a physical constraint for the interpretation of radar prospections: the accurate definition of wall thickness and local geometry makes it possible to reduce uncertainty in estimating wave propagation velocity and to limit artefacts linked to simplifying assumptions (parallel planes, constant thicknesses) (Zhang et al., 2023; Alaia et al., 2026). The limitation of this approach lies in the need for access to both sides of the structure for metric calibration, which makes it applicable only to above-ground structures and unsuitable for subsurface investigations. Finally, the third trajectory concerns the translation of evidence into information models and immersive environments: the integration of above-ground surveys and underground data within HBIM is proposed as a privileged pathway for managing stratifications and uncertainties, by articulating levels of reliability and linking geophysical anomalies to semantic entities (Barrile et al., 2022; Clini et al., 2025). In this case, integration is achieved by anchoring GPR data to the TLS/HBIM model and visualising them as information layers and interactive sections, which merge more semantically than physically. What emerges from this framework is that, although TLS-GPR integration is now a relatively consolidated practice, it still presents constraints that must be critically considered: the 3D representation of data requires an inherently multi-level and multidisciplinary interpretation, based on spatial alignment of information that often demands specific processing procedures, as well as on the reading of evidence that may at times be ambiguous. It is within this research domain that the present study is situated, developed through the case study of the Early Christian Baptistery of Nocera Superiore (Italy). The investigation aims to construct an integrated knowledge framework based on range-based data (primarily TLS point clouds) for defining the geometry of above-ground structures, and on GPR data for characterising internal and subsurface structures. The overall objective is the creation of a unified knowledge system, in which information derived from different instruments is connected on a metric, spatial, and interpretative basis, according to a logic of multi-source integration, to mitigate interpretative uncertainties. Attention is also devoted to the definition of a semi-automatic data-processing workflow for their integration within 3D space, with the aim of better supporting integrated information visualisation through the use of simplified tools.

3. Case study and Methodology

The Early Christian Baptistery of Santa Maria Maggiore, located in Nocera Superiore (SA, Italy), constitutes a paradigmatic example for testing the validity of the research premises. Among the most significant Late Antique complexes in southern Italy, the baptistery presents morphological and stratigraphic conditions that require documentation capable of recognising the 3D relationships between the above-ground and sub-surface components. The architectural layout consists of two concentric spaces: a central area covered by a post-imperial dome, dominated by the baptismal font, which in size is the second largest in Italy after that of San Giovanni in Laterano in Rome; and an annular ambulatory with a barrel vault articulated by fifteen pairs of Roman columns (Stettler, 1940). The complex is also characterised by a subsoil environment connected with underground pathways and hypogeal structures related both to earlier Roman buildings upon which the baptistery was erected in the sixth century, and to the medieval hospital that stands in the garden behind the baptistery. The documentation campaign was organised in four main phases: (i) a 3D morpho-metric survey of the entire complex using TLS, conducted with a FARO Focus 3D CAM2 system; (ii) aerial photogrammetric survey using a DJI mini 4 pro, to capture the roofing systems and their relationship with the surroundings; (iii) detailed terrestrial photogrammetric survey performed with a Sony Alpha7 iv camera for certain architectural features of note in the baptistery; (iv) GPR prospections using several devices (GSSI UtilityScan RLT3, 200 HS MHz, and FlexNX systems) for non-invasive archaeological investigations and heritage assessment in selected areas: the interior of the baptistery, the Terra Santa garden, and the hospital hypogea. Furthermore, the GPR technique was applied to the roof of the baptistery to verify its structural stratigraphy (Figure 1). In this context, the actions carried out within the interior spaces of the baptistery are illustrated, focusing on the integrated application of TLS and GPR techniques, while outlining the survey strategies, data acquisition procedures, and interpretative outcomes contributing to the understanding of the architectural and stratigraphic characteristics of the monument.

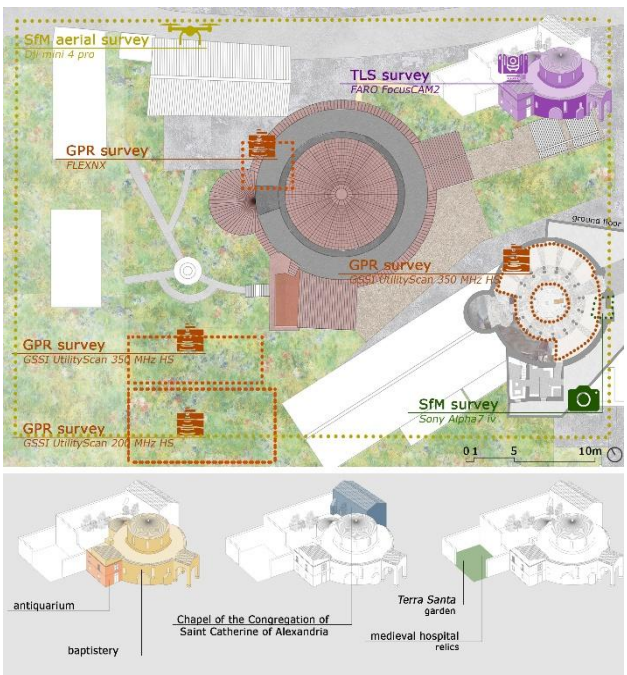

Figure 1. Main areas of the baptistery and surveyed areas map.

