Ontological Patterns for Modeling the Validity of Spatiotemporal Statements

Nicola Carboni[1,*]

[1] University of Geneva, 1, Rue De-Candolle 1205 Geneva, Switzerland

Abstract

The paper tackles the challenges of modeling the temporal validity of artist-recorded locations documented in exhibition catalogs by proposing three different ontological patterns for recording spatiotemporal-constrained assertions in a Knowledge Graph. Exhibition catalogs often document the address used by an artist participating in exhibitions. The analysis of this information offers significant insights into the spatial and social dimensions of artistic communities. However, the time-bound nature of the information makes its modeling particularly challenging, as it requires a framework that can effectively represent both the spatial and temporal dimensions of the relation. The connection between artist and their address, in fact, is always a fluent, a relationships that only hold within specific time intervals. To overcome this challenge, the study proposes and compares three ontological patterns designed for formalizing the dynamic aspects of spatial, temporal, and conceptual characteristics of an entity. The three approaches, based on RDF-star, CIDOC-CRM, and 4D ontology, are evaluated based on their ability to accurately represent the diachronic nature of addresses, formalizing (whether implicitly or explicitly) the validity of spatiotemporal assertions, and facilitating data integration and analysis.

Keywords

Ontological Patterns, Knowledge Graphs, CIDOC-CRM, Fluents, Spacetime, Temporal Validity, Digital Humanities, Digital History, Digital Art History, Cultural Heritage, Exhibition Data

CEUR-WS.org/Vol-3809/paper5.pdf

1. Introduction

With the increasing availability of digital museum resources researchers can now potentially explore the circulation of art and artists across the globe. Catalogues raisonnés and exhibition catalogs provide crucial information on what has been produced by an artist as well as where and when it was exhibited. Studying exhibitions means examining the history of forms through the circulation of ideas, images, and artworks. They are a primary medium for the construction [1] and reinforcement [2] of the artistic canon. Datafication of exhibition information opens the door to the use of novel methodologies (e.g., distant reading) to investigate the phenomenon at scale, specifically in relation to important topics such as the geography of art, or the development of the artistic communities. For such a study to take place, however, it is important to integrate multiple datasets, which still present some ontological challenges, specifically with respect to the modeling of the temporal and spatial dimensions. Interconnecting scattered collections through a shared conceptualization is, therefore, a matter of urgency, specifically given the rise of data-driven studies in fields such as Digital Art History and Cultural Analytics [3, 4].

This paper tackles the challenge of accurately modeling the temporal validity of artist locations recorded in exhibition catalogs, where multiple addresses for the same artist are often documented over time. Analyzing this data can reveal the extent to which the lives of these artists were interconnected, offering insights into the social and urban dimensions of the community. However, it also raises the question of how to model the temporal validity of spatial and diachronic information, such as an address. This paper seeks to answer this question by developing ontological patterns that can effectively model the temporal validity of artist-recorded addresses in a way that supports both detailed historical analysis and large-scale data integration. To develop the answer, the contribution

---

SWODCH'24: International Workshop on Semantic Web and Ontology Design for Cultural Heritage, October 30–31, 2024, Tours, France

[*] Corresponding author.

nicola.carboni@unige.ch (N. Carboni)
0000-0003-4912-4947 (N. Carboni)

© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

begins with an initial literature review (section 2), followed by an examination of the use case and its information requirements (section 3, 4). Section 5 then presents three potential ontological patterns for modeling the temporal validity of address. Section 5.2 examines how RDF-star can help express the temporal validity of rdf statements. Section 5.3 introduces a CIDOC-CRM spatiotemporal pattern (3D+1 modeling) for bounding the documentation of spatial features in time. Section 5.4 presents a pattern that instantiates each historical trace of the address as a distinct spatiotemporal region (4D modeling). An analysis of the advantages/disadvantages of these patterns is provided in section 6.

2. Related Work

The annotation of temporal dimensions in RDF has been extensively studied. A recent survey [5] offers a comprehensive review. A series of articles have been published on the use of CIDOC-CRM, and its extension CRMgeo, for recording spatial regions [6, 7, 8, 9]. The modeling of time-based maps has been discussed in [10]. The authors used a combination of the web annotation model, together with GeoSPARQL and OWL-Time to annotate textual sources about Australian explorations with geographical and temporal information. The problem of describing spatial diachronic entities has been addressed by the SONADUS project [11]. They employ BFO for spatiotemporally structure a dataset about administrative units of Switzerland, tracking geographical changes and assigning new identity and geometry over any documented change. Another important research that has touched upon the topic of this work has been developed as part of the FinnONTOproject. The work of Hyvönen et al [12] propose a 4D ontology for the annotation of spatio-temporal regions, where each spatiotemporal slice representing the region is recorded as an individual. The collection of spatiotemporal slices defines the identity of a region. The ontology has been used to describe the municipalities of Finland in the period 1865–2010. More recently, the French project Social Dynamics in Urban Context has created a geohistorical knowledge graph based on named entities extracted from 19th-century Parisian trade directories [13]. The project uses the LOCN ontology[1] to annotate each address, and PAV[2] to link each statement to the trade directory that documents it. Charles and Hernandez [14] have recently proposed an ontology for representing hierarchical historical territories (HHT) and their evolution, division, and changes.

