From Versioned Cities to Versioned Documents: Extending CityRDF for Modelling Urban Documentation over Time

Diego Vinasco-Alvarez[1], John Samuel[1,2], Pierre-Antoine Chiron[1], Sylvie Servigne[1] and Gilles Gesquière[1]

[1] LIRIS, UMR 5205, Universite Claude Bernard Lyon 1, Université Lumière Lyon 2, CNRS, INSA Lyon, Lyon, France
[2] CPE Lyon, Villeurbanne, France
diego.vinasco-alvarez@cnrs.fr; pierre-antoine.chiron@univ-lyon2.fr; john.samuel@cpe.fr;
sylvie.servigne@insa-lyon.fr; gilles.gesquiere@univ-lyon2.fr

Keywords: Versioned cities, versioned documents, urban documentation, temporal evolution, multimedia corpora, urban analysis, cultural heritage documentation, historical urban studies

Abstract

Cities are documented by large and heterogeneous multimedia corpora, including photographs, maps, plans, reports, videos, and web resources. While versioned urban models have been studied for representing the temporal evolution of city objects, the documents describing these objects are still often treated as static attachments. This paper investigates how versioned cities and versioned documents can be modelled together by treating urban documentation as principal evolving entities connected to evolving urban features. We propose a CityRDF extension as an OWL ontology, interoperable with CityGML 3.0 that represents documents, their qualified links to urban objects, and their associated temporal semantics within the same framework. The key idea is to separate the identity of a document from the interpretative relation that connects it to the city model, making it possible to distinguish document history, depiction time, and city-object history while preserving the version history of urban objects. The proposed model treats documents as core urban elements and links diverse, versioned records to equally versioned cities across time and scale for heritage analysis and urban exploration.

1. Introduction

Cities generate large and heterogeneous collections of multimedia documents, including photographs, maps, reports, videos, and web pages. These resources capture different facets of urban space and provide essential evidence for understanding how cities evolve over time. Previous research has explored the versioning of urban models (Chaturvedi et al., 2017, Kutzner et al., 2022, Samuel et al., 2020, Idris et al., 2024), making it possible to represent successive states of city objects such as buildings, streets, and infrastructure. By contrast, far less attention has been given to the temporal evolution of the documents themselves and to the ways in which they relate to evolving urban objects.

This gap matters because documents are not merely static attachments to a city model. They are evidence-bearing resources that may themselves evolve through digitisation, restoration, annotation, translation, republication, or metadata enrichment. Ignoring document version histories risks losing critical information about provenance, interpretation, and reuse, thereby limiting their value for urban analysis, cultural heritage documentation, and historical studies (Samuel et al., 2016, Münster et al., 2017).

In this paper, we extend the concept of versioned cities (Chaturvedi et al., 2017, Kutzner et al., 2022, Samuel et al., 2020) toward versioned urban documentation. We investigate how documents linked to urban objects can be modeled as evolving entities with their own version histories, including revisions, digitization processes, annotations, and derived representations. Each document may be associated with one or more urban objects at different spatial scales, from individual buildings to districts or entire cities, creating complex spatial and semantic relationships.

We propose a framework that connects two complementary temporal structures: the version history of city objects and the version history of documents describing or depicting them. This dual versioning model (two coupled version histories) enables the representation of partial correspondences between document versions and urban object states, acknowledging that documents may capture only fragments of urban change or represent interpretations produced at different moments.

By linking versioned documents with evolving urban objects, the proposed approach facilitates the exploration of large multimedia corpora in their spatial and temporal context. Such a representation opens new possibilities for urban analysis, cultural heritage documentation, and historical urban studies, enabling users to navigate urban environments through their associated documentation (Gaspari et al., 2023, Mota et al., 2023) while preserving the provenance and evolution of both city models and documents. Here, urban analysis refers mainly to document-supported historical reconstruction, comparison of documentary evidence across city-object states, provenance-aware inspection of model changes, and spatio-temporal retrieval of records attached to buildings or districts.

The contributions of this paper are threefold. First, we identify the need for a versioned document model that captures the evolution of urban documentation alongside the evolution of urban objects. Second, we propose a CityRDF extension (OGC Incubator, 2026) as an OWL ontology, interoperable with CityGML 3.0, that provides a conceptual framework for jointly modelling versioned cities and versioned documents. Third, we illustrate the approach through a case study showing how documentary and urban histories can be connected and queried in practice.

The remainder of the paper is organised as follows. Section 2 presents the modelling requirements, Section 3 reviews related work, Section 4 introduces the proposed versioned document model, Section 5 presents the case study, Section 6 discusses the main implications and limitations, and Section 7 concludes the paper.

2. Modelling Requirements

Research on versioned city models (Chaturvedi et al., 2017, Kutzner et al., 2022, Samuel et al., 2020, Idris et al., 2024) has mainly focused on successive states of urban objects: buildings are constructed, renovated, subdivided or demolished, and 3D city models are updated to represent these changes (Biljecki et al., 2015, Eriksson and Harrie, 2021, Vitalis et al., 2022). Urban knowledge, however, is also preserved in photographs, maps, plans, reports, videos, web pages and archival records that document the city from different viewpoints (Huffman et al., 2017, Morlighem et al., 2022, Moncla et al., 2017, di Lenardo et al., 2025).