3.1 TLS survey: data acquisition and processing

The laser scans were arranged in a network of 99 stations distributed along the interior spaces and the external perimeter, with a variable acquisition sampling step in relation to the working distance (2-10 meters) and the 3D scanner set-up (3-6mm@10mt; with and without color information) (Figure 2).

Figure 2. TLS survey: scheme and point cloud.

The alignment between scans has been carried out using ICP, as the presence of several similar columns with the same architectural background prevented the application of automatic orientation procedures, resulting in a cloud-to-cloud accuracy of ±4 mm.
The resulting point clouds were filtered to remove noise and decimated, preserving areas of greater formal complexity. The resulting model enabled the generation of metric orthoprojections and horizontal sections, useful for subsequent correlation with radar data.

3.2 GPR acquisition campaign and data processing

The GPR surveys at the Nocera Inferiore complex investigated four areas, as previously mentioned: the interior of the Baptistery, the "Terrasanta" garden, the hypogeum, and the roof. To optimize results, three GSSI systems were deployed: the 200 MHz antenna for deep penetration, the UtilityScan RLT3 (specifically the 350 HS for the Baptistery) for best combination of resolution and penetration, and the FlexNX for maximum resolution.
These choices balance the inverse relationship between frequency and penetration: lower frequencies reach deeper layers, while higher frequencies provide the detail necessary to resolve thin layers or small objects (Conyers, 2016). In the Baptistery, the survey was particularly detailed. Profiles were acquired using a radial grid with a spacing of 0.5m at the starting base, ensuring nearly total coverage of the flooring (Figure 3). The UtilityScan 350 HS system was configured with a time window of 138 nsec, an inline trace distance of 2.5cm, and 512 samples per scan.
A critical aspect of this survey was the use of a "benchmark" structure: as will be explained in detail in the following paragraph, a wall visible in the excavated section was observed continuing at a depth of 50cm beneath the current pavement. By analyzing the signal reflection of this known object, the team established a reference for interpreting other anomalies throughout the site.
The GPR data were processed using Geolitix software to transform raw data into interpretable images. The processing sequence included: time calibration for zero-time correction; a band-pass frequency filter to remove noise outside the antenna's 350 MHz central frequency; background removal to eliminate horizontal ringing; and manual gain to amplify deeper signals. Finally, migration and the Hilbert transform were applied to collapse diffraction hyperbolas and calculate the signal envelope. The resulting depth slices were visualized using a JET color scale, which highlights amplitude variations through a high-contrast spectrum from blue (low) to red (high).

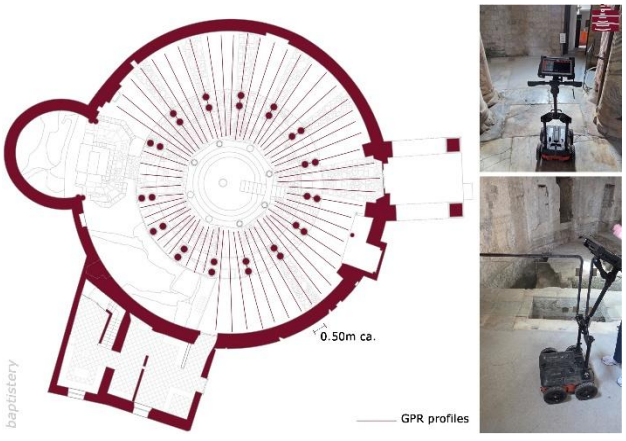

Figure 3. GPR survey: radial grid of acquired profiles.

4. TLS and GPR data integration for the 3D interior of the Baptistery: results and discussion

The data collected during the integrated survey campaign were analyzed to obtain information on the hidden and stratified architecture of the Baptistery, following a multi-layer approach. From a methodological perspective, the process initially involved a more traditional analysis of the GPR data, namely one based on 2D representations, referenced to precise coordinates derived directly from laser-scanning-based spatial model. Subsequently, in order to improve interpretative reasoning and visualizing anomalies in more intelligible way, a workflow was set-up to translate information on the dielectric behavior in 3D volumes, coherently positioned with respect to the TLS point-cloud.

4.1 Preliminary analysis of GPR data

An initial analysis of the subsurface data was based on the cross-correlation between horizontal depth slices and individual vertical radar sections, allowing for a simultaneous assessment of the data through both 2D-only representations and three-dimensional ones, understood as the positioning of 2D slices within a 3D space. In both modes of observation, a fundamental step consisted in registering the slices according to the relative orientation system defined by the TLS survey, made possible by adopting feature elements of the floor in both datasets. More specifically, the transformation of the slices within the 2D plane was carried out with the geometric support of a plan representation at a depth of z = 0.00, generated on the basis of the corresponding orthomosaic derived from the point cloud. The orthomosaic extracted by 3D TLS data constitutes a metric and locally referenced planimetric representation of the investigated area, providing a reliable spatial reference for the orientation and scale processes. Starting from the orthomosaic, a CAD plan was produced, defining a common reference framework for both datasets. The interpretative phase of the anomalies was based, as a key aspect, on the presence of an outcropping wall, visible in the 50 cm deep excavation, as a benchmark. During the data acquisition phase it was observed in the on-site pre-visualisation of the radargrams a discontinuity in the profiles in the area adjacent to the exposed excavation zone and aligned with the entrance to the Baptistery from the museum area. This discontinuity occurred in a portion of soil that was optically aligned with the aforementioned emerging wall and consistent with the corresponding depth level. This approach follows established methodological protocols for ground-truth calibration in complex archaeological layers (Leucci, 2019). By characterizing the electromagnetic signature of this "known object," the research team established a reference for correlating physical geometries with specific reflection patterns (Piro, 2000). It was essential to reduce interpretive ambiguity and reliably, identifying similar buried structures in the unexcavated areas of the Baptistery (Figure 4). In addition, during the survey, specific architectural features of the church were distinguished between true archaeological anomalies and interference features, such as reflections caused by the church's geometry or air waves (surface reflections) that create noise in the data. The survey, which covered the entire surface down to a depth of 3 m, highlighted many geometric anomalies at different depth levels, representing distinct construction phases. These were made clearly readable by positioning the previously oriented 2D slices within the 3D space (Figure 5). Several anomalies seem to be aligned with the current positions of the columns, suggesting their presence in the original buried bases. Key archaeological features were identified in the slices at depths of 1.90 m and 2.80 m, confirming the presence of buried remains beneath the floor.