3. Use Case

Exhibition data can be modeled theoretically (*top-down*) or by analyzing how they are described in existing information structures (*bottom-up*) [15]. The latter is more effective because it adheres to the concreteness of the data and encourages a culture of reuse. In line with this approach, selecting a dataset that captures a wide variety of exhibition information was crucial. With this objective in mind, BasArt—a comprehensive, collaborative database of exhibition catalogs—was chosen as the guiding model. Developed under the aegis of the Artl@s project[3] [16], BasArt provide researchers access to a wealth of information. At the time of this writing, BasArt documents 5653 exhibitions from the 19th century to the present day. Exhibition information is extracted from digitized catalogs using semi-automatic methods, and subsequently datified, normalized, and inserted in a SQL database. Every record is registered in the base using a comprehensive set of descriptors including temporal information about the exhibition, artwork exhibited and by whom, as well as many details of the artists involved. The recorded data focus mainly on four entities, the exhibition itself, its participants (artists), the artwork exhibited and the source used for the description (catalog). Given its global focus and the presence of diverse types of historical and contemporary exhibitions, BasArt has been the perfect use case to develop for model integrating and examining exhibition data. Due to its popularity within memory institutions, the data has been modeled using CIDOC-CRM, a bottom-up, event-centric, domain

---

[1] https://semiceu.github.io/Core-Location-Vocabulary/

[2] https://pav-ontology.github.io/pav/

[3] https://artlas.huma-num.fr

ontology created under the aegis of ICOM (International Council of Museums). CRM aims to provide a way to formalize statements about human activities, and specifically their interactions with heritage objects and practices [17]. It purposefully uses a neutral formalism, in order to be implemented using diverse encoding (e.g., database schemas, rdf, neo4j schemas). CRM is actively developed through the CRM-SIG (Special Interests Group), and the last version available is 7.3. Aside from the core classes and properties, CRM has been extended to cover multiple domains. Particularly useful to the context of this paper is the extension developed for documentation of spatiotemporal phenomena [18].

4. Context

Due to space constraints and the complexity of the subject, this paper focuses on a specific aspect of modeling, expressing the temporal validity of diachronic spatio-temporal features. Specifically, we will focus on one, seemingly simple, ontological pattern: the address. Exhibition catalogs offer valuable historical data about the artistic community, detailing who participated in an exhibition, when it took place, which artworks were displayed and, most importantly, where the artist may have resided at the time (see example in Fig. 1). These catalogs serve as historical records, indicating that within a specific period, an artist was associated with a particular address. However, the context of this address remains unclear—it could represent the artist's studio, residence (often the same as the studio), or simply a correspondence address. Despite this ambiguity, each artist's entry gives us key data on the spatial and temporal distribution of the artistic community within a city.

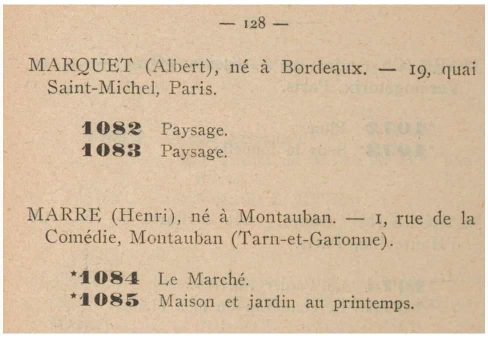

— 128 —

MARQUET (Albert), né à Bordeaux. — 19, quai Saint-Michel, Paris.

1082 Paysage.
1083 Paysage.

MARRE (Henri), né à Montauban. — 1, rue de la Comédie, Montauban (Tarn-et-Garonne).

***1084** Le Marché.
***1085** Maison et jardin au printemps.

Figure 1: Société des artistes indépendants. 25, Catalogue de la 25e exposition 1909.source Gallica, BNF

Nevertheless, this information has a significant limitation, as the addresses recorded are valid only at the time of the exhibition's opening. As addresses change over time, a later exhibition might associate the artist with a different address. To better understand the problem we can use the example (**EX1**) of the artist Albert Marquet. The famous French painter has participated in a large number of exhibitions, and by studying exhibition catalogs we are able to retrace the addresses he used over the years. In the catalog of the 19e Société des Indépendants, published in 1903, he provided as address '62, rue Bargue, Paris'. The next year, in 1904, the catalog of the 20e Société des Indépendants revealed that his address changed to '211 bis avenue de Versailles, Paris'. According to the catalog of the 23e Société des Indépendants, held in 1907, his address changed again, this time in '29, place Dauphine, Paris'. Finally, in the catalog of the 24e Société des Indépendants, held in 1908, Albert Marquet's address was '19 Quai Saint-Michel, Paris'. If some catalogs are rich in details, others do not report any address information. It is the case of the Katalog der Fünfzehnten Ausstellung der Berliner Secession, Berlin in 1908 and in the International Exhibition of Modern Art Association (Armory Show) in 1913. Marquet appears in both catalogs but without any indications regarding his current residence, although we know from the catalog of the 24e Société des Indépendants that in 1908 he was living in 19 Quai Saint-Michel in Paris. Simply recording in a database the list of addresses used by an artist during their career is surely useful, but will provide us with quite limited insights into their activities, as we would not be able to explore the temporal dimension of the artist's movement. The key information to record is when each address was in use by an artist. Although this information is not explicitly stated in the catalog's artist entry, we can infer it from the exhibition itself. Each entry, in fact, can only be considered accurate and valid within the timeframe of the exhibition, as we cannot determine if the presence of an artist at an address is true or false outside of their initial statement. Hence, the relationship between address and artist is temporal, it exists in time. To complicate the matter further, the address itself is also a diachronic entity, it comes into existence and eventually ceases to exist. Streets can be reclassified, or become part of a new municipality or borough. For instance, (**EX2**) the artist Théodore Béthume, according to the 'Salon des Artistes Français' of April 1852, lived in 'rue de La Villette, 4' in Belleville, a municipality that only a few years later would become part of Paris. Another artist, (**EX3**) Simone Perrin, in the catalog of the '39e exposition de la société des artistes indépendants' in 1928 indicated the address '38, avenue du Parc-Montsouris'. While it was its correct name at the time, it was later renamed 'Avenue René-Coty'. Given this level of information richness and complexity, it is essential to develop a model capable of sustaining information about historical changes over physical entities, applicable to the many contexts that work with temporally constrained facts, specifically facts about dynamic physical and spatial entities.

