Visualization of Cultural Heritage Models with Semantic Region-Level Narratives Using glTF Metadata Extensions

Bura Adem Atasoy [1], Alper Tunga Akın [1], Fatih Terzi [2], Ziya Usta [3]

[1] Karadeniz Technical University, Faculty of Engineering, Department of Geomatics Engineering, 61080 Ortahisar/Trabzon, Türkiye – bugratasoy@ktu.edu.tr; alpertunga@ktu.edu.tr

[2] Karadeniz Technical University, Graduate School of Natural and Applied Sciences, 61080 Ortahisar/Trabzon, Türkiye – fatihterzi@ktu.edu.tr

[3] Department of Geomatics Engineering, Artvin Coruh University, Artvin 08100, Türkiye – ziyausta@artvin.edu.tr

Keywords: Cultural Heritage, glTF, Structural Metadata, Semantic Regions, 3D Web Visualization.

Abstract

This study presents a web-based workflow for linking predefined semantic regions of a 3D cultural heritage model with externally managed historical, material, architectural, and restoration-related records. Saint Michael Church in Akçaabat, Trabzon, was used as the test asset. Its photogrammetric model was converted to glTF and enriched with stable triangle identifiers through EXT_structural_metadata. The model was rendered in a Three.js-based viewer, while semantic and narrative records were maintained in MongoDB. When a user selected a surface, raycasting identified the intersected triangle and sent its identifier to the application programming interface. MongoDB then returned the semantic-region document whose children_ids array contained the selected triangle identifier. The retrieved information was displayed in the interface, and all triangles belonging to the matched region were highlighted on the model. Using one triangle identifier instead of three vertex identifiers reduced the identifier payload and required database references by approximately 66.7%. The prototype provided building-level access to architectural-plan and restoration-history records and region-level access to the large window, north entrance, apse, and nave dome. All tested components were successfully retrieved and highlighted, with displayed vertex counts ranging from 57,642 to 302,955. The results demonstrate the functional feasibility of connecting browser-based 3D interaction with separately managed and updateable cultural heritage narratives.

1. Introduction

Three-dimensional (3D) modelling and visualization techniques are widely used for the digital documentation, interpretation, and presentation of cultural heritage assets. These techniques make heritage objects accessible in digital environments and support the examination of their geometric and visual characteristics (Andújar et al., 2026). However, cultural heritage assets commonly contain heterogeneous materials, irregular geometries, different construction phases, and architectural components with distinct historical meanings. These characteristics make the semantic organization and interpretation of 3D heritage models challenging (Grilli and Remondino, 2019; Mazzacca et al., 2025).

Metadata, paradata, and narrative structures have therefore been incorporated into digital heritage environments to connect geometric representations with historical, material, and cultural information. Metadata and paradata have been used to document the sources and interpretation processes underlying immersive heritage experiences (Ammirati et al., 2025). Semantic metadata and narrative structures have also been combined with 3D visualization and storytelling approaches to improve the contextual interpretation of heritage assets (Barzaghi et al., 2025). These studies indicate that the documentary value of a 3D model depends partly on the ability to associate its geometry with structured descriptive information.

Different information-management approaches have been developed for this purpose. Web-based Geographic Information System architectures have been used to connect 3D heritage models with Linked Open Data and external web resources (Nishanbaev et al., 2021). Heritage Building Information Modeling (HBIM) approaches have enabled geometric and non-geometric information to be organized together for documentation, conservation, and management purposes (López et al., 2018; Quattrini et al., 2017). Semantic-web technologies have also been used to make structured HBIM information queryable through web environments (Quattrini et al., 2017). However, the creation of BIM or HBIM models from existing survey data can require considerable modelling, conversion, and maintenance effort (Volk et al., 2014).

Current approaches address different stages and requirements of the digital-heritage workflow. HBIM is primarily concerned with the structured creation and management of building-related information, whereas browser-based dissemination requires geometric representations that can be transmitted, rendered, and interacted with efficiently. For applications that only require users to select predefined regions and access the related descriptive records, constructing a comprehensive HBIM model may be unnecessary. In such cases, the 3D model can be used solely as the visual interaction layer, while the associated documentary information is stored and managed separately. Similar web-based heritage platforms have used semantically segmented 3D geometry as an interaction and navigation layer while retrieving the corresponding information from external databases (Auer et al., 2014). Within such a decoupled architecture, a lightweight 3D delivery format can represent the geometric layer, while stable identifiers maintain the connection between model regions and externally managed semantic records. Considering modern web solutions on 3D object delivery, the GL Transmission Format (glTF) provides an extensible format for transmitting and loading 3D assets in web-based environments (Khronos Group, 2021).