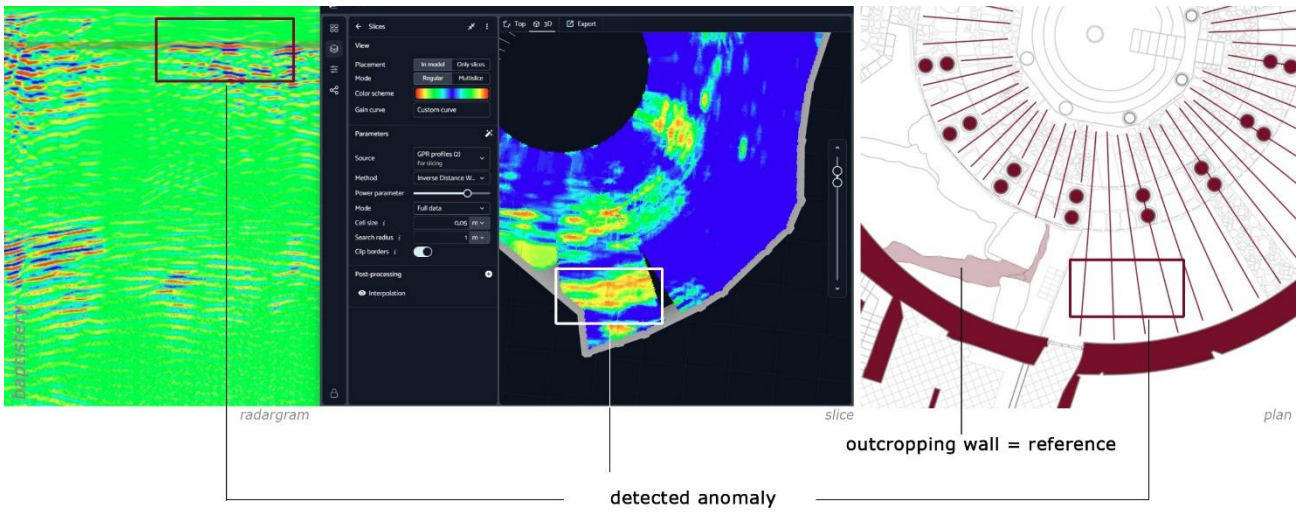

Figure 4. Detected anomaly in the radargram and its relationship with the electromagnetic signature of the "know object".

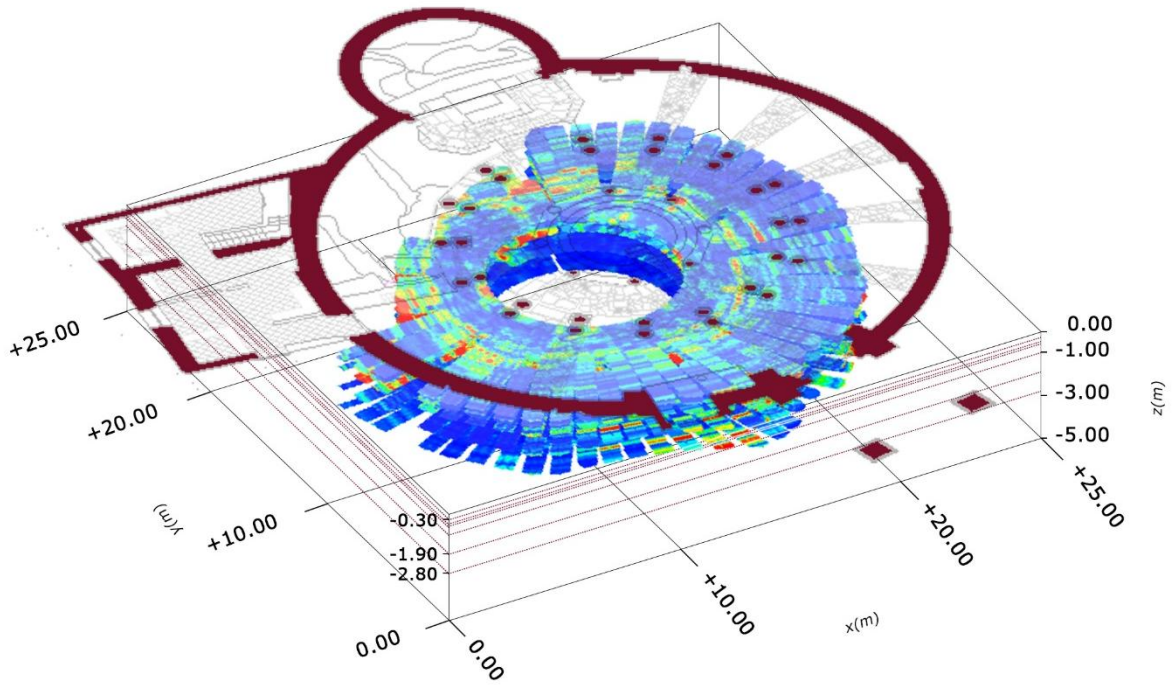

Figure 5. Time-depth slices in axonometric view related to the baptistery planimetric orientation.