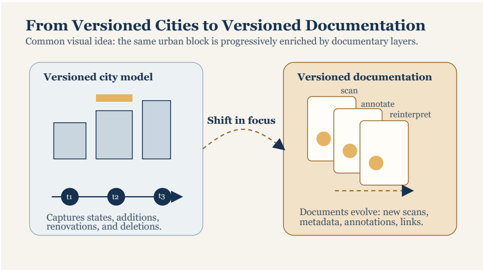

Figure 1. R1 – From versioned cities to versioned documentation: the same urban scene is no longer only represented through successive city states, but also through evolving documentary layers.

This calls for modelling the documentation about changing city objects, not only the city model states themselves. Figure 1 introduces this requirement: the urban timeline is accompanied by the timeline of the records through which it is described. These records are not static attachments. They may evolve through digitisation, restoration, annotation, translation, republication or metadata enrichment. For example, a historical image may exist as an original print, a scanned surrogate, a derivative and a semantically annotated graph resource (Matzen and Snavely, 2014, Janowicz et al., 2022, di Lenardo et al., 2025). If these transformations are ignored, provenance, interpretation and reuse are only partially represented.

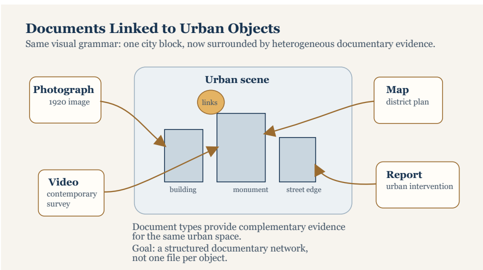

Figure 2. R2 – Heterogeneous documents linked to the same urban scene: photographs, maps, videos and reports provide complementary evidence for urban objects.

Figure 2 shows that heterogeneous records may converge on the same urban object while preserving different media, viewpoints and documentary roles. Their histories are related to, but distinct from, the lifecycles of city objects. A document produced today may describe a city state from decades earlier (Morlighem et al., 2022); a recent annotation may also reinterpret an older image using newly available urban data. These temporal layers must therefore be explicit when documents support historical reconstruction, model validation or exploration of an object's documentary context.

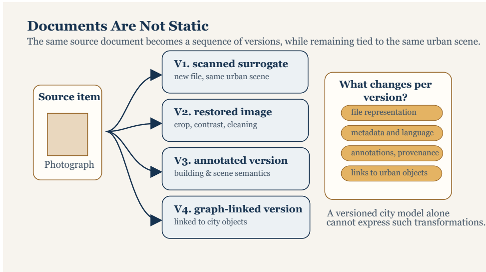

Figure 3. R3 – Documents are themselves versioned objects: scans, restorations, annotations and semantic enrichments create a documentary history distinct from the urban history being described.

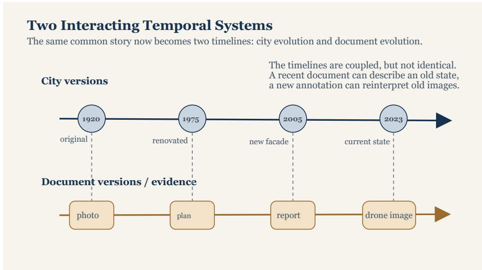

Figure 4. R4 – Two interacting temporal systems: city versions and document versions evolve separately but remain connected through evidence and interpretation.

Figures 3 and 4 make this distinction concrete: document revisions intersect only partially with the history of the objects they describe. A further requirement is multi-scale, many-to-many linking. A record may refer to a facade, building, street or district; conversely, one building may be associated with photographs, plans, reports and web resources produced at different dates and for different purposes. A simple document-to-object link is therefore insufficient. The relation must be qualified by temporal scope, interpretative purpose, provenance and, where needed, uncertainty, especially in cultural heritage and historical urban studies (Huffman et al., 2017, Matzen and Snavely, 2014).

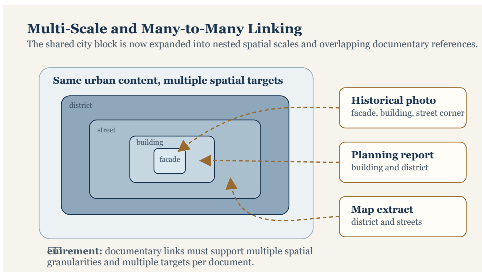

Figure 5. R5 – Documents may target multiple spatial scales simultaneously, from facade to district, which requires qualified links rather than simple attachments.

Figure 5 summarises the final requirement: documentary links must remain qualified across spatial granularities rather than attach a file to a single object in a flat way. The goal is thus to analyse urban documentation and urban-object evolution together while preserving their distinct temporalities, enabling richer navigation in documentary corpora and tighter integration between urban knowledge graphs (Janowicz et al., 2022, Ning and Liu, 2023) and documentary resources.

3. Related Work

The development of versioned documents for urban knowledge graphs lies at the intersection of four research areas: urban modelling standards and their extension mechanisms, versioning of urban models and documented cities, general-purpose version control systems, and versioning of RDF and knowledge graphs. Each of these aspects contributes part of the solution, but none directly addresses the joint evolution of urban objects and the documentary resources that describe them.

