

From Fragmented Repertoires to a Property Graph: A Queryable Atlas of Byzantine Geometric Mosaic Bands for Architectural Heritage

Giulia Flenghi¹, Violette Abergel^{2,3}, Michele Russo¹

¹ History, Representation and Restauration Dept., Sapienza University, Rome, Italy,
(giulia.flenghi; m.russo)@uniroma1.it

² CNRS, MAP UPR 2002, Marseille, France

³ CNRS, INSA Lyon, Universite Claude Bernard Lyon 1, LIRIS, UMR5205, Villeurbanne, France
violette.abergel@map.cnrs.fr

Keywords: Byzantine mosaic bands; Photogrammetric survey; Property graph; Typological normalization.

Abstract

Research often focuses on architecture as a whole, seeking to explore its various scales as thoroughly as possible. However, some aspects, such as decorative friezes, play a decisive role not only in terms of decoration but also in understanding architecture as a whole. Besides, understanding an entire system of ornamental decorations is linked to resolving specific complexities, which make it difficult to access in its entirety. These include the study of types, their recurrence and comparison, the possibility of inferring connections from common elements, and the accessibility of all these aspects to experts in the field. In fact, comparing historical-artistic ornamental repertoires across sites and periods is often hindered by fragmented documentation and inconsistent terminology, which make systematic verification of recurrences and variants difficult. This paper proposes a methodological framework for constructing an interrogable Atlas of architectural ornament based on a property-graph model, designed to reconnect visual evidence, typological classification, contextual data, and bibliographic sources within a single relational structure. The workflow integrates image-based acquisition and orthometric products, vector reference drawings for typological definition, a controlled thesaurus for terminological normalization, and a repeatable ETL pipeline for data ingestion into a graph database. Within this framework, decorative instances, classes, and morphological families are explicitly linked to architectural context, chronology, images, and sources, enabling evidence-first navigation and controlled comparative analysis. To assess the approach, the method is applied to a case study on geometric mosaic bands, comprising 34 parietal and pavement classes on planar surfaces (5th–12th centuries) in monuments from Ravenna and Rome. The resulting Atlas supports cross-site and cross-period comparison by making typological proximity, variants, and attestations directly inspectable together with their supporting documentation. Rather than a generic digital archive, the contribution defines an operational structure for comparative research on architectural ornament, offering a scalable and verifiable framework for cumulative knowledge building and future extensions.

1. Introduction

Research often focuses on architecture as a whole, seeking to explore its various scales as thoroughly as possible. However, some aspects, such as decorative friezes, play a decisive role not only in terms of decoration but also in understanding architecture as a whole. Besides, understanding an entire system of ornamental decorations is linked to resolving specific complexities, which make it difficult to access in its entirety. These include the study of types, their recurrence and comparison, the possibility of inferring connections from common elements, and the accessibility of all these aspects to experts in the field. In fact, the comparison of historical-artistic repertoires of architectural ornament on a territorial and diachronic scale is often hindered by two main factors: the fragmentation of documentation (images, drawings, heterogeneous records) and the terminological discontinuity with which formally related motifs are described across different contexts and publications. The result is a rich body of data that is difficult to interrogate systematically, in which recurrences, co-occurrences, and variants remain scarcely visible and hard to verify through a transparent link between descriptions, images, and sources.

This paper proposes an interrogable atlas based on a property graph, conceived to recombine fragmented repertoires into a coherent relational structure. The atlas explicitly connects instances, typologies, and morphological families with

architectural contexts, chronologies, images, and bibliographic references, making the evidential chain that supports each relationship explicit.

The objectives are twofold: (i) to make formally related decorations comparable through normalized criteria and metadata; and (ii) to provide a graph-driven environment for exploration and consultation, enabling the verification of recurrences and variants through visual navigation and direct access to images and sources. The workflow integrates image-based acquisition and orthometric products, vector reference drawings for classification, an ETL pipeline for controlled data ingestion, and user-oriented navigation tools. To assess the approach, it is applied to a case study focusing on mosaic decorative bands: 34 parietal and pavement classes on planar surfaces (5th–12th centuries) in monuments in Ravenna and Rome.

2. State of the Art

The construction of digital repertoires for historical ornament currently develops along two complementary directions: (i) the production of visual and metric documentation reliable enough to support comparison, and (ii) the semantic structuring of content, required to make relationships, recurrences, and variants queryable beyond the single case. In the specific domain of mosaics, this dual need is particularly evident: the object is both strongly image-dependent and inherently comparative, yet documentation, labels, and references often

remain distributed across heterogeneous outputs and publication contexts. Recent research has therefore focused on robust acquisition and representation strategies that can sustain cross-case inspection and reuse.