4.2 Semi-automatic workflow for 3D visualization of GPR data

The workflow aim is to obtain a volumetric spatial visualisation of the anomalies highlighted in the GPR slices, by transforming the sequence of planimetric maps at different depths into a set of coherent sections along the depth axis and three-dimensional geometries reconstructed through loft procedure in a CAD environment.

Each slice is interpreted as a "section" of subsurface elements and the 3D reconstruction derives from the identification of regions that are spatially compatible - i.e., persistent - across successive levels. This approach is consistent with the use of time/depth slices for interpreting GPR anomalies in archaeological contexts. The process was structured as a sequence of consecutive operations which can be grouped into four main operational phases and sub-actions, starting from the raster images representing the time/depth slices: (i) colour-based image segmentation to isolate red-toned areas and green-toned areas; (ii) retention of areas that remain coherent across slices; (iii) conversion of raster regions into vector outlines; (iv) three-dimensional modelling of the sections at different z-heights. To speed up the workflow, the first three macro-activities were carried out automatically using a Large Language Model (LLM), which, through a dedicated prompt, was tasked with executing the various processes and sub-processes up to the export, in .dxf format, of the different vector regions.

All actions performed by the LLM were iteratively checked, refining both the choice of approaches adopted and the configuration of processing settings. The starting dataset consists of the previously oriented slices, from which the metric scale of 1 px = 0.333 cm was derived, and therefore the pixel-to-metric conversion factor, used both to define matching tolerances and for CAD export.

4.2.1 Colour-based segmentation: the execution of the first task started by the selection of the colour areas to be modelled, according to their persistence across the different slices representing the various stages of depth progression (Figure 6a). With regard to the definition of the colours to be selected, and considering that in a GPR slice red and green areas indicate dielectric behaviours that differ from the surrounding background, namely anomalies, only the areas with maximum saturation were taken into account, in order to eliminate computational artefacts and potential noise.

(a) *Geometric pre-processing: co-registration and common reference space*
The persistence criterion presupposes that the (x, y) coordinates of a colour area in a slice at a given depth level correspond to the same planimetric position plan-view position) in the other slices. Even a variation of only 1–2 pixels between one raster slice and the subsequent one would render the analysis of this vertical continuity inconsistent, introducing a shift that breaks the intersection and causes the apparent loss of persistence. To ensure this verification, all images were brought into the same shared raster extent. Operationally, the LLM was instructed to convert the slices into RGBA format, cropping to the minimum shared height and width. As all slices belong to the same dataset and had previously been registered according to the same reference derived from the TLS data. This process ensured exact pixel correspondence, without the need for further image registration, as no residual translations or rotations were present.

(b) *Chromatic segmentation: extraction of "saturated green" and "saturated red"*
For colour-based segmentation, we operated in the HSV colour space, which separates the chromatic component (Hue) from intensity (Value) and Saturation. This choice is more stable with images characterised by gradients, as HSV reduces sensitivity to luminance variations compared to direct RGB thresholding and allows the concept of "saturated green/red" to be formalised as constraints on Hue and Saturation. Accordingly, for each slice I_i, the computation of a boolean mask $G_i(x, y)$ was required, taking the value 1 when a pixel falls within a specified range of green/red hues and exhibits high saturation, together with a minimum brightness value, to avoid false positives in dark areas. The mask was further cleaned by excluding pixels close to white and black (legends, borders, background). The same operation was performed for areas falling within a red-toned interval.

(c) *Noise reduction: connected-component analysis and area-based filtering*
The mask G_i may include speckle and micro-regions, resulting from noise or processing artefacts, which are not meaningful for matching and subsequent modelling. To transform the mask into a set of interpretable elements (a kind of "blob" representation), connected-component labelling with 8-connectivity was performed, yielding geometric statistics for each component (most notably area). This operation was implemented using connectedComponentsWithStats (OpenCV), a widely used "morphological–statistical" filter in binary denoising tasks, which returns labels and statistics for each connected region. Components with an area smaller than a threshold $A_{min} = 50$ pixels were then removed. This value was empirically defined in relation to the pixel-to-metre conversion scale and was considered conservative with respect to the risk of introducing bias in the investigated archaeological context (i.e. small but real features potentially being discarded), thus ensuring the preservation of structurally significant anomalies while eliminating isolated or irrelevant reflections. In addition, this parameter setting is also motivated by the fact that excessively small regions tend to generate unstable contours for subsequent surface extraction through lofting. The threshold was empirically determined based on the spatial resolution of the GPR data (3 px = 1 cm) and the minimum expected size of archaeological structures of interest, ensuring the retention of structurally meaningful anomalies and the removal of isolated or non-significant reflections.