The first relevant aspect is the CityGML standard and its Application Domain Extension (ADE) mechanism. CityGML 3.0 provides a conceptual model (Kolbe et al., 2021) for representing 3D city objects and their temporal properties through the `core:AbstractFeatureWithLifespan` class, which associates features with creation and termination dates and supports multiple representational states. The standard also defines an ADE mechanism that allows domain-specific packages to extend the CityGML object model without modifying its core. Several ADEs have been developed for energy simulation, noise mapping, indoor navigation, and documentary enrichment. In particular, earlier work already proposed a CityGML 2.0-based *DocumentADE* to associate historical documents with urban objects and support their visualization in documented city models (Samuel et al., 2016). In parallel, the ongoing CityRDF initiative develops an RDF ontology for CityGML and CityJSON, providing an important bridge between 3D city modelling standards and graph-native semantic representations (OGC Incubator, 2026). CityRDF is especially relevant here because it shows how urban-model concepts can be formalised in RDF/OWL rather than remaining tied only to exchange encodings. Our proposal follows this semantic-extension trajectory, but focuses more specifically on documentary resources, qualified references, and their explicit temporal versioning.

The second aspect concerns the versioning of urban and geospatial models and the integration of documents with 3D cities. Earlier work on documented or time-aware 3D cities showed that documents can be connected to spatial objects and explored in a 3D environment (Samuel et al., 2016, Chagnaud et al., 2016, Jaillot et al., 2021, Huffman et al., 2017, Morlighem et al., 2022, Moncla et al., 2017). More generally, urban model versioning has investigated how alternative states, planning scenarios, and temporal transitions can be represented for city objects (Langran and Chrisman, 1988, Worboys, 2005, Biljecki et al., 2015, Eriksson and Harrie, 2021, Vitalis et al., 2022). In this line, UrbanCo2Fab proposed a Git-inspired perspective on concurrent viewpoints of urban fabric (Samuel et al., 2018), while more recent work such as ConVer-G addresses concurrent versioning of urban knowledge graphs (Puget Gil et al., 2024). These contributions are important because they demonstrate that urban data are not static and that version-aware data management is essential for collaboration and historical reasoning. However, even when document links are already present, the main focus remains the evolution of the urban dataset itself. Documents are generally integrated as linked resources for consultation and interpretation, not as entities with explicit internal version histories and qualified temporal references.

The third aspect comes from general-purpose version control and structured-data versioning. Software configuration management systems such as CVS and Git established the conceptual vocabulary of commits, branches, merges, and collaborative history management. Their success has inspired the adaptation of versioning principles to geospatial and tabular data, as illustrated by systems such as GeoGig (GeoGig Contributors, 2024), and DoltHub (DoltHub Inc., 2024). These systems show that version control is not restricted to source code: it is also valuable for datasets, schemas, and collaborative editing of structured information. Nevertheless, they usually operate at the level of files, rows, or repository states. They do not directly model the semantic distinction between a document, its successive representations, and its qualified references to urban objects.

The fourth aspect concerns RDF and knowledge-graph versioning, which is particularly relevant to the semantic representation targeted in this paper. Early systems such as SemVersion (Volkel et al., 2005), R&Wbase (Vander Sande et al., 2013), Quit Store (Arndt et al., 2016), and R43ples (Graube et al., 2014) explored Git-inspired or revision-oriented management of RDF graphs. Other works addressed RDF archiving and time-travel querying more explicitly, including RDF triple-store version management (Im et al., 2012), surveys of linked-data archiving systems (Papakonstantinou et al., 2016), and archive-oriented storage engines such as OSTRICH (Taelman et al., 2018) and its underlying storage model for random-access versioned querying (Taelman et al., 2019). These approaches make clear that graph versioning is a mature problem in the Semantic Web community. They provide strong foundations for storing and querying evolving RDF data, but they generally remain agnostic to the specific semantics of urban documentation.

Our work builds on these four aspects while addressing a gap between them. The CityGML ADE mechanism already provides a valid extension point for linking documents to urban objects, as illustrated by DocumentADE (Samuel et al., 2016), but documentary versioning itself remains largely unmodelled. Existing work on documented cities shows how multimedia documents can be linked to 3D urban environments, yet without an explicit account of document version histories, curatorial revisions, and temporally qualified references. Version-control systems provide collaborative history management for files and datasets, but not semantic qualification of documentary relations. RDF versioning research provides efficient graph archiving and querying, but remains agnostic to urban and documentary semantics. What is still missing is a model that treats urban documents themselves as principal evolving entities and connects their version histories to the evolving urban objects they describe, depict, or interpret (Janowicz et al., 2022, Ning and Liu, 2023). This is the space targeted by the versioned document model proposed in this paper.

The proposal is also related to archival and bibliographic description models. Dublin Core, ISAD(G), EAD, FRBRoo and BIBFRAME provide established ways to describe documentary resources, agents, editions, manifestations and archival context, while PROV-O provides a generic vocabulary for provenance chains involving digitisation, restoration or annotation activities. Our model does not replace these vocabularies. Instead, it defines the urban-spatial reference layer that connects documentary resources to CityGML-aligned features. In practice, document metadata and agent provenance can be attached through existing vocabularies, whereas `doc:Reference` captures the temporally qualified relation between a document state and an urban feature.