In the proposed implementation, the glTF model was used as the geometric interaction layer. Semantic regions were represented through vertex-index sets and corresponding triangle indices as EXT_structural_metadata, while the corresponding historical, material, architectural, and restoration-related records were maintained in an external MongoDB database. Attribute-driven delivery approaches have also been examined in other web-based 3D environments, including urban digital-twin visualization (Usta et al., 2026).

The proposed workflow has a limited relationship with HBIM. It does not reproduce the broader information-management capacity of an HBIM environment, which may include monitoring data, structural information, historical documents, and conservation states. Its scope is focused on the browser-based presentation of predefined semantic regions and the retrieval of related narrative records. The external database structure used in the prototype also does not constitute an ontology-driven Linked Data implementation.

Accordingly, this study aims to implement and examine a web-based workflow that links predefined semantic regions of a 3D cultural heritage model with externally managed narrative records. Saint Michael Church in Akçaabat, Trabzon, was selected as the test heritage asset. The contribution of the study consists of a vertex-indexed method for identifying semantic regions, the connection of stable region identifiers with an external database, and the demonstration of building-level and region-level narrative access through a browser-based viewer.

2. Methodology

2.1 Study Asset and Photogrammetric Model

Saint Michael Church, situated in the Akçaabat district of Trabzon, was selected as the test heritage asset. The structure was originally built during the 13th–14th centuries and was expanded during a later construction phase dated to 1846 (Atasoy et al., 2026). These construction phases and the presence of architecturally distinct components made the church suitable for examining semantic-region-level access. The principal construction phases and architectural components used to contextualize the semantic regions are illustrated in Figure 1.

The high-resolution 3D model was produced through a Structure from Motion photogrammetry workflow described by Atasoy et al. (2026). Data acquisition was conducted using an Unmanned Aerial Vehicle flight at an altitude of 56.4 m. A total of 178 aerial images were captured with 80% forward overlap and 70% side overlap. The flight configuration yielded a Ground Sampling Distance of 8.74 mm/pixel and a point density of 2.79 points/cm². The reported reconstruction outputs included 140,625 tie points and a reprojection error of 0.448 pixels. Dense-cloud processing, arbitrary mesh generation, and mosaic texturing were applied to produce the final textured model.

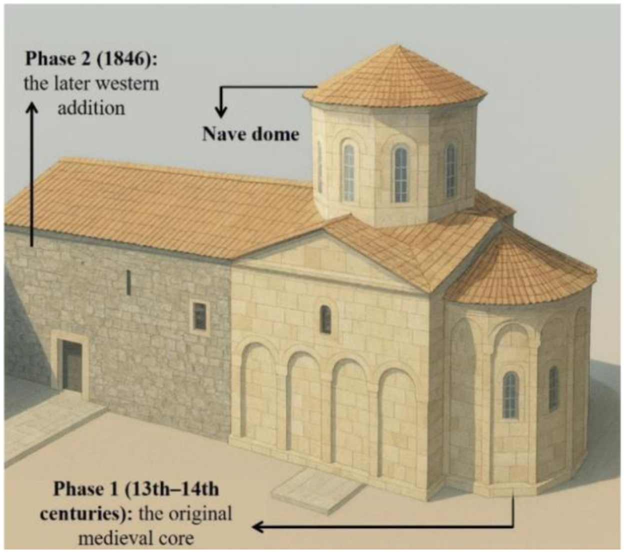

Figure 1. Historical construction phases and principal architectural components of Saint Michael Church (Atasoy et al., 2026).

2.2 Semantic-Region and Database Organization

After the photogrammetric model had been prepared for web visualization, architecturally and historically meaningful components were predefined as semantic regions. The regional organization included openings, entrances, the apse, the nave dome, material zones, and surfaces associated with different construction phases, using a custom object labeller over the glTF model. Besides, the glTF model has been enriched with triangle indices in addition to the existing vertex indices, thanks to the EXT_structural_metadata extensibility.

On MongoDB side, each semantic region was represented by a unique identifier and a corresponding set of triangle indices referring to the geometry loaded by the web viewer. Examples of these identifiers included large_window, north_entrance, apse, and nave_dome. The identifier functioned as the persistent link between the selected geometric region and its descriptive database record. Region names intended for human interpretation were maintained separately from these machine-readable identifiers. Each semantic region, which are defined by the identifiers, has following fields:

- Historical Phases
- Material Composition
- Art-Historical Narrative
- Construction Phases
- Restoration Logs
- Children IDs

The fields were designed to describe complementary aspects of each semantic region. Historical Phases situate the region within the broader chronological development of the building, whereas Construction Phases document the specific stages during which the region was built, modified, or extended. Material Composition identifies the principal materials associated with the region. The Art-Historical Narrative provides a descriptive and interpretative account of its architectural form, ornamentation, spatial role, and historical significance. Restoration Logs record documented conservation, repair, or reconstruction interventions, including their dates and available explanatory information. Finally, Children IDs contain the triangle indices assigned to the semantic region and therefore define its geometric extent within the model. Through this common structure, different architectural components can be associated with their own chronological, material, narrative, restoration, and geometric information. Empty fields were retained where relevant information was unavailable, thereby preserving a consistent database structure across all semantic-region records. An example of the database record used for this association is shown in Figure 2.

