Establishing a Dialogue Between Textual Sources and 3D-Spatialised Annotations: A First Experiment Within Notre-Dame's Scientific Data Corpus

Pierre Arese [1], Roxane Roussel [1], Livio De Luca [1]

[1] MAP UPR 2002 CNRS, 31 chemin Joseph Aiguier, 13402 Marseille, France – (pierre.arese, roxane.roussel, livio.deluca)@map.cnrs.fr

Keywords: 3D Annotation, Semantic Annotation, Notre-Dame de Paris, Knowledge Graph, Cultural Heritage.

Abstract

Structured 3D annotations offer a powerful means of documenting and localizing scientific and professional observations within spatial models. However, the production and transmission of knowledge in heritage science ultimately relies on narratives—discursive forms that explain how and why these observations were made. Reconciling the precision of structured, queryable data with the contextual depth of scholarly reasoning remains a central challenge in digital humanities and cultural heritage research. How can these two descriptive trajectories—analytical and narrative—be interconnected? How can we enable researchers to move fluidly between the what / where / when and the how / why within a single exploration environment? This article presents an experimental framework developed in response to that challenge, situated within the unprecedented multidisciplinary effort launched after the fire at Notre-Dame de Paris on April 15, 2019. Capitalizing on the extensive research activity that followed, we sought to interconnect two complementary corpora: (1) a rich body of spatialized 2D/3D annotations created on the Aïoli platform, and (2) textual sources—including diagnostic reports and academic publications—produced by researchers and heritage professionals. The paper introduces automated mechanisms for linking these two forms of documentation and describes an interactive visualization system that allows users to explore and assess these connections both semantically and spatially.

1. Introduction

The day after the fire of Notre-Dame de Paris, the CNRS and the French Ministry of Culture took the initiative of creating a scientific worksite structured into working groups (Dillmann et al., 2023). Intended to allow scientists from various fields (architecture, archaeology, art history, computer science, anthropology, chemistry, etc.) to study and collect data on the cathedral, the Scientific Project has enabled the development of numerous projects. In this context, the “digital data working group” set itself the objective of creating a “digital ecosystem” with the aim of integrating, creating, disseminating and archiving the recovered data. The digital platform developed within this framework allowed the gathering of different layers of data coming from involved stakeholders (Néroulidis et al., 2024).

Within this platform a central role was played by the 3D annotation. Aïoli (Manuel and Abergel, 2018), a web service designed for reality-based 3D annotation and multi-user collaborative documentation of cultural and heritage artefacts platform allows us to stake several annotation layers around a common denominator (Roussel and De Luca, 2023). It was developed by the MAP laboratory (Models and Simulations for Architecture and Cultural Heritage) with the support of the CNRS and the French Ministry of Culture.

If structured annotations represent an excellent way to document and represent scientific and professional observations and analysis within a spatialised representation, the production and dissemination of knowledge always rise from narratives. Which means, on the one hand, we describe facts structured into entities and observations that can be spatially represented and characterized by quantitative and qualitative attributes, and on the other hand, to elucidate the reasoning behind the processes of analysis and interpretation of these facts, we have to rely them on narratives. These narratives do not manipulate the same structured representations but instead integrate these elements within logical sequences that reflect the intellectual approach followed during the study. Which means that, at some point, we have to reconcile an encyclopedic approach - suitable for database queries - with the necessity of contextualizing these data and integrating them into knowledge production trajectories.

This leads us to the following questions: How can these two descriptive trajectories be reconciled? How can they be interconnected? How can the strengths of analytical description and narrative be combined within a data exploration framework?

2. State of the Art: Bridging Structured Data and Narrative in Cultural Heritage Research

The articulation between structured data and narrative discourse remains a key challenge in the digital humanities and heritage sciences. While structured annotations enable the precise localization and characterization of information within three-dimensional models, they often fall short of conveying the intellectual paths and interpretative processes behind scholarly observations. Conversely, narrative texts offer rich and contextualized accounts of these processes but lack direct spatial anchoring or formalized representation. This section reviews current approaches and technological frameworks that aim to address this gap, thereby setting the stage for our own contribution.

2.1 Structured Spatialised Annotations in Heritage Science

Structured annotations have become essential tools in the documentation and analysis of cultural heritage. Solutions such as Arches (Peytavie et al., 2009), Recogito (Simon et al., 2015), and Aïoli (Manuel and Abergel, 2018) provide researchers with environments to annotate spatial data, enabling multi-user collaboration and the layering of interpretations on top of representations in 2D and 3D.. These systems typically rely on structured models that associate entities, attributes, and values with spatial coordinates or geometric primitives, thus making the data computable, queryable, and shareable. These efforts are often supported by semantic frameworks such as CIDOC-CRM (Doerr, 2003) or Linked Open Data (Bizer et al., 2017), which promote interoperability across domains and institutions. However, despite their technical rigor (or descriptive flexibility), these annotation systems tend to abstract away the interpretive richness that characterizes heritage research. They are not well-suited for expressing epistemic uncertainty, competing hypotheses, or the reasoning behind interpretative choices. In other words, while they excel at describing **what**, **where**, and **when**, they struggle to convey **how** and **why**. This is particularly true in spatial annotation systems, where location and time can be precisely encoded, but the logic and context behind interpretative actions remain implicit.

2.2 Textual Sources and Concept Extraction in Humanities Research

Textual sources remain the primary medium for conveying scholarly narratives in the humanities. These texts—excavation reports, architectural analysis, scientific publications—encode rich interpretative knowledge often resistant to structured representation. Recent advances in Natural Language Processing (NLP) now enable the extraction of entities, concepts, and relationships from unstructured text using techniques such as Named Entity Recognition (NER) (Keraghel et al., 2024), topic modeling (e.g., LDA) (Drissi, 2024), and contextual embeddings (e.g., BERT) (Illina and Fohr, 2022). In the heritage domain, such methods face challenges: specialized vocabulary, historical language shifts, and ambiguous references require tailored tools and domain-specific models. Extracted concepts also rarely include spatial references in geometric terms, making alignment with structured datasets—like annotations or 3D models—difficult, though such alignment is crucial to grasping the scholarly logic behind data production and interpretation.