On the documentation side, image-based and photogrammetric approaches are increasingly adopted to record decorated surfaces and to generate analytical outputs such as orthophotos and 3D models. Applications show how such workflows can deliver usable, metrically grounded documentation even under operational constraints, enabling coherent visual data for analysis and comparison (Adami et al., 2018; Doria & Picchio, 2020). Related studies also emphasize how photogrammetry supports conservation-oriented readings by mapping materials, decorative patterns, and past restorations (Fioretti et al., 2020). These efforts connect with broader initiatives aimed at building reusable digital collections for preservation-oriented access and consultation (Manferdini, 2010; Manferdini & Cipriani, 2011). In parallel, more “systemic” proposals have highlighted the potential of integrated environments for navigation and comparative access—such as the virtual mosaic museum introduced by Maino and Visparelli (2003)—anticipating issues that remain central today: user-oriented exploration, structured access, and the linkage between documentation and interpretation.

The second direction concerns semantic modelling and the formalization of cultural-heritage knowledge. Conceptual models and ontologies provide a coherent framework for describing cultural objects and intangible components, and for making visual and iconographic representations computationally addressable and queryable (Carboni & De Luca, 2016; Carboni & De Luca, 2019). Within this perspective, reference frameworks such as the CIDOC CRM provide a widely adopted conceptual backbone for cultural-heritage interoperability (Doerr, 2003), and its formalization as an international standard further supports stable, exchangeable knowledge structures (ISO, 2014). In operational terms, semantic platforms have also been proposed to connect architectural evidence, analysis outputs, and interpretive statements within unified environments designed for querying and reuse (De Luca et al., 2011). Complementarily, research on ontology and knowledge-structure visualization provides methodological grounding for making such semantic relations explorable through interactive interfaces (Katifori et al., 2007). At the infrastructure level, recent work stresses the need for integrated ecosystems that couple data production, knowledge modelling, and long-term access within coherent paradigms for heritage science (Thomas et al., 2025). Within ornament studies, controlled vocabularies and classification schemes remain a key prerequisite for comparability. Foundational graphic repertoires—such as the *Répertoire graphique du décor géométrique dans la mosaïque antique* (Blanchard-Lemée et al., 1973) and the later *Le décor géométrique de la mosaïque romaine* (Balmelle et al., 1985)—exemplify the role of stabilized nomenclature and typological criteria in supporting identification and comparison across heterogeneous corpora. More recent experiments in atlas-making have shown how morphology and proportion-based criteria can be operationalized to enable navigable typological readings of ornamental configurations (Arena, 2020). In a related direction, collaborative image-based annotation platforms provide practical means to link visual evidence to structured descriptors and controlled vocabularies, supporting cumulative enrichment and reuse at scale (Roussel & De Luca, 2018).

Despite these advances, a recurring operational difficulty remains: integrating heterogeneous documentation (images, metric products, drawings), typological classification, and contextual links (site, chronology, bibliography) within a single

interrogable infrastructure, while keeping the evidential chain between statements and sources explicit and verifiable. The present contribution addresses this gap by proposing a graph-based model in which such relations become navigable and checkable, enabling transversal exploration and comparative control through explicit links between descriptions, visual evidence, and references.

3. Case Study

To assess the proposed approach on an ornament repertoire that is historically stratified and documentarily fragmented, the methodology is applied to geometric mosaic decorative bands—on walls and floors—understood as linear units that articulate and measure architectural space. These bands structure the perception of surfaces: they regulate rhythm and mark edges and thresholds in relation to figurative fields and architectural articulation. Despite their pervasiveness in Late Antique and medieval decorative programs, documentation is often episodic and described through non-aligned nomenclatures, making systematic comparison and the reconstruction of recurrences, associations, and variants difficult. On this basis, the case study brings together 34 classes of decorations on planar surfaces, dating from the 5th to the 12th century, selected from monumental contexts in Ravenna and Rome (Figure 1).

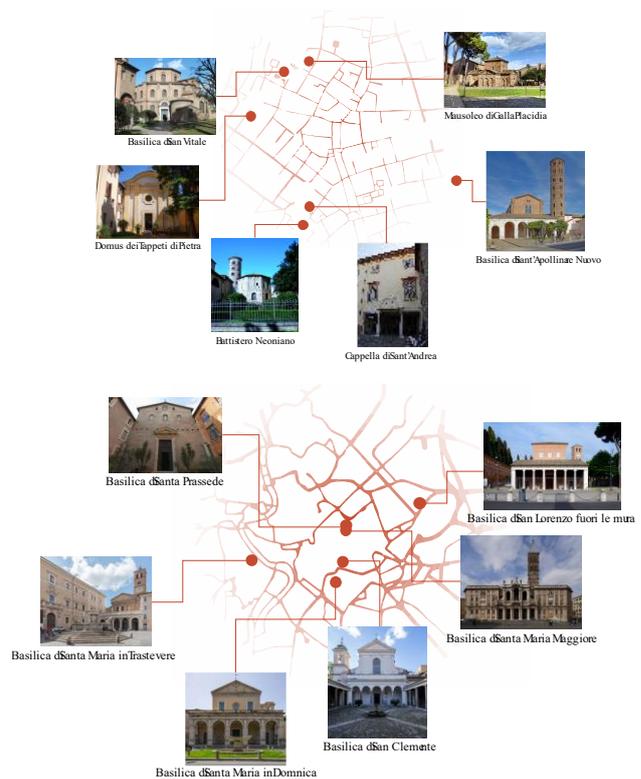


Figure 1. Location of the surveyed monumental sites in Ravenna and Rome included in the study.