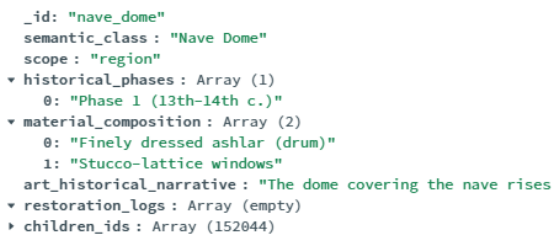

Figure 2. Example of a MongoDB record containing region-level metadata and narrative information.

2.3 System Architecture

The prototype was implemented as a decoupled web-based architecture in which the geometric representation, semantic-region records, and interaction logic were managed as separate but connected components. The principal components of the system were a textured glTF model, a Three.js-based client-side viewer, an application programming interface, and a MongoDB Cloud database. The glTF asset contained the geometry, textures, and triangle-level identifiers required for interaction, whereas MongoDB stored the descriptive and semantic information associated with the architectural regions. The application programming interface connected these components by receiving a triangle identifier from the viewer, locating the corresponding semantic region document in the database, and returning that document to the client. Figure 3 summarizes the system architecture.

When the application was opened, the textured glTF model was loaded into the browser and rendered through the Three.js-based viewer. Each selectable triangle in the model was associated with a stable triangle identifier made available through the structural metadata of the glTF asset using EXT_structural_metadata. The interaction workflow began when the user clicked a visible surface of the model. The viewer performed a ray-based intersection test to determine which triangle had been selected. Instead of extracting and transmitting the identifiers of all three vertices defining the intersected triangle, the viewer retrieved the single triangle identifier associated with that face. This identifier was then submitted through the application programming interface as the database lookup value.

On the server side, MongoDB was queried to find the semantic region document whose children_ids array contained the submitted triangle identifier. Conceptually, this operation corresponded to a document lookup of the form children_ids: “Selected triangle index”. Because each triangle identifier belonged to the children_ids array of its associated semantic region, the database could determine which architectural component contained the selected face without requiring the client to know the region identifier in advance. The semantic classification was therefore resolved dynamically through the database rather than through a complete region-to-geometry mapping stored in the browser.

Following a successful match, the complete MongoDB document associated with the selected region was returned to the client. Each document contained a machine-readable region identifier, a human-readable semantic class, a scope definition, and the descriptive fields defined for the prototype. These fields included Historical Phases, Material Composition, Art-Historical Narrative, Construction Phases, Restoration Logs, and Children IDs.

The returned document served two purposes on the client side. First, its descriptive fields were formatted and displayed in the information panel. Second, its children_ids array was used to reconstruct the complete geometric extent of the selected semantic region. For every triangle identifier listed in the array, the viewer located the corresponding triangle in the loaded glTF geometry. The vertex indices defining these triangles were then retrieved from the geometry index buffer and combined to produce the highlighted region. The selected component was therefore highlighted as a complete architectural unit rather than as an individual clicked triangle.