2.3 Towards Integrated Narratives: Connecting Structured Data and Discourse

Various attempts have been made to bridge the gap between structured datasets and narrative expression. Research in digital storytelling and narrative databases has explored how structured information can be integrated into story-based interfaces, particularly in museum contexts and educational applications. StoryMaps, timeline visualizations, and projects such as Eureka3D-XR (https://eureka3d.eu/eureka3d-xr/) have illustrated the potential of combining spatial data with interactive narratives. This convergence raises theoretical questions about the nature of knowledge representation. Should digital heritage systems favor document-centric approaches that reflect the original sources, or knowledge-centric approaches that prioritize formalized entities and relationships? Concepts such as data storytelling, hermeneutic loops, and thick description provide useful lenses through which to understand this tension. What is needed is a way to let narratives and structured data coexist—each informing and enriching the other.

2.4 A Need for a Dual-Trajectory Approach

This state of the art reveals a persistent disjunction between structured data models and interpretive narratives. On one hand, annotations enable rigorous, spatialized documentation, but they lack narrative depth. On the other hand, textual narratives provide interpretative insight but are rarely structured or linked to 2D or 3D spatial data. Our experiment approach seeks to reconcile these two trajectories by developing a framework in which concepts extracted from narrative texts can be semantically and visually connected to structured 3D annotations. This enables the construction of spatialized narratives grounded in data, offering a novel way to explore, interpret, and communicate cultural heritage knowledge.

3. Bridging Two Descriptive Domains: Research Writings and Spatial Annotations

To address the challenge of reconciling structured analytical representations with narrative reasoning, our approach relies on the parallel mobilization of two complementary corpora. Each of these corpora reflects a distinct epistemic trajectory: one rooted in the textual expression of research, the other in the spatialized documentation of observations.

On one side, scientific writings—such as academic articles and project reports—offer discursive windows into the processes, methods, and interpretative frameworks underlying research practices (see section 4). These texts, particularly when written by those directly involved in field investigations, carry valuable insights into the conceptual strategies and epistemological positions adopted by researchers. By examining the language used in these writings, we aim to trace the emergence of disciplinary perspectives and uncover patterns in knowledge production. To this end, we developed an exploratory pipeline for extracting and classifying relevant terms from scientific publications using Large Language Models (LLMs), in order to construct a structured view of conceptual domains emerging from the Notre-Dame scientific restoration project.

On the other side, the 3D annotation corpus developed within the Aïoli platform offers a spatial and semantic structuring of observations made on the cathedral's fabric (see section 5). Created collaboratively by architects, scientists, and heritage professionals, these annotations are grounded in photogrammetric models and linked to contextual data, forming a georeferenced knowledge base. Their structured format enables quantitative querying, but also reflects the physical and spatial dimension of research activity.

The combination of these two corpora—narrative and spatial—forms the basis of our experimental framework. While each originates from different epistemic logics, their convergence opens new possibilities for understanding how knowledge is constructed, represented, and interconnected. The following section presents the methodologies developed for the semi-automated analysis of research writings and the organization of spatial annotations, before exploring how these can be connected to support new modes of scientific documentation and exploration (see section 6) (Fig. 1).

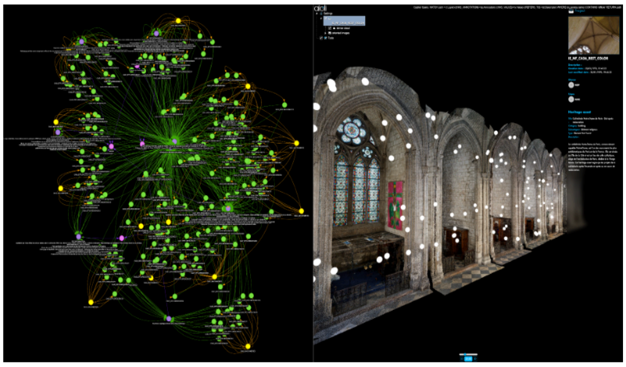

Figure 1. Integrated visualization of a set of spatialized annotations on the conservation state of the chapels and the graph of associated semantic relationships

4. An Exploratory Approach to Scientific Documentation through the Analysis of Research Writings

4.1 A semi-automatic extraction of textual data

This work is part of a scientific documentation effort that aims to better explore and understand research dynamics through the study of texts produced by researchers themselves. Scientific documentation plays a key role in understanding research practices, especially in interdisciplinary contexts. Instead of conducting direct observations or interviews, our approach is based on the analysis of scientific writings, as reflective and discursive representations of research work. By focusing on publishing texts – such as those from the Journal of Cultural Heritage (1) – we assume that the words chosen by the authors can serve as traces of the forms of thinking involved in knowledge production.

Our analysis relies on the semi-automatic extraction of textual data using freeware artificial intelligence tools that are easy to manipulate without requiring advanced technical skills. This extraction process does not rely on traditional statistical models but instead uses optimized prompts and lightweight linguistic processing techniques, allowing for minimum human intervention while maintaining adaptability. Each text is analysed in terms of the relative positions of terms within a single document, as well as how the texts relate to one another.

4.2 Workflow

Between the summer of 2023 and January 2024, an exploratory study was conducted to develop a semi-automated method for extracting and classifying specialized terms from a scientific corpus using Large Language models (LLMs), particularly ChatGPT (versions 3,5 and 4), Llama, and Bard. The objective was to assess the capacity of these models to identify, classify and, when possible, count terms relevant to a predefined categorization. This categorization was structured into four domains: Material Thinking (including alterations, architectural composition, materials and architectural elements), Knowledge Thinking (research domains, methodologies and disciplinary areas), Digital thinking (digital tools, representations and metadata), and Events (including temporal markers, professional activities, and actors).

Initial tests with full chapters showed partial success: while relevant terms were often detected, classification inconsistencies and difficulties with numerical data – especially temporal data – were recurrent. In particular, the subcategory D2 "Spatio Temporal Framework" revealed major gaps: from forty-four manually identified dates, GPT-4 retrieved only sixteen, and Llama introduced hallucinated dates. Restricting the task to category D and using visual cues (e.g., bold text) slightly improved results, but remained insufficient.