The choice of the two centers responds to criteria of representativeness and comparability: Ravenna offers a sequence of exceptional quality and continuity for Late Antique and sixth-century production; Rome allows the observation of processes of reception and re-elaboration of Eastern models within a Western context characterized by strong continuity in patronage and decorative programs. The corpus is organized into two groups according to placement—wall and floor—in order to distinguish different visual functions and documentation

conditions, and to make placement a queryable and navigable attribute in the subsequent phases of graph-based representation. For each decorative instance, architectural context, chronology, bibliographic references, and links to iconographic documentation are recorded, ensuring traceability and cross-site comparison. The broad chronological span (5th–12th centuries) was adopted to obtain a sufficiently diverse corpus in terms of formal solutions and contexts of use. To ensure comparability and reduce geometric ambiguity, the investigation excludes decorations on curved surfaces (domes, apses, and other non-planar intradoses), since perspective deformations and local scale variations can alter the perception of modules and introduce uncertainty in normalization; such cases are deferred to future developments, which will require dedicated correction and sampling procedures. The corpus includes, among others, jeweled frames, meanders/Greek keys, braids, sinusoids, peltae, interlaced circles, and related modular combinations (Figure 2). For each class, whenever possible, placement and position within the architectural system, chronology, relationships with the iconographic context, and notes on borderline cases (variants, hybrids, spatial adaptations) are recorded, so as to make explicit the conditions of comparability and the interpretive margins.

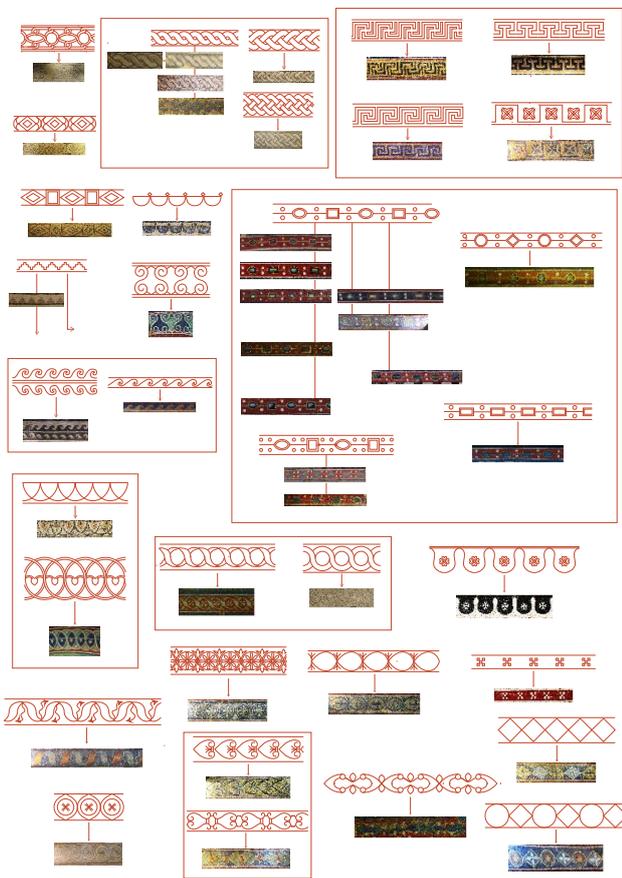


Figure 2. Main geometric band types included in the corpus and their vector reference templates

4. Methodology

The methodology (Figure 3) treats geometric mosaic ornament as an operational domain in which measurements extracted by survey activity, geometric and pattern recurrences analysis, construction of learning models and visualization networking can be brought in a common platform, which allow interconnections and knowledge inferences.

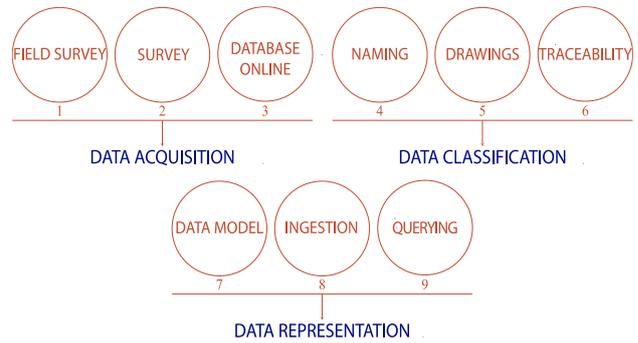


Figure 3. Methodology workflow. Traceable pipeline from data acquisition to classification and graph-based representation.

The workflow is articulated as a sequence of traceable phases, where each decorative instance is linked to its iconographic documentation, typological classification, and reference sources, culminating in an interrogable structure:

1. 3D data acquisition: photogrammetric documentation of surfaces (and, where necessary, integration from archives) to obtain dataset of images from different point of views and ortho-images products suitable for morphological reading and comparison.
2. Data classification: identification and extraction of decorative instances on planar surfaces; definition of classes and families through reference vector drawings and a unified nomenclature; recording and management of borderline cases (variants, hybrids, spatial adaptations).
3. Data organization: structuring essential metadata (chronology, location, denominations, bibliography) into a structured and exportable schema (CSV/JSON), with explicit provenance to ensure verifiability and updatability.
4. Data representation: modeling as a property graph implemented in ArangoDB, linking images, instances, classes/families, contexts, and sources, enabling exploratory navigation and comparative paths (e.g., cross-site recurrences and relationships between motifs and placements), with direct pointers back to documentary evidence.