4. Development of the Versioned Document Model

The proposed model (Figure 6) is conceived as a CityGML 3.0–aligned RDF/OWL extension dedicated to urban documentation (Kolbe et al., 2021). The proposal is not a formal CityGML ADE with a UML package and GML encoding. It is an RDF/OWL extension of CityRDF that follows ADE principles at the conceptual level by extending CityGML-aligned classes without modifying the urban core. More precisely, it extends the conceptual layer imported through the urban core ontology and introduces a document package whose main classes are `doc:Document`, `doc:Reference`, `doc:Tag`, and `doc:PurposeTypeValue`, together with explicit temporal entities from OWL-Time (Cox and Little, 2022). In the OWL formalisation, `doc:Document` is declared as a subclass of `core:AbstractFeatureWithLifespan` (Kolbe et al., 2021). This choice is fundamental: documents are no longer treated as external attachments to the city model, but as urban knowledge objects that can participate in the same semantic and temporal ecosystem as CityGML entities. In this sense, the model responds directly to the first requirement: it shifts the focus from versioning only the city model to versioning both the city and the documentation associated with it.

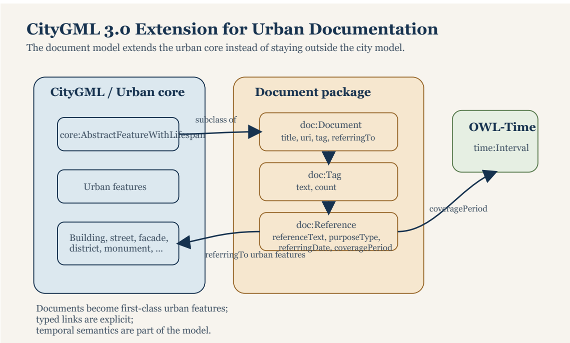

Figure 6. Proposed CityGML 3.0 documentation extension. The document package extends the urban core, introduces document-specific classes, and connects documentary semantics to both urban features and explicit temporal intervals.

The extension is intentionally lightweight. A `doc:Document` carries stable identity information through properties such as `doc:Document.title` and `doc:Document.uri`, and it may be associated with one or more tags through `doc:Document.tag`. The corresponding ontology axioms specify that document titles are strings, URIs are typed as `xsd:anyURI`, and tags are separate objects of type `doc:Tag`. An additional `rdfs:seeAlso` property can be used on a document individual to point to its provenance page (for example, a Wikimedia Commons revision URL), keeping the canonical media URI stable while providing a durable link to the curatorial event that produced or described it (R2). Urban documentation is heterogeneous: photographs, maps, reports, web pages or scans may all need to be represented under a common abstraction while preserving typeable metadata and multilingual labels. The `doc:Document` class provides that common abstraction, while the tag mechanism creates a first indexing layer for retrieval and thematic organisation.

The most important modelling decision is the introduction of `doc:Reference` as an intermediate semantic object between a document and the city model. In the ontology, `doc:Document.referringTo` does not point directly to a building, street or district; instead, it points to a `doc:Reference` (R3 and R4). A reference can carry a textual interpretation through `doc:Reference.referenceText`, a documented purpose through `doc:Reference.purpose` or `doc:Reference.purposeType`, the date on which the reference was asserted through `doc:Reference.referringDate`, and the historical period covered by the document content through `doc:Reference.coveragePeriod`. In other words, the model does not equate "this file exists" with "this file documents this urban object at this time for this reason". These are represented separately and explicitly.

This separation is what makes document versioning compatible with city versioning. A document may have several successive digital or semantic states while still referring to the same urban scene; conversely, the same urban object may be associated with different documentary states over time. By treating the document as one evolving feature and the reference as another, more interpretative layer, the model can represent two coupled but distinct temporal trajectories. The lifespan of the urban object remains handled by the CityGML-side model. The period depicted by the document is handled by `doc:Reference.coveragePeriod`. The date of curation, digitisation, alignment, or annotation is handled by `doc:Reference.referringDate`. This distinction is essential in historical and heritage workflows, where a modern digital surrogate often documents a much older urban state. It is precisely the situation highlighted by the dual-timeline illustration.

For multi-scale and many-to-many links, the reference target is `core:AbstractFeatureWithLifespan`, not a single thematic CityGML class such as `bldg:Building`. A document reference may therefore target any urban feature that participates in the CityGML 3.0 object space: a facade, a building, a street object, a district-level abstraction, or any other feature modelled in the urban graph. This range (Kolbe et al., 2021) preserves compatibility with the multi-scale organisation of CityGML rather than hard-coding the extension to one application domain.

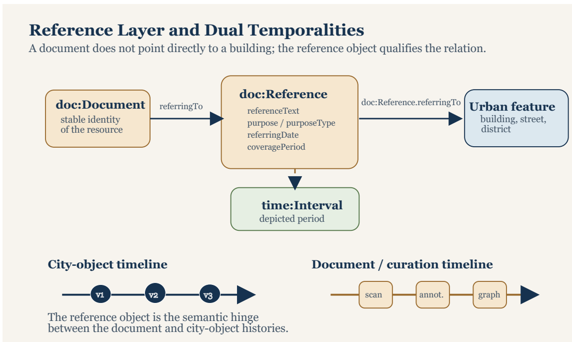

Figure 7. Reference layer of the proposed model. A document is linked to an urban feature through a qualified reference object that carries interpretation, purpose, assertion date, and historical coverage period. This layer separates city-object evolution from document-curation evolution.