A more effective strategy emerged by analysing the text paragraph by paragraph. This reduced hallucinations and improved term detection, especially for temporal data. Capitalizing key terms in the input further enhanced extraction accuracy and prompted the model to infer additional relevant concepts.

Key conclusions of this first phase include the need to restrict input length, avoid asking for term counts, match the prompt language to the text, and minimize model interpretation. Furthermore, the models should not be granted interpretative freedom; prompts must enforce rigid, explicit instructions. Also, the use of visual markers and pre-identified keywords significantly improves performance, mostly when a non-English text is preceded.

4.3 First experiment: The Journal of Cultural Heritage

Building on the previously tested methodology, the full text-mining process was applied to all articles in Volume 65 of The Journal of Cultural Heritage, a special issue entirely dedicated to the scientific restoration project of Notre-Dame de Paris. The thirty-one articles analysed reflect the diversity and technical complexity of the research conducted by various interdisciplinary working groups involved in the project.

The objective was to extract and categorize all relevant terms from these texts using the predefined conceptual framework, without relying on statistical models that typically focus on abstracts and metadata or that reduce analysis to frequency counts. Instead, we employed ChatGPT 3,5 – a generative model trained on a vast corpus – chosen specifically for its ability to provide nuanced outputs guided by carefully formulated prompts. The use of a conversational interface further allowed for iterative refinement or results through dialogue.

To operationalize this process, each article was systematically structured in an Excel table. The texts were segmented into their constituent parts (e.g., introduction, chapter sections, conclusion), with each section placed into a dedicated cell corresponding to its structural position in the article (Fig. 2). ChatGPT was then instructed to process the text paragraph by paragraph, extracting and classifying key terms according to the predefined categories. The model was further asked to return the results in a structured JSON format, which was subsequently converted to CSV for integration into the Excel table. JSON format was chosen for its hierarchical and machine-readable structure, allowing for clear mapping between terms and categories while facilitating reliable data transformation and export across different platforms. The extracted terms were then manually recorded in the relevant section of the Excel table, preserving both the semantic structure and contextual integrity of each article. This structured and interactive approach enabled

(1) Journal of Cultural Heritage, *Notre-Dame de Paris : a multidisciplinary scientific site,* Vol.65 - January/February 2024,pp.1-240. (https://www.sciencedirect.com/journal/journal-of-cultural-heritage/vol/65/)

a comprehensive and fine-grained term extraction process, tailored to the disciplinary and technical complexity of the corpus.

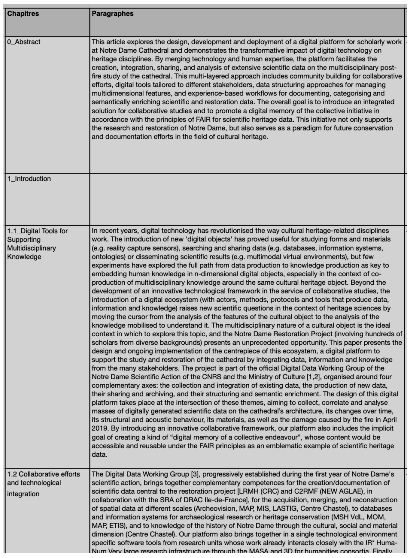

Chapitres	Paragraphes
0_Abstract	This article explores the design, development and deployment of a digital platform for scholarly work at Notre Dame Cathedral and demonstrates the transformative impact of digital technology on heritage disciplines. By merging technology and human expertise, the platform facilitates the creation, integration, sharing, and analysis of extensive scientific data on the multidisciplinary post-fire study of the cathedral. This multi-layered approach includes community building for collaborative efforts, digital tools tailored to different stakeholders, data structuring approaches for managing multidimensional features, and experience-based workflows for documenting, categorising and semantically enriching scientific and restoration data. The overall goal is to introduce an integrated solution for collaborative studies and to promote a digital memory of the collective initiative in accordance with the principles of FAIR for scientific heritage data. This initiative not only supports the research and restoration of Notre Dame, but also serves as a paradigm for future conservation and documentation efforts in the field of cultural heritage.
1_Introduction	
1.1_Digital Tools for Supporting Multidisciplinary Knowledge	In recent years, digital technology has revolutionised the way cultural heritage-related disciplines work. The introduction of new 'digital objects' has proved useful for studying forms and materials (e.g. reality capture sensors), searching and sharing data (e.g. databases, information systems, ontologies) or disseminating scientific results (e.g. multimodal virtual environments), but few experiments have explored the full path from data production to knowledge production as key to embedding human knowledge in n-dimensional digital objects, especially in the context of co-production of multidisciplinary knowledge around the same cultural heritage object. Beyond the development of an innovative technological framework in the service of collaborative studies, the introduction of a digital ecosystem (with actors, methods, protocols and tools that produce data, information and knowledge) raises new scientific questions in the context of heritage sciences by moving the cursor from the analysis of the features of the cultural object to the analysis of the knowledge mobilised to understand it. The multidisciplinary nature of a cultural object is the ideal context in which to explore this topic, and the Notre Dame Restoration Project (involving hundreds of scholars from diverse backgrounds) presents an unprecedented opportunity. This paper presents the design and ongoing implementation of the centrepiece of this ecosystem, a digital platform to support the study and restoration of the cathedral by integrating data, information and knowledge from the many stakeholders. The project is part of the official Digital Data Working Group of the Notre Dame Scientific Action of the CNRS and the Ministry of Culture [1,2], organised around four complementary axes: the collection and integration of existing data, the production of new data, their sharing and archiving, and their structuring and semantic enrichment. The design of this digital platform takes place at the intersection of these themes, aiming to collect, correlate and analyse masses of digitally generated scientific data on the cathedral's architecture, its changes over time, its structural and acoustic behaviour, its materials, as well as the damage caused by the fire in April 2019. By introducing an innovative collaborative framework, our platform also includes the implicit goal of creating a kind of “digital memory of a collective endeavour”, whose content would be accessible and reusable under the FAIR principles as an emblematic example of scientific heritage data.
1.2 Collaborative efforts and technological integration	The Digital Data Working Group [3], progressively established during the first year of Notre Dame's scientific action, brings together complementary competences for the creation/documentation of scientific data central to the restoration project [LRMH (CRC) and C2RMF (NEW AGLAE), in collaboration with the SRA of DRAC Ile-de-France], for the acquisition, merging, and reconstruction of spatial data at different scales (Archeovision, MAP, MIS, LASTIG, Centre Chastel), to databases and information systems for archaeological research or heritage conservation (MSH VdL, MOM, MAP, ETIS), and to knowledge of the history of Notre Dame through the cultural, social and material dimension (Centre Chastel). Our platform also brings together in a single technological environment specific software tools from research units whose work already interacts closely with the IR* Huma-Num Very large research infrastructure through the MASA and 3D for humanities consortia. Finally,