5. Modeling the Validity of Spatiotemporal Statements

5.1. Requirements and Objectives

As described in the section above, addresses are complex entities that are inherently paired with the person who inhabits them, as they are used not as mere identifiers of a place, but to identify the person who lives in that place. For such reason, they are always a fluent, i.e. a relationship between a person and place that holds for a specific time interval [19]. A model for documenting addresses should be structured in order to be capable of formalizing statements about a person's temporary presence within a specific spatiotemporal segment. Consequently, we should be able to document the temporary presence E_{x} of an artist A_{x} at a specific location P_{x} during the time $(T_0 - T_1)$, recording both the location's current and/or historical name of P_{x} $(N_{\mathrm{x}}(P_{\mathrm{x}}, T_{0\ldots\mathrm{n}}))$ as needed. Since we are working with historical sources (S_{x}), it is crucial to reference the origin of the statement, in this case, the catalog $(C_{\mathrm{x}}(EX))$ that documents the exhibition $(EX_{\mathrm{t}_{0..\mathrm{n}}})$.

With these requirements in mind, we can establish three challenges: (**CH1**) Formalize the temporal validity of an address provided by an artist in an exhibition catalog; (**CH2**) Formalize the different names given to an address in time; (**CH3**) Facilitate the extraction of a list of artists who indicated the same address in the same time period. We will examine three ontological patterns representing three different methods to tackle these challenges. Section 5.2 uses binary relationships to express temporal validity with the aid of RDF-star; section 5.3 presents a partial endurantist modeling, relying on CIDOC-CRM and n-ary construct; section 5.4 present a full endurantist model applied to CRM.

5.2. LOCN and Tempo

It is possible to encode the address and its temporal validity using two ontologies, LOCN and TempO. TempO or Temporal Ontology[4], formalizes a set of classes and properties for documenting valid time and decision time for bitemporal and tritemporal databases. LOCN is an ontology for the description of the fundamental characteristics of a location. It is part of the Core Vocabularies defined in the context of the European Union. By combining TempO and LOCN we can express the temporal validity of a recorded address. However, it is important to note that what is considered valid is not the address itself, but the relationship between a person and that address. This type of statement requires more than two predicates $(\mathit{Valid}(\mathit{Person}, \mathit{Address}, \mathit{Time}))$, and since RDF 1.1 does not support the representation of higher arity predicates, we must define a method to express this statement using one of the various available approaches [20, 21, 22, 23]. Among the different methods available, the one chosen is RDF-star, as it is (although differently) largely supported by rdf graph databases, it has been chosen as a base for RDF 1.2 and, it provides a less cumbersome approach to reification as well as an

---

[4] http://purl.org/tempo

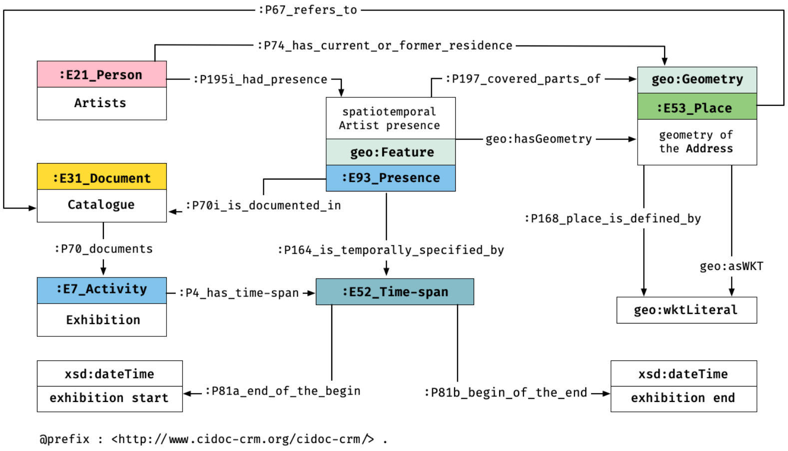

Figure 2: The diagram illustrates the modeling of an artist's temporary address. The model demonstrates how a spacetime segment can be used to map the artist's presence at the address in both space and time, moreover connecting this statement with the relevant source(s). The colors indicate the domain of each class within the CRM framework. Prefixes used are listed on the image or in prefix.cc.

easy-to-understand serialization. RDF-star provides a compact alternative to standard RDF reification, extending the model with a new construct, quotedTriple. A quoted triple is a triple used as the subject or object of another triple. Therefore, an artist's address can be represented with a traditional triple `<s><p><o>`. Using RDF-star we use such triple as the subject of (quotedTriple) of other triples. By doing so, we can further specify the validity of the original statement using a predicate from TempO, such as `tempo:validFrom` or `tempo:validTill` (see example in Listing 1). Using the same approach we can assign a temporal frame to street names and toponyms[5]. This approach captures the time-bound nature of spatial features.

```
<< <ex:Marquet> cv:domicile <ex:addressMarquet> >> tempo:validFrom
    ”1903-04-01”^^xsd:date >> cito:isDocumentedBy bas:catalog894 .
<ex:addressMarquet> locn:fullAddress ”62, rue Bargue, Paris” .
```

Listing 1: Example of RDF-star. The statements constrain the validity of a Marquet's address to a time period. Prefixes used are listed on prefix.cc.

Using LOCN and TempO in combination with RDF-star we can express the temporal validity of an address (**CH1**), the various names it has been given over time (**CH2**), and, using SPARQL-star, retrieve information about artists who have shared the same address (**CH3**). From a modeling perspective, RDF-star offers tremendous potential for streamlining the expression of historical changes in an easy-to-understand manner.

5.3. Fluents and CIDOC-CRM

Starting from version 5.1.2, released in October 2013, CRM introduced the concept of `E92 Spacetime Volume`, shifting from a fully endurantist perspective to a partial perdurantist one. The change made it possible to link physical entities with spatiotemporal volumes. Using these classes and properties we can formalize statements about a person's temporary presence within a specific spatiotemporal segment. The proposed model (Fig. 2) showcases how we can use CRM to document the temporal presence of an artist within a specific address. To understand the possibilities and limitations of the model, and its possible uses for the formalization of dynamic spatio-temporal features, we can examine closely the choices about its diverse components.