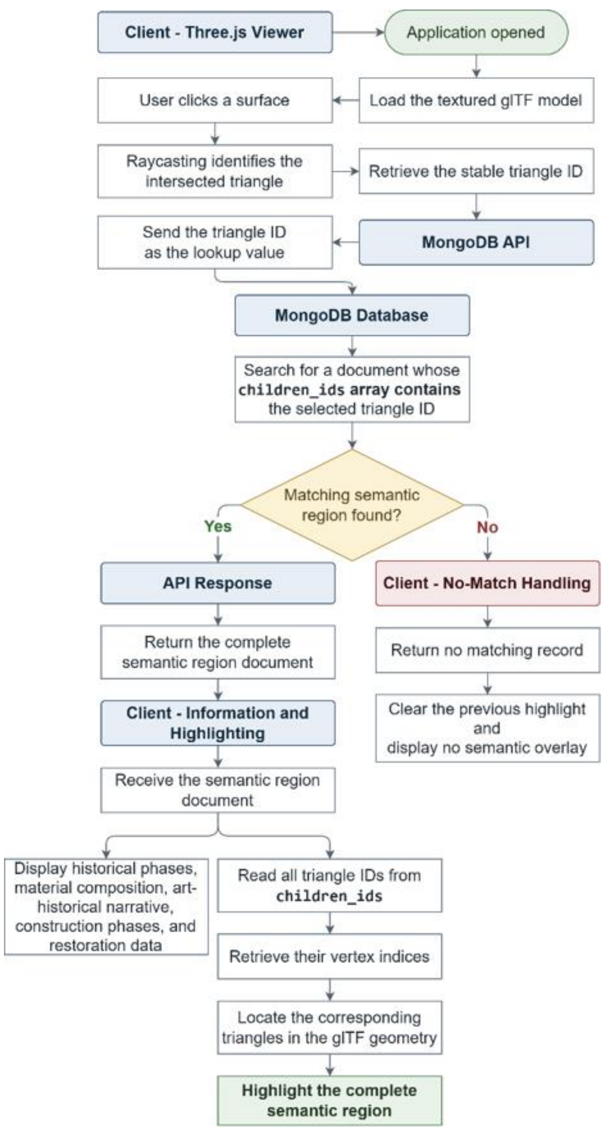

Figure 3. Data flow of the triangle-index-based semantic retrieval and region-highlighting workflow.

This architecture differs from an approach in which semantic-region definitions are fully loaded into the browser during initialization. The client initially requires only the glTF asset and the triangle-level metadata necessary to identify the clicked face. The relationship between that face and a larger semantic region is resolved through the MongoDB query. As a result, semantic records can be modified, extended, or corrected without altering the vertex coordinates, model topology, texture coordinates, or image textures of the glTF asset. Similarly, the visual model can be updated independently, provided that the triangle identifiers used in the structural metadata and the corresponding children_ids values in the database remain synchronized.

The use of triangle identifiers also reduces the amount of geometric identification data required by the interaction workflow. A triangle is normally described geometrically by three vertex indices. A lookup based on such a representation would require the client to transmit a three-value vertex tuple and the database to store or compare the corresponding vertex references. In the proposed implementation, the same face is represented by a single triangle identifier. The number of identifier values transmitted for each selection is therefore reduced from three to one, corresponding to a 66.7% reduction in the identifier payload.

The same principle applies to database storage. When the geometry of a semantic region contains (n) triangles, storing the complete three-vertex tuple of every triangle would require up to (3n) integer references. Storing one identifier for each triangle requires only (n) integer references. Thus, compared with a structure that stores three vertex identifiers for every face, the children_ids representation reduces the number of stored identifier values to approximately one third. It also simplifies the lookup condition because the database searches for a single scalar value inside an array rather than matching a three-element vertex combination.

This reduction should be understood primarily as a reduction in identifier transmission, storage, and matching workload. The identifier-related part of the lookup involves approximately one third of the values required by a three-vertex representation. However, the complete end-to-end response time also depends on network latency, API processing, database indexing, document size, and client-side rendering. Therefore, the architecture provides a clear theoretical and structural reduction in query input size, although an exact threefold reduction in total response time would require a separate performance benchmark.

3. Results

3.1 Building-Level Thematic Retrieval

The implemented viewer provided two building-level thematic access modes: Restoration History and Architectural Plan. These modes were activated through dedicated controls in the information panel and did not require the selection of a geometric region. In both cases, the textured glTF model remained visible in the scene while the corresponding building-level document was retrieved and displayed.

The Restoration History mode returned a general narrative describing the historical development and documented interventions associated with Saint Michael Church. The displayed narrative referred to the original construction of the church during the 13th–14th centuries, its enlargement and restoration in 1846, and the restoration programme conducted between 2018 and 2021. The narrative was followed by individual restoration-log entries presented as structured records.

The first displayed record, identified as rest_1846, documented the enlargement and restoration dated 20 May 1846. The interface presented the intervention type, responsible actor, documentary source, and a short description. The second record, identified as rest_2018_2021, represented the restoration programme initiated in 2018 and completed in 2021. Its displayed fields included the start and end dates, intervention type, and available descriptive information. The retrieved narrative and restoration logs are shown in Figure 4.

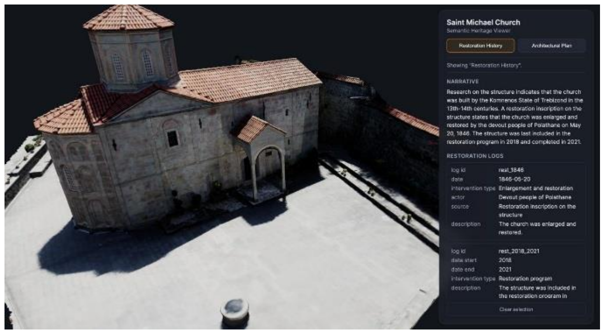

Figure 4. Building-level retrieval and presentation of the restoration-history narrative and structured intervention records.