The ontology further formalises this design through a set of cardinality and typing constraints. A `doc:Reference` (Figure 7) must be associated with one coverage period, one assertion date, and at least one purpose type. A `doc:Tag` must carry textual content and may carry a count. These constraints are not just implementation details: they encode minimal documentation requirements that make the resulting graph queryable and auditable. For example, requiring a coverage interval forces the modeller to distinguish the temporal scope of what is depicted from the date at which the graph statement was entered. Requiring a purpose type makes the relation interpretable as an aerial view, administrative source, renovation plan, or another controlled semantic category. The `doc:PurposeTypeValue` class is declared as a subclass of `skos:Concept`, so that purpose values can be organised into SKOS concept schemes, labelled in multiple languages, and aligned with external vocabularies for documentary typology or archival description. In this respect, the extension answers the requirement for richer qualified links rather than simple document attachments.

From a CityGML 3.0 perspective (Kolbe et al., 2021), the model should therefore be read as an ADE-like conceptual layer for documentation-aware urban knowledge. It preserves the object-centred semantics of CityGML while adding a documentary layer that is itself temporal, typed and queryable. The `doc:Document` class answers the need to represent documents as principal urban resources. The `doc:Reference` class answers the need to qualify documentary relations and to keep interpretation separate from file identity. The use of OWL-Time answers the need to distinguish depiction time from curation time. The range based on `core:AbstractFeatureWithLifespan` answers the need for multi-scale linking across urban feature types. Taken together, these choices provide a direct response to the five requirements introduced earlier: they move from city versioning to document versioning, integrate heterogeneous document types, make document histories explicit, connect documentary and urban temporalities, and support multi-scale urban referencing in a single semantic framework (Table 1).

Although the instantiated example uses a photograph, the model treats the media URI and purpose type generically; videos, maps or reports would require different metadata and preview interfaces, but not a different reference mechanism.

5. Case Study

5.1 Dataset and Urban Context

The case study uses the Gratte-Ciel district of Villeurbanne, a housing estate built in the 1930s near Lyon, France. The urban data are drawn from the UD-Demo-DocGraph repository (VCityTeam, 2026a), which reuses UD-Graph ontologies and datasets (VCityTeam, 2025) to provide a CityGML 3.0-compatible knowledge graph of the district in RDF/OWL. The proposed extension targets a documentary corpus linked to this urban graph; for compactness, the aligned sample instantiates it with one representative photograph and its building-level references.

5.2 Graph Structure

The instance graph contains four building individuals (`data:69266BO56_0` to `data:69266BO56_3`) typed as `bldg:Building`. On the documentary side, `data:document_1` represents a 1936 aerial photograph hosted on Wikimedia Commons. The document is described with two explicit `doc:Tag` individuals, one `doc:PurposeTypeValue` linked to a small SKOS concept scheme, a direct media URI, and an `rdfs:seeAlso` link to the source revision page.

The document is linked to the city model through four reference individuals (`data:reference_1` to `data:reference_4`), each connecting the image to one building. All references share the same 1936 coverage interval and store the assertion date (March 11, 2020) through `doc:Reference.referringDate`. This structure keeps the strict ontology unchanged while preserving building-level provenance.

5.3 Querying Versioned Documentary Links

A key advantage of representing urban documentation as a knowledge graph is that documentary and urban histories become jointly queryable. At the current stage of the model, this already allows queries over qualified documentary links: one can retrieve which documents depict which buildings, for which historical period, and when those links were asserted in the graph. The SPARQL query below retrieves all documents together with their linked buildings and the historical period they depict, enabling temporal navigation across the documentary corpus:

```
PREFIX doc: <.../document#>
PREFIX time: <http://www.w3.org/2006/time#>

SELECT ?doc ?title ?building ?start ?end WHERE {
  ?doc  a doc:Document ;
        doc:Document.title ?title ;
        doc:Document.referringTo ?ref .
  ?ref  doc:Reference.referringTo ?building ;
        doc:Reference.coveragePeriod ?period .
  ?period  time:hasBeginning/time:inXSDgYearMonth ?start ;
           time:hasEnd/time:inXSDgYearMonth ?end .
}
```

Applied to the Gratte-Ciel dataset, this query returns one row per building per document, recovering the four buildings covered by `data:document_1` together with the 1936 interval. With additional documents, the same pattern would support retrieval of all images depicting a given building within a chosen historical window.

The same logic extends to versioned documents. If successive states are represented as explicit document versions linked to a stable document identity, one can query both the documentary timeline and the urban timeline together. The following abstract SPARQL pattern illustrates this case:

```
PREFIX doc: <.../document#>
PREFIX ex:  <http://example.org/versioning#>
PREFIX time: <http://www.w3.org/2006/time#>

SELECT ?doc ?version ?building ?assertionDate WHERE {
 ?doc     a doc:Document .
 ?version ex:isVersionOf ?doc ;
          doc:Document.referringTo ?ref .
  ?ref    doc:Reference.referringTo ?building ;
          doc:Reference.referringDate ?assertionDate .
}
ORDER BY ?doc ?assertionDate
```

Conceptually, this query asks what each document version referred to and when that interpretation was asserted. It supports comparisons between original scans, restored images, and later annotations through three anchors: stable document identity, an explicit reference layer, and separate depiction and curation times.