5. Data Acquisition

The acquisition strategy was planned to produce geometrically consistent and morphologically readable images of the bands, with a particular attention over the following aspects:

- (i) spatial resolution, expressed as Ground Sample Distance (GSD) on the bands;
- (ii) optical conditions affecting the recognizability of modules (specular reflections on glass tesserae, colour casts from mixed lighting, non-orthogonal shots and the resulting local scale variations).

The survey approach has been adapted to the different external conditions. In the church where natural or artificial lighting allowed sufficient surface visibility, a photogrammetric approach has been carried out. It was applicable in the Ravenna case studies and in the Basilica of Santa Prassede in Rome. In the other Roman sites—whose direct acquisition was not feasible within project-compatible timeframes and conditions—archival documentation and digital resources were integrated. Such material was included only when it met minimum quality requirements (module readability, limited obliquity, and absence of degrading compression).

The campaigns were preceded by site inspections to optimize acquisition set-up (acquisition point of view, height and distance constraints), lighting conditions (natural light and existing artificial ones), and physical obstacles (walkways, parapets, confined

spaces). Across all sites, a three-level protocol was adopted: (1) general views to situate the band within the architectural system; (2) photogrammetric strip sequences with a minimum 50% overlap (longitudinal and transversal), ensuring robust orientation even when acquisition were not perfectly nadiral to the band; and (3) close-up detail acquisition to document tesseræ, borders, and interlace junctions, as well as micro-irregularities and execution variants.

To reduce specular reflections, no directed artificial lighting was used; when necessary, long exposure times and stable supports (tripod/telescopic pole) were employed (Figures 4-5). Camera settings were oriented toward preserving sharpness and micro-contrast, favouring mid-to-small apertures ($f/2.5$ – $f/14$) for adequate depth of field and low ISO values as lighting permitted (ISO 100–1600); where useful, longer focal lengths (up to 129 mm, or ~217 mm in 35 mm equivalent, depending on the device) were adopted to mitigate perspective effects in non-orthogonal views.



Figure 4. Photographic survey carried out with the camera mounted on a tripod.



Figure 5. Photographic survey carried out with the camera mounted on a telescopic pole.

GSD was derived a priori from acquisition parameters and validated a posteriori by GSD extraction from ortho-images. An estimated GSD range of approximately 0.07 mm to 2 mm, modulated according to the acquisition distance and used as an operational indicator to distinguish: detail scale (reading minimal units), intermediate scale (reading module and borders), and context scale (localization and continuity along the band). Besides, other important metadata have been recorded and reported in a common list (aperture, shutter speed, ISO). Images were processed in Agisoft Metashape following a standard workflow (orientation images → homologous points extraction → interpolated dense cloud definition → mesh construction → texture extraction → orthomosaic projection), with settings oriented toward reconstruction quality and the production of ortho-images supporting normalization and vector redrawing. Consistently with the paper's aim—focused on morphological comparison and evidence traceability—absolute metric scaling was not systematically enforced through targets/scalebars: comparability is ensured by the local geometric consistency of the orthometric products and by transparent recording of acquisition conditions and provenance. No colour analysis and colour checker has been used, focusing the attention on the extraction of the geometric pattern. This is part of the possible future activities which can implement the actual dataset.

6. Data Classification

Manual classification is the necessary control step to transform a heterogeneous repertoire of decorative bands into a comparable and queryable set: it defines units, classes, and morphological families, stabilizes nomenclature, and produces the attributes that will later be modeled in the property graph. The process combines: (i) terminological normalization; (ii) typological definition through reference drawing; and (iii) a traceable link between classes and documented instances, with explicit handling of variants and borderline cases.

6.1 Terminological normalization and an operational thesaurus

Specialist literature employs vocabularies that are not always aligned (across different disciplines, translations, varying levels of detail, and polysemy). To prevent such discontinuities from undermining cross-site comparisons, a one-to-one operational thesaurus was built (one normalized term per class), following a three-level scheme:

1. Primary reference (Balmelle, 1985): adopt Balmelle's denomination as the preferred label whenever the motif is attested in *Le décor géométrique de la mosaïque romaine*.
2. Secondary reference (Blanchard-Lemée et al., 1973): when the motif is not covered by Balmelle, use the denomination from *Le décor géométrique de la mosaïque antique*.
3. Undocumented cases (expert consultation): when the motif is not documented in either repertoire, introduce a controlled descriptive denomination based on consultation with Prof. Maria Cristina Carile (Department of Cultural Heritage, Alma Mater Studiorum – University of Bologna), whose work on Byzantine geometric frames includes *Forms and Ideas in the Fifth-Century Mediterranean* (2019) and *Visualità del celeste* (2020); the provenance of the choice is explicitly recorded.