The Architectural Plan mode returned information on the construction phases, plan configuration, masonry, and principal architectural elements of the church. Two construction phases were represented in the retrieved record.

Phase 1 corresponded to the medieval structure dating to the 13th–14th centuries. The record identified the Komnenos State of Trebizond as the builder and described a square-plan church with a single nave and apse, surmounted by a dome resting on a high drum. Finely dressed ashlar masonry was associated with the original structure. The record also listed principal ashlar elements, including the bearing walls, apse, dome drum, north entrance, and triangular pediments.

Phase 2 represented the western addition dated to 1846. This record described the westward extension of the church and identified river-smoothed stone, rubble, and roughly hewn ashlar among the materials used in the later section. The retrieved construction-phase information is shown in Figure 5.

The two building-level modes therefore returned different document structures through the same interface. Restoration History combined a general narrative with multiple intervention logs, whereas Architectural Plan combined architectural information with structured construction-phase records.

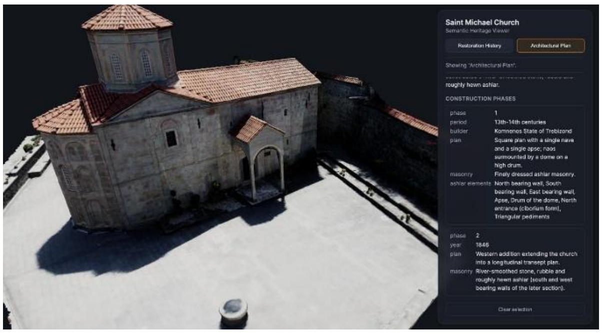

Figure 5. Building-level retrieval and presentation of the architectural-plan record and associated construction phases.

3.2 Semantic-Region Retrieval and Highlighting

Region-level interaction was examined using four predefined architectural components: the large window on the north façade, the north entrance, the apse, and the nave dome. The components differed in their position, geometric form, and spatial extent.

For each case, the user selected a visible surface of the glTF model. The Three.js raycasting operation identified the intersected triangle, and the stable identifier assigned to that triangle was submitted to the application programming interface. The database was then queried for the document whose children_ids array contained the selected triangle identifier.

Following a successful match, the corresponding semantic-region document was returned to the viewer. Its descriptive fields were displayed in the information panel. The complete children_ids array was subsequently used to locate all triangles belonging to the matched region. The vertex indices defining those triangles were retrieved from the glTF geometry and used to generate the orange highlight.

The interface also displayed the three vertex indices defining the initially clicked face as diagnostic information. These displayed values documented the selected geometric face, but the database lookup itself was performed using the single stable triangle identifier.

Selection of a triangle within the large window resulted in retrieval of the document identified as large_window. The interface displayed the semantic class Large Window (North Façade) and classified the object as a region.

The component was assigned to Phase 1, corresponding to the 13th–14th centuries. The displayed materials included polychrome ashlar in green, brown, and yellow tones, tooth-mould ornamentation, and a rope-mould arch. The narrative described the vertically oriented opening, shallow round arch, surrounding mouldings, ornamental inner frame, and coloured ashlar arrangement.

The returned children_ids array was used to highlight the complete window region rather than only the selected triangle. The viewer reported 57,642 highlighted vertices. The resulting geometric highlight and descriptive record are shown in Figure 6.

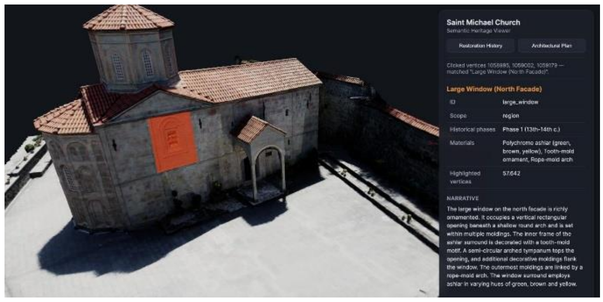

Figure 6. Retrieval and highlighting of the large window on the north façade, with the associated historical, material, and narrative information displayed in the viewer.

The north entrance was retrieved by selecting a triangle belonging to the projecting entrance structure. The returned document was identified as north_entrance, and the semantic class was displayed as North Entrance.

The region was assigned to Phase 1 and associated with finely dressed ashlar. The narrative described the entrance opening, recessed lintel, relieving arch, projecting canopy, supporting piers, engaged pilasters, pitched roof, and triangular pediment.