Figure 2. Example of a text segmented into its constituent parts. On the left side, the titles and subtitles; on the right side, the corresponding texts

Once all extracted terms had been entered into their respective sections of the Excel table, a critical data curation phase was undertaken. Although the extraction process using ChatGPT proved effective overall, it frequently produced classification inaccuracies. Each category was therefore systematically reviewed to ensure that every term was correctly assigned. Misclassified terms were relocated to their appropriate categories, and duplicate entries were removed to eliminate redundancies and preserve the clarity and relevance of the dataset. This step was essential to ensuring the reliability and coherence of the final term inventory.

In total, thirty-one articles from the Volume 65 of The Journal of Cultural Heritage were analysed using the structured, prompt-driven extraction method. This process yielded three hundred and five individual CSV files and a corpus of five thousand three hundred seventy classified terms (Fig. 3).

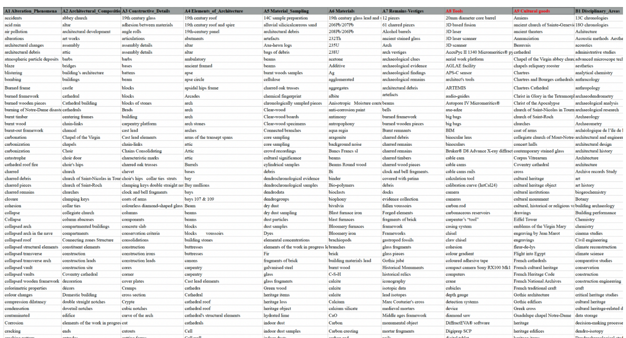

A1 Alteration_Phenomena	A2 Architectural_Compositio	A3 Constructive_Details	A4 Elements_of_Architecture	A5 Material_Sampling	A6 Materials	A7 Remains-Vestiges	A8 Tools	A9 Cultural goods	B1 Disciplinary_Areas
accidents	abbey church	19th century glass	19th century roof	14C sample preparation	19th century glass lead and s	12 pieces	20mm diameter core barrel	Amiens	13C chronologies
acid rain	altar	adhesion between materials	19th century roof and spire	alluvial silicicalcareous sand	206Pb/207Pb	61 charred pieces	3D based fusion	ancient church of Sainte-Genevi	18O chronologies
air pollution	architectural development	angle rolls	19th-century panel	architectural debris	208Pb/206Pb	Alcohol barrels	3D laser	ancient theaters	Achitecture
alterations	art works	articulations	abutments	artefacts	232Th	ancient stained glass	3D laser scanner	Annunciation	Acoustic methods Aesthe
architectural changes	assembly	assembly details	altar	Axe-hewn logs	235U	Arch	3D scanner	Beauvais	acoustics
architectural debris	attic	assembly details	altar	bags of debris	238U	arch vestiges	AccuPyc II 1340 Micromeritics® py	cathedral	administrative studies
atmospheric particle deposits	barbs	barbs	ambulatory	beams	acetone	archaeological clues	aerial work platform	Chapel of the Virgin abbey churc	advanced microscope tec
blaze	bridges	bases	ancient framed	beams	Additive	archaeological evidence	AGLAE facility	chapels reliquary rooster	aesthetics
blistering	building's architecture	battens	apse	burnt woods samples	Ag	archaeological findings	APS-C sensor	Chartres	analytical chemistry
bombing	buildings	beam	apse circle	cellulose	agglomerated	archaeological remains	architect's tools	Chartres and Bourges cathedrals	anthracology
Burned frame	castle	blocks	apsidal hips frame	charred oak trusses	aggregates	architectural debris	ARTEMIS	Chartres Cathedral	anthropology
burned framework	cathedral	blocks	Arcades	chemical fingerprint	albite	artefacts	audio-guides	Christ in Glory in the Tetramorp	archaeodendrometry
burned wooden pieces	Cathedral building	blocks of stones	arch	chronologically sampled pieces	Anisotropic Moisture conte	beams	Autopore IV Micromeritics®	Christ of the Apocalypse	archaeological analysis
burning of Notre-Dame disaste	cathedrals	Brads	arch	Clear-wood	anti-corrosion paint	bells	axe-adze	church of Saint-Nicolas in Tourn	archaeological research
burnt timber	centering frames	building	arch	Clear-wood boards	antimony	burned framework	big bags	church of Saint-Roch	Archaeology
burnt wood	chain-links	carpentry platform	arch stones	Clear-wood specimens	antropophony	burned wooden pieces	big bags	churches	Archaeometry
burnt-out framework	chancel	cast lead	arches	Connected branches	aqua regia	Burnt remnants	BIM	coat of arms	archéologique de l'île de l
carbonation	Chapel of the Virgin	Cast lead elements	arms of the transept spans	core sampling	aragonite	charred debris	binocular lens	collegiate church of Mont-Notre-	architectural and engineer
carbonization	chapels	chain-links	attic	core sampling	background noise	charred remains	binoculars	concert halls	architectural design
carbonization	Choir	Chains Consolidating	Attic	crowd recordings	Bancs Francs sl	charred remains	Bruker® D8 Advance X-ray diffract	contemporary stained glass	architectural history
catastrophe	choir door	characteristic marks	attic	cultural significance	beams	charred timbers	cable cam	Corpus Vitrearum	Architecture
cathedral roof fire	choir's hips	charred oak trusses	Barrels	cylindrical samples	Beams Round wood	charred wood pieces	cable cams	Coventry cathedral	architecture
charred	church	chevet	bases	debris	Bi	clock and bell fragments.	cable cams rails	cross	Archive records Study
charred debris	church of Saint-Nicolas in Tour	choir's hips collar ties struts	bay	dendrochronological evidence	binder	covered with patina	calculation tool	cultural heritage	art
charred pieces	church of Saint-Roch	clamping keys double straight no	Bay mullions	dendrochronological samples	Bio-polymers	debris	calibration curve (IntCal24)	cultural heritage object	art history
charred remains	churches	clock and bell fragments	bays	dendrodata	bioclasts	docks	camera	cultural institutions	biogeochemistry
closure	clamping keys	coats of arms	bays 107 & 109	dendrogroups	biophony	evidence collection	cameras	cultural monument	Botany
cohesion	collar ties	colourless diamond-shaped glass	Beam	dry dust	bivalvia	fallen voussoirs	carbon rod	cultural, historical or religious va	building archaeology
collapse	collegiate church	columns	beams	dry dust sampling	Blast furnace iron	Forged elements	carbonaceous reservoirs	drawings	Building performance
Collapse	column abacuses	components	beams	dust particles	blast furnaces	fragments of brick	carpenter's “tool”	Eiffel Tower	Chemistry
collapsed arch	compartmented buildings	concrete slab	blocks	dust samples	Bloomery furnaces	framework	casing system	emblems of the Virgin Mary	chemistry
collapsed arch in the nave	compartments	conservation criteria	blocks voussoirs	Dyes	Bloomery iron	Frameworks	chisel	engraving by Jean Marot	cinema studies
collapsed roof	Connecting zones Structure	consolidations	building stones	elemental concentrations	brachiopods	gastropod fossils	claw chisel	engravings	Civil engineering
collapsed structural elements	constituent elements	construction	buttresses	elements of the work in progress le	branches	glass fragments	cohesion	fleur-de-lys	climate reconstruction
collapsed transverse	construction	construction irons	buttresses	Fir	brick	glass pieces	colour gradient	Flight into Egypt	climate science
collapsed transverse arch	construction leads	construction leads	canons	fragments of brick	building materials lead	Gothic jubé	coloured adhesive tape	French cathedrals	comparative studies
collapsed vault	construction site	cores	carpentry	galvanised steel	burnt wood	Historical Monuments	compact camera Sony RX100 Mk1	French cultural heritage	conservation
collapsed vaults	Coventry cathedral	corner	carpentry	glass	C-S-H	historical relics	computers	French Heritage Code	construction
collapsed wooden framework	decoration	cover plates	Cast lead elements	glass fragments	calcite	iconography	crane	French National Archives	construction engineering
colorimetric properties	décors	Cramps	cathedra	Green wood	calcite	isotopic data	cubicles	French traditional craft	craft
colour changes	Domestic building	cross section	Cathedral	heritage items	calcite	lead isotopes	depth gauge	Gothic architecture	critical heritage studies
compression dilatancy	double straight notches	Crypte	cathedral roof	heritage loss	Calcium	Marc Couturier's cross	detection systems	Gothic edifices	cultural heritage
condensation	dovetail notches	cubic notches	cathedral roof	heritage object	calcium silicate	medieval mortars	device	Greek cross	cultural heritage-related d
contaminated	edifice	curve of the arch	cathedral's structural elements	hydrated lime	CaO	Middle ages framework	diamond saw	Guadalupe chapel Notre-Dame	data storage
Corrosion	elements of the work in progres	cut	cathedrals	indoor dust	Carbon	monumental object	DiffractEVA® software	heritage	decision-making processe
cracking	ends	cutouts	Cell	indoor dust samples	Carbon cresting	mortar fragments	Digiprep SCP	heritage edifices	dendro-isotopy
cracking pattern	extrados	cutting forms	Cell wall	indoor dusts	carbon rod	nails	digital tablet	heritage items	Dendrochronological stud