In the aligned sample, tags are represented by explicit `doc:Tag` individuals, the documentary purpose by a `doc:PurposeTypeValue` in a SKOS scheme, and each building-level link by its own `doc:Reference`. The result is a compact pattern that preserves qualified references while leaving room for future revision, provenance, and uncertainty annotations.

https://doi.org/10.5194/isprs-annals-XII-4-W1-2026-301-2026 | © Author(s) 2026. CC BY 4.0 License.

Requirement	Model construct	Response
R1 – From versioned cities to versioned documents	`doc:Document` as subclass of `core:AbstractFeatureWithLifespan`	`data:document_1` is modelled as a key urban feature rather than as an external file attachment.
R2 – Heterogeneous document resources	`doc:Document`, `doc:Tag`, `doc:PurposeTypeValue`	The sample photograph combines URI, title, explicit tags, and a typed aerial-view concept within one common documentary structure.
R3 – Version histories of documents	Separation between `doc:Document` and `doc:Reference`	`data:document_1` remains stable while its four references provide an evolvable layer for interpretation, annotation, and linkage.
R4 – Interaction between city and document timelines	`doc:Reference.referringDate` and `doc:Reference.coveragePeriod` (Cox and Little, 2022)	The sample distinguishes the depicted 1936 interval from the March 11, 2020 assertion date, enabling dual temporal analysis.
R5 – Multi-scale and many-to-many linking	`doc:Reference.referringTo` with range `core:AbstractFeatureWithLifespan` (Kolbe et al., 2021)	Four building-specific references show how one document can point to several urban features without losing atomic provenance.

Table 1. How the proposed CityGML 3.0 document extension responds to the five requirements (R1–R5) through specific modelling constructs.

5.4 Demonstration

The proposed model is illustrated through an interactive web demonstration distributed through the same repository (Figure 8). It combines the Gratte-Ciel 3D city model, the documentary graph, and a semantic query interface in a Docker-based application with a UD-Viz frontend from UD-SV (VCityTeam, 2026b) and a Blazegraph RDF store served through SPARQL.

The demonstration supports two complementary interactions. First, users select a building in the 3D city model and retrieve documents whose references point to that building within a chosen temporal range. Figure 8 shows two such states, differing in highlighted documents and camera position. Second, a workspace module visualises the RDF graph itself, making the connections between `doc:Document`, `doc:Reference`, the buildings, and the OWL-Time interval directly inspectable. The paper illustrates this with one aligned photograph, but the same query pattern and interface are designed for a larger documentary corpus. The graph view also makes the dual temporal structure visible by showing the 1936 coverage period alongside the 2020 assertion date.

The demonstration confirms that the model is not only formally expressible in OWL but also operational as a browser-accessible, queryable knowledge graph linked to a 3D urban scene.

6. Discussion

The proposed model shows that extending CityGML 3.0 with a document layer is a viable way to treat urban documentation as evolving knowledge rather than as static attachments. Its main contribution is the separation between document identity and the interpretative reference that links a document to urban features, allowing documents, links, and city objects to coexist in one semantic framework while preserving their different temporalities.

The Gratte-Ciel example confirms that the strict ontology can be instantiated without ad hoc exceptions. Although the aligned sample is minimal, the model is not restricted to a single image: the same SPARQL pattern scales to photographs, maps, drawings, and other documentary resources linked to the same urban objects. The distinction between the 1936 coverage interval and the March 11, 2020 assertion date shows why qualified references are preferable to flat attachments, while the sample also shows that controlled vocabularies, atomic references, and correct temporal typing remain compatible with compact, queryable RDF.

This opens several research directions. The proposed conceptual model should be integrated with versioned graph backends such as those surveyed in the Related Work to support longitudinal querying over urban and documentary histories at corpus scale, including finer-grained exploration of city states, document states, and the transactions that connect their revisions. Finally, the usability of the model should be assessed in real interfaces where historians, planners, or heritage experts interact with city objects through their associated documentary corpora.

7. Conclusion

This paper addressed the problem of representing urban documents as evolving entities connected to evolving city objects. Starting from the observation that versioned cities and documented cities have often been studied separately, we proposed a CityGML 3.0-aligned RDF/OWL extension that introduces a document package centred on `doc:Document`, `doc:Reference`, `doc:Tag`, and `doc:PurposeTypeValue`, together with OWL-Time intervals. The key idea is to represent documentary links through qualified reference objects rather than through direct attachments, thereby separating the identity of a document from the semantics, temporality, and provenance of its relation to urban features.

Using the Gratte-Ciel dataset, we showed that this model can represent documentary resources, their building-level references, their coverage intervals, and their curation dates within the same semantic graph and make them jointly queryable via SPARQL. The paper illustrates this with one 1936 aerial photograph linked to four building individuals, but the contribution

(a) View centred on the SPARQL widget and the 3D scene, with one set of highlighted documents and a first point of view over the study area.

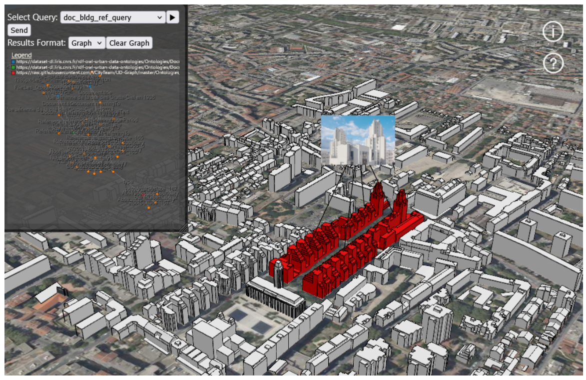

(b) Companion view of the same interface showing a different document selection and an alternative point of view over the 3D buildings.