4.2.2 Persistence computation and preservation of consistent areas between slices: between adjacent slices, the same anomaly may exhibit slight variations in extent and/or position due to changes in the signal with depth, as well as to smoothing introduced by GPR processing. Requiring perfect coincidence would therefore be overly restrictive. For this reason, the analysis of area matching between slices considered a spatial tolerance r, allowing for limited positional and geometric variability while still preserving the criterion of persistence across depth levels.

(a) *Definition of spatial tolerance r and local, non "pixel-perfect" persistence*
The value of the parameter r (in pixels), namely the maximum tolerance of "proximity", was estimated by asking the LLM to measure, on consecutive slices, how much a region needs to be "dilated" in order to achieve a minimum degree of overlap with the corresponding region. In the analysed dataset, this estimation yielded r = 4 px, corresponding to approximately 1.33 cm for green regions, and r = 2 px for red regions. This parameter represents a sensitive hyperparameter, as values that are too small would prevent the identification of groups of persistent sections, whereas values that are too large would lead to the unification of distinct sections. The computed values demonstrate an order of magnitude consistent with the spatial resolution of the dataset (in pixels) and with the expected drift between slices, examined in 2D during Phase 4.1.
A further fundamental aspect for the retention of persistent regions is the definition of the persistence criterion, denoted as K. A real target may be visible only within a limited depth interval; therefore, a global criterion such as "present in all slices" would eliminate plausible targets. In practical GPR slice interpretation, coherence is primarily local along the depth axis, that is, between neighbouring slices. By analysing the dataset, the persistence criterion K was set as local persistence with K = 2 adjacent slices, meaning that a region is considered plausible if it appears in at least two consecutive slices (or, in any case, if it exhibits an extended area of "support" in one of the two adjacent slices). This choice clearly represents a compromise, as it reduces false positives (such as single-slice noise) without eliminating targets that exist only in a limited number of slices. Operationally, adjacency was formalised such that, for each slice i, the LLM was asked to compute a "locally persistent" mask of the form:

$$M_i := G_i \cap \text{dilate}(G_{i-1} \cup G_{i+1}, r) \quad (1)$$

(at the extremes of the sequence, only the available slice is used). This computation accounts for the function dilate(·, r), that is, morphological dilation, understood as a mathematical operation that expands a region according to a structuring element which, in this case, was set equal to the tolerance radius r. In this way, correspondence within a well-defined spatial range is formally established (Figure 6b).

(b) *Tracking of persistent regions (chains)*
Starting from the mask M_i, similarity between regions is defined since the overlap metric "Intersection over Union" (IoU) / Jaccard index, a standard metric for quantifying the overlap

between binary masks. Accordingly, given that, in general terms, the similarity between two regions A and B is measured as:

$$IoU(A, B) = |A \cap B| / |A \cup B| \quad (2)$$

the metric was used to incorporate spatial tolerance r.
Consequently, the comparison was performed by the LLM between A and a dilated version of B, namely:

$$IoU(A, dilate(B, r)) \quad (3)$$

In practical terms, a region in slice i is considered compatible with a region in slice i + 1 if a significant portion of the former falls within a neighbourhood of r pixels around the latter.
Since dilation may generate weakly coherent overlaps in cases of numerous or densely distributed regions, an additional constraint was introduced in the form of a maximum distance between region centroids (15 pixels ≈ 5 cm), expressed in centimetres using the pixel-to-metre conversion scale. This constraint reduces physically implausible matches from an archaeological standpoint. Therefore, between consecutive slices i and i + 1, for each region in i, the LLM was required to select the best candidate in i + 1 by maximising the IoU value, with a minimum threshold set at 30% to avoid accidental partial overlaps due to noise or unreliable geometric continuity, thus pursuing a conservative approach (Figure 6c). As the final action of this sub-phase, correspondences were required to be concatenated into "chains", understood as ideal groups of regions representing the same phenomenon/object across multiple slices and satisfying the minimum condition K = 2 (presence in at least two adjacent slices). This condition is also necessary for creating loft surfaces, which requires at least two sections. Given the parameters adopted, the procedure applied to the dataset produced a total of 63 chains for green regions and 15 chains for red regions.

4.2.3 Conversion of raster regions into vector boundaries: The conversion of persistent regions into vector contours involved a series of operations performed after their direct generation, including geometry cleaning. These operations were necessary to facilitate subsequent 3D modelling, rather than merely performing a direct raster vectorisation.

(a) *Contour extraction and boundary complexity control*
For each previously identified region, the contour was extracted, again using the LLM, through the 'findContours' function in OpenCV, which performs reliably on binary images and returns, for each object, a sequence of (x, y) points along the boundary.
The contour retrieval mode was set to RETR_EXTERNAL, to obtain only the outer boundaries. This choice was motivated by the intention to ignore any internal holes, also considering the scale of the dataset and of the computed regions. In addition, the CHAIN_APPROX_NONE option was selected to retain all extracted points prior to applying a subsequent simplification step (contour simplification / vertex reduction).

(b) *Mild simplification with Douglas–Peucker and topology cleaning*
Pixel-wise contours may consist of thousands of vertices, making CAD modelling unstable or inefficient. It was therefore decided to apply a polygonal approximation using 'approxPolyDP' in OpenCV, which implements the Douglas–Peucker algorithm and ensures that the maximum distance between the original and the approximated contour remains below a tolerance value ϵ. This parameter controls the maximum geometric error in pixels and acts as a generalisation parameter. Within the pipeline, ϵ was kept low (ϵ = 0.8, sub-pixel) to preserve correspondence between slices.