Figure 3. Sample of the final file presenting the classification of the five thousand three hundred seventy classified terms

All data were consolidated into a master Excel database, where a final round of curation ensured the elimination of residual duplicates and the correction of any remaining classification errors. To enhance semantic granularity, a final step of sub-categorization was carried out using Mistral AI Le Chat – selected after targeted testing for its improved performance over ChatGPT 3,5 in hierarchical classification tasks. Unlike earlier stages driven by top-down categorization, this phase adopted a bottom-up approach: the model was asked to infer meaningful sub-categories from the curated term lists associated with each class. This operation significantly improved the semantic resolution of the dataset and enabled the construction of accurate, domain-specific knowledge graphs. Overall, the methodology demonstrated the potential of combining generative LLMs with controlled input structuring and iterative curation to support advanced forms of knowledge modelling in complex, interdisciplinary research contexts.

5. 2D/3D annotations in the aïoli platform

5.1 The annotation process in Aïoli

A 2D/3D annotation corpus was established on the Aïoli platform, developed by the MAP laboratory with support from the CNRS and the French Ministry of Culture. This collaborative web tool enables semantic enrichment of heritage data via photogrammetric inputs, processed either through an integrated Micmac pipeline or imported from Metashape using a conversion script (https://page.hn/s9anv5). To ensure spatial and metric consistency across projects, photogrammetric datasets are aligned with topographic references, facilitating integration with web viewers (ThreeJS, PotreeJS) and other tools within the project's digital ecosystem.

Once projects are created on the platform, the point cloud serves as a shared geometric framework for propagating and correlating 2D annotations across the entire image dataset. Users initiate the process by creating a 2D region on an image, which is then projected across the dataset and corresponding 3D point cloud using a 2D–3D–2D projection service, deployed via Docker-based web services. This process operates asynchronously and can be repeated in parallel. Annotations are structured into groups and layers, enriched with semantic descriptors (e.g., text, dates, numerical values, lists, URLs) and associated media (documents, images, videos, audio). Additionally, computed 3D descriptors provide morphological information, including basic geometry, color, normals, roughness, and ambient occlusion.

5.2 The annotation process in Aïoli

The first and most substantial dataset integrated into the Aïoli platform was the condition report produced between 2019 and 2021 by architects Philippe Villeneuve, Rémi Fromont, Pascal Prunet, and their teams. Based on orthophotographs and annotated plans, this documentation was supplemented with detailed photographs to support spatial localization and interpretation. Covering a broad spectrum of topics and areas of interest, it remains the largest annotated corpus on the platform, comprising 10,836 annotations, approximately 78% of the total (Fig. 4).