6.2 Class definition: morphological–geometric criteria and reference drawing

Starting from the case-study corpus, the classification identified 34 classes (Figure 6) from 47 exemplars. Discrimination is based on morphological–geometric criteria: generating module,

rhythm, borders, concatenating rules, and relationships among recurring elements, prioritizing structural properties that are observable and repeatable.



Figure 6. The classification of the case-study corpus resulted in the identification of 34 decorative classes.

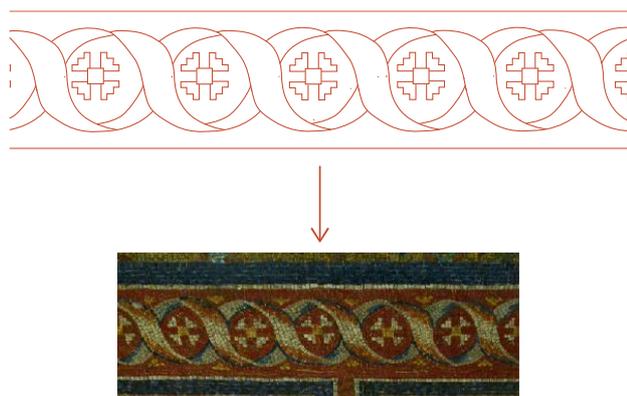


Figure 7. Simplified vector drawing and corresponding documented instance.

For each class, a simplified vector drawing was produced as a reference template (Figure 7): a normalization device that makes explicit repeated units, symmetries, and attachment

points between elements (interlace nodes, meander returns, peltae concatenations, etc.). When the same motif is attested in multiple contexts, the reference template was anchored to the earliest securely documented exemplar, treating it as a baseline typological formulation.

6.3 Families, variants, and borderline cases

Classes were also grouped into families (Figure 8) based on structural affinities (e.g., interlace logics, key/return patterns, sinusoidal structures), supporting both navigation and higher-level queries within the graph. Variants are handled with a conservative criterion: minor differences that can be traced back to the same geometric rule are treated as variants of the same class, whereas changes that affect the compositional rule or the generating module result in distinct classes.

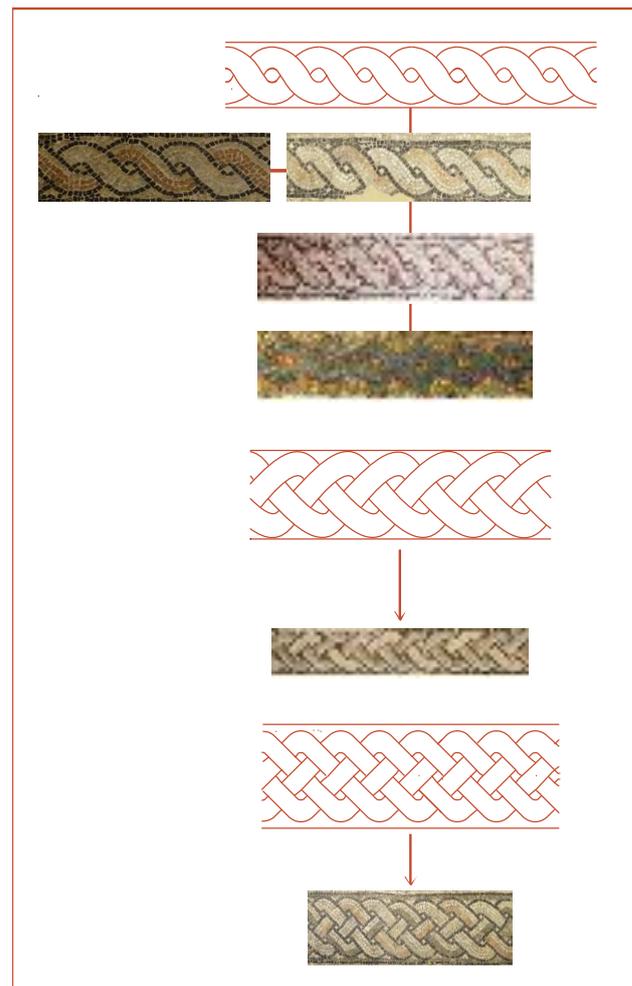


Figure 8. Example of a decorative family based on interlace structural logic.

6.4 Data schema and alignment with the property graph

The outcome of the classification was formalized into a structured data schema, maintained in a working sheet and exportable to CSV/JSON, designed to ensure a one-to-one correspondence between: class ↔ reference drawing ↔ documented instances. The schema is organized to support both editorial/descriptive use and subsequent modeling in the graph database.

For each class and each instance, at minimum the following are recorded:

- stable identifiers (persistent IDs and/or slugs) for class and instance, and family membership;
- normalized name (preferred label) and, when useful, terminological variants/aliases with source attribution;
- placement (wall/floor) and position within the architectural system when determinable;
- site and city, and chronology with an indication of uncertainty when applicable;
- essential bibliographic references and notes (variants, adaptations, lacunae, ambiguities);
- image links/IDs.