The complete entrance ensemble was highlighted in orange. The highlighted geometry included the projecting canopy, its supporting elements, roof surfaces, and the other triangles assigned to the same semantic region. The viewer reported 245,740 highlighted vertices. The result is shown in Figure 7.

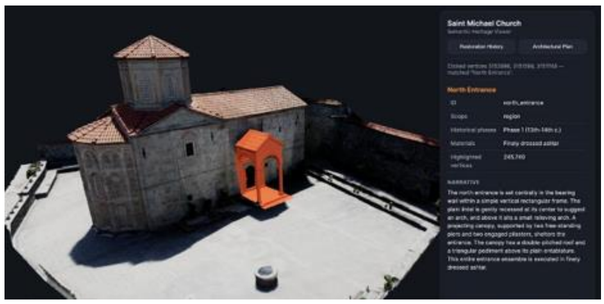

Figure 7. Retrieval and highlighting of the north entrance, including the projecting canopy and supporting architectural elements.

Selection of a triangle on the exterior apse surface returned the document identified as apse. The region was assigned to Phase 1 and associated with finely dressed ashlar, plaster bands, stucco lattice window screens, and convex rope mouldings.

The displayed narrative described the semicircular interior and heptagonal exterior form of the apse. It also documented the division of the exterior into seven vertical facets and three horizontal registers, together with the arrangement of niches, arches, windows, and decorative mouldings.

The orange overlay covered the complete exterior apse region, including its polygonal wall surfaces and roof. The viewer reported 302,955 highlighted vertices, the largest reported value among the four demonstrated semantic regions. The result is shown in Figure 8.

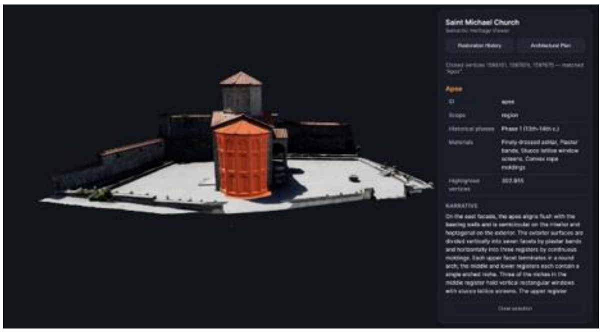

Figure 8. Retrieval and highlighting of the apse, with its complete polygonal exterior and roof represented as a single semantic region.

Selection of a triangle on the upper structure returned the document identified as nave_dome. The semantic class was displayed as Nave Dome, and the region was assigned to Phase 1.

The material information included finely dressed ashlar associated with the drum and stucco-lattice windows. The narrative described the octagonal drum, conical upper form, recessed rectangular panels, window openings, horizontal mouldings, and the transition between the drum and the upper roof structure.

The returned children_ids array defined the complete dome region above the nave. The viewer resolved the corresponding triangles against the glTF geometry and highlighted the component in orange. A total of 152,044 highlighted vertices was reported. The result is shown in Figure 9. The principal outputs of the four region-level cases are summarized in Table 1.

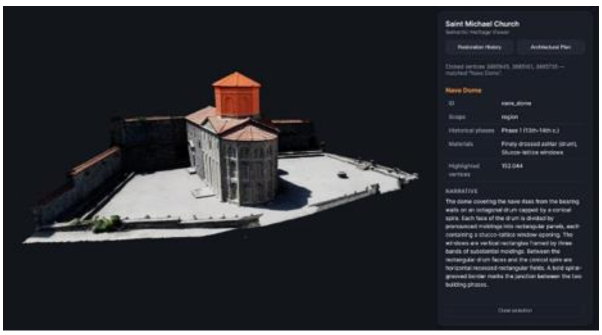

Figure 9. Retrieval and highlighting of the nave dome, including the drum and associated upper architectural surfaces.

The highlighted-vertex values refer to the vertices used by the viewer to render the complete orange overlays. They do not represent the number of triangle identifiers stored in the corresponding children_ids arrays. The database stored one identifier for each triangle, whereas the vertex indices required for rendering were derived from the glTF geometry after retrieval of the matched document.

Functional verification covered six principal interface cases: two building-level thematic retrievals and four region-level selections. The building-level cases examined whether the corresponding thematic control returned and displayed the intended document. The region-level cases examined the complete workflow from surface selection to document retrieval, information-panel presentation, and complete-region highlighting.