In addition, to avoid failures in modelling through the profile lofting due to self-intersecting polylines, a geometric validation step was applied. If a polygon was found to be invalid, the heuristic procedure 'buffer(0)' in Shapely was used to "clean" self-intersections (e.g. bow-tie geometries / self-crossing shapes). Shapely documentation (Shapely, 2025) explicitly states that 'buffer(0)' may, in some cases, "clean" self-crossing polygons, while warning that the outcome may vary and that very small buffer values may behave differently. The adopted strategy therefore represents a heuristic repair choice, whose potential influence on geometry was verified a posteriori, at the end of processing, by checking that the area did not change significantly and that the contour did not invert (flip orientation / self-reverse).

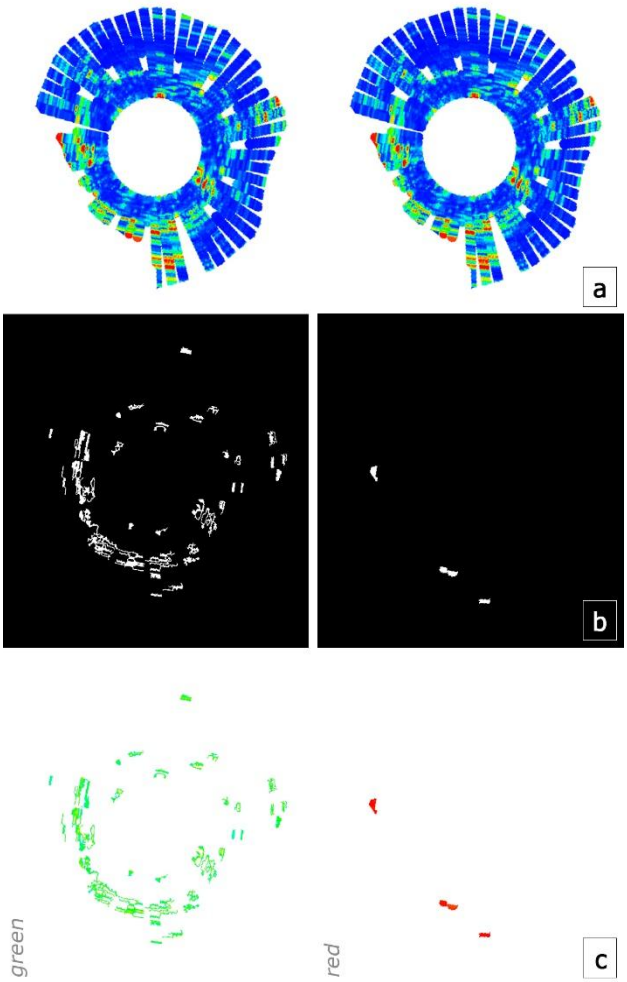

Figure 6. (a) GPR Slice -0.50m; (b) binary mask/spatial tolerance radius; (c) persistent regions, on the for green, on the right for red.

(c) *DXF output*
Once the contours had been extracted and cleaned, the LLM was instructed to export them in DXF format for use in a CAD environment, after performing a pixel-to-centimetre conversion, given the known scale.
Using the scale (1 px = 0.333 cm), the coordinates were converted as follows: $x_{(cm)} = x_{(px)} / 3$; $y_{(cm)} = y_{(px)} / 3$. Each slice was also assigned a z-coordinate: $z_{(cm)} = 100 \cdot depth_{(m)}$. In this way, each contour became a 3D polyline at the corresponding depth level. Finally, within the DXF file, layers were structured by assigning different layers to each chain and its corresponding depth, thus enabling, within the CAD modelling environment, the

isolation of a chain of coherent sections and the execution of the loft operation.

4.2.4 Volumetric modelling of GPR anomaly chains: the final phase of the workflow consisted in the volumetric reconstruction of the anomalies identified within the GPR slices. Modelling was carried out in a CAD environment using the loft function applied to sets of coherent planimetric sections, defined as chains. Each chain represents a sequence of vector contours derived from anomalous areas that persist across adjacent slices and are interpreted as successive manifestations of the same physical phenomenon along the depth axis.
The choice to adopt lofting is motivated by the fact that, in the GPR context, the original data are intrinsically volumetric but are often analysed through stratified planimetric representations, which may be unclear or difficult to interpret for non-experts. The transition from the raster domain to the vector domain and, ultimately, to CAD solids therefore enables a more intuitive spatial interpretation and facilitates integration with other geometric or design-related data.
From an operational perspective, the use of chains as input for the loft operation responds to a fundamental methodological requirement: avoiding the arbitrary pairing of contours belonging to different anomalies. Since the loft command requires geometrically and semantically coherent sections, the prior classification of areas into chains ensures that each generated solid represents a single interpretative entity.
In this way, modelling process is based on a simple geometric interpolation, incorporating information on the continuity and spatial stability of anomalies. The outcome can also support subsequent analyses (such as comparison with archaeological or structural data, or integration into HBIM models).