The annotation corpus was then enlarged with the integration of the studies carried out by the various working groups of the scientific worksite, such as : the catalog of lapidary signs throughout the cathedral ; the inventory of the upper stained glass windows' lead sealing, staples, and gutter walls ; stratigraphic analysis of the western façade and analysis of stone samples ; condition report and dendrochronological analysis of the timber remains ; cartographies of tool traces on the stones of the central nave's interior elevations and spatialization of sound recordings made in the cathedral.

The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Volume XLVIII-M-9-2025
30th CIPA Symposium "Heritage Conservation from Bits:
From Digital Documentation to Data-driven Heritage Conservation", 25–29 August 2025, Seoul, Republic of Korea

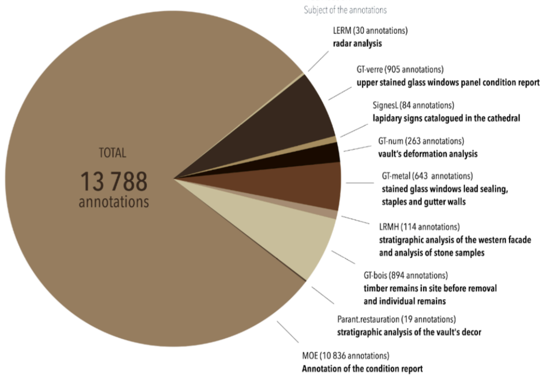

Figure 4. Graph presenting the number of annotations and subjects integrated in the aïoli platform, per aïoli account

The corpus currently comprises 13,788 annotations, distributed across nine Aïoli accounts and thematic domains, excluding datasets still undergoing integration. Annotation density varies by subject, from 19 to 894 for scientific worksite studies, and up to 10,836 for the condition reports. The scale and diversity of the corpus, along with its structured organization, provide a robust foundation for experimentation in data visualization, cross-referencing, and large-scale querying.

6. Towards the Reconciliation of Textual and Spatial Trajectories

To explore how scientific narratives and spatial annotations can be effectively interconnected, we conducted a focused case study on the chapels of Notre-Dame de Paris. This reduced-scale experiment allowed us to test methods for linking textual reasoning with annotated 2D/3D data. After presenting the two components of the corpus—diagnostic texts and spatial annotations—we detail the implementation of a computational framework designed to interrelate them. This includes the development of tools such as the smart paragraph, which enables semantic linking between descriptive passages and spatial entities.

6.1 The chapels case study: Diagnostic study texts

Following the large-scale semantic extraction performed on the Journal of Cultural Heritage, the methodological focus shifted to a new, domain-specific corpus: Volume 3 of the post-fire diagnostic study dedicated to the chapels of Notre-Dame de Paris (Villeneuve et al., 2019). Authored by Philippe Villeneuve, Rémi Fromont, and Pascal Prunet – Chief Architects for Historic Monuments – this document represents a critical component of the scientific and heritage response following the 2019 fire. The volume is structured into four main sections: 1. A historical and chronological overview of the creation and evolution of the cathedral's chapels; 2. A detailed architectural description 3. A pathological assessment of the damage caused by the fire; and 4. The projected restoration interventions (Fig. 5). The complexity and richness of this corpus made it an ideal candidate for further testing and application of the text mining pipeline previously developed.

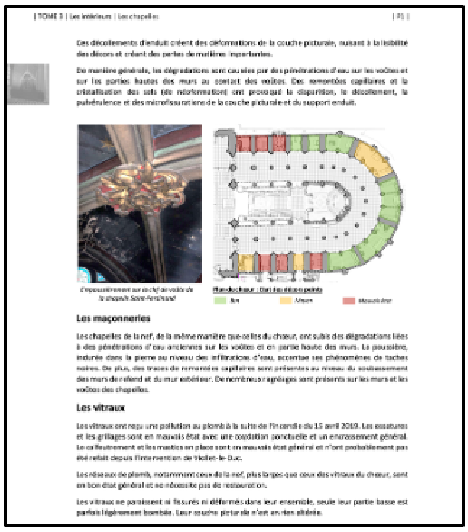
Les maçonneries

Les vitraux

Figure 5. Page forty-four of the condition report, presenting the distribution of painted decors across the chapels, and their description - Cathedral of Notre-Dame de Paris – Diagnostic Study Following the Fire of April 15, 2019. Villeneuve, P., Fromont, R., Prunet, P., Architectes en Chef des Monuments Historiques, dec 2020

Using the same multi-stage method – combining fine-tuned prompting, paragraph-level analysis, structured extraction in JSON format, and manual curation – all significant terms in the report were extracted, cleaned, and categorized. For this task, the LLM Mistral AI Le Chat was employed, selected for its high performance on French-language corpora and its ability to maintain lexical precision in highly specialized architectural discourse. The process allowed not only for the refinement of semantic categorization within a focused architectural and historical dataset but also tested the generalizability of the approach beyond peer-reviewed academic publications. This phase reinforced the relevance of LLMs for processing heterogeneous heritage documentation, provided that strict structuring, iterative refinement, and expert supervision are maintained throughout.

6.2 The chapels case study: 3D annotations

The annotations related to the condition reports of the choir and nave chapels were based on 32 dedicated Aïoli projects, documenting the state of the cathedral after the fire and prior to restoration. Each project follows a consistent structure, distinguishing a general information group—with layers concerning furnishings, statuary, earlier restorations, and painted decorations—from a diagnostics and interventions group, which includes layers addressing the condition of masonry, stained glass, painted decorations, and electrical networks.