Figure 8. Interactive demonstration of the proposed model. Both screenshots show the SPARQL widget coupled with the 3D building view, while differing in the highlighted documents, the camera position over the zone and time period (1932 and 2020).

targets documentary collections rather than an isolated image. We also showed that the model responds directly to the five requirements (R1–R5): it promotes documents to principal urban features, supports heterogeneous documentary resources, creates a basis for document version histories, separates city and document temporalities, and enables multi-scale documentary links. The enhanced sample data further shows that the strict ontology can be instantiated with structured tags, controlled documentary purposes, atomic references, and correct OWL-Time typing in a single coherent example.

The main result of this work is therefore a conceptual bridge from versioned cities to versioned documents. Rather than treating multimedia documentation as contextual metadata around a city model, the proposed extension makes it part of the urban information space itself. Future work will need to expand the aligned dataset, develop richer SKOS-based controlled vocabularies for documentary purposes, and integrate the conceptual layer into a complete operational framework for documentary revision management, concurrent graph versioning, and detailed exploration of city and document versions together with their version transactions. A further direction is to examine how large language models could assist metadata extraction, document summarisation and candidate link generation, while keeping the asserted `doc:Reference` relations auditable and provenance-qualified.

Acknowledgements

The authors would like to thank the University Lumière Lyon 2 and INSA Lyon for funding respectively the Ph.D. Thesis and the post-doctoral position of Diego Vinasco-Alvarez. This work has been done within the LIRIS laboratory as a part of the VCity project (VCity Project-LIRIS, 2026).

References

Arndt, N., Radtke, N., Martin, M., 2016. Distributed collaboration on rdf datasets using git: Towards the quit store. *Proceedings of the 12th International Conference on Semantic Systems*, 25–32.

Biljecki, F., Stoter, J., Ledoux, H., Zlatanova, S., Çöltekin, A., 2015. Applications of 3D City Models: State of the Art Review. *ISPRS International Journal of Geo-Information*, 4(4), 2842–2889.

Chagnaud, C., Samuel, J. S., Servigne, S., Gesquière, G., 2016. Visualization of Documented 3D Cities. *Proceedings*, Proceedings, Liège, Belgium, 87–93.

Chaturvedi, K., Smyth, C. S., Gesquière, G., Kutzner, T., Kolbe, T. H., 2017. Managing versions and history within semantic 3d city models for the next generation of citygml. A. Abdul-Rahman (ed.), *Advances in 3D Geoinformation*, Springer International Publishing, Cham, 191–206.

Cox, S. J. D., Little, C. C., 2022. Time Ontology in OWL. *W3C Recommendation*. https://www.w3.org/TR/owl-time/.

di Lenardo, I., Vaienti, B., Guhennec, P., Ubelmann, Y., Bianchi, I., Kaplan, F., 2025. Reconstructing Saint Mark's Square in Venice: A 4D Model with Point Cloud Integration for Analyzing Historiographical Hypotheses. *Heritage*, 8(2). https://www.mdpi.com/2571-9408/8/2/75.

DoltHub Inc., 2024. DoltHub: Git for data. Software project. `https://www.dolthub.com/`.

Eriksson, H., Harrie, L., 2021. Versioning of 3D City Models for Municipality Applications: Needs, Obstacles and Recommendations. *ISPRS International Journal of Geo-Information*, 10(2), 55.

Gaspari, F., Ioli, F., Barbieri, F., Rivieri, C., Dondi, M., Pinto, L., 2023. Rediscovering cultural heritage sites by interactive 3D exploration: A practical review of open-source WebGL tools. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-M-2-2023, 661–668. https://isprs-archives.copernicus.org/articles/XLVIII-M-2-2023/661/2023/.

GeoGig Contributors, 2024. GeoGig: Distributed version control for geospatial data. Software project. `https://geogig.org/`.

Graube, M., Hensel, S., Urbas, L., 2014. R43ples: Revisions for triples. *Proceedings of the 1st Workshop on Linked Data Quality co-located with 10th International Conference on Semantic Systems (SEMANTiCS 2014)*.

Huffman, K., Giordano, A., Bruzelius, C. (eds), 2017. *Visualizing Venice: Mapping and Modeling Time and Change in a City*. Routledge Research in Digital Humanities, Routledge.

Idris, N. H., Azri, S., Ujang, U., 2024. Reviewing the versioning in 3D city model to track building history and progress. *IOP Conference Series: Earth and Environmental Science*, 1412(1), 012035. https://iopscience.iop.org/article/10.1088/1755-1315/1412/1/012035.

Im, D.-H., Lee, S.-W., Kim, H.-J., 2012. A version management framework for RDF triple stores. *International Journal of Software Engineering and Knowledge Engineering*, 22(01), 85–106.

Jaillot, V., Rigolle, V., Servigne, S., Samuel, J. S., Gesquière, G., 2021. Integrating multimedia documents and time-evolving 3D city models for web visualization and navigation. *Transactions in GIS*, 25(3), 1419–1438. https://hal.science/hal-03178005. Publisher: Wiley.