Semantic region	**Historical phase**	**Material information displayed**	**Highlighted vertices**
Large Window, North Façade	Phase 1, 13th–14th centuries	Polychrome ashlar; tooth-mould ornament; rope-mould arch	57,642
North Entrance	Phase 1, 13th–14th centuries	Finely dressed ashlar	245,740
Apse	Phase 1, 13th–14th centuries	Finely dressed ashlar; plaster bands; stucco lattice window screens; convex rope mouldings	302,955
Nave Dome	Phase 1, 13th–14th centuries	Finely dressed ashlar; stucco-lattice windows	152,044

Table 1. Semantic-region records retrieved and highlighted in the web-based viewer

3.3 Functional Verification Results

Both building-level controls returned the intended records. Restoration History displayed the general historical narrative and two structured restoration logs. Architectural Plan displayed the construction characteristics of the church together with two construction-phase records.

All four region-level selections returned the intended semantic-region documents. In each case, the triangle selected through raycasting was associated with the correct MongoDB document through membership in its children_ids array. The returned identifier, semantic class, scope, historical phase, material information, narrative, and highlighted-vertex count were subsequently displayed in the information panel.

Complete-region highlighting was also produced for all four selected components. Although the large window, north entrance, apse, and nave dome had different geometric forms and extents, the same retrieval and highlighting procedure was applied in every case. The initial selection identified one triangle, while the document returned from MongoDB defined the complete semantic region.

The verification focused on the functional correctness of the implemented workflow. Model-loading time, API response time, MongoDB query duration, network latency, frame rate, memory consumption, and cross-browser performance were not measured. The reported results therefore demonstrate successful retrieval, presentation, and highlighting rather than quantitative system performance.

4. Discussion

The results show that building-level thematic records and localized architectural narratives can be accessed through the same web-based environment. Building-level access provides a general interpretation of construction and restoration history, while region-level access connects individual architectural components with their historical, material, and descriptive information. These two access levels complement each other by allowing the heritage asset to be examined at different semantic scales.

The separation of the geometric asset from the descriptive database allows narrative records to be managed without embedding their complete contents in the 3D geometry. The model-related structure retains the identifiers required for interaction, while longer and potentially changeable descriptive records are maintained externally. This arrangement supports the retrieval and updating of narrative records without requiring the underlying photogrammetric geometry to be reconstructed.

The implementation shares the broader objective of combining 3D visualization with metadata and narrative structures, as examined in previous cultural heritage studies (Ammirati et al., 2025; Andújar et al., 2026; Barzaghi et al., 2025). Its information-management scope remains narrower than that of HBIM environments, which may integrate a wider range of geometric, structural, monitoring, and conservation records (López et al., 2018; Quattrini et al., 2017). The current prototype is focused on browser-based visualization and access to selected narrative records.

The use of vertex indices as regional references introduces a model-version dependency. Vertex order may change when a mesh is simplified, re-triangulated, re-indexed, or exported again using different processing parameters. Such a change could invalidate an existing region definition even when the visible geometry remains similar. The glTF asset and its associated semantic-region definitions must therefore be versioned and validated together. For a production-scale implementation, stable feature identifiers represented through a standardized feature-metadata mechanism would provide a more interoperable alternative to direct vertex-index references.
A second limitation concerns the regional definition process. In the current prototype, semantic regions were predefined for a selected group of architecturally meaningful components. The study did not evaluate automatic segmentation, overlapping semantic regions, hierarchical component relationships, or regions whose boundaries change according to historical interpretation. Future work should examine whether the same retrieval structure can support larger region sets, multiple levels of semantic granularity, and alternative interpretations of the same geometry.

The presented results establish functional feasibility for a single cultural heritage asset and four demonstrated semantic regions. Performance, scalability, interoperability across different software environments, and user-oriented accessibility were not quantitatively evaluated. The current MongoDB and API structure also does not provide a fully semantic interoperability or formal ontology integration. Future evaluations may include additional heritage assets, larger region sets, cross-viewer compatibility tests, response-time measurements, and user-based assessments.

5. Conclusion

A web-based workflow was implemented to connect predefined semantic regions of a 3D cultural heritage model with narrative records stored in an external database. The workflow was demonstrated using a photogrammetric model of Saint Michael Church. Building-level access was provided for architectural and restoration information, while region-level access was demonstrated for the large window on the north façade, north entrance, apse, and nave dome.

When a region was selected, the corresponding geometry was highlighted and its identifier, historical phase, material information, vertex count, and narrative description were displayed. The demonstrated regions contained between 57,642 and 302,955 highlighted vertices. These results confirm the functional operation of region matching, visual highlighting, and narrative retrieval for the tested components.