This formalization ensures terminological coherence, explicit provenance, and direct alignment with the property-graph model, where classes, families, and instances are represented as connected entities through navigable and inspectable relations (e.g., belongs to family, instance of, attested in, dated to, documented by, cited in).

7. Ingestion: A reproducible pipeline from the classification schema

To operationalize this alignment between the classification schema and the property-graph model, the structured records are ingested through a reproducible pipeline that enables their transformation into an interactive graph environment. We propose an integrated pipeline for easily transforming structured tabular data into an interactive three-dimensional graph representation. Made of three layers, its architecture combines server-side data processing (data ingestion and processing layer), ArangoDB database (property graph layer), and client-side web-based visualization (rendering layer), all deployed within a single Dockerized environment. Authentication and user identity management are delegated to Keycloak, an external identity provider (<https://github.com/keycloak/keycloak>).

First, the data ingestion mechanism aims to populate a property graph from a CSV file, uploaded through the web interface. To this end, each CSV file is parsed on the server side and converted into a structured representation suitable for further processing. During this parsing phase, the data is curated. It undergoes systematic normalization to ensure consistency across heterogeneous sources: empty values are standardized, numeric fields are correctly typed, and textual fields are cleaned to remove extraneous whitespace and formatting artifacts. When multiple values are encoded within a single cell, typically through line breaks, they are explicitly transformed into arrays. This normalization step guarantees that subsequent graph construction operates on well-defined and homogeneous data structures.

Due to the major importance of visual documentation in historical ornament digital repositories, a particular emphasis is placed on the handling of image references embedded within the CSV data. Rather than relying on predefined column names, the system automatically detects potential image paths by identifying string values that correspond to known image file extensions. This strategy makes the pipeline resilient to schema variability and allows the ingestion of datasets with evolving or inconsistent column naming conventions. All detected paths are subjected to a curation process that removes filesystem-dependent prefixes, rejects absolute paths or directory traversal patterns, and normalizes separators. This approach prevents invalid or unsafe references while preserving portability across operating systems and deployment environments. The web interface proposes an option allowing users to upload their

corresponding image datasets alongside the CSV file. Then, image paths are resolved into public URLs that can be accessed by the rendering layer. The reference dataset presented in this paper rely on a permanent storage area exposed through a stable URL namespace, whereas user-uploaded datasets are mapped to temporary, user-scoped directories. This distinction allows the system to persist authoritative datasets while isolating ephemeral data generated during exploratory sessions. The resolved image URLs are stored directly as node attributes, enabling the visualization layer to load textures without direct access to the underlying filesystem.

The graph is populated in ArangoDB, a multi-model database, in which nodes and relations are stored as JSON documents in nodes and edges collections. We decided not to use a dataset-driven approach, with one node and one edge collection per dataset, to prevent the proliferation of collections and the subsequent querying issues. For this reason, each graph element is associated with its dataset unique identifier, allowing multiple datasets to coexist within the same collections, while enabling efficient filtering, querying, and cleaning operations. Each row of the normalized CSV file is saved as a graph node, preserving all its curated attributes as node properties. To ensure global connectivity, a blank root node is introduced when necessary. Edges are created when parent/child relationships are explicitly specified in the data (e.g. through an optional ParentID or BELONGS_TO_FAMILY attribute) (Figure 9).

7.1 Data representation

The Atlas of Byzantine geometric mosaic bands is used as a reference dataset, which is automatically imported during the container startup if no prior instance exists, ensuring that users are always presented with a baseline graph upon accessing the application. In contrast, authenticated user datasets are designed to be temporary. They are replaced upon subsequent uploads by the same user, may be deleted explicitly through the interface, and are periodically cleaned up by a time-based expiration mechanism. Although somewhat limited, this strategy has the advantage of preventing an uncontrolled resource accumulation while supporting iterative experimentation.

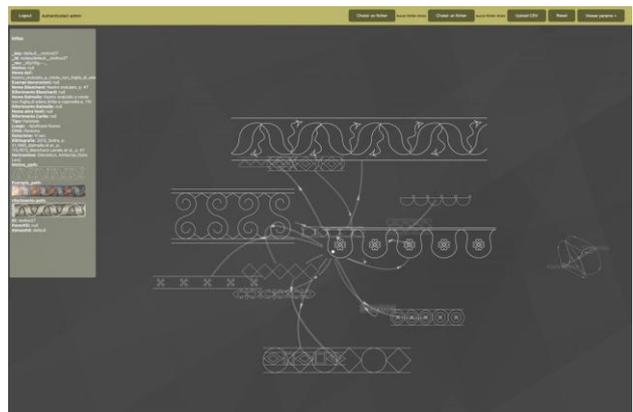


Figure 9. ArangoDB graph viewer after CSV ingestion: each row becomes a node (JSON properties) and explicit links (e.g., ParentID, BELONGS_TO_FAMILY) generate edges; elements are scoped by dataset_id within shared collections.

When a dataset is requested for visualization, the backend executes an AQL query to extract all nodes and edges documents related to the specified dataset identifier, and transmits the result to the client as a structured JSON object. The visualization layer renders the graph using 3D-force-graph

(<https://github.com/vasturiano/3d-force-graph>) a JavaScript graph rendering library built on top of Three.js (<https://github.com/mrdoob/three.js>). Nodes are displayed as sprites and may be textured with images when associated to consistent paths, while edges are rendered either as directional arrows or as animated particles, depending on user preference. The force-based layout is computed dynamically to adapt the representation to user inputs.