Janowicz, K., Hitzler, P., Li, W., Rehberger, D., Schildhauer, M., Zhu, R., Shimizu, C. et al., 2022. Know, Know Where, KnowWhereGraph: A Densely Connected, Cross-Domain Knowledge Graph and Geo-Enrichment Service Stack for Applications in Environmental Intelligence. *AI Magazine*, 43(1), 30–39.

Kolbe, T. H., Kutzner, T., Smyth, C. S., Nagel, C., Roensdorf, C., Heazel, C., 2021. OGC city geography markup language (CityGML) part 1: Conceptual model standard. Technical Report OGC 20-010, Open Geospatial Consortium.

Kutzner, T., Chaturvedi, K., Kolbe, T. H., 2022. CityGML 3.0: New Functions Open Up New Applications. *PFG – Journal of Photogrammetry, Remote Sensing and Geoinformation Science*, 88(1), pp.43–61.

Langran, G., Chrisman, N. R., 1988. A Framework for Temporal Geographic Information. *Cartographica: The International Journal for Geographic Information and Geovisualization*, 25(3), 1–14.

Matzen, K., Snavely, N., 2014. Scene chronology. *Proceedings of the European Conference on Computer Vision (ECCV)*, Lecture Notes in Computer Science, 8592, Springer, 615–630.

Moncla, L., Gaio, M., Joliveau, T., Lay, Y.-F. L., 2017. Automated geoparsing of paris street names in 19th century novels. *Proceedings of the 1st ACM SIGSPATIAL Workshop on Geospatial Humanities*, GeoHumanities '17, Association for Computing Machinery, New York, NY, USA, 1–8.

Morlighem, C., Ledoux, H., Commandeur, T., 2022. Reconstructing Historical 3D City Models. *Urban Informatics*, 1(1), 11.

Mota, R., Ferreira, N., Silva, J. D., Horga, M., Lage, M., Ceferino, L., Alim, U., Sharlin, E., Miranda, F., 2023. A Comparison of Spatiotemporal Visualizations for 3D Urban Analytics. *IEEE Transactions on Visualization and Computer Graphics*, 29(1), 1277-1287.

Münster, S., Friedrichs, K., Kröber, C., Bruschke, J., Henze, F., Maiwald, F., Niebling, F., 2017. Urban history in 4 dimensions–Supporting research and education. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLII-2/W5, 525–530. https://isprs-archives.copernicus.org/articles/XLII-2-W5/525/2017/.

Ning, Y., Liu, Y., 2023. UrbanKG: An Urban Knowledge Graph System. *ACM Transactions on Intelligent Systems and Technology*, 14(4).

OGC Incubator, 2026. CityRDF. GitHub repository. `https://github.com/ogcincubator/cityrdf`.

Papakonstantinou, V., Flouris, G., Fundulaki, I., Stefanidis, K., Roussakis, G., 2016. Versioning for Linked Data: Archiving Systems and Benchmarks. *BLINK@ ISWC*, 1700.

Puget Gil, J., Coquery, E., Samuel, J., Gesquière, G., 2024. ConVer-G: Concurrent versioning of knowledge graphs. *40ème conférence sur la gestion des données (BDA)*, Orléans, France.

Samuel, J., Périnaud, C., Servigne, S., Gay, G., Gesquière, G., 2016. Representation and Visualization of Urban Fabric through Historical Documents. C. E. Catalano, L. D. Luca (eds), *Eurographics Workshop on Graphics and Cultural Heritage*, The Eurographics Association.

Samuel, J., Servigne, S., Gesquiere, G., 2018. UrbanCo2Fab: Comprehension of Concurrent Viewpoints of Urban Fabric Based on Git. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 4, 65–72.

Samuel, J., Servigne, S., Gesquière, G., 2020. Representation of concurrent points of view of urban changes for city models. *Journal of Geographical Systems*. http://link.springer.com/10.1007/s10109-020-00319-1.

Taelman, R., Vander Sande, M., Van Herwegen, J., Mannens, E., Verborgh, R., 2019. Triple storage for random-access versioned querying of RDF archives. *Journal of Web Semantics*, 54, 4–28.

Taelman, R., Vander Sande, M., Verborgh, R., 2018. Ostrich: versioned random-access triple store. *Companion Proceedings of the The Web Conference 2018*, 127–130.

Vander Sande, M., Colpaert, P., Verborgh, R., Coppens, S., Mannens, E., Van de Walle, R., 2013. R&Wbase: git for triples. *LDOW*, 996.

VCity Project-LIRIS, 2026. VCity. Research project website. `https://projet.liris.cnrs.fr/vcity/`.

VCityTeam, 2025. UD-Graph. GitHub repository. `https://github.com/VCityTeam/UD-Graph`.

VCityTeam, 2026a. UD-Demo-DocGraph. GitHub repository. `https://github.com/VCityTeam/UD-Demo-DocGraph`.

VCityTeam, 2026b. UD-SV. Software project. `https://github.com/VCityTeam/UD-SV`.

Vitalis, S., Arroyo Ohori, K., Stoter, J. E., 2022. Applying versioning to multi-lod 3d city models. *ISPRS International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-4/W4-2022, 177–182.

Volkel, M., Winkler, W., Sure, Y., Kruk, S. R., Synak, M., 2005. Semversion: A versioning system for rdf and ontologies. *Proc. of ESWC*.

Worboys, M. F., 2005. Event-Oriented Approaches to Geographic Phenomena. *International Journal of Geographical Information Science*, 19(1), 1–28.