In addition to visual annotations, each project includes a series of descriptive fields—referred to as "description sheets"—that allow for the inclusion of supplementary textual information, effectively forming a legend derived from the original reports produced by the architects. Attached files, such as historical photographs, drawings of previous decorations, and original condition reports, have been integrated to further contextualize and substantiate the observations (Fig. 6).

https://doi.org/10.5194/isprs-archives-XLVIII-M-9-2025-49-2025 |

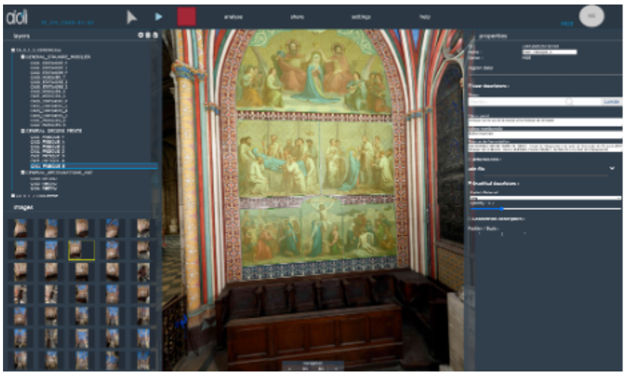

Figure 6. Semantic annotations of the general informations of the chapels painting in the Aïoli platform

6.3 Implementing a Framework for Linking Text and Space

To move from the conceptual ambition of reconciling narrative reasoning with spatial observation to a working prototype, we built a lightweight three-layer framework that acts as a semantic “bridge” between written paragraphs and 2D/3D annotations stored on the Aïoli platform. The guiding principle was separation of concerns: textual discourse should keep its rhetorical structure, and annotations should keep their metric anchorage, yet the two domains must be able to discover one another on-demand. All communication therefore travels through an intermediate knowledge layer—implemented as a Neo4j graph—that stores the vocabulary mined from both corpora and exposes it through a simple REST autocomplete service and a multi-concept search endpoint. The **Smart Paragraph** is the user-facing pivot of this architecture. Every time a scholar pastes or types a paragraph into the narrative editor (Editor JS), the Smart Paragraph block normalises the text, generates n-grams, and queries the autocomplete service against a controlled taxonomy extracted from the full Aïoli corpus (annotation labels, groups, layers, field names and values). Returned concepts are rendered as colour-coded chips inside the paragraph: solid tones for direct textual matches, faded tones for indirect suggestions, and a star icon to mark a concept as central. With a click the author can promote, demote, or highlight any chip; these decisions are stored in the block’s JSON so that semantic intent stays close to the prose itself (Fig. 7).

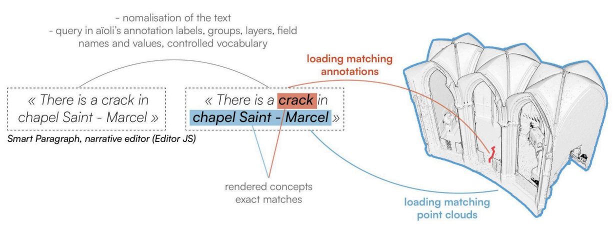

Figure 7. Results generated following the input of the sentence “There is a crack in the Saint-Marcel Chapel.” During text entry, the system automatically detects textual entities corresponding to pre-existing annotations in the database, in this case “Saint-Marcel Chapel” and “crack.” Once the query is submitted, the identified entities are rendered within the 3D viewer, enabling their precise spatial localization. Selecting an annotation triggers the display, in a lateral panel, of an associated detail image, accompanied by the full set of descriptive metadata linked to that entity

At the end of the tag extraction process, pressing **“query”** converts the active chips into a progressively relaxed Cypher query. If a central concept is present, all returned annotations include it; other concepts are matched with decreasing strictness until results are found. Up to 1,000 annotations—preserving client performance—are sent to the Potree viewer via an iframe, where their spatial footprints are immediately highlighted, allowing researchers to explore geometry, imagery, and diagnostic metadata *in situ*. The loop is reciprocal: clicking an annotation in the viewer triggers an event that scrolls relevant Smart Paragraphs into view and highlights the corresponding chips, linking what is said to where it was observed. Interoperability relies on minimalist design. All paragraph metadata—text, chip states, centrality—are stored as plain JSON in the Editor, with no external database. On the spatial side, the Aïoli viewer remains unchanged: it listens for narrativeSearchResults messages and returns annotation IDs. Since both modalities use shared concept IDs, linkage is robust, bidirectional, and easily extended to new datasets or vocabularies.

6.4 Limitations

The current prototype is admittedly exploratory. It has yet to address large-scale performance (cache layers and batch queries will be necessary), explicit temporal reasoning (time-range filters could be added to the Cypher templates), and inline authoring workflows that would let scholars create new annotations when no match exists. Nonetheless, the early results confirm that a fine-grained, live correspondence between narrative explanation and spatialised evidence is technically feasible with modest resources. By letting researchers move fluidly between the *how* and *why* of their interpretation and the *what / where / when* of their observations, the framework opens new possibilities for transparent, traceable, and richly contextualised heritage documentation.

7. Perspectives and Future Work

The prototype described above proves that *Smart Paragraphs* can already knit together text, taxonomic concepts and spatial annotations in a single query. Yet its most exciting promise lies ahead: turning a collection of isolated paragraph–annotation links into full **mnémogrammes** (2)— coherent, spatio-narrative trajectories that scholars can compose, replay and share (Fig. 8).

(2) The term *Mnémogrammes* is a contraction of *mnemo* (from the Greek *mnēmē*, meaning memory) and *graphein* (from the Greek *graphein*, meaning to write). It refers to coherent spatio-narrative trajectories that researchers can compose, replay, and share.

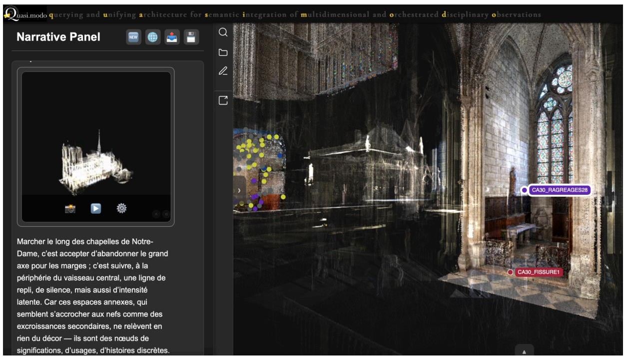

Figure 8. The mnémogramme is constructed through continuous writing, enriched by contextual integration of 3D annotations or viewpoint captures. These elements, extracted directly from the Viewer, are linked to the narrative in real time. The Smart Paragraph module enables the dynamic retrieval and spatialization of relevant annotations. Figure X displays a chapel view associated with a segment of the text focusing on the chapels of Notre-Dame.