Interaction is exploratory and visual: the user selects a graph node (motif/instance) directly within the network and retrieves the associated information through a side information panel. Selection provides descriptive metadata (normalized label, placement, chronology, bibliographic references), as well as links to images and to the reference drawings used as evidence. When a motif is connected to a class or to a morphological family, the viewer also reveals the related nodes, enabling users to explore recurrences and variants within the repertory through the graph's relational structure (Figure 10). The user can configure rendering parameters (background color, edge style, etc.), and export high-resolution snapshots of the current view. His configuration settings are stored locally to provide him with a consistent experience across sessions.

The Atlas is currently deployed within the ESPADON digital ecosystem, dedicated to heritage sciences (R. Thomas et al., 2025). Long-term operation requires on the one hand versioning procedures, change tracking, and maintenance of the relationship vocabulary and identifiers, in order to preserve data cumulability and traceability, and on the other hand, the support for collaborative or multi-user workflows.

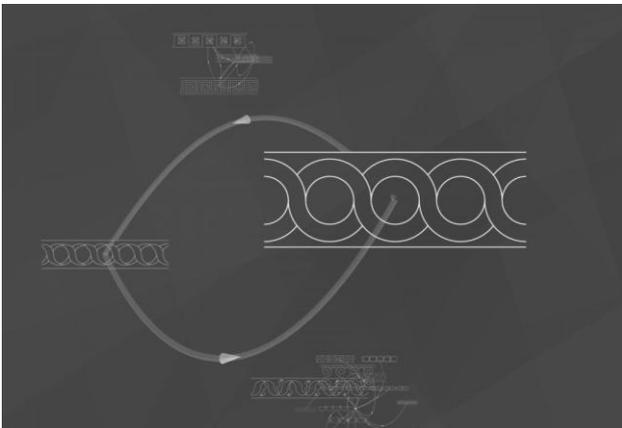


Figure 10. Graph viewer: selecting a motif reveals its links to the corresponding class and morphological family, enabling exploration of recurrences and variants across the repertory.

7.2 Target Users and Practical Value.

The Atlas is primarily intended for scholars and practitioners (art historians, archaeologists, architectural representation specialists, conservators, and cataloguers), while remaining readable for non-specialist audiences when consultation starts from images and from guided paths by classes and families. For researchers, the graph structure makes comparisons feasible that would otherwise be costly within a fragmented literature: it enables rapid verification of attestations, variants, and co-occurrences across motifs, placements, and chronologies, while keeping sources and supporting visual documentation continuously traceable. For cataloguing and conservation, the Atlas provides an operational vocabulary and a reusable data model, helping reduce terminological ambiguity, consistently record borderline cases (hybrids, spatial adaptations, lacunae), and update the repertory incrementally without losing the

history of relationships. Looking ahead, the infrastructure can serve as a shared baseline to integrate new sites, expand the bibliography, and build thematic consultation paths (by family, architectural context, or chronological range), thereby supporting interdisciplinary communication and the replicability of analyses.

8. Conclusions and Future Work

This paper outlines a practical workflow for turning fragmented ornamental documentation into an evidence-linked, queryable Atlas, where each record remains explicitly connected to its visual sources, contextual information, and bibliography. Applied to a testbed of 34 classes of geometric mosaic bands on planar surfaces (5th–12th centuries) from Ravenna and Rome, the approach supports controlled comparison across sites and periods by making typological proximity, attestations, and variants inspectable together with the documents that justify them.

The main contribution is therefore an operational structure for comparative reading: an Atlas in which navigation follows typological relationships and returns, at each step, to the underlying evidence (images, reference drawings, and cited sources). This framing is especially useful for research and professional practice, where interpretive claims often depend on dispersed documentation and on terminology that shifts across publications.

Several limitations define a coherent roadmap. The current corpus remains bounded (two centres, planar surfaces), while extending the method to curved or non-planar areas will require dedicated procedures for geometric normalisation and sampling. In addition, long-term sustainability depends on robust data governance—stable identifiers, duplicate control, change tracking, and the maintenance of a controlled relation vocabulary—particularly given that, at the time of submission, the Atlas is deployed online within a controlled-access environment.

Future work will focus on: (1) expanding the corpus (sites, typological coverage) and addressing non-planar surfaces with tailored correction and sampling strategies, keeping acquisition conditions as explicit, queryable attributes; (2) strengthening infrastructure and governance through a consolidated read-only web viewer, controlled access, backup/replication procedures, and a maintained relation vocabulary (e.g., `INSTANCE_OF`, `BELONGS_TO_FAMILY`, `VARIANT_OF`, `LOCATED_IN`, `CITED_IN`); (3) improving interoperability by progressively aligning selected attributes and relations with reference conceptual/ontological models; (4) enabling collaborative curation via proposal–review–validation workflows (e.g., through integration with platforms such as Aïoli), while preserving the Atlas as the convergence layer for disciplinary judgement and data infrastructure; and (5) adding optional text search and filters (location, site, chronology) as complementary access paths, without replacing the evidence-first, graph-driven navigation.