4.3 3D integration between GPR and TLS data

The concluding phase of the work is devoted to the integration, within a shared reference space, of the 3D solids derived from GPR anomalies and the TLS point cloud, moving towards a multi-layer interpretation of the survey data. This integration requires the alignment of two datasets characterised by different nature, resolution, and acquisition methods. The adopted linking element is the reference CAD plan, drawn at the conventional elevation z = 0.0, derived from the orthomosaic extracted from the TLS point cloud.
This reference CAD plan was used both for inserting and visualising the TLS point cloud and for the prior orientation of the GPR slices. The modelling of GPR anomalies, extracted and transformed into three-dimensional solids through lofting, therefore implicitly inherited this spatial alignment.
The joint insertion of the TLS point cloud and the GPR solids within the same reference system enabled an integrated volumetric visualisation, in which geophysical anomalies were analysed in direct relation to above-ground geometries (emergent structures) and to the built context (Figure 7).
In this way, interpretations of GPR data initially carried out through traditional analysis of 2D slices were validated in a more informed manner (with greater spatial awareness), reducing spatial ambiguity and laying the groundwork for subsequent comparative analyses.
Indeed, the spatial reconstruction of the evidence made it possible to analyse in greater detail the mutual relationships between different stratigraphic depth levels, offering a dynamic reading of the subsurface as an architectural palimpsest.
Anomalies closest to the floor level show a clear correspondence with the current arrangement of the columns, reinforcing the hypothesis that they may refer to original bases or earlier foundation systems, later obliterated by remodelling and repaving interventions.
Conversely, the concentration and continuity of reflectors in the deeper layers attest to the high density of buried structures, almost certainly indicating intense building activity and a long chronological articulation of the phases of use and transformation of the Baptistery.

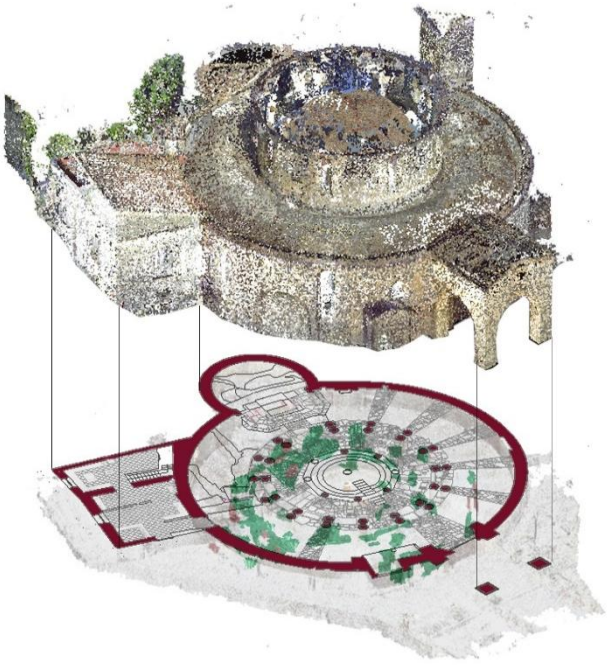

Figure 7. Integrated 3D visualization of TLS point cloud and GPR anomalies.

5. Overview of preliminary results for other areas under investigation

Preliminary analyses were also carried out for the remaining areas investigated with GPR surveys. Initial interpretations yielded significant results, both in terms of discovering new structural elements and validating the integrated approach. Regarding the analyses performed on the roof, the time-slices revealed a series of high-intensity parallel reflections distributed along the entire ring of the ambulatory.
The comparative analysis with the TLS model enabled these reflections to be located at the level immediately below the visible roof, suggesting the presence of a reinforced-concrete secondary vault, which is not documented in historical surveys and is attributable to consolidation interventions carried out during the twentieth century.
This element profoundly alters both the structural and historical interpretation of the ambulatory roof, raising new questions about previous restoration campaigns and confirming the indispensable role of integrated diagnostics in the non-destructive investigation of complex structures.
Concerning the prospection conducted in the hospital hypogea, the GPR data made it possible to identify a sub-rectangular structure, oriented transversely to the axis of the main known hypogeal chamber. The radar reflections display sharp boundaries and arched horizontal surfaces, indicating the presence of a second funerary room, probably contemporary with the original core. In addition, the direct correlation between the radar data, also acquired in the Terra Santa area, and the 3D model made it possible to propose a preliminary plan for the

reconstruction of the volume, suggesting the presence of further hypogeal compartments precisely in the Terra Santa area.

6. Conclusions

The research conducted on the Baptistery of Nocera Superiore confirms that the understanding of complex historic heritage sites requires the integration of multiple data sources. The combined use of TLS and GPR proves methodologically essential in re-establishing spatial and interpretative continuity between above-ground and subsurface components, enabling the detection of hidden features and buried structures that would otherwise remain inaccessible and supporting both historical research and conservation activities.

From a methodological standpoint, the proposed workflow accelerates the spatial visualisation of GPR data by leveraging LLMs, while ensuring control and reproducibility through clearly defined mathematical operations and parameters. Key settings include HSV thresholds for anomaly detection, area and IoU filtering criteria, spatial tolerance r, centroid-distance limits, and geometric precision for contour simplification. The adoption of $K = 2$ adjacent slices balances the reduction of noise with the preservation of meaningful targets and supports 3D modelling via lofting. The workflow is fully implementable in Python using standard libraries and produces both raster and vector outputs.