Spatio-temporal segment(E_x). The spatio-temporal segment representing the person using an address is encoded as an instance of the class `crm:E93_Presence`. This class formalizes a snapshot of

---

[5] A full example is given in https://gist.github.com/ncarboni/ba04250bbc9a2ed0c5953a156d5ccdee

Table 1
Places and encoded Place Serialization Mappings between CRMgeo and the GeoSPARQL Ontology. Prefixes used are listed on prefix.cc.

Element(from)	Relation	Element(to)
crmgeo:SP2_Phenomenal_Place	rdfs :subClassOf	crm:E53_Place
crmgeo:SP2_Phenomenal_Place	rdfs :subClassOf	geo:Feature
crmgeo:SP6_Declarative_Place	rdfs :subClassOf	crm:E53_Place
crmgeo:SP6_Declarative_Place	rdfs :subClassOf	geo:Geometry
crmgeo:SP15_Geometry	rdfs :subClassOf	geo:Geometry
crmgeo:gmlLiteral	rdfs :subClassOf	crmgeo:SP5_Geometric_Place_Expression
crmgeo:wktLiteral	rdfs :subClassOf	crmgeo:SP5_Geometric_Place_Expression

a `crm:E92_Spacetime_volume`, a restricted part of a larger volume chosen to determine and record the extent of a phenomenon. In the modeling, instances of `crm:E93_Presence` are linked to both spatial and temporal projections, respectively the address of the artist and the exhibition's timespan. Since addresses serve as contact points rather than definitive indicators of presence, it is uncertain whether the artist was actually at the address during the exhibition. Consequently, the temporal projection only records the partial and temporary presence (`crm:P197_covered_parts_of`) of an actor at a given address (`crm:E53_Place`). Each address used by an artist can be modeled as an instance of `crm:E93_Presence`

Source(S_x). The connection to the source is crucial in a historical context. For such reason, each spatiotemporal segment representing an artist living in a specific address should be linked (`crm:P70i_is_documented_in`) with the document(s) (`crm:E31_Document`) which serves as the source for the recorded statement.

Temporal Projection($T_0 - T_1$). The temporal information of the presence (`crm:E93_Presence`), i.e. the segment representing the person using an address, is specified using `crm:E52_Time-span`. In this instance, since our knowledge of the address is based solely on the details provided by the exhibition (`crm:E7_Activity`) catalog, the timespan associated with the spatiotemporal projection always aligns with that of the exhibition.

Spatial Projection(P_x). The spatial projection is modeled as an instance of `crm:E53_Place`. To make use of the capabilities of GeoSPARQL-compliant Graph Stores, the data have been further aligned to the GeoSPARQL ontology[6]. There exists already a mapping between CIDOC-CRM and GeoSPARQL in the CRMgeo extension [24, 18]. A more detailed representation of the interconnection between the two is given in Table 1. The author believes there are a few steps still missing for completing an integration between GeoSPARQL and CRM. A few properties for geometric serialization are not mapped with GeoSPARQL (see Table 2), and an important serialization (`geo:asGeoJSON`) is not yet present. Additionally, the mapping lacks a foundation in CRM itself. The class `crm:E53_Place`, for instance, is not mapped yet, and within the community, there is a lack of clarity about its ontological status (see the different mappings CRM-GeoSPARQL in [7] and [25]). Due to these reasons, the model in Fig. 2 formalizes the geometry of the address using GeoSPARQL. By instantiating the spatial projection as both `crm:E53_Place` and `geo:Feature` it is possible to use the GeoSPARQL ontological pattern for the definition of a geometrical serialization.

The diagram in Fig. 2, presents a CRM-way to formalize spatiotemporal presence. Each time a catalog documents an artist residing at a particular address we can instantiate it as a `crm:E92_Presence`, which is going to be further linked to a spatial and temporal dimension. The latter represented as `crm:E52_Time-span`, provides us with an easy mechanism for distinguishing when artists occupied the same space. The spatial dimension represents the address, formalized as an instance of `crm:E53_Place` with an associated WKT geometry (Point). Each address, defined by its street name and street number, is encoded as an individual entity with a unique IRI, thereby facilitating the identification of situations

---

[6] Documentation at https://docs.ogc.org/is/22-047r1/22-047r1.html

Table 2
Additional mappings between CRM, CRMgeo, and the GeoSPARQL Ontology. Prefixes used are listed on prefix.cc

Element(from)	Suggested Relation	Element(to)
crmgeo:Q11i_is_approximated_by	rdfs :subPropertyOf	geo:hasGeometry
crm:E53_Place	rdfs :subClassOf	geo:Geometry
crm:E92_Spacetime_Volume	rdfs :subClassOf	geo:Feature
crm:E18_Physical_Thing	rdfs :subClassOf	geo:Feature
crmgeo:Q10i_place_is_defined_by	rdfs :subPropertyOf	geo: hasSerialization
crmgeo:asWKT	rdfs :subPropertyOf	geo:asWKT
crm:P168_place_is_defined_by	rdfs :subPropertyOf	geo:asWKT
crmgeo:asGML	rdfs :subPropertyOf	geo: hasSerialization

where two artists share the same address. A limitation of this approach is the need for constant data curation, which is not always feasible, especially as new catalogs are frequently added to BasArt. However, by utilizing the GeoSPARQL ontology, we can effectively apply its functions at the query level. For instance, the `geof:distance` function allows us to retrieve artists who have used the same address by calculating the distance between documented geometries, while the `geof:nearby` function returns all geometries within a specified radius of a given point. Moreover, when the dataset is curated, the `geof:sfContains` function can be employed to query artists residing within specific districts, streets, or any defined bounding box.

Using this model, we can systematically document the various addresses associated with an artist throughout their lifetime. To enhance data usability, the model in Fig. 2 integrates a simple and complex layer. The basic exhibition information is organized within a 'simple layer,' an ontological pattern that provides users seamless access to the data. This layer links the exhibition, its timeframe, the address, the residing artist, and the historical source, allowing retrieval without requiring knowledge of the underlying spatiotemporal modeling. For instance, using the property `crm:P74_has_current_or_former_residence` we can easily retrieve the list of documented addresses associated with an artist without having to further navigate the graph. However, by using the class `crm:E93_Presence` we can change the shape in which the same information is linked together, providing a way to query the spatiotemporal presence of an artist at each documented address. The two modeling solutions help provide easy access to CRM data while maintaining an information-rich environment for domain specialists who may want to extract as many details as possible from the dataset.