Our immediate roadmap is therefore two-fold. First, we will extend the Smart Paragraph query pipeline so that it can retrieve not only discrete annotations but also their contextual assets: high-density point-cloud subsets, oriented photographs, and time-stamped annotations. By packaging those assets with each paragraph, a researcher will be able to “drop” multiple Smart Paragraphs into the editor and watch a composite storyline emerge, each step seamlessly blending prose, concepts and 3-D evidence inside the viewer.

Second, we are formalising the notion of a mnémogramme as a first-class object in the platform. A mnémogramme records, in order, the paragraph blocks selected by the author, the queries they generated, the fine-tuned selections of the annotations (curated manually by the user), and the camera paths or point-cloud regions activated in the 3Dviewer. Because every element is saved as JSON, the entire trajectory can be replayed, annotated further, versioned, or forked by another scholar—turning narrative construction itself into a shareable, inspectable research artefact.

Beyond these short-term goals, several research avenues remain open. Adaptive prompting and incremental concept learning could make tag detection progressively smarter as new projects enrich the taxonomy. Bidirectional editing would allow users to create or refine Aïoli annotations from within a paragraph when no satisfactory match exists, closing the documentation loop. Finally, user-centred evaluations—both within the Notre-Dame worksite and in other heritage contexts—are planned to assess how mnémogrammes support collaborative interpretation, knowledge transfer, and long-term conservation decision-making.

By evolving from single paragraph queries to articulated mnémogrammes, the framework aims to offer heritage scholars a novel workspace where narrative reasoning, semantic structure and spatial reference are no longer parallel tracks, but tightly interwoven threads of the same intellectual fabric.

Acknowledgments

This work was funded since 2019 by the CNRS and the French Ministry of Culture, within the framework of the national scientific action Notre-Dame de Paris, and since 2022 by by the European Research Council (ERC Advanced Grant nDame_Heritage : n-Dimensional analysis and memorization ecosystem for building cathedrals of knowledge in Heritage Science, Project ID: 101055423). Authors wish to acknowledge the help and collaborative support from: the chief architects of historical monuments in charge, Philippe Villeneuve, Pascal Prunet and Rémi Fromont, the Établissement public chargé de la conservation et de la restauration de la cathédrale Notre-Dame de Paris (RNDP), and the heritage conservators. The authors thank the numerous scientific partners and collaborators, with special recognition of the scientific worksite working groups, and with particular acknowledgments to our interlocutors at Archéovision, MIS, MOM, LRMH, LCE, ArScAn, IRAMAT and Héritages : Culture/s, Patrimoine/s, Création/s laboratory.

References

Bizer Ch., Vidal M.-E., Skaf-Molli H.. Linked Open Data. *Encyclopedia of Database Systems*, 2017. https://doi.org/10.1007/978-1-4899-7993-3_80603-2

Dillmann, P., Liévaux, P., De Luca, L., Magnien, A., Regert, M., The CNRS/MC Notre-Dame scientific worksite: an extraordinary interdisciplinary adventure. Journal of Cultural Heritage, 2024, SI: Notre-Dame de Paris, 65, pp.2-4. https://doi.org/10.1016/j.culher.2024.02.004

Drissi A., Khemiri A., Sassi S., Tissaoui A., Chbeir R., et al.. LDA+: An Extended LDA Model for Topic Hierarchy and Discovery. 1716, Springer Nature Singapore, pp.14-26, 2022, Communications in Computer and Information Science, https://dx.doi.org/10.1007/978-981-19-8234-7_2

Illina I., Fohr D.. BERT Semantic Context Model for Efficient Speech Recognition. *ICCAS 2022 - International Conference on Cognitive Aircraft Systems*, ISAE-SUPAERO, Jun 2022, Toulouse, France. http://dx.doi.org/10.5220/0011948200003622

Keraghel I., Morbieu S., Nadif M.. Recent Advances in Named Entity Recognition: A Comprehensive Survey and Comparative Study. 2024. https://doi.org/10.48550/arXiv.2401.10825

Doerr, M. (2003). The CIDOC Conceptual Reference Module: An Ontological Approach to Semantic Interoperability of Metadata. AI Magazine, 24(3), 75. https://doi.org/10.1609/aimag.v24i3.1720

Manuel, A., Abergel, V., Aïoli, a Reality-Based Annotation Cloud Platform for the Collaborative Documentation of Cultural Heritage Artefacts. *Un patrimoine pour l'avenir, un avenir pour le patrimoine*, 2022. http://dx.doi.org/10.1016/j.daach.2023.e00285

Néroulidis, A., Pouyet, T., Tournon, S., Abergel, V., Rousset, M. et al., A digital platform for the centralization and long-term preservation of multidisciplinary scientific data belonging to the Notre Dame de Paris scientific action. Journal of Cultural Heritage, 2024, SI: Notre-Dame de Paris, 65, pp.210-220. https://doi.org/10.1016/j.culher.2023.09.016

Peytavie A., Galin E., Grosjean J., Mérillou S., Arches: a Framework formodelling Complex Terrains. Computer Graphics Forum, 2009, 28, pp.457-467 http://dx.doi.org/10.1111/j.1467-8659.2009.01385.x

Roussel, R., De Luca, L., An approach to build a complete digital report of the Notre-Dame Cathedral after the fire, using the Aioli platform. 29th CIPA Symposium “Documenting, Understanding, Preserving Cultural Heritage: Humanities and Digital Technologies for Shaping the Future”, Jun 2023, Florence, Italy. pp.1359-1365, https://doi.org/10.5194/isprs-archives-XLVIII-M-2-2023-1359-2023

Simon R., Barker E., Isaksen L., de Soto Cañamares P., Linking early geospatial documents, one place at a time: annotation of geographic documents with Recogito, 2015, e-Perimetron. 10.2, pp.49-59 (https://oro.open.ac.uk/43613/1/Simon_et_al.pdf)

Villeneuve Ph., Fromont R., Prunet P., Cathédrale Notre-Dame de Paris : Étude de diagnostic à la suite de l’incendie du 15 avril 2019, pp.29-53