The implementation further showed that vertex-indexed region definitions can provide a lightweight connection between a detailed photogrammetric mesh and externally maintained cultural-heritage records. However, the dependence of these definitions on the vertex ordering of a particular model version limits their direct portability. Future development will therefore focus on stable feature identifiers, standardized glTF feature-metadata representations, larger semantic-region inventories, and quantitative evaluation of interaction performance and usability.

The study provides a focused proof of concept. Its findings are limited to one heritage asset and selected semantic regions, and they should not be interpreted as evidence of large-scale performance, cross-platform interoperability, or user-oriented effectiveness. Nevertheless, the demonstrated workflow provides a basis for connecting browser-based inspection of complex cultural heritage geometry with separately managed and updateable narrative records.

Acknowledgements

The authors would like to thank the Scientific and Technological Research Council of Türkiye (TÜBİTAK) for the financial support provided under the BİDEB 2224-A Travel Fellowship for International Scientific Events Program (Grant number: 1919B022606085).

References

Ammirati, L., Watrous, B., Ftaita, A., Rigauts, T., de Gruchy, M., 2025. Dive into heritage: Paradata and metadata in an immersive digital heritage experience. In: Ioannides, M., Baker, D., Agapiou, A., Siegkas, P. (Eds.), *3D Research Challenges in Cultural Heritage V: Paradata, Metadata and Data in Digitisation. Lecture Notes in Computer Science*, 15190, Springer, Cham, 36-51. doi.org/10.1007/978-3-031-78590-0_4.

Andújar, C., Cayuela, B., Lorés Otzet, I., Bosch, C., Guàrdia, M., Muñoz-Pandiella, I., 2026. Revealing the layers of history: Enriching medieval cultural heritage 4D models with interactive narratives. *ACM Journal on Computing and Cultural Heritage*, 19(1), Article 5, 1-22. doi.org/10.1145/3715157.

Atasoy, B. A., Akın, A. T., Terzi, F., Yuka, A., 2026. A digital layer of memory: 3D documentation, art historical interpretation, and web-based presentation of the Saint Michael Church. *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci.*, XLVIII-4/W18-2025, 57-61. doi.org/10.5194/isprs-archives-XLVIII-4-W18-2025-57-2026.

Auer, M., Agugiaro, G., Billen, N., Loos, L., Zipf, A., 2014. Web-based Visualization and Query of semantically segmented multiresolution 3D Models in the Field of Cultural Heritage. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, II-5, 33–39. doi.org/10.5194/isprsannals-ii-5-33-2014

Barzaghi, S., Colitti, S., Moretti, A., Renda, G., 2025. From metadata to storytelling: A framework for 3D cultural heritage visualization on RDF data. arXiv preprint, arXiv:2505.14328. doi.org/10.48550/arXiv.2505.14328.

Grilli, E., Remondino, F., 2019. Classification of 3D digital heritage. *Remote Sensing*, 11(7), 847. doi.org/10.3390/rs11070847.

Khronos Group, 2021. glTF 2.0 Specification. https://registry.khronos.org/glTF/specs/2.0/glTF-2.0.html (13 July 2026).

López, F. J., Lerones, P. M., Llamas, J., Gómez-García-Bermejo, J., Zalama, E., 2018. A review of Heritage Building Information Modeling (H-BIM). *Multimodal Technologies and Interaction*, 2(2), 21. doi.org/10.3390/mti2020021.

Mazzacca, G., Sterpin, A., Rigon, S., Medici, M., Remondino, F., 2025. A Roman carved tale modelled in 3D and interpreted with AI. *ISPRS Ann. Photogramm. Remote Sens. Spatial Inf. Sci.*, X-M-2-2025, 253-260. doi.org/10.5194/isprs-annals-X-M-2-2025-253-2025.

Nishanbaev, I., Champion, E., McMeekin, D. A., 2021. A Web GIS-based integration of 3D digital models with Linked Open Data for cultural heritage exploration. *ISPRS International Journal of Geo-Information*, 10(10), 684. doi.org/10.3390/ijgi10100684.

Quattrini, R., Pierdicca, R., Morbidoni, C., 2017. Knowledge-based data enrichment for HBIM: Exploring high-quality models using the semantic-web. *Journal of Cultural Heritage*, 28, 129-139. doi.org/10.1016/j.culher.2017.05.004.

Usta, Z., Akın, A. T., Arroyo Ohori, K., Stoter, J., 2026. Visualization of urban digital twins on the web with attribute-driven adaptive tiling. *Environmental Modelling & Software*, 197, 106863. doi.org/10.1016/j.envsoft.2026.106863.

Volk, R., Stengel, J., Schultmann, F., 2014. Building Information Modeling (BIM) for existing buildings: Literature review and future needs. *Automation in Construction*, 38, 109-127. doi.org/10.1016/j.autcon.2013.10.023.