The modeling in Fig 2, while satisfying the requirements of **CH1** and **CH3**, needs more details for **CH2**. In fact, the address is modeled only as a static entity. Given the examples in **EX2** and **EX3** it is essential to develop a model capable of recording the dynamicity of an address. The mapping in Fig 2 is complemented by the one in Fig. 3, which illustrates how to use CRM spatiotemporal classes to formalize the temporal validity of street names and municipalities. The model formalizes the address as the geometrical projection of a spatiotemporal segment. This makes it possible to record and link its different phases, instantiated as `crm:E93_Presence`, specifying for each phase its temporal attributes as well as the name used during the phase (as in **EX3**). CRM formalizes a series of properties for expressing time relationships (similarly to Allen operators), but only for the class of temporal entities (`crm:E2_Temporal_Entity`). As spatiotemporal entities and temporal entities are classified under separate branches of the taxonomy, we cannot use these properties to express relationships between spatiotemporal segments. Due to the issue, it is not possible to express with CRM that one address was used earlier than another. Since such relationships are not defined in CRM, we formalize them using OWL-Time[7].

We can still rely on CRM to express the relationship between an instance representing the geometry of the address and an instance representing the municipality it belongs to (**CH2**). As done for the address, we can also document its different phases, instantiated as `crm:E93_Presence`, and specify

---

[7] https://www.w3.org/TR/owl-time/

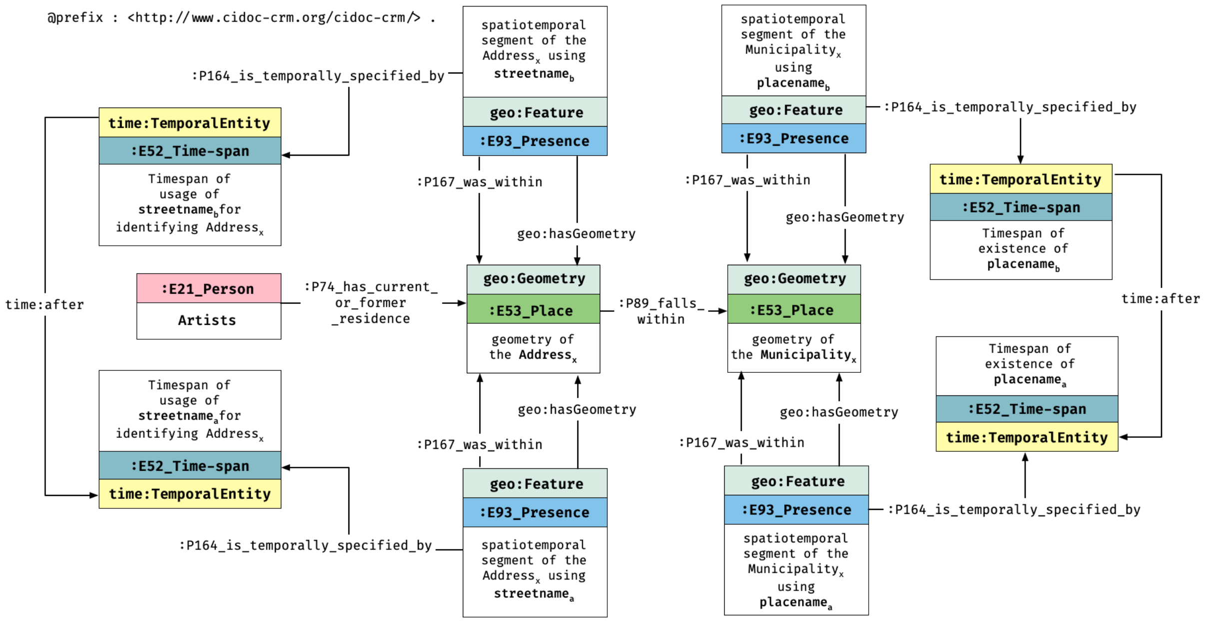

Figure 3: The diagram illustrates the modeling of temporal and spatial relationships of interconnected geometries representing a street and its corresponding municipality. This approach enables the formalization of the temporal validity of the different names used to identify an entity. Prefixes used are listed on prefix.cc

for each of them its corresponding label and temporal attributes. This makes it possible to describe **EX2**, documenting that an address at time T belongs to a municipality that, at the same time T, was identified with a specific label.

5.4. 4D Fluents

The perdurantist logic of CIDOC-CRM differs slightly from other 4D ontologies [26, 19]. Unlike full perdurantist 4D ontologies, not all entities in the CRM are of a temporal nature. Physical entities ”*defines a spacetime volume unique to it*” [27] but as a matter of identity they are still considered instances of `E77 Persistent item`, endurants, disjoint from perdurants, which are formalized as instances of the class `E2 Temporal Entity`. In the current incarnation of the ontology, version 7.3, endurants (and persons) preserve their identity in time. Their identities are linked with a spacetime but they are not a spacetime themselves. In perdurantist ontologies entities are understood as perdurants, i.e. 4-d ’slice’ of a larger hypervolume of 4-dimensional space [28]. We do not document statements about the entity, but about the entity’s temporal part. Given their time-related nature, perdurantist ontologies appear to be quite promising for the description of historical phenomena. To investigate their use, this section presents a hybrid approach combining the 4D ontology[8] with CRM.