Overall, the work proposes a feasible route from heterogeneous documentation and naming practices to a cumulative, inspectable repertory, designed to grow incrementally while preserving traceability and reuse.

Acknowledgements

The authors would like to thank Daouda Ngom for his contribution to making the property graph accessible online, and Prof. Maria Cristina Carile (University of Bologna) for her expert consultations on nomenclature and terminological choices.

This experiment benefited from State aid managed by the Agence Nationale de la Recherche (French National Research Agency) under the Investment for the Future Program integrated into France 2030, bearing the reference ANR 21-ESRE-0050 EquipEx+ ESPADON. Giulia Flenghi drafted Sections 1–4 and 6, and contributed to Section 8. Violette Abergel drafted Sections 7 and contributed to Sections 1 and 8. Michele Russo drafted Section 5 and contributed to Sections 1 and 8.

References

- Adami, A., Fassi, F., Fregonese, L., Piana, M., 2018: Image-based techniques for the survey of mosaics in the St Mark's Basilica in Venice. *Virtual Archaeology Review*, 9(19), 1.
- Arena, M., 2020. Connessioni geometriche: Per una catalogazione "fantastica" dei pattern bizantini. In: *Connettere – Un disegno per annodare e tessere · Connecting – Drawing for Weaving Relationships*. FrancoAngeli srl.
- Balmelle, C., Prudhomme, R. (ill.), Morricone Matini, M.-L. (trad.), 1985. *Le décor géométrique de la mosaïque romaine: Répertoire graphique et descriptif des compositions linéaires et isotropes*, Vol. 1. Picard, Paris.
- Balmelle, C., Prudhomme, R. (ill.), Morricone Matini, M.-L. (trad.), 1985. *Répertoire graphique et descriptif des décors centrés*, Vol. 2. Picard, Paris.
- Blanchard-Lemée, M., Balmelle, C., Prudhomme, R. (Eds.), 1973. *Le décor géométrique de la mosaïque antique. Bulletin de l'Association internationale pour l'étude de la mosaïque antique*, IV fascicule, Mai 1973. Paris.
- Carboni, N., De Luca, L., 2016: Towards a conceptual foundation for documenting tangible and intangible elements of a cultural object. *Digital Applications in Archaeology and Cultural Heritage*, 3(4), 108–116.
- Carboni, N., De Luca, L., 2019: An ontological approach to the description of visual and iconographical representations. *Heritage*, 2(2), 1191–1210.
- De Luca, L., Véron, P., Florenzano, M., 2011. A semantic-based platform for the digital analysis of architectural heritage. *Computers & Graphics*, 35(2), 227–241.
- Doerr, M., 2003. The CIDOC Conceptual Reference Model. *AI Magazine*, 24(3), 75–92.
- Doria, E., Picchio, F., 2020: Techniques for mosaics documentation through photogrammetry data acquisition: The Byzantine mosaics of the Nativity Church. *ISPRS Ann. Photogramm. Remote Sens. Spatial Inf. Sci.*, V-2-2020, 965–972.
- Fioretti, G., Acquafredda, P., Calò, S., Cinelli, M., Germanò, G., Laera, A., Moccia, A. (2020). Study and conservation of the St. Nicola's Basilica mosaics (Bari, Italy) by photogrammetric survey: Mapping of polychrome marbles, decorative patterns and past restorations. *Studies in Conservation*, 65(3), 160–171.
- International Organization for Standardization (ISO), 2014. ISO 21127:2014: Information and documentation — A reference ontology for the interchange of cultural heritage information. *International Organization for Standardization*, Geneva, Switzerland.
- Katifori, A., Halatsis, C., Lepouras, G., Vassilakis, C. & Giannopoulou, E., 2007. Ontology visualization methods — a survey. *ACM Computing Surveys (CSUR)*, 39(4), Article No. 10.
- Maino, G., Visparelli, D., 2003. Il museo virtuale dei mosaici. In: *Atti del Convegno Contesti Virtuali e Fruizione dei Beni Culturali* (Napoli, 22–23 maggio 2003), 1–6. ENEA – Centro Ricerche "E. Clementel"; Certosa e Museo di San Martino, Napoli.
- Manferdini, A. M. (2010). Digital survey of ancient mosaics of Ravenna. In *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XXXVIII-5, 387–392.
- Manferdini, A.M., Cipriani, L., 2011. Digital 3D collections of mosaics. In: *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci.*, XXXVIII-5/W16, 201–206.
- Roussel, D., De Luca, L., 2018. Aïoli: a web-based platform for the collaborative annotation of large image datasets. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, IV-2, 209–216.
- Thomas, R., Detalle, V., Abergel, V., Anheim, E., Bernet, Q., De Luca, L., Petitcol, R., Wilkie-Chancelier, N., 2025. The augmented heritage object: An integrated paradigm for multi-disciplinary data and knowledge production in heritage science. *Journal of Cultural Heritage*. Submitted, under review.