In terms of results, the vertical and planimetric distribution of coherent anomalies within the baptistery hall suggests the presence of masonry structures and foundations predating the current layout. In other surveyed areas, the identification of a reinforced-concrete counter-vault and a transverse hypogeal space confirms the method's ability to generate new knowledge unattainable through traditional surveys, offering insights into twentieth-century restoration practices and opening new perspectives on the historical development of the baptistery complex.

Acknowledgements

This article is the result of the joint research efforts of the authors. The specific written contributions of the authors are attributed as follows: M.R. wrote the paragraphs 1 and 6; V.C. wrote paragraphs 2, 4, 4.2, 4.3, 5; M.C. wrote paragraphs 3 and 3.1; M.P. wrote paragraphs 3.2 and 4.1 The authors would like to thank Don Fabio Senatore and Raffaele Alfano for their support and availability, and Massimo Capasso, who was involved in the study for his degree thesis in Architecture at Federico II University.

References

Adamopoulos, E., Colombero, C., Comina, C., Rinaudo, F., Volinia, M., Girotto, M., Ardissono, L. 2021. Integrating multiband photogrammetry, scanning, and GPR for built heritage surveys: the façades of Castello del Valentino. *ISPRS Ann. Photogramm. Remote Sens. Spatial Inf. Sci*, VIII-M-1, 1-8. doi.org/10.5194/isprs-annals-VIII-M-1-2021-1-2021.

Alaia, G., Ercoli, M., Brigante, R., Marconi, L., Cavalagli, N., Radicioni, F. 2026. Combining Ground Penetrating Radar and a Terrestrial Laser Scanner to Constrain EM Velocity: A Novel Approach for Masonry Wall Characterization in Cultural Heritage Applications. *Remote Sensing*, 18(1), 15. doi.org/10.3390/rs18010015.

Barrile, V., Bilotta, G., Meduri, G. M., De Carlo, D., Nunnari, A. 2017. Laser Scanner Technology, Ground-Penetrating Radar and Augmented Reality for the Survey and Recovery of Artistic, Archaeological and Cultural Heritage. *ISPRS Ann. Photogramm. Remote Sens. Spatial Inf. Sci*., IV-4/W4:123–127. doi.org/10.5194/isprs-annals-IV-4-W4-123-2017.

Barrile, V., Bernardo, E., Fotia, A., Bilotta, G. 2022. Integration of Laser Scanner, Ground-Penetrating Radar, 3D Models and Mixed Reality for Artistic, Archaeological and Cultural Heritage Dissemination. *Heritage*, 5(3), 1529-1550. doi.org/10.3390/heritage5030080.

Clini, P., Nespeca, R., Angeloni, R., Ferretti, U. 2025. Above-Ground and Underground Architectural Heritage Documentation: A Scan-to-HBIM-to-XR Approach for Historic Centres. *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci*., XLVIII-M-9, 345-351. doi.org/10.5194/isprs-archives-XLVIII-M-9-2025-345-2025.

Conyers, L. B., 2016. Ground-penetrating radar for geoarchaeology. John Wiley & Sons, 45-60.

Forte, E., Mocnik, A., Basso, P., Casagrande, G., Martinucci, D., Pillon, S., Possamai, M., Zambrini, R., 2021. Optimised Extraction of Archaeological Features from Full 3-D GPR Data. *Applied Sciences*, 11(18), 8517. doi.org/10.3390/app11188517.

Leucci, G., 2019. Nondestructive Test for Archaeological Heritage. Springer International Publishing, 88-112.

Mele, M., Degli Esposti, M., Giudici, M., Comunian, A., Al Tamimi, A. M., Al Aufi, A. S., Zerboni, A. 2025. Evaluating Subsurface Risk for Archaeological Heritage Through Ground-Penetrating Radar Surveys: The Case Study of Bisya and Salūt Archaeological Site (Sultanate of Oman). *Heritage*, *8*(10), 399. doi.org/10.3390/heritage8100399.

Pansini, R., Guzel, S., Morelli, G., Barsuglia, F., Penno, G., Catanzariti, G., Merico, A., Carpentiero, M. G., Camporeale, S., Campana, S. 2024. Multi-modal/multi-resolution 3D data acquisition and processing for a new understanding of the historical city of Siena (Italy). *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci*., XLVIII-2/W4-2024, 341-347. doi.org/10.5194/isprs-archives-XLVIII-2-W4-2024-341-2024, 2024.

Pieraccini, M., Miccinesi, L., Conti, A., Fiorini, L., Tucci, G., Pieri, I., Corazzini, S. 2023. Integration of GPR and TLS for investigating the floor of the 'Salone dei Cinquecento' in Palazzo Vecchio, Florence, Italy. *Archaeological Prospection*, 30(19), 27-32. doi.org/10.1002/arp.1788.

Piro, S., 2000. GPR prospecting to stack out the continuity of buried archaeological structures. *Archaeological Prospection*, 7(2), 103-113.

Shapely, 2025. Python package for computational geometry, Version 2.12. shapely.readthedocs.io (19 January 2026).

Stettler, M. 1940. Das Baptisterium zu Nocera Superiore. *Rivista di Archeologia Cristiana*, fasc.1-2, Roma 1940.

Zhang, D., Jia, D., Ren, L., Li, J., Lu, Y., Xu, H. 2023. Multisensor and Multiscale Data Integration Method of TLS and GPR for Three-Dimensional Detailed Virtual Reconstruction. *Sensors*, *23*(24),9826. doi.org/10.3390/s23249826.