The 4D ontology is encoded to be a lightweight top ontology, and it formalizes only a few classes to encode temporal slices of entities and their extent in time. It takes the perdurantist stance that entities have a diachronic identity (they are different at every documented time), and assertions about entities should be always framed in time. Each temporal part of an entity is encoded as an individual with a time extent, and each statement encoded with the ontology is a temporal statement [26]. We can combine this approach with CRM, by classifying the top classes in the CRM taxonomy (`crm:E53_Place`, `crm:E77_Persistent_Item` and `crm:E2_Temporal_Entity`) and the top properties as, respectively, subclass of `4d:TemporalPart` and subproperties of `4d:temporalProperty`. In such a way, we formalize all the mentions of the same address found in the exhibition catalogs as temporal parts of the same entity, separate individuals that are associated with one or multiple artists (see example in Fig. 4). Artists are also space-time worms. A record in a catalog in 1901 mentioning that ”Marquet live in 38

---

[8] https://www.emse.fr/ zimmermann/ndfluents.html

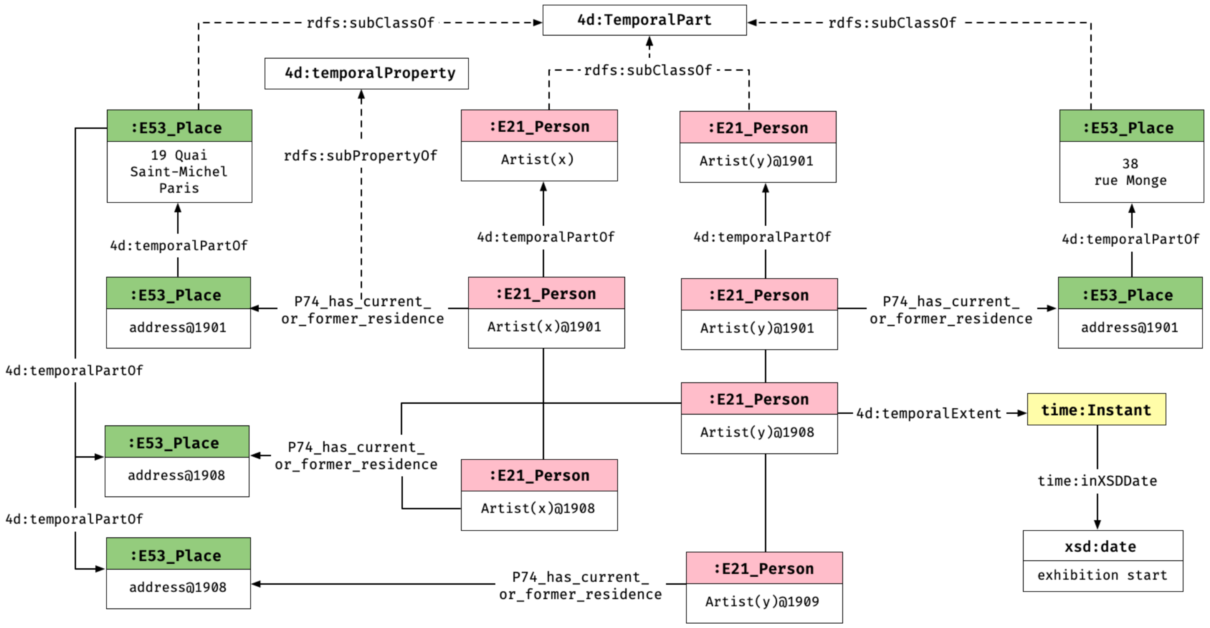

Figure 4: The modeling in the picture formalizes addresses and artists as temporal parts (4d:temporalPartOf) of their respective entities. Each known address (same can be said for the artist) is formalized as an individual with clear temporal boundaries, expressed with 4d:temporalExtent (although not all temporal parts in the diagram, due to limited space, are linked with their own temporal extent). The notation '@year' is used to simplify the illustration and highlight the identity differences between temporal parts of the same entity.

rue Monge, Paris” should refer not to Marquet himself, but to the ”Marquet living in 1901”. In the same way, the statement would not refer to the entity representing ’38 rue Monge, Paris’, but its 4D slice, 38 rue Monge existing in 1901. The entity Marquet, for instance, would be constituted by the collection of all the recorded facts about him in the base. In our example, all the statements about the addresses associated with him would be data points that constitute its identity: Marquet1903(62, rue Bargue), Marquet1904(211 bis avenue de Versailles), Marquet1907(29, place Dauphine), Marquet1907(19 Quai Saint-Michel). The context of each of these statements is always temporally bound, as the information recorded is limited to the documented timeframe, reflecting a source-first approach. Using this modeling voids the validity problem (**CH1**), as the statements regarding the presence of each artist are framed in time. The property `crm:P74_has_current_or_former_residence`, in fact, connects a version of the artist existing in Tx with a version of the address also existing in Tx. The documentation of diachronic places (**CH2**), such as in (**EX2**) and (**EX3**), would simply document different temporal segments of such places, annotating them with the corresponding historical information valid at the time. The same can be said for the representation of artists sharing the same address (**CH3**). With these models defined, the next section will compare their effectiveness for the presented use case.

6. Analysis

The modelings proposed above present different ontological patterns for the representation of temporally constrained facts. Following a classification by Hayes [28], they can be classified into three patterns: the one in section 5.2 uses a 3D modeling pattern, the one in section 5.3 a 3D+1, and the one in section 5.4 a 4D. Among the three different modeling proposed above, only the one in section 5.2 express explicitly the validity of a recorded address. The CRM modeling does annotate the validity, although implicitly. The recorded propositions are, in fact, contained, and thus valid, within a determined spatiotemporal frame. The modeling in section 5.4 does not explicitly annotate the validity of the assertions, since the documented entities referred to already possess temporal boundaries. Any of the along modeling successfully retrieves the address and its temporal boundaries, however with some differences. The ontological pattern proposed in section 5.2 does express the validity of the statement but with very little semantics. The richness provided by the myriad of properties and classes in CIDOC-CRM is, of course, incomparable with the basic vocabulary used here. However, such richness is not always necessary, and its usefulness depends on the scope of the data modeling. If the primary objective is analytical rather than documentative, this solution is particularly effective, as it preserves the structure of the data while ensuring a lean and efficient dataset. Moreover, this type of pattern is considered to be the most user-friendly and easy to adopt [29]. The only two limitations to be considered are (i) the presence of few RML mappers that support RDF-STAR and (ii) the possible complexity in querying when dealing with highly nested triples, specifically for non-computer scientists (e.g., art history/heritage researchers). The 4D perdurantist logic proposed in section 5.4 is a perfect solution for documenting historical knowledge, as it takes into account the dynamicity of entities that represent historical objects and events. However, there are some drawbacks at the application level. The method can produce a very high proliferation of time slices [30], and therefore the creation of a very large quantity of triples that may hinder both reasoning and retrieval of the data. Another problem of the pattern proposed in section 5.4 is that is regarded, by a qualitative study on data modelers [29] as the most complicated ontological pattern to use. The CRM solution proposed in section 5.3 strikes a middle ground, as it does present the possibility to map the data with spatiotemporal properties without creating a massive proliferation of time slices as in the 4D approach. The result is far from lean, but it has the great advantage of working alongside the CRM universe, immediately aligned with the rest of the model, and many other cultural heritage linked data resources. A great advantage of CRM is that it better conveys the nuances in historical statements, due to its connection with humanist practitioners and scholars. The distinction between the properties `crm:P197_covered_parts_of`, which describes the partial or complete overlap of a time slice with a spatial extent, and `P161_has_spatial_projection`, which refers to the coverage of a spatial extent by a spacetime volume, exemplifies the nuanced distinctions embedded in the ontology. Such subtleties may appear less relevant in the case of large-scale data-driven applications, but they are highly valued by humanists. The presence of a large number of nuanced properties, while highly regarded by data owners, can create quite difficult models to query and fully comprehend. It is particularly difficult, specifically for new users on CRM, to strike the balance between expressiveness and practicality. That being said CRM, due to its ontological richness and widespread adoption in the cultural heritage domain, it will increase the potential for data reuse and integration across different digital humanities sources. For such reason, it is the preferred modeling choice for BasArt, although more tests are necessary to understand how its integration with a full 4D framework would help data representation.

7. Conclusion

This paper presented a framework for modeling the temporal validity of artist-recorded addresses as documented in exhibition catalogs. Through the comparative analysis of three ontological patterns—RDF-star, CIDOC-CRM, and a 4D ontology— the contribution has demonstrated various approaches that effectively capture the diachronic nature of addresses. Each method offers distinct advantages. RDF-star provides an easy-to-understand, lean, and efficient solution that is ideal for analytical tasks. CIDOC-CRM offers a balanced approach that aligns well with existing cultural heritage resources, offering a rich semantic depth and numerous nuanced distinctions that are particularly valuable in humanities research. The 4D ontology, while potentially adding complexity to data management, offers comprehensive representational power and flexibility and, moreover, it aligns closely with the underlying thinking of historical research. The three patterns are conceived to be used for structuring and retrieving information about historical changes over physical entities and can be applied to the many contexts that work with temporally constrained facts, specifically facts about dynamic physical and spatial entities. Further research should investigate the computational differences, in terms of transformation and analysis between the three models, providing quantitative (e.g. number of triples produced, query performance, efficiency, reasoning time) and qualitative metrics (e.g., user studies on querying 4D data vs 3D vs 3D+1) to guide data modelers in selecting the most appropriate approach for their case study.

The CIDOC-CRM pattern discussed in section 5.3 has been used to transform a large dataset of artists' addresses from the BasArt database into RDF. This work has made it possible to extract temporal information about each artist's residency along with any address documented in the data, which in turn enables us to examine the temporal and spatial distribution of the artistic communities in 20th century Paris.

8. Bibliography

References

[1] F. Vogel, On the canon of exhibition history, in: Re-Envisioning the Contemporary Art Canon, Routledge, London, 2016, pp. 205–218.

[2] B. Altshuler, Salon to Biennial - Exhibitions that Made Art History, Volume 1: 1863-1959, Phaidon Press, London ; New York, 2007.

[3] D. S. Greenwald, Painting by Numbers: Data-driven Histories of Nineteenth-Century Art, Princeton University Press, Princeton, 2021.

[4] L. Manovich, Cultural Analytics, Mit Press, Cambridge, Massachusetts, 2020.

[5] D. Wu, H.-T. Wang, A. U. Tansel, A survey for managing temporal data in RDF, Information Systems 122 (2024) 102368. doi:10.1016/j.is.2024.102368.

[6] S. Migliorini, Enhancing CIDOC-CRM Models for GeoSPARQL Processing with MapReduce, in: Proceedings of the 2nd Workshop On Computing Techniques For Spatio-Temporal Data in Archaeology And Cultural Heritage Co-Located with 10th International Conference on Geographical Information Science (GIScience 2018), volume 2230, CEUR-WS, Melbourne, Australia, 2018.

[7] B. Ranjgar, A. Sadeghi-Niaraki, M. Shakeri, S.-M. Choi, An ontological data model for points of interest (POI) in a cultural heritage site, Heritage Science 10 (2022) 13. doi:10.1186/s40494-021-00635-9.

[8] G.-A. Nys, M. Van Ruymbeke, R. Billen, Spatio-temporal reasoning in CIDOC CRM: An hybrid ontology with GeoSPARQL and OWL-Time, in: Proceedings of the 2nd Workshop On Computing Techniques For Spatio-Temporal Data in Archaeology And Cultural Heritage Co-Located with 10th International Conference on Geographical Information Science (GIScience 2018), volume 2230, CEUR-WS, Melbourne, Australia, 2018.

[9] V. Bartalesi, G. Coro, E. Lenzi, N. Pratelli, P. Pagano, M. Moretti, G. Brunori, A Semantic Knowledge Graph of European Mountain Value Chains, Scientific Data 11 (2024) 978. doi:10.1038/s41597-024-03760-9.

[10] F. Yu, E. Champion, D. McMeekin, Exploring Historical Australian Expeditions with Time-Layered Cultural Maps, ISPRS International Journal of Geo-Information 12 (2023) 104. doi:10.3390/ijgi12030104.

[11] F. Gantner, B. Waldvogel, R. Meile, P. Laube, The Basic Formal Ontology as a Reference Framework for Modeling the Evolution of Administrative Units, Transactions in GIS 17 (2013) 206–226. doi:10.1111/j.1467-9671.2012.01356.x.

[12] E. Hyvönen, J. Tuominen, T. Kauppinen, J. Väätäinen, Representing and Utilizing Changing Historical Places as an Ontology Time Series, in: N. Ashish, A. P. Sheth (Eds.), Geospatial Semantics and the Semantic Web: Foundations, Algorithms, and Applications, Springer US, Boston, MA, 2011, pp. 1–25. doi:10.1007/978-1-4419-9446-2_1.

[13] S. Tual, N. Abadie, B. Duménieu, J. Chazalon, E. Carlinet, Création d'un graphe de connaissances géohistorique à partir d'annuaires du commerce parisien du 19 ème siècle: Application aux métiers de la photographie, in: 4es Journées Francophones d'Ingénierie Des Connaissances (IC 2023) @ Plate-Forme Intelligence Artificielle (PFIA 2023),, Strasbourg, France, 2023.

[14] W. Charles, N. Hernandez, HHT: An Ontology to Represent Territorial Dynamics for Digital

Humanities, in: Formal Ontology in Information Systems, IOS Press, Sherbrooke, QC, Canada, 2023, pp. 257–271. doi:10.3233/FAIA231133.

[15] M. Fernández-López, A. Gómez-Pérez, N. Juristo, Methontology: From ontological art towards ontological engineering, in: AAAI-97 Spring Symposium Series, volume Proceedings of the Ontological Engineering AAAI-97 Spring Symposium Series, American Association for Artificial Intelligence, 1997.

[16] B. Joyeux-Prunel, O. Marcel, Exhibition catalogues in the globalization of art. A source for social and spatial art history, Artl@s Bulletin 4 (2015) 8.

[17] G. Bruseker, N. Carboni, A. Guillem, Cultural Heritage Data Management: The Role of Formal Ontology and CIDOC CRM, in: Heritage and Archaeology in the Digital Age, Quantitative Methods in the Humanities and Social Sciences, Springer, Cham, 2017, pp. 93–131.

[18] G. Hiebel, M. Doerr, Ø. Eide, CRMgeo: A spatiotemporal extension of CIDOC-CRM, International Journal on Digital Libraries 18 (2017) 271–279. doi:10.1007/s00799-016-0192-4.

[19] C. Welty, R. Fikes, A Reusable Ontology for Fluents in OWL, in: Formal Ontology in Information Systems, IOS Press, 2006, pp. 226–236.

[20] D. Hernández, A. Hogan, M. Krötzsch, Reifying RDF: What works well with wikidata, in: Proceedings of the 11th International Workshop on Scalable Semantic Web Knowledge Base Systems, volume 1457, CEUR-WS, Bethlehem, PA, 2015, pp. 32–47.

[21] O. Hartig, Foundations of RDF* and SPARQL* : (An Alternative Approach to Statement-Level Metadata in RDF), in: AMW 2017 11th Alberto Mendelzon International Workshop on Foundations of Data Management and the Web, Montevideo, Uruguay, June 7-9, 2017., volume 1912, Juan Reutter, Divesh Srivastava, 2017.

[22] V. Nguyen, O. Bodenreider, A. Sheth, Don't like RDF reification?, in: 23rd International World Wide Web Conference (WWW 2014), volume Proceedings of the 23rd international conference on World wide web, Association for Computing Machinery, 2014, pp. 759–770. doi:10.1145/2566486.2567973.

[23] J. J. Carroll, C. Bizer, P. Hayes, P. Stickler, Named graphs, provenance and trust, in: Proceedings of the 14th International Conference on World Wide Web, WWW '05, Association for Computing Machinery, New York, NY, USA, 2005, pp. 613–622. doi:10.1145/1060745.1060835.

[24] M. Doerr, G. Hiebel, CRMgeo: Linking the CIDOC CRM to GeoSPARQL through a Spatiotemporal Refinement, Technical Report 435, ICS-FORTH, 2013.

[25] V. Ducatteeuw, Developing an urban gazetteer : A semantic web database for humanities data, in: GeoHumanities'21 : Proceedings of the 5th ACM SIGSPATIAL International Workshop on Geospatial Humanities, Association for Computing Machinery, Beijing, China, 2021, pp. 36–39. doi:10.1145/3486187.3490204.

[26] J. M. Giménez-García, A. Zimmermann, P. Maret, NdFluents: An Ontology for Annotated Statements with Inference Preservation, in: E. Blomqvist, D. Maynard, A. Gangemi, R. Hoekstra, P. Hitzler, O. Hartig (Eds.), The Semantic Web, Springer International Publishing, Cham, 2017, pp. 638–654. doi:10.1007/978-3-319-58068-5_39.

[27] C. Bekiari, G. Bruseker, C. Erin, M. Doerr, P. Michon, C. E. Ore, S. Stead, A. Velios, CIDOC Conceptual Reference Model, 2024.

[28] P. Hayes, Formal Unifying Standards for the Representation of Spatiotemporal Knowledge, Technical Report 02TA4-SP1-RT1, Florida Institute for Human & Machine Cognition, 2004.

[29] A. Scheuermann, E. Motta, P. Mulholland, A. Gangemi, V. Presutti, An empirical perspective on representing time, in: Proceedings of the Seventh International Conference on Knowledge Capture, K-CAP '13, Association for Computing Machinery, New York, NY, USA, 2013, pp. 89–96. doi:10.1145/2479832.2479854.

[30] V. Zamborlini, G. Guizzardi, On the Representation of Temporally Changing Information in OWL, in: 2010 14th IEEE International Enterprise Distributed Object Computing Conference Workshops, 2010, pp. 283–292. doi:10.1109/EDOCW.